



Microsoft
GOLD CERTIFIED
Partner

IBM XIV[®] Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution



Tested with: ESRP – Storage Version 3.0
Tested Date: October 28, 2009
Authors: Aviad Offerⁱ, David Hartmanⁱⁱ
Contributors: Tom Quiereⁱⁱⁱ
Reviewers: David West^{iv}
Document version: 3.8



Content

IBM XIV® Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution..... 1

Content..... 2

Overview 3

Disclaimer 3

Features..... 3

Solution Description 4

 XIV Storage System Model 2810 4

 IBM System x3650 10

Solution Description 11

 Targeted Customer Profile 16

 Tested Deployment 16

 Simulated Exchange Configuration: 16

 Storage Hardware 17

 Storage Software 18

 Storage Disk Configuration (Mailbox Store Disks)..... 18

 Storage Disk Configuration (Transactional Log Disks) Utilizes same disks as Mailbox Store Database 18

Best Practices 19

 Storage - Mailbox Resiliency 19

 Backup strategy 20

Contact for Additional Information..... 20

Test Result Summary 20

 Reliability 20

 Storage Performance Results 21

 Database Backup/Recovery Performance 22

 Database Read-only Performance 22

 Transaction Log Recovery/Replay Performance 22

Conclusion 23

Appendix A. Test Results 24

 Stress Testing 24

 Performance Testing 28

 Streaming Backup Testing 37

 Soft Recovery Testing 42



Overview

This document provides information on the IBM XIV 2810 Storage System storage solution for Microsoft Exchange Server, based the *Microsoft Exchange Solution Reviewed Program (ESRP) – Storage* program*. For any questions or comments regarding the contents of this document, see [Contact for Additional Information](#).

*The *ESRP – Storage* program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide information on their storage solutions for Microsoft Exchange Server software. For more details on the *Microsoft ESRP – Storage* program, please click:

<http://www.microsoft.com/technet/prodtechnol/exchange/2007/esrp.msp>

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.

Please refer to the [IBM disclaimer](#), located at the end of this document for additional information.

Features

This solution implements the new Mailbox Resiliency features in Exchange Server 2010; specifically the database availability group feature. A database availability group (DAG) is a set of up to 16 Microsoft Exchange Server 2010 Mailbox servers that work together to provide automatic database-level recovery from a database, server, or network failure. Mailbox servers in a DAG monitor each other for failures. Solutions that implement this feature delivery a very robust Exchange environment.

More information on this feature can be found by visiting:

<http://technet.microsoft.com/en-us/library/dd351172%28EXCHG.140%29.aspx>

A single DAG solution comprising four-mailbox servers was created that supported 40,000 mailboxes with a mailbox size of 1GB. Each server hosted 10,000 users, and had five active databases, with 2,000 users per database. Within the DAG, there were two copies of every database; one local, and one on another server connected to a second XIV storage array. This configuration can provide for both high-availability, and disaster-recovery scenarios. Database replication is handled by the mailbox servers over Ethernet. Two XIV frames were used, and the databases and copies were equally distributed across them. A failure of any database will cause the database to mount and become active on whichever server is hosting the database copy, in a manner transparent to the end users.

IBM XIV® Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution	Page 3	Aviad Offer, David Hartman, Tom Quiere, David West ©Copyright IBM Corp. 2009
---	--------	--

A graphical representation of the DAG architecture is represented in figure 1 below.

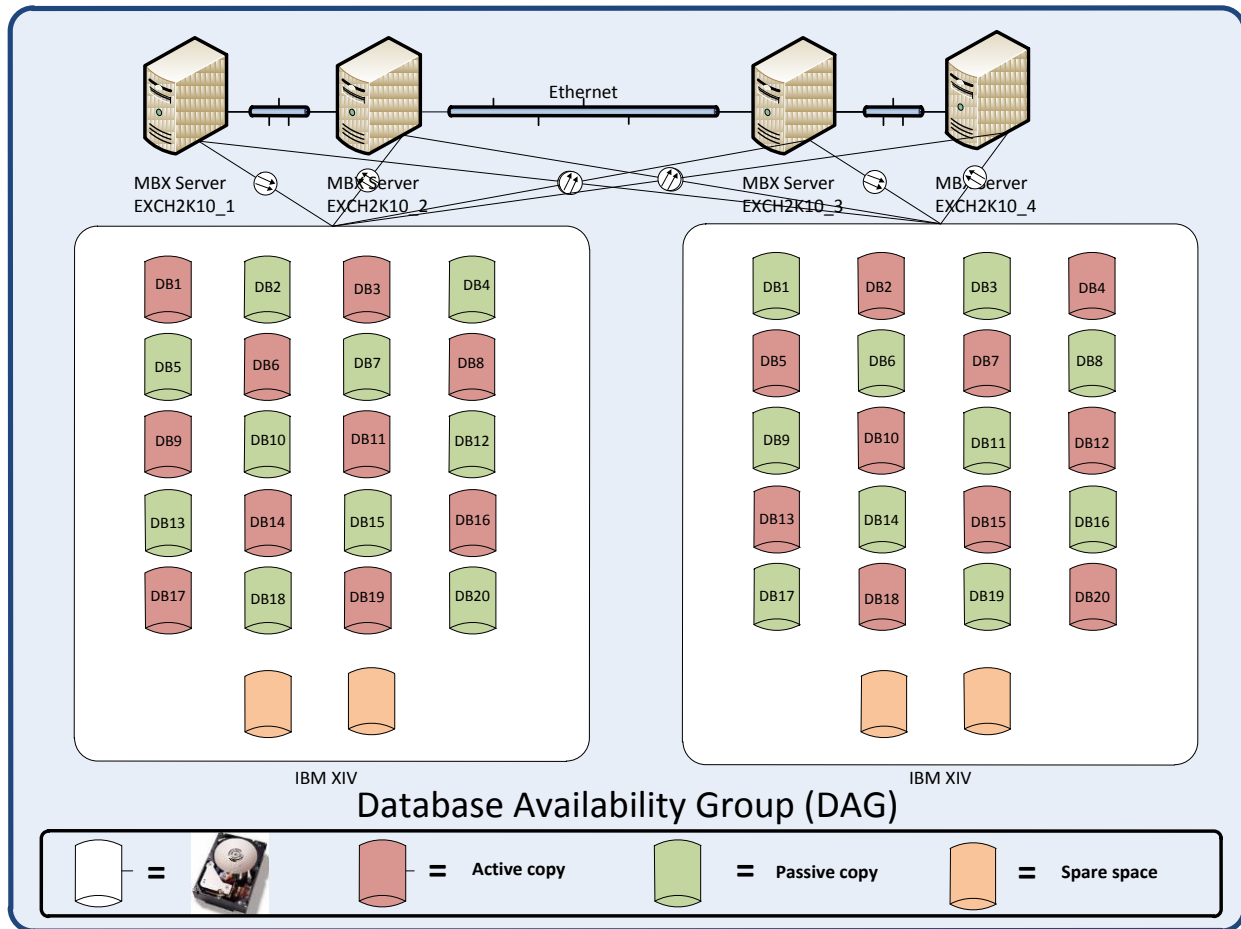


Figure 1) Exchange Server 2010 DAG configuration.

The solution in this paper features the IBM XIV Storage System together with IBM System x3650 Intel based servers.

The following sections describe the actual hardware components used in this solution.

Solution Description

XIV Storage System Model 2810

As one of the newest members of IBM's family of storage systems, the XIV Storage System is an enterprise-class array offering superior performance, solid-protection, and ease of manageability.

Exclusive to this storage system is the inherent self-optimizing grid-based architecture which provides massive parallelism. Massive parallelism distributes



data in parallel across all available storage resources, by harnessing all the processing power available to the system. Due to this design, hot spots (heavily saturated storage components) normally associated with poor storage configuration and/or workload variances, is eliminated. The XIV architecture delivers continuous optimal performance and eliminates the need for dubious tuning activities on the storage.

Built using industry-standard Intel/Linux components -- a fully equipped XIV Model 2810/A14 Storage System is supplied with fifteen data modules, 6 out of those are interface modules. A data module consists of one Intel Xeon® Quad-core 64-bit CPU, 8 GB fully buffered RAM, and twelve 1 TB SATA II 7,200 RPM hard disks, an interface module will have an additional Quad-core 64bit CPU

In comparison to Fibre-Channel based disks, the XIV has more disk density allowing for a smaller physical storage footprint. The unique design of the architecture can also yield similar or better performance.

Hardware Highlights	
Total Storage Central Processing Units	21
Combined Physical Memory	120 GB
Raw Storage	180 TB
Disk to Cache Bandwidth	240 Gbps
Internal Switching Capacity	160 Gbps
Fibre Channel Ports (Speed: 1 Gbs, 2 Gbs, & 4 Gbs)	24
iSCSI Gigabit Ethernet Ports	6
UPS	3
Gigabit Ethernet Switch	2
Disk Hot Spares	3
Data Module Hot Spare	1
Max Snapshots	16,000



Figure 2 below represents a fully configured XIV Storage System:



Figure 2) XIV front and rear view (full rack)

Reliability/Protection

By design, the XIV Storage System architecture is self-healing as data in the background is automatically checked for integrity and mirrored to provide data redundancy. Should the XIV lose a 1 TB disk, the system triggers an event alerting IBM support while rebuilding the faulty disk. The disk rebuild process is transparently handled without user intervention and can finish in as few as 30 minutes when the XIV's capacity is 100% utilized.

Consistent performance is maintained even during the loss of a hard disk. The rebuild time of losing one disk does not degrade performance, unlike other competing storage systems which absorb a significant degradation in performance.

Should a storage sub-component fail on monolithic storage systems, about 50% (CPU and read/write cache) of performance is lost. This not only impacts productivity, but could have serious implications affecting application performance.



Losing an entire data module¹ on the XIV Storage System would temporarily account for only 1/15th reduced performance.

To maintain data redundancy, all data on the storage and memory cache is mirrored. The cache data is never stored in the same data module. This not only provides data cache protection, but eliminates loss of data in case a data module's power is lost or interrupted.

Each data module is supplied with two field replaceable hot-swappable power supplies, and three UPS devices provide interim power in the event of interrupted or lost main power. Enough power is provided to allow safe cache destaging and a graceful shutdown of the XIV Storage System.

The XIV Storage System also supports synchronous and asynchronous Remote Mirroring, allowing for replication of data to a remote XIV Storage System.

Thin Provisioning

This feature permits the logical allocation of storage to appear larger than what is physically committed. This provides the ability to resize storage pools and volumes even after defining initial storage sizes.

This feature cannot only reduce upfront storage costs by avoiding purchasing needless unused storage, but reduces the hardware and environmental footprint. A larger footprint is typically more expensive because it requires more hard disks, electrical power, cooling requirements, disk expansions enclosures, floor space, and management overhead.

Thin provisioning reduces or eliminates the impact of future storage capacity issues when current storage limits have been reached. This occurs when storage has quickly outgrown original capacities. It can be a difficult task to predict how data will grow as storage usage is often highly dynamic in nature. In an effort to circumvent future capacity issues, storage can be appropriately sized beforehand or resized later.

¹ Loss of a data module compared to a data interface module has lesser impact in performance.



Figure 3 below depicts how to easily resize an XIV volume.

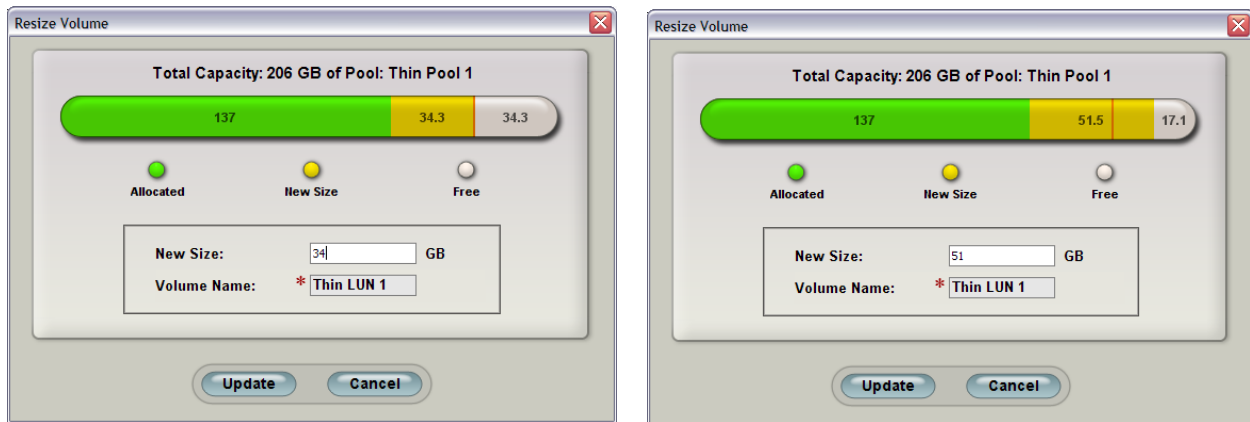


Figure 3) Resizing an XIV volume.

Important: While thin provisioning enables allocated storage to appear larger than what is physically committed, constant monitoring of current available physical storage space is vital to prevent capacity issues. Storage writes destined to logical capacity not allocated with enough physical storage can hinder performance, delete snapshots, and cause storage operations to fail due to lack of space. Prior to reaching this limit, storage threshold alerts should be configured and enabled to warn administrators of the condition. Regular monitoring of storage events and proactive action can prevent these events from occurring.

Another innovative feature is the ability to consume disk space as only actual data is written to the volume. This means that while a LUN may be allocated 100GB of capacity, if only 10GB is in actual use, only 10GB is of physically allocated & reported. This concept is referred to as a volume's hard or physical utilization. This can result in a significant savings in both space and cost. This differs from traditional storage systems, in which disk space is flagged as 0's and marked as used storage. Traditional storage systems quickly outgrow space by immediately pre-allocating storage due to the inefficient method of marking free space as reserved. Whereas on the XIV Storage System it is tracked by the system as free space, and can be used for other hosts, thus maximizing overall storage utilization.

Note: It is important to realize that this concept differs from that of thin provisioning. Also, this feature does not apply to free space within the EDB file(s).

Ease of Management

Managing the XIV Storage System compared to other storage is markedly different. Typical storage management tasks are eliminated and handled automatically by the XIV Storage System including automatic provisioning, self-tuning of data placement, read/write cache tuning, self-healing, automatic rebuilding of failed disk drives, and automatic phasing-out or fencing of failed components. The modular, redundant architecture dramatically reduces impact to service levels from such events.



In general, storage management can be a daunting task, requiring extensive storage experience and specialized training. These barriers have been removed, as the XIV Storage System provides a highly intuitive GUI called XIV Storage Management, with a near zero learning curve.

A less experienced storage administrator with no XIV Storage System experience can allocate storage and map volumes to hosts in less than 15 minutes*. Storage management tasks take less time, allowing storage to be utilized faster.

Figure 4 below is a sample screen shot of the XIV Storage Management interface.



Figure 4) The highly intuitive XIV Storage Management interface

A full-featured command-line management interface, XCLI is also provided for scripting tasks, management, configuration, and monitoring.

Lower Power Consumption

The IBM XIV Storage System is equipped with high-capacity SATA II disks. SATA II disk drives provide significantly better GB/\$ value than Fibre-Channel (FC) disks.

* We assume the host attached to the IBM XIV is already configured with the vendor specific SAN drivers and is properly zoned in the SAN fabric or a point-to-point connection is used.



SATA II disks also require less power consumption than FC disks. Since SATA II disk drives are generally much higher in capacity than FC disks – this translates to requiring less hard drives, cooling, and power for a given capacity.

Snapshots

A key component of the XIV Storage System is the inherent snapshot capabilities which deliver near-instantaneous volume replicas^{**}. A total of 16,000 snapshots are supported. A snapshot is simply a point-in-time image of data used for quick data retrieval purposes. Transportable snapshots are often use for offloaded backups of Exchange databases, which remove the backup load from the production servers.

The snapshot method employed by the XIV Storage System uses an efficient method called *Redirect-on-Write*. The *Redirect-on-Write* snapshot-based method addresses the typical performance overhead associated with the *Copy-on-Write* or *COW* snapshot-based method. For example, during a *Copy-on-Write* snapshot operation, new writes destined to target storage must wait until the original data is copied to the snapshot area. This step is eliminated using *Redirect-on-Write* as the new writes are written immediately to the snapshot area resulting in faster snapshots. The *Redirect-on-Write* snapshot method does not impact the source volumes as do *Copy-on-Write* based snapshots.

For more information on the XIV Storage System visit the IBM website at:

<http://www.xivstorage.com>

IBM System x3650

The System x3650 a rack-mount server designed for critical and high-performance computing.

For this particular solution, two System x3650 model 7979AC1 systems containing 2 quad core e5450 3.0 GHz Intel Xeon CPUs were used. The System x3650 was equipped with 16 GB of internal memory.

^{**} Actual snapshot operation times are measured on the IBM XIV. Monitoring of an ad-hoc or scheduled snapshot operation from the TSM or IBM Data Protection interface known as the VSS requester may appear longer as there is additional communication occurring between the requester client and server.



Figure 5 below is a view of the System x3650



Figure 5) System x3650 server

Product specifications below list the actual configuration used in this solution.

Product Description

Description:	IBM System x3650
Form Factor/Height	Rack/2U
Machine Type/Model:	3650
Part Number:	7979MC1
BIOS	Build ID: IBM GGE141AUS
	Released: 10/8/2008
	Revision: 1.12
Processor	Dual 3.0 GHz Intel Xeon Quad-Core processors
Physical Memory	16 GB
On-board Network	Broadcom Gigabit Ethernet:
	Model: BCM5708C NetXtreme II GigE
	Driver Version: 4.4.15.0
Fibre Channel adapter	Qlogic® 4Gb dual port Fiber HBA
	Expansion (PCIe):
	HBA Driver Version:
	Storport 9.1.7.18
	HBA Bios: 2.02
	HBA Firmware Version: 4.0.4
	Part Number: QLE2562

For more information on System x3650, please visit the IBM website:

<http://www-03.ibm.com/systems/x/hardware/rack/x3650/glance.html>

Solution Description

This solution demonstrates the performance capabilities utilizing two XIV Storage System and four System x3650 servers supporting a total of 40,000 mailboxes using a target .8 IOP profile (1 IOP - with 20% overhead - was tested). The mailbox size is 1 GB.



Important: Though this solution was tested using 40,000 mailboxes, additional storage to accommodate snapshots or frequent mailbox move operations may be required for larger transaction log volumes. Carefully review these parameters before finalizing your final target limits. Also, in a production system, we would likely recommend adding two additional servers (one per array) and redistributing the number of mailboxes and databases to remain under Microsoft's recommendation of no more than 10,000 mailboxes/server in the event of a database failover. The additional servers would allow for multiple database failovers without exceeding that threshold. As configured though, the system did provide adequate headroom for database failover.

Storage hardware:

- 2 x XIV Storage Systems Model 2810/A14
- XIV Storage System Software Version: 10.1.0.a
- 8 GB physical memory/data module
- Twelve 1 TB drives per data module
- 180 SATA II drives (1 TB, 7,200 RPM)
- Form Factor: Rack-mount
- Rack Height: 42u

Exchange servers:

- 4 x System x3650 servers
 - 2 quad-core e5450 3Ghz Intel Xeon CPUs
 - 16 GB RAM/Server
- Windows Server 2008 (x64) – Enterprise Edition, SP 2
 - .NET Framework 2.0
 - Microsoft Multipath I/O
- Qlogic 4Gb dual-port Fibre HBA Expansion (PCIe):
 - HBA Driver Version Storport 9.1.7.18
 - HBA Bios: 2.02
 - HBA Firmware Version: 4.0.4
- XIV Storage System Host Attachment Package for Windows (x64) – Version 1.0.4

Fibre-Channel network:

- 2 –8 Gb Brocade Model BR-48000 Fibre Channel switches
- 8 –Fibre-Channel Ports (Tested Speed: 4 Gb)
- 16 - FC LC Fibre Channel 50/125 micrometer shortwave cables
- 16 - Multi-mode cables (20 meters)

Figure 6 below depicts a high-level representation of this solution.

Note: The FC switches are not shown as to keep the overall diagram simplified.

The brackets are to imply each FC switch as described in the section above.

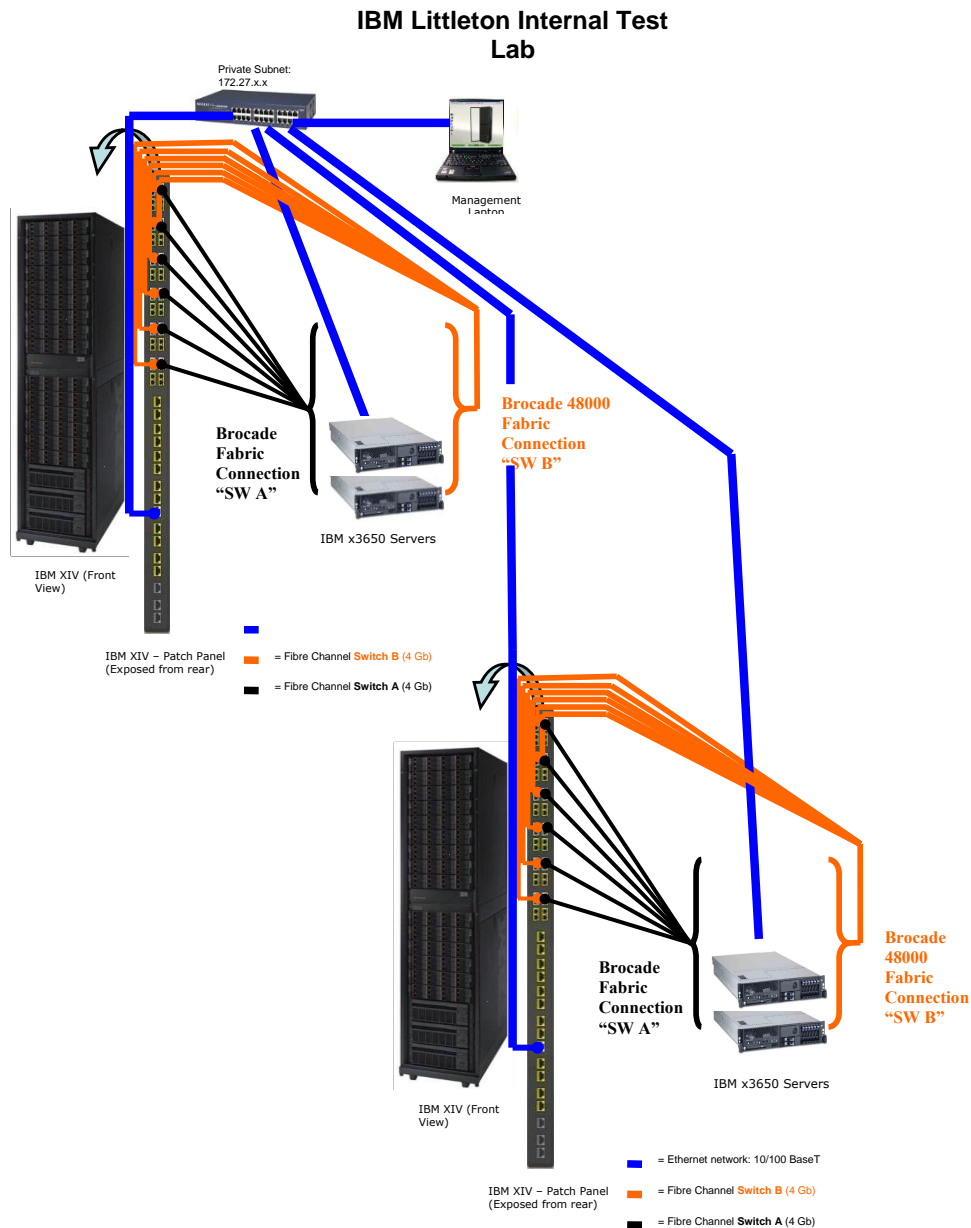


Figure 6) ESRP physical configuration

IBM XIV® Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution	Page 13	Aviad Offer, David Hartman, Tom Quiere, David West ©Copyright IBM Corp. 2009
---	---------	--



As mentioned in the previous sections, the XIV Storage System provides SAN connectivity via 24 FC ports or 6 iSCSI Ethernet ports. For this solution, only the Fibre Channel ports are utilized.

The rear of the XIV Storage System's patch panel reveals access to all six host interface modules, each of which contains FC port designations 1, 2, 3, and 4. For each Exchange server, a minimum of 6 paths were used to provide SAN connectivity to the XIV Storage System.

Two FC switches are used to establish 12 physical connections to the XIV Storage System. Each switch has six FC ports connected. Each FC switch port is connected to each interface module's FC port #3 or FC port #1. This provides an FC connection to all FC Interface Modules on the XIV Storage System.

This configuration provides multiple redundant paths to the XIV Storage System and guarantees each host has a path to each host interface module. This allows the hosts to evenly distribute I/O workload across all host interface modules to obtain the best possible performance.

For each host to achieve optimal performance, each host FC adapter's max queue depth was set to 128.

Each FC port on each interface module concurrently supports up to 1,400 IOPs. There are a total of 4 FC ports on each interface module. Two of the four are available for target FC connections. The remaining two are reserved for remote replication.

To calculate the preferred max queue depth for each host, divide total number of hosts defined in your zone by 1,400.

Zoning Layout

The FC ports on each host are connected via both Fibre channel switches.

To help provide a starting reference for this test environment refer to how XIV host interface modules are connected in each SAN A and B



Note: Zone XIV_CAPRICA contains all XIV host interface modules.

Zone Name	Switch A - Zone Members	Switch B - Zone Members
Caprica_XIV_HOST_30	XIV30 HBA 1 Port 0 HBA 1 Port 1 XIV_CAPRICA(port 3)	XIV30 HBA 2 Port 0 HBA 2 Port 1 XIV_CAPRICA (port 1)
Captria_XIV_HOST_31	XIV31 HBA 1 Port 0 HBA 1 Port 1 XIV_CAPRICA(port 3)	XIV31 HBA 2 Port 0 HBA 2 Port 1 XIV_CAPRICA (port 1)

Table 1) FC zoning configuration

Important: Before deploying FC switches into production carefully run 24 hour tests and closely monitor the FC activity for errors. If errors are observed, take action by ensuring that zoning parameters are correctly defined. Not following these guidelines may not guarantee reliable operation of your storage environment and could result in data loss.

The XIV 2810 System Storage array can be found in the Windows Server Catalog here:

<http://www.windowservercatalog.com/item.aspx?idItem=bd1edc8c-dbe9-7364-58b6-0f9456c01025&bCatID=1282>

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. Other factors which affect the server scalability are: 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, and client usage profiles. All these factors are beyond the scope for ESRP-Storage. Therefore, the number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployments.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please refer to Microsoft's *Troubleshooting Microsoft Exchange Server Performance*, available at:

<http://go.microsoft.com/fwlink/?LinkId=23454>.



Targeted Customer Profile

This overall solution was designed toward medium to large organizations that require high-availability and disaster-recovery features inherent in a Mailbox Resiliency configuration, while minimizing administration complexity. The System x3650 and XIV Storage System offer scale-out flexibility to expand the solution as the Exchange requirements increase in number of users, mailbox size, or both. The targeted profile for this solution includes:

- 4 Exchange 2010 servers (2 tested, simulating 4 for DB copies)
- 2 backend storage subsystems (1 tested)
- 40,000 mailboxes
- 1.0 IOPS user profile (.80 tested plus for 20% growth)
- 1GB mailbox quota
- 2TB maximum database size tested
- Mailbox Resiliency provides high availability and is the primary data protection mechanism. XIV RAID provides for data protection against physical loss/failure
- 24x7 Background Database Maintenance is utilized

Tested Deployment

The following tables summarize the testing environment:

Simulated Exchange Configuration:

Number of Exchange mailboxes simulated	40,000
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	4
Number of active mailboxes/server	10,000
Number of databases/host	5
Number of copies/database	2
Number of mailboxes/database	2,000
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.80 (1.0 tested)
Database LUN size	5TB
Log LUN size	750GB
Total database size for performance testing	20TB
% storage capacity used by Exchange database**	40%

**Storage performance characteristics change based on the percentage utilization of the individual disks. Tests that use a small percentage of the storage (~25%)

IBM XIV® Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution	Page 16	Aviad Offer, David Hartman, Tom Quiere, David West ©Copyright IBM Corp. 2009
---	---------	---



may exhibit reduced throughput if the storage capacity utilization is significantly increased beyond what is tested in this paper. Actual (hard) capacity used will vary due to the way the XIV Storage System physically allocates disk space. See the previous section for further information on how XIV allocates & tracks physical disk space.

Storage Hardware

Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI)	SATA II
Storage model and OS/firmware revision	XIV Storage System Model:2810 XIV Storage System Code Version 10.1.0.a Link to the WHQL website of this storage model.
Storage cache	120GB
Number of storage controllers	6
Number of storage ports	24
Maximum bandwidth of storage connectivity to host	48Gbit (12 X 4Gb)
Switch type/model/firmware revision	2 Brocade 48000 (64 Ports) (8Gb) Switch Firmware Version: 6.1.0
HBA model and firmware	Storage: LSI Corporation 4GB FC PCI Express Adapter: - Model: FC949E Fibre Channel - Operation Speed: 1, 2 & 4Gbps) HBA Firmware Version: - MPT FC FCode Version: 1.00.41 (2006.04.13) - Active firmware version: 01031800 (1.03.24) - Firmware image's version: LSIFC949E-1.03.24 (6/3/2008) Host: Qlogic QLE2562 8Gb dual-port FC HBA (PCIe)
Number of HBA's/host	2
Host server type	4 x IBM x3650 servers each with two quad-core 3.0 GHz Intel Xeon processors
Total number of disks tested in solution	180
Maximum number of spindles can be hosted in the storage	180



Storage Software

HBA driver	Qlogic Storport 9.1.7.18
HBA QueueTarget Setting	Default
HBA QueueDepth Setting	128
Multi-Pathing	Microsoft Multipath I/O (Native to Windows Server 2008), round robin mode
Host OS	Windows Server 2008 – Enterprise Edition Service Pack 2 (6.0.6002.131072)
ESE.dll file version	14.00.0639.011
Replication solution name/version	N/A

Storage Disk Configuration (Mailbox Store Disks)

Disk type, speed and firmware revision	SATA II, 7,200 RPM
Raw capacity per disk (GB)	1.0TB
Number of physical disks in test	180
Total raw storage capacity (GB)	Actual: 180TB Net: 79,1TB
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	180
Raid level	RAID X (IBM XIV RAID)
Total formatted capacity	50TB
Storage capacity utilization	63.3%
Database capacity utilization	40%

Storage Disk Configuration (Transactional Log Disks) Utilizes same disks as Mailbox Store Database

Disk type, speed and firmware revision	SATA II, 7,200 RPM
Raw capacity per disk (GB)	1.0TB
Number of Spindles in test	180
Total raw storage capacity (GB)	Actual: 180TB Net: 79,1TB
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	180
Raid level	RAID X (IBM XIV RAID)
Total formatted capacity	5TB



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 2010 guidance on storage design, please visit:

<http://technet.microsoft.com/en-us/library/dd346703.aspx>

Storage - Mailbox Resiliency

1. To accurately size final formatted volume size on the XIV Storage System, stage the target volume size prior to implementing. By default, the minimum volume extent size is created in blocks of 17 GB. To get close to target volume size, divide the target size by extent size of 17 GB. For example, if the target volume size is 650 GB, divide $650/17 = 38$ Extents (equals to about 646 GB). After partitioning, and formatting volume, the final size will be reduced even more. Factor in another 17 GB, to compensate for the lost free space. $646\text{ GB} + 17\text{ GB}$ will provide 663 GB unformatted drive space. Format the newly created volume and check the available or free space if this meets target needs. Otherwise, increase or reduce the volume to get a more accurate target size.

Note: "*" 1GB=1000MB, 1MB=1000KB, 1KB=1000 bytes"

663GB translated is actually $663 \times 1000 \times 1000 \times 1000 / 1024 / 1024 / 1024 = 617\text{GB}$ on windows.

2. For optimal performance, format the database volumes using a 64k Windows allocation unit size. The volumes used to store log data can be formatted using default allocation unit size, as there aren't significant gains increasing this from the default value.
3. Increasing the Queue Depth setting on the HBA Fibre Controller may increase overall IOPS. Refer to your manufacturer's HBA documentation for detailed support on how this value can be modified.
4. Increasing the number of FC HBA's can provide additional IOPS and resiliency options. Make sure to not define more than 24 paths to the storage, as this the current limit at this time.
5. We strongly recommend implementing Mailbox Resiliency. This allows for using a fewer number of larger LUNS for your databases and yields significantly increased storage performance. Designing your Exchange XIV storage configuration with fewer and larger LUNs is often preferable, and we have observed performance and capacity increases on the order of 25% when configured in this manner.
6. Troubleshooting any latency on the XIV Storage System can be monitored via standard tools such as Microsoft Performance Monitor or the XIV Storage System monitoring tools. These tools can help determine where and when



the heaviest disk I/O patterns are occurring, and what the data transfer rates are, which helps identify and isolate potential bottlenecks.

7. Assuming all storage components are functioning, balanced, and configured correctly, most Exchange Server 2010 storage latency issues are caused by an insufficient number of disks for a given I/O load. The number of IOPS hitting a given drive can be measured using the Performance Monitor Logical Disk object, Disk Transfers/sec counter. The average disk latency should be less than 20ms, with the maximum value no higher than 100ms. Disk latency can be measured using the Performance Monitor Logical Disk object, Avg. Disk sec/Read. Since the XIV Storage System uses all 180 disks for I/O, and is dynamically self-tuning, this concern is all but eliminated versus conventional storage systems.
8. During our testing, the storage subsystems were dedicated to the Exchange Server workload. Sharing the storage with other applications may negatively impact Exchange I/O performance.
9. For performance tips and best-practice guidelines on the XIV Storage System, please visit IBM [Techdocs](#).
10. For detailed technical information on the XIV Storage System, please visit IBM [Redbooks](#).
11. Planning and monitoring your Exchange storage solution based on these guidelines will help ensure optimal client response times, allow room for growth, and avoid potential performance issues.

Backup strategy

N/A

Contact for Additional Information

For additional information regarding the IBM components used in this solution, please visit:

<https://www-912.ibm.com/ssg/tsfeedback.nsf/feedback?OpenForm>

Test Result Summary

This section provides a high level summary of the test data from ESRP and the links to the detailed html reports which are generated by ESRP testing framework. Please click on the underlined headings below to view the html report for each test.

Reliability

A number of tests in the framework are designed to check storage system reliability; these tests run for 24 hours. The goal is to verify the storage can handle high I/O load for an extended period of time. Both log and database files are analyzed for integrity after the stress test to ensure that no database and/or log corruption has occurred.

IBM XIV® Storage System Model 2810 40,000 Mailbox Resiliency Exchange 2010 Storage Solution	Page 20	Aviad Offer, David Hartman, Tom Quiere, David West ©Copyright IBM Corp. 2009
---	---------	--



The following list provides an overview: (clicking on the underlined word will show the html report generated after the reliability test runs)

- There were no errors reported in the saved eventlog file.
- There were no errors reported during the [database](#) and [log](#) checksum process.

Storage [Performance](#) Results

The primary storage performance testing is designed to exercise the storage with the maximum sustainable Exchange type of I/O for 2 hours. The test is designed 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/O's and average of all the logical disks I/O latency in the 2 hour test duration. Each server is listed separately and the aggregate numbers across all servers are listed as well.

Individual Server Metrics:

The sum of I/O's across Storage Groups and the average latency across all Storage Groups on a per server basis.

Server 1: XIV30

Database I/O	
Database Disks Transfers/sec	10,380.945
Database Disks Reads/sec	7,169.811
Database Disks Writes/sec	3,211.134
Average Database Disk Read Latency (ms)	16.9454
Average Database Disk Write Latency (ms)	7.3104
Transaction Log I/O	
Log Disks Writes/sec	922.777
Average Log Disk Write Latency (ms)	2.3194

Server 2: XIV31

Database I/O	
Database Disks Transfers/sec	10,174.773
Database Disks Reads/sec	7,034.888
Database Disks Writes/sec	3,139.885
Average Database Disk Read Latency (ms)	16.9656
Average Database Disk Write Latency (ms)	7.1802
Transaction Log I/O	
Log Disks Writes/sec	907.067
Average Log Disk Write Latency (ms)	2.3154



Aggregate Performance across all servers Metrics:

The sum of I/O's across servers in solution and the average latency across all servers in solution.

Database I/O	
Database Disks Transfers/sec	20,555.72
Database Disks Reads/sec	14,204.70
Database Disks Writes/sec	6,351.02
Average Database Disk Read Latency (ms)	16.9555
Average Database Disk Write Latency (ms)	7.2453
Transaction Log I/O	
Log Disks Writes/sec	1829.844
Average Log Disk Write Latency (ms)	2.3174

Database Backup/Recovery Performance

There are two tests reports in this section. The first one is to measure the sequential read rate of the database files, and the second is to measure the recovery/replay performance (playing transaction logs into the database).

Database Read-only Performance

The test is to measure the maximum rate at which databases could be backed up via VSS. The following table shows the average rate for a single database file.

MB read/sec per database	86.02
MB read/sec total per server	430.10

Transaction Log Recovery/Replay Performance

The test is to measure 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 1 MB in size.

Average time to play one Log file (sec)	6.683069
---	----------



Conclusion

This testing validates this Exchange ESRP – Storage submission for the IBM XIV Storage System under the stated configuration parameters, and demonstrates that this storage array is an ideal repository for Exchange Server 2010 data.

It is imperative to note that this particular solution was tested using an exceptionally heavy user profile that was over seven times larger than what Microsoft considers a typical heavy profile under Exchange Server 2010. This was done to compare the I/O behavior of Exchange Server 2010 under Exchange Server 2007 user profiles. The observed results, when configured in Mailbox Resiliency mode, was a much more efficient I/O profile, that resulted in the storage system being able to contain a larger number of mailboxes than under Exchange Server 2007. Given the updated, significantly lower user profile metrics for Exchange Server 2010, the XIV Storage System would be able to contain a significantly higher number of mailboxes (dependent upon mailbox size limits) when configured with Exchange Server 2010 user profiles (i.e. 0.11 IOPs versus 0.80 IOPs for a heavy user).

The incredible ease of management, advanced data distribution technology, self-tuning capability, and automatic (and extremely quick) drive rebuild offer Exchange administrators and storage architects new tools to deliver a robust, yet easily managed Exchange storage solution. Advanced features such as copy services, thin provisioning, full MPIO & VSS support offer additional flexibility and support in terms of performance, availability, quick recovery, and ensure optimal, efficient usage of the storage subsystem.

This document was developed by storage solution providers, and reviewed by the Microsoft Exchange Product team. The test results/data presented in this document are based on the tests introduced in the ESRP test framework. Customers should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the requisite 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 obtaining 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 differing vendor solutions.



Appendix A. Test Results

This section contains Jetstress reports for Stress, Performance, Streaming Backup, and Soft Recovery. All server test results were reviewed by Microsoft, and had similar performance results. We include the results from only one of the servers here, to make this section more readable.

Stress Testing

Server 1: XIV30

Microsoft Exchange Server Jetstress

Test Result Report

Checksum Statistics - All

Database	Seen pages	Bad pages	Correctable pages	Wrong page-number pages	File length / seconds taken
E:\Jetstress001001.edb	65830786	0	0	0	2057212 MBytes / 15869 sec
F:\Jetstress002001.edb	65829506	0	0	0	2057172 MBytes / 15537 sec
g:\Jetstress003001.edb	65827970	0	0	0	2057124 MBytes / 15522 sec
h:\Jetstress004001.edb	65829762	0	0	0	2057180 MBytes / 15482 sec
i:\Jetstress005001.edb	65830274	0	0	0	2057196 MBytes / 15482 sec
(Sum)	329148298	0	0	0	10285884 MBytes / 15869 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
E:	0.028	0.000	2073.962	0.000	65536.000
F:	0.028	0.000	2118.429	0.000	65536.000
g:	0.027	0.000	2120.418	0.000	65536.000
h:	0.027	0.000	2126.078	0.000	65536.000
i:	0.027	0.000	2126.140	0.000	65536.000



Memory System Performance (of checksum)

Counter	Average	Minimum	Maximum
% Processor Time	13.688	2.266	16.032
Available MBytes	14721.670	14680.000	14742.000
Free System Page Table Entries	33563044.466	33560695.000	33563437.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	61617105.455	61591552.000	61734912.000
Pool Paged Bytes	154749021.091	153411584.000	199933952.000

Test Log 10/27/2009 9:19:34 AM -- Jetstress testing begins ...
 10/27/2009 9:19:34 AM -- Prepare testing begins ...
 10/27/2009 9:19:40 AM -- Attaching databases ...
 10/27/2009 9:19:40 AM -- Prepare testing ends.
 10/27/2009 9:19:40 AM -- Dispatching transactions begins ...
 10/27/2009 9:19:40 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 10/27/2009 9:19:40 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 10/27/2009 9:19:46 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
 10/27/2009 9:19:46 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
 10/27/2009 9:19:48 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 10/27/2009 9:19:48 AM -- Performance logging begins (interval: 15000 ms).
 10/27/2009 9:19:48 AM -- Attaining prerequisites:
 10/27/2009 9:21:05 AM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1211498000.0 (lower bound: 1207960000.0, upper bound: none)
 10/28/2009 9:21:06 AM -- Performance logging ends.
 10/28/2009 9:21:06 AM -- JetInterop batch transaction stats: 2811025, 2811286, 2812367, 2809529, and 2813546.
 10/28/2009 9:21:06 AM -- Dispatching transactions ends.
 10/28/2009 9:21:06 AM -- Shutting down databases ...
 10/28/2009 9:21:08 AM -- Instance2556.1 (complete), Instance2556.2 (complete), Instance2556.3 (complete), Instance2556.4 (complete), and Instance2556.5 (complete)
 10/28/2009 9:21:09 AM -- Performance logging begins (interval: 30000 ms).
 10/28/2009 9:21:09 AM -- Verifying database checksums ...
 10/28/2009 1:45:38 PM -- E: (100% processed), G: (100% processed), H: (100% processed), and I: (100% processed)
 10/28/2009 1:45:38 PM -- Performance logging ends.
 10/28/2009 1:45:38 PM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_28_9_21_8.blg has 528 samples.



Server 2: XIV31

Microsoft Exchange Server Jetstress

Test Result Report

Checksum Statistics - All

Database	Seen pages	Bad pages	Correctable pages	Wrong page-number pages	File length / seconds taken
E:\Jetstress001001.edb	68445826	0	0	0	2138932 MBytes / 16351 sec
F:\Jetstress002001.edb	68449154	0	0	0	2139036 MBytes / 16031 sec
G:\Jetstress003001.edb	68440962	0	0	0	2138780 MBytes / 16024 sec
H:\Jetstress004001.edb	68448642	0	0	0	2139020 MBytes / 16222 sec
I:\Jetstress005001.edb	68449410	0	0	0	2139044 MBytes / 16222 sec
(Sum)	342233994	0	0	0	10694812 MBytes / 16351 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
E:	0.028	0.000	2091.993	0.000	65536.000
F:	0.027	0.000	2134.510	0.000	65536.000
G:	0.027	0.000	2135.319	0.000	65536.000
H:	0.028	0.000	2109.433	0.000	65536.000
I:	0.027	0.000	2109.431	0.000	65536.000

Memory System Performance (of checksum)

Counter	Average	Minimum	Maximum
% Processor Time	13.948	3.093	19.237
Available MBytes	10996.849	10896.000	11013.000
Free System Page Table Entries	33562022.619	33559418.000	33562429.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	63651636.706	63610880.000	63680512.000
Pool Paged Bytes	124527653.647	122896384.000	164315136.000

Test Log 10/26/2009 11:09:23 AM -- Jetstress testing begins ...

10/26/2009 11:09:23 AM -- Prepare testing begins ...

10/26/2009 11:09:29 AM -- Attaching databases ...

10/26/2009 11:09:29 AM -- Prepare testing ends.

10/26/2009 11:09:29 AM -- Dispatching transactions begins ...

10/26/2009 11:09:29 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)

10/26/2009 11:09:29 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)



10/26/2009 11:09:35 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/26/2009 11:09:35 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/26/2009 11:09:37 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/26/2009 11:09:37 AM -- Performance logging begins (interval: 15000 ms).

10/26/2009 11:09:37 AM -- Attaining prerequisites:

10/26/2009 11:09:45 AM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 233312300.0 (lower bound: 1207960000.0, upper bound: none)

10/26/2009 11:09:45 AM -- Performance logging ends.

10/26/2009 11:09:45 AM -- JetInterop batch transaction stats: 114, 87, 87, 111, and 115.

10/26/2009 11:09:46 AM -- Dispatching transactions ends.

10/26/2009 11:09:46 AM -- Shutting down databases ...

10/26/2009 11:09:46 AM -- Instance1804.1 (complete), Instance1804.2 (complete), Instance1804.3 (complete), Instance1804.4 (complete), and Instance1804.5 (complete)

10/26/2009 11:09:46 AM -- Jetstress testing ends.

10/26/2009 11:10:50 AM -- Jetstress testing begins ...

10/26/2009 11:10:50 AM -- Prepare testing begins ...

10/26/2009 11:10:56 AM -- Attaching databases ...

10/26/2009 11:10:56 AM -- Prepare testing ends.

10/26/2009 11:10:56 AM -- Dispatching transactions begins ...

10/26/2009 11:10:56 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)

10/26/2009 11:10:56 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)

10/26/2009 11:11:02 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/26/2009 11:11:02 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/26/2009 11:11:04 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/26/2009 11:11:04 AM -- Performance logging begins (interval: 15000 ms).

10/26/2009 11:11:04 AM -- Attaining prerequisites:

10/26/2009 11:11:04 AM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 166682600.0 (lower bound: 1207960000.0, upper bound: none)

10/26/2009 11:11:04 AM -- Performance logging ends.

10/26/2009 11:11:04 AM -- JetInterop batch transaction stats: 56, 69, 45, 40, and 48.

10/26/2009 11:11:05 AM -- Dispatching transactions ends.

10/26/2009 11:11:05 AM -- Shutting down databases ...

10/26/2009 11:11:05 AM -- Instance1804.1 (complete), Instance1804.2 (complete), Instance1804.3 (complete), Instance1804.4 (complete), and Instance1804.5 (complete)

10/26/2009 11:11:05 AM -- Jetstress testing ends.

10/27/2009 9:19:38 AM -- Jetstress testing begins ...

10/27/2009 9:19:38 AM -- Prepare testing begins ...

10/27/2009 9:19:43 AM -- Attaching databases ...

10/27/2009 9:19:43 AM -- Prepare testing ends.

10/27/2009 9:19:43 AM -- Dispatching transactions begins ...

10/27/2009 9:19:43 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)

10/27/2009 9:19:43 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)

10/27/2009 9:19:49 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

10/27/2009 9:19:49 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

10/27/2009 9:19:51 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

10/27/2009 9:19:51 AM -- Performance logging begins (interval: 15000 ms).

10/27/2009 9:19:51 AM -- Attaining prerequisites:



10/27/2009 9:21:08 AM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1218388000.0 (lower bound: 1207960000.0, upper bound: none)
 10/28/2009 9:21:09 AM -- Performance logging ends.
 10/28/2009 9:21:09 AM -- JetInterop batch transaction stats: 2579747, 2578281, 2575525, 2577396, and 2575772.
 10/28/2009 9:21:09 AM -- Dispatching transactions ends.
 10/28/2009 9:21:09 AM -- Shutting down databases ...
 10/28/2009 9:21:12 AM -- Instance1804.1 (complete), Instance1804.2 (complete), Instance1804.3 (complete), Instance1804.4 (complete), and Instance1804.5 (complete)
 10/28/2009 9:21:12 AM -- Performance logging begins (interval: 30000 ms).
 10/28/2009 9:21:12 AM -- Verifying database checksums ...
 10/28/2009 1:53:44 PM -- E: (100% processed), F: (100% processed), G: (100% processed), H: (100% processed), and I: (100% processed)
 10/28/2009 1:53:44 PM -- Performance logging ends.
 10/28/2009 1:53:44 PM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_28_9_21_12.blg has 544 samples.

Performance Testing

Server 1: XIV30

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name XIV30

Test Description

Test Start Time 10/23/2009 7:57:27 PM

Test End Time 10/23/2009 9:58:55 PM

Jetstress Version 14.00.0639.012

Ese Version 14.00.0639.011

Operating System Windows Server (R) 2008 Enterprise Service Pack 2 (6.0.6002.131072)

Performance Log C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_23_19_57_38.blg
 C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_23_21_58_55.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 10380.944
Target Transactional I/O per Second 10000
Initial Database Size (bytes) 10612715094016
Final Database Size (bytes) 10636496797696
Database Files (Count) 5



Jetstress System Parameters

Thread Count	64 (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

Instance5160.1 Log Path: J:\1
 Database: E:\Jetstress001001.edb

Instance5160.2 Log Path: j:\2
 Database: F:\Jetstress002001.edb

Instance5160.3 Log Path: j:\3
 Database: g:\Jetstress003001.edb

Instance5160.4 Log Path: j:\4
 Database: h:\Jetstress004001.edb

Instance5160.5 Log Path: j:\5
 Database: i:\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 Average Bytes	I/O Database 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
Instance5160.1	17.376	7.750	1433.855	641.962	33012.670	34839.569	0.000	2.324	0.000	184.423	0.000	12051.129
Instance5160.2	16.850	7.597	1433.561	642.274	33020.953	34813.958	0.000	2.320	0.000	184.603	0.000	11987.824
Instance5160.3	16.866	7.396	1433.931	642.650	33022.400	34803.989	0.000	2.319	0.000	184.968	0.000	11963.726
Instance5160.4	16.810	7.004	1430.128	639.693	33014.397	34798.152	0.000	2.323	0.000	183.729	0.000	11950.455
Instance5160.5	16.825	6.805	1438.336	644.555	33020.926	34793.426	0.000	2.311	0.000	185.054	0.000	11938.460

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance5160.1	29.875	261992.338
Instance5160.2	28.079	261984.418
Instance5160.3	28.074	261966.032
Instance5160.4	28.105	261938.089
Instance5160.5	28.106	261974.123

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance5160.1	9.438	232555.901
Instance5160.2	9.403	232553.913
Instance5160.3	9.405	232555.575
Instance5160.4	9.329	232555.170
Instance5160.5	9.385	232554.733



Total I/O Performance

MSExchange Database ==> Instances	I/O Data base Read s Average Latency (msec)	I/O Data base Write s 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
Instance 5160.1	17.376	7.750	1463.730	641.962	37686.243	34839.569	15.246	2.324	9.438	184.423	23255.901	12051.129
Instance 5160.2	16.850	7.597	1461.639	642.274	37419.410	34813.958	15.189	2.320	9.403	184.603	23255.3913	11987.824
Instance 5160.3	16.866	7.396	1462.004	642.650	37418.595	34803.989	15.244	2.319	9.405	184.968	23255.5575	11963.726
Instance 5160.4	16.810	7.004	1458.233	639.693	37426.508	34798.152	15.358	2.323	9.329	183.729	23255.5170	11950.455
Instance 5160.5	16.825	6.805	1466.442	644.555	37409.056	34793.426	15.612	2.311	9.385	185.054	23255.4733	11938.460

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	17.001	3.920	21.223
Available MBytes	13349.700	13324.000	13418.000
Free System Page Table Entries	33560749.415	33560314.000	33562201.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	61277175.467	61272064.000	61288448.000
Pool Paged Bytes	167644552.533	167575552.000	167735296.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/23/2009 7:57:27 PM -- Jetstress testing begins ...

10/23/2009 7:57:27 PM -- Prepare testing begins ...
 10/23/2009 7:57:32 PM -- Attaching databases ...
 10/23/2009 7:57:32 PM -- Prepare testing ends.
 10/23/2009 7:57:32 PM -- Dispatching transactions begins ...
 10/23/2009 7:57:32 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 10/23/2009 7:57:32 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 10/23/2009 7:57:38 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 50 msec/read).
 10/23/2009 7:57:38 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 50 msec/write).
 10/23/2009 7:57:40 PM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 10/23/2009 7:57:40 PM -- Performance logging begins (interval: 15000 ms).
 10/23/2009 7:57:40 PM -- Attaining prerequisites:



10/23/2009 7:58:51 PM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1224880000.0 (lower bound: 1207960000.0, upper bound: none)
10/23/2009 9:58:52 PM -- Performance logging ends.
10/23/2009 9:58:52 PM -- JetInterop batch transaction stats: 223268, 223335, 223240, 222880, and 223551.
10/23/2009 9:58:53 PM -- Dispatching transactions ends.
10/23/2009 9:58:53 PM -- Shutting down databases ...
10/23/2009 9:58:55 PM -- Instance5160.1 (complete), Instance5160.2 (complete), Instance5160.3 (complete), Instance5160.4 (complete), and Instance5160.5 (complete)
10/23/2009 9:58:55 PM -- Performance logging begins (interval: 30000 ms).
10/23/2009 9:58:55 PM -- Verifying database checksums ...
10/24/2009 2:16:17 AM -- E: (100% processed), F: (100% processed), g: (100% processed), h: (100% processed), and i: (100% processed)
10/24/2009 2:16:17 AM -- Performance logging ends.
10/24/2009 2:16:17 AM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_23_21_58_55.blg has 514 samples.
10/24/2009 2:16:28 AM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_23_21_58_55.html is saved.
10/24/2009 2:16:28 AM -- Verifying log checksums ...
10/24/2009 2:16:30 AM -- J:\1 (24 log(s) processed), j:\2 (24 log(s) processed), j:\3 (25 log(s) processed), j:\4 (24 log(s) processed), and j:\5 (23 log(s) processed)
10/24/2009 2:16:30 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_23_19_57_38.blg has 484 samples.
10/24/2009 2:16:30 AM -- Creating test report ...
10/24/2009 2:16:35 AM -- Instance5160.1 has 17.4 for I/O Database Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.1 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.1 has 2.3 for I/O Log Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.2 has 16.8 for I/O Database Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.2 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.2 has 2.3 for I/O Log Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.3 has 16.9 for I/O Database Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.3 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.3 has 2.3 for I/O Log Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.4 has 16.8 for I/O Database Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.4 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.4 has 2.3 for I/O Log Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.5 has 16.8 for I/O Database Reads Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.5 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:35 AM -- Instance5160.5 has 2.3 for I/O Log Reads Average Latency.
10/24/2009 2:16:35 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/24/2009 2:16:35 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/24/2009 2:16:35 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_23_19_57_38.xml has 479 samples queried.



Server 2: XIV31

Performance Test Result Report

Test Summary

Overall Test Result **Pass**

Machine Name XIV31

Test Description

Test Start Time 10/23/2009 7:57:32 PM

Test End Time 10/23/2009 9:59:16 PM

Jetstress Version 14.00.0639.012

Ese Version 14.00.0639.011

Operating System Windows Server (R) 2008 Enterprise Service Pack 2 (6.0.6002.131072)

Performance Log C:\Program Files\Exchange
Jetstress\ESRP\Performance_2009_10_23_19_57_44.blg
C:\Program Files\Exchange
Jetstress\ESRP\DBChecksum_2009_10_23_21_59_16.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 10174.773
Target Transactional I/O per Second 10000
Initial Database Size (bytes) 10607430270976
Final Database Size (bytes) 10630985482240
Database Files (Count) 5

Jetstress System Parameters

Thread Count 64 (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

Instance2932.1 Log Path: J:\1
Database: E:\Jetstress001001.edb

Instance2932.2 Log Path: J:\2
Database: F:\Jetstress002001.edb

Instance2932.3 Log Path: J:\3
Database: G:\Jetstress003001.edb

Instance2932.4 Log Path: j:\4
Database: H:\Jetstress004001.edb

Instance2932.5 Log Path: j:\5
Database: I:\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
Instance2932.1	17.40	7.695	1406.321	626.849	33241.361	34884.229	0.00	2.31	0.000	181.0	0.00	12000.643
Instance2932.2	16.86	7.397	1406.559	627.668	33240.170	34876.003	0.00	2.31	0.000	181.6	0.00	12043.133
Instance2932.3	16.86	7.218	1407.714	627.741	33230.243	34866.123	0.00	2.31	0.000	181.4	0.00	12010.513
Instance2932.4	16.84	6.967	1407.229	628.574	33244.868	34851.959	0.00	2.31	0.000	181.3	0.00	11981.009
Instance2932.5	16.84	6.624	1407.065	629.053	33238.005	34863.523	0.00	2.31	0.000	181.5	0.00	12044.245



Background Database Maintenance I/O Performance

MSEExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2932.1	29.875	261962.436
Instance2932.2	28.058	262001.030
Instance2932.3	28.085	261963.597
Instance2932.4	28.115	262018.125
Instance2932.5	28.124	261997.817

Log Replication I/O Performance

MSEExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2932.1	9.274	232558.229
Instance2932.2	9.339	232558.264
Instance2932.3	9.303	232553.711
Instance2932.4	9.279	232556.639
Instance2932.5	9.336	232555.376

Total I/O Performance

MSEExchange 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
Instance 2932.1	17.402	7.695	1436.196	626.849	37999.091	34884.229	15.707	2.316	9.274	181.032	232558.229	12000.643
Instance 2932.2	16.866	7.397	1434.617	627.668	37714.192	34876.003	15.860	2.315	9.339	181.676	232558.264	12043.133
Instance 2932.3	16.868	7.218	1435.799	627.741	37704.338	34866.123	15.799	2.315	9.303	181.477	232553.711	12010.513
Instance 2932.4	16.844	6.967	1435.344	628.574	37726.056	34851.959	15.973	2.315	9.279	181.343	232556.639	11981.009
Instance 2932.5	16.848	6.624	1435.190	629.053	37720.829	34863.523	15.630	2.316	9.336	181.539	232555.376	12044.245



Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	16.746	2.882	23.433
Available MBytes	9605.400	9587.000	9724.000
Free System Page Table Entries	33559967.165	33559514.000	33561517.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	72571724.800	72318976.000	73269248.000
Pool Paged Bytes	168278024.533	168194048.000	168394752.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/23/2009 7:57:32 PM -- Jetstress testing begins ...
10/23/2009 7:57:32 PM -- Prepare testing begins ...
10/23/2009 7:57:38 PM -- Attaching databases ...
10/23/2009 7:57:38 PM -- Prepare testing ends.
10/23/2009 7:57:38 PM -- Dispatching transactions begins ...
10/23/2009 7:57:38 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
10/23/2009 7:57:38 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
10/23/2009 7:57:44 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 50 msec/read).
10/23/2009 7:57:44 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 50 msec/write).
10/23/2009 7:57:46 PM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/23/2009 7:57:46 PM -- Performance logging begins (interval: 15000 ms).
10/23/2009 7:57:46 PM -- Attaining prerequisites:
10/23/2009 7:59:13 PM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1216131000.0 (lower bound: 1207960000.0, upper bound: none)
10/23/2009 9:59:13 PM -- Performance logging ends.
10/23/2009 9:59:13 PM -- JetInterop batch transaction stats: 218405, 219075, 219030, 218693, and 218990.
10/23/2009 9:59:14 PM -- Dispatching transactions ends.
10/23/2009 9:59:14 PM -- Shutting down databases ...
10/23/2009 9:59:16 PM -- Instance2932.1 (complete), Instance2932.2 (complete), Instance2932.3 (complete), Instance2932.4 (complete), and Instance2932.5 (complete)
10/23/2009 9:59:17 PM -- Performance logging begins (interval: 30000 ms).
10/23/2009 9:59:17 PM -- Verifying database checksums ...
10/24/2009 2:16:11 AM -- E: (100% processed), F: (100% processed), G: (100% processed), H: (100% processed), and I: (100% processed)
10/24/2009 2:16:11 AM -- Performance logging ends.
10/24/2009 2:16:11 AM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_23_21_59_16.blg has 513 samples.
10/24/2009 2:16:22 AM -- C:\Program Files\Exchange Jetstress\ESRP\DBChecksum_2009_10_23_21_59_16.html is saved.
10/24/2009 2:16:22 AM -- Verifying log checksums ...
10/24/2009 2:16:23 AM -- J:\1 (25 log(s) processed), J:\2 (26 log(s) processed), J:\3 (26 log(s) processed), j:\4 (25 log(s) processed), and j:\5 (22 log(s) processed)
10/24/2009 2:16:23 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_23_19_57_44.blg has 485 samples.
10/24/2009 2:16:23 AM -- Creating test report ...
10/24/2009 2:16:29 AM -- Instance2932.1 has 17.4 for I/O Database Reads Average Latency.
10/24/2009 2:16:29 AM -- Instance2932.1 has 2.3 for I/O Log Writes Average Latency.
10/24/2009 2:16:29 AM -- Instance2932.1 has 2.3 for I/O Log Reads Average Latency.



10/24/2009 2:16:29 AM -- Instance2932.2 has 16.9 for I/O Database Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.2 has 2.3 for I/O Log Writes Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.2 has 2.3 for I/O Log Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.3 has 16.9 for I/O Database Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.3 has 2.3 for I/O Log Writes Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.3 has 2.3 for I/O Log Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.4 has 16.8 for I/O Database Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.4 has 2.3 for I/O Log Writes Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.4 has 2.3 for I/O Log Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.5 has 16.8 for I/O Database Reads Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.5 has 2.3 for I/O Log Writes Average Latency.
 10/24/2009 2:16:29 AM -- Instance2932.5 has 2.3 for I/O Log Reads Average Latency.
 10/24/2009 2:16:29 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 10/24/2009 2:16:29 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
 10/24/2009 2:16:29 AM -- C:\Program Files\Exchange
 Jetstress\ESRP\Performance_2009_10_23_19_57_44.xml has 479 samples queried.

Streaming Backup Testing

Server 1: XIV30

Microsoft Exchange Server Jetstress

Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance4776.1	2110268.09	06:48:46	86.04
Instance4776.2	2110372.09	06:51:00	85.58
Instance4776.3	2110204.09	06:51:20	85.50
Instance4776.4	2110412.09	06:51:13	85.53
Instance4776.5	2110276.09	06:51:44	85.42

Jetstress System Parameters

Thread Count 64 (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

Instance4776.1 Log Path: j:\1
Database: e:\Jetstress001001.edb

Instance4776.2 Log Path: j:\2
Database: f:\Jetstress002001.edb

Instance4776.3 Log Path: j:\3
Database: g:\Jetstress003001.edb

Instance4776.4 Log Path: j:\4
Database: h:\Jetstress004001.edb

Instance4776.5 Log Path: j:\5
Database: i:\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
Instance4776.1	4.782	0.000	344.200	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance4776.2	4.785	0.000	342.314	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance4776.3	4.790	0.000	341.980	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance4776.4	4.790	0.000	342.123	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance4776.5	4.790	0.000	341.649	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00



Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	4.167	2.104	7.442
Available MBytes	14804.038	14764.000	14811.000
Free System Page Table Entries	33563421.588	33562826.000	33563495.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	62443006.754	62398464.000	62509056.000
Pool Paged Bytes	163591604.010	162455552.000	208773120.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/29/2009 8:57:20 PM -- Jetstress testing begins ...
10/29/2009 8:57:20 PM -- Prepare testing begins ...
10/29/2009 8:57:25 PM -- Attaching databases ...
10/29/2009 8:57:25 PM -- Prepare testing ends.
10/29/2009 8:57:32 PM -- Performance logging begins (interval: 30000 ms).
10/29/2009 8:57:32 PM -- Backing up databases ...
10/30/2009 3:49:16 AM -- Performance logging ends.
10/30/2009 3:49:16 AM -- Instance4776.1 (100% processed), Instance4776.2 (100% processed), Instance4776.3 (100% processed), Instance4776.4 (100% processed), and Instance4776.5 (100% processed)
10/30/2009 3:49:16 AM -- C:\Program Files\Exchange Jetstress\ESRP\DatabaseBackup_2009_10_29_20_57_26.blg has 822 samples.
10/30/2009 3:49:16 AM -- Creating test report ...



Server 2: XIV31

Microsoft Exchange Server Jetstress

Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance2084.1	2192044.09	07:02:05	86.55
Instance2084.2	2192132.09	07:03:09	86.34
Instance2084.3	2191892.09	07:02:38	86.43
Instance2084.4	2192220.09	07:03:03	86.37
Instance2084.5	2192140.09	07:02:44	86.42

Jetstress System Parameters

Thread Count 64 (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

Instance2084.1 Log Path: j:\1
Database: e:\Jetstress001001.edb

Instance2084.2 Log Path: j:\2
Database: f:\Jetstress002001.edb

Instance2084.3 Log Path: j:\3
Database: g:\Jetstress003001.edb

Instance2084.4 Log Path: j:\4
Database: h:\Jetstress004001.edb

Instance2084.5 Log Path: j:\5
Database: i:\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 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
Instance2084.1	4.751	0.000	346.204	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2084.2	4.754	0.000	345.294	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2084.3	4.755	0.000	345.711	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2084.4	4.749	0.000	345.448	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2084.5	4.753	0.000	345.641	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	4.206	2.929	6.364
Available MBytes	11045.922	11009.000	11048.000
Free System Page Table Entries	33563005.664	33562506.000	33563079.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	64585863.725	64565248.000	64647168.000
Pool Paged Bytes	134642301.425	133402624.000	174411776.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log 10/29/2009 8:57:25 PM -- Jetstress testing begins ...

10/29/2009 8:57:25 PM -- Prepare testing begins ...

10/29/2009 8:57:30 PM -- Attaching databases ...

10/29/2009 8:57:30 PM -- Prepare testing ends.

10/29/2009 8:57:37 PM -- Performance logging begins (interval: 30000 ms).

10/29/2009 8:57:37 PM -- Backing up databases ...

10/30/2009 4:00:47 AM -- Performance logging ends.

10/30/2009 4:00:47 AM -- Instance2084.1 (100% processed), Instance2084.2 (100% processed), Instance2084.3 (100% processed), Instance2084.4 (100% processed), and Instance2084.5 (100% processed)

10/30/2009 4:00:47 AM -- C:\Program Files\Exchange

Jetstress\ESRP\DatabaseBackup_2009_10_29_20_57_30.blg has 845 samples.

10/30/2009 4:00:47 AM -- Creating test report ...



Soft Recovery Testing

Server 1: XIV30

Microsoft Exchange Server Jetstress

SoftRecovery Test Result Report

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance5568.1	512	365.4635427
Instance5568.2	513	375.0108039
Instance5568.3	507	360.6899121
Instance5568.4	516	355.3858781
Instance5568.5	500	356.9770883

Database Configuration

Instance5568.1 Log Path: j:\1

Database: e:Jetstress001001.edb

Instance5568.2 Log Path: j:\2

Database: f:Jetstress002001.edb

Instance5568.3 Log Path: j:\3

Database: g:Jetstress003001.edb

Instance5568.4 Log Path: j:\4

Database: h:Jetstress004001.edb

Instance5568.5 Log Path: j:\5

Database: i:Jetstress005001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Data base Read Averages Latency (msec)	I/O Data base Write Averages Latency (msec)	I/O Database Reads /sec	I/O Database Writes /sec	I/O Database Average Bytes	I/O Database 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
Instance5568.1	13.155	2.340	1249.818	8.425	36052.310	32399.820	10.969	0.000	12.638	0.000	229902.398	0.000
Instance5568.2	13.406	2.284	1163.885	8.264	35378.047	32768.000	12.003	0.000	12.396	0.000	232577.978	0.000
Instance5568.5	13.84	2.280	1179.	8.538	35438	32768	12.0	0.00	12.80	0.000	232505	0.00





568.3	2	701	.794	.000	81	0	7	.576	0			
Instance5	13.41	2.291	1232.	8.807	36037	32391	11.9	0.00	13.21	0.000	229890	0.00
568.4	6	082	.534	.356	07	0	1	.002	0			
Instance5	13.33	2.297	1234.	8.499	36038	32391	12.5	0.00	12.74	0.000	229837	0.00
568.5	2	206	.311	.356	69	0	9	.251	0			

Background Database Maintenance I/O Performance

MSEExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance5568.1	30.241	261997.966
Instance5568.2	29.982	262010.156
Instance5568.3	30.162	261945.445
Instance5568.4	30.226	261959.083
Instance5568.5	30.137	262088.268

Total I/O Performance

MSEExchange Database ==> Instances	I/O Database Reads /sec	I/O Database Writes /sec	I/O Database Average Bytes	I/O Database Average Latency (msec)	I/O Database Reads /sec	I/O Database Writes /sec	I/O Database Average Bytes	I/O Database Average Latency (msec)	I/O Log Reads /sec	I/O Log Writes /sec	I/O Log Average Bytes	I/O Log Average Bytes
Instance5	13.15	2.340	1280.	8.425	41390	32399	10.9	0.00	12.63	0.000	229902	0.00
568.1	5	059	.292	.820	69	0	8	.398	0			
Instance5	13.40	2.284	1193.	8.264	41069	32768	12.0	0.00	12.39	0.000	232577	0.00
568.2	6	867	.503	.000	03	0	6	.978	0			
Instance5	13.84	2.280	1209.	8.538	41085	32768	12.0	0.00	12.80	0.000	232505	0.00
568.3	2	863	.557	.000	81	0	7	.576	0			
Instance5	13.41	2.291	1262.	8.807	41447	32391	11.9	0.00	13.21	0.000	229890	0.00
568.4	6	307	.163	.356	07	0	1	.002	0			
Instance5	13.33	2.297	1264.	8.499	41426	32391	12.5	0.00	12.74	0.000	229837	0.00
568.5	2	343	.468	.356	69	0	9	.251	0			

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	6.448	1.550	13.808
Available MBytes	13410.430	13329.000	14579.000
Free System Page Table Entries	33562833.903	33561435.000	33563408.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	64408631.054	62427136.000	65032192.000



Pool Paged Bytes	163951428.817	162467840.000	164552704.000
Database Page Fault Stalls/sec	0.005	0.000	0.497

Test Log 10/30/2009 9:34:28 AM -- Jetstress testing begins ...
10/30/2009 9:34:28 AM -- Prepare testing begins ...
10/30/2009 9:34:33 AM -- Attaching databases ...
10/30/2009 9:34:33 AM -- Prepare testing ends.
10/30/2009 9:34:33 AM -- Dispatching transactions begins ...
10/30/2009 9:34:33 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
10/30/2009 9:34:33 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
10/30/2009 9:34:39 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 50 msec/read).
10/30/2009 9:34:39 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 50 msec/write).
10/30/2009 9:34:40 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/30/2009 9:34:40 AM -- Performance logging begins (interval: 15000 ms).
10/30/2009 9:34:40 AM -- Generating log files ...
10/30/2009 9:46:40 AM -- j:\1 (102.4% generated), j:\2 (102.8% generated), j:\3 (101.6% generated), j:\4 (103.4% generated), and j:\5 (100.2% generated)
10/30/2009 9:46:40 AM -- Performance logging ends.
10/30/2009 9:46:40 AM -- JetInterop batch transaction stats: 15411, 15547, 15271, 15388, and 15268.
10/30/2009 9:46:40 AM -- Dispatching transactions ends.
10/30/2009 9:46:40 AM -- Shutting down databases ...
10/30/2009 9:46:42 AM -- Instance5568.1 (complete), Instance5568.2 (complete), Instance5568.3 (complete), Instance5568.4 (complete), and Instance5568.5 (complete)
10/30/2009 9:46:42 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_39.blg has 47 samples.
10/30/2009 9:46:42 AM -- Creating test report ...
10/30/2009 9:46:43 AM -- Instance5568.1 has 11.0 for I/O Database Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.1 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.1 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.2 has 10.8 for I/O Database Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.2 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.2 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.3 has 10.7 for I/O Database Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.3 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.3 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.4 has 10.7 for I/O Database Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.4 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.4 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.5 has 10.7 for I/O Database Reads Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.5 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:46:43 AM -- Instance5568.5 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:46:43 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/30/2009 9:46:43 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/30/2009 9:46:43 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_39.xml has 46 samples queried.
10/30/2009 9:46:43 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_39.html is saved.
10/30/2009 9:46:44 AM -- Performance logging begins (interval: 4000 ms).
10/30/2009 9:46:44 AM -- Recovering databases ...
10/30/2009 9:52:59 AM -- Performance logging ends.
10/30/2009 9:52:59 AM -- Instance5568.1 (365.4635427), Instance5568.2 (375.0108039),



Instance5568.3 (360.6899121), Instance5568.4 (355.3858781), and Instance5568.5 (356.9770883)
 10/30/2009 9:53:00 AM -- C:\Program Files\Exchange
 Jetstress\ESRP\SoftRecovery_2009_10_30_9_46_43.blg has 93 samples.
 10/30/2009 9:53:00 AM -- Creating test report ...

Server 2: XIV31

Microsoft Exchange Server Jetstress
 SoftRecovery Test Result Report
 Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance992.1	500	326.198091
Instance992.2	510	322.7504689
Instance992.3	514	316.1204264
Instance992.4	507	316.1204264
Instance992.5	512	307.633972

Database Configuration

Instance992.1 Log Path: j:\1
 Database: e:Jetstress001001.edb

Instance992.2 Log Path: j:\2
 Database: f:Jetstress002001.edb

Instance992.3 Log Path: j:\3
 Database: g:Jetstress003001.edb

Instance992.4 Log Path: j:\4
 Database: h:Jetstress004001.edb

Instance992.5 Log Path: j:\5
 Database: i:Jetstress005001.edb

Transactional I/O Performance

MSExchange Instance	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
Instance 992.1	12.288	2.047	1274.782	9.298	35486.459	32768.000	9.431	0.038	13.963	0.019	232479.251	6.400



Instance 992.2	11.49	2.076	1372.124	9.586	35334	32768	9.50	0.00	14.37	0.000	232474	0.00
	6				.272	.000	6	0	8		.193	0
Instance 992.3	11.83	2.127	1352.839	9.815	35345	32768	10.9	0.00	14.72	0.000	232535	0.00
	7				.634	.000	29	0	3		.822	0
Instance 992.4	11.99	2.098	1331.135	9.699	35326	32768	10.3	0.00	14.54	0.000	232517	0.00
	4				.828	.000	33	0	9		.293	0
Instance 992.5	12.27	2.191	1330.853	10.07	35479	32768	10.8	0.00	15.11	0.000	232582	0.00
	6				.188	.000	89	0	6		.784	0

Background Database Maintenance I/O Performance

MSEExchange Database Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance992.1	30.858	261928.234
Instance992.2	31.094	261878.361
Instance992.3	31.076	261865.134
Instance992.4	30.969	262076.690
Instance992.5	31.030	261977.170

Total I/O Performance

MSEExchange 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
Instance 992.1	12.28	2.047	1305.640	9.298	40838	32768	9.43	0.03	13.96	0.019	232479	6.40
	8				.202	.000	1	8	3		.251	0
Instance 992.2	11.49	2.076	1403.218	9.586	40354	32768	9.50	0.00	14.37	0.000	232474	0.00
	6				.292	.000	6	0	8		.193	0
Instance 992.3	11.83	2.127	1383.915	9.815	40432	32768	10.9	0.00	14.72	0.000	232535	0.00
	7				.107	.000	29	0	3		.822	0
Instance 992.4	11.99	2.098	1362.104	9.699	40482	32768	10.3	0.00	14.54	0.000	232517	0.00
	4				.273	.000	33	0	9		.293	0
Instance 992.5	12.27	2.191	1361.883	10.07	40639	32768	10.8	0.00	15.11	0.000	232582	0.00
	6				.828	.000	89	0	6		.784	0

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	7.267	2.859	12.839
Available MBytes	9647.716	9589.000	10817.000



Free System Page Table Entries	33562035.210	33561169.000	33562497.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	68644914.568	64688128.000	69050368.000
Pool Paged Bytes	135223826.963	133816320.000	135286784.000
Database Page Fault Stalls/sec	0.009	0.000	0.496

Test Log 10/30/2009 9:34:33 AM -- Jetstress testing begins ...
10/30/2009 9:34:33 AM -- Prepare testing begins ...
10/30/2009 9:34:38 AM -- Attaching databases ...
10/30/2009 9:34:38 AM -- Prepare testing ends.
10/30/2009 9:34:38 AM -- Dispatching transactions begins ...
10/30/2009 9:34:38 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
10/30/2009 9:34:38 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
10/30/2009 9:34:44 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 50 msec/read).
10/30/2009 9:34:44 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 50 msec/write).
10/30/2009 9:34:46 AM -- Operation mix: Sessions 64, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
10/30/2009 9:34:46 AM -- Performance logging begins (interval: 15000 ms).
10/30/2009 9:34:46 AM -- Generating log files ...
10/30/2009 9:49:49 AM -- j:\1 (100.2% generated), j:\2 (102.0% generated), j:\3 (102.8% generated), j:\4 (101.4% generated), and j:\5 (102.4% generated)
10/30/2009 9:49:49 AM -- Performance logging ends.
10/30/2009 9:49:49 AM -- JetInterop batch transaction stats: 14857, 14963, 14881, 14898, and 14945.
10/30/2009 9:49:49 AM -- Dispatching transactions ends.
10/30/2009 9:49:49 AM -- Shutting down databases ...
10/30/2009 9:49:51 AM -- Instance992.1 (complete), Instance992.2 (complete), Instance992.3 (complete), Instance992.4 (complete), and Instance992.5 (complete)
10/30/2009 9:49:51 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_44.blg has 60 samples.
10/30/2009 9:49:51 AM -- Creating test report ...
10/30/2009 9:49:51 AM -- Instance992.1 has 11.9 for I/O Database Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.1 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:49:51 AM -- Instance992.1 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.2 has 11.5 for I/O Database Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.2 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:49:51 AM -- Instance992.2 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.3 has 11.6 for I/O Database Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.3 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:49:51 AM -- Instance992.3 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.4 has 11.5 for I/O Database Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.4 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:49:51 AM -- Instance992.4 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.5 has 11.6 for I/O Database Reads Average Latency.
10/30/2009 9:49:51 AM -- Instance992.5 has 1.5 for I/O Log Writes Average Latency.
10/30/2009 9:49:51 AM -- Instance992.5 has 1.5 for I/O Log Reads Average Latency.
10/30/2009 9:49:51 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
10/30/2009 9:49:51 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
10/30/2009 9:49:51 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_44.xml has 59 samples queried.
10/30/2009 9:49:52 AM -- C:\Program Files\Exchange Jetstress\ESRP\Performance_2009_10_30_9_34_44.html is saved.



10/30/2009 9:49:53 AM -- Performance logging begins (interval: 4000 ms).
10/30/2009 9:49:53 AM -- Recovering databases ...
10/30/2009 9:55:19 AM -- Performance logging ends.
10/30/2009 9:55:19 AM -- Instance992.1 (326.198091), Instance992.2 (322.7504689),
Instance992.3 (316.1204264), Instance992.4 (316.1204264), and Instance992.5 (307.633972)
10/30/2009 9:55:19 AM -- C:\Program Files\Exchange
Jetstress\ESRP\SoftRecovery_2009_10_30_9_49_52.blg has 81 samples.
10/30/2009 9:55:19 AM -- Creating test report ...



Copyright © 2009 by International Business Machines Corporation.



This document could include technical inaccuracies or typographical errors. IBM may make changes, improvements or alterations to the products, programs and services described in this document, including termination of such products, programs and services, at any time and without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only, and IBM shall have no responsibility to update such information.

Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing any such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS-IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM program or product in this document is not intended to state or imply that only that program or product may be used. Any functionally equivalent program or product, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

The provision of the information contained herein is not intended to, and does not grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

IBM, the IBM logo, System x, and System Storage are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Microsoft, Windows, and Exchange Server are trademarks of Microsoft Corporation in the United States, other countries or both.

ⁱ IBM Systems & Technology Group, Kirkland, WA, tquiere@us.ibm.com

ⁱⁱ IBM Systems & Technology Group, Raleigh, NC, dhartma@us.ibm.com

ⁱⁱⁱ IBM Systems & Technology Group, Littleton, MA, aviad@us.ibm.com

^{iv} IBM Systems & Technology Group, Kirkland, WA, davwest@us.ibm.com