Nimble Storage Exchange 2010 40,000-Mailbox Resiliency Storage Solution

Tested with: ESRP Storage Version 3.0
Test date: July 10, 2012

Overview

This document provides information on Nimble Storage's 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 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

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 Nimble 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 Nimble Storage, and Nimble Storage cannot guarantee the accuracy of any information presented after the date of publication.

Features

This document describes the testing of a 40,000 mailbox resiliency (2-copy) database availability group (DAG) configuration for Exchange Server 2010 with a Nimble Storage CS460G-X2 storage system. Nimble Storage has developed an entirely new approach to data storage that converges primary and backup storage into one array, dramatically cutting costs and complexity. To better understand why existing storage systems cannot solve today's IT problems, it helps to step back a few years and examine several key storage technologies that are still in use today.

Meeting the Needs of the Modern Datacenter

The growth of virtualization and server consolidation has compounded the need for storage to simultaneously deliver superior performance, increased capacity, and simplified data protection. Traditional storage architectures are unable to comprehensively meet these demands.

Nimble's patented Cache Accelerated Sequential Layout (CASL™) is the first storage architecture designed from the ground up to seamlessly integrate SSDs with high-capacity disks, delivering affordable performance, integrated backup and disaster recovery, and pain-free operations.

Performance and Capacity Efficiency

Nimble's CASL architecture coalesces random writes into a single RAID stripe, improving write performance by 10x -100x compared to traditional systems. Hot data is served from flash which responds to workload changes within milliseconds instead of hours or days.

Nimble's variable block-based compression technology delivers capacity savings of 30%-75% with no impact on system performance. Combined with highly efficient snapshots, thin provisioning, and zero-copy clones, Nimble offers the industry's best primary storage density, all in a 3U form factor.

Instant Backups and Restores

Nimble eliminates the constraints of traditional backup solutions. Thousands of snapshots can be stored with minimal space consumption, eliminating backup windows, allowing for more frequent recovery points, and enabling quick and efficient replication over a WAN.

Nimble integration simplifies application-consistent backups for Microsoft, Oracle, and VMware. Backups and restores can be completed in seconds, improving application-level recovery point objectives and recovery time objectives.

Stress-Free Operations

Nimble arrays are easy to purchase, set up, and use. A single SKU includes all hardware and software features; there are no extra-cost options. The arrays are designed to be managed by an IT generalist. An intuitive UI simplifies management and eliminates configuration and administration complexities. Built-in application profiles and protection templates make it easy to tune storage systems to optimize performance and data protection for individual workloads.

It is easy to keep Nimble arrays up and running. Firmware upgrades are completely non-disruptive. Built-in real-time monitoring and phone-home support flag problems before failures occur. Remote support access enables quick triage, troubleshooting, and remediation without the need for onsite access.

Solution Description

The following section outlines the Exchange 2010 mailbox resiliency solution that Nimble Storage implemented to run the ESRP tests.

Storage:

- Nimble Storage CS460G-X2 array
- Dual Storage Controllers
- (2) 10 Gbps Ethernet ports per controller with Jumbo Frames
- Nimble Operating System 1.4
- 48 TB of Usable Storage Capacity
- 4.8 TB of Usable Flash Cache
- Link to CS460G-X2 Windows Server Catalog listing: http://windowsservercatalog.com/item.aspx?idItem=4e47eb71-ac19-bba5-7ba1-9f8b9ff2a1d0&bCatID=1282

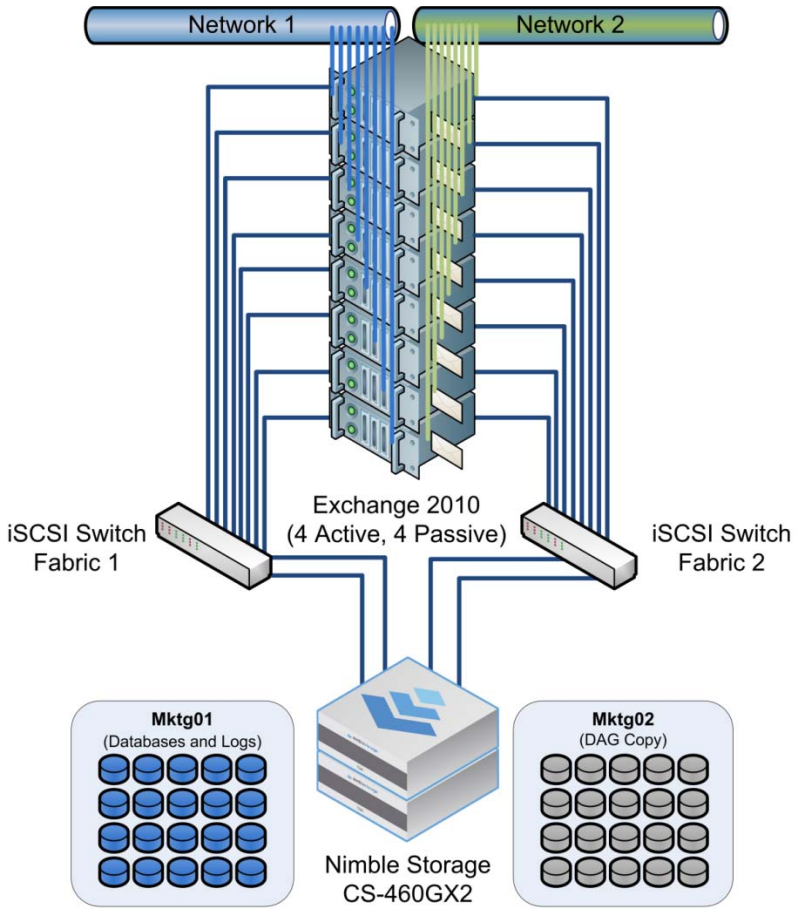

The primary copy storage is described in this document. The secondary copy storage is configured identically (from the host on down to the spindle including brand, model, firmware, drivers, etc.).

The Exchange DAG storage was configured as follows:

- 4 Active Mailbox Servers and 4 Passive Mailbox Servers.
- 2 Nimble Storage CS460G-X2 Arrays.
- 20 Active Databases, 2,000 mailboxes per database.
- Nimble Storage arrays use double-parity to protect against drive failure.
- The Nimble Storage CASL file system uses a block checksum to ensure data integrity.
- 2 copies of each database.

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 deployment.

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://technet.microsoft.com/en-us/library/dd335215.aspx.

Targeted Customer Profile

The Nimble Storage solution for a mid to large-sized Enterprise Exchange environment is designed to support:

- 40,000 Mailboxes
- 8 Servers (4 Active, 4 Passive)
- 0.15 IOPS / Mailbox (0.189 IOPS / Mailbox achieved for 26% headroom)
- 1,000 MB per mailbox
- 24 x 7 background database maintenance enabled
- Providing controller capacity for replication and disk reconstruction
- Permitting space for backups and restores using Nimble Storage Snapshot Backup technology
- Mailbox Resiliency (2-Copy)

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)	2
Number of servers/DAG	4
Number of active mailboxes/server	10,000
Number of databases/host	5 (20 Total)
Number of copies/database	2
Number of mailboxes/database	2,000
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	0.15 IOPS / Mailbox (0.189 IOPS / Mailbox achieved for 26% headroom)
Database LUN size	2.08 TB
Log LUN size	250 GB
Total database size for performance testing	41.6 TB
% storage capacity used by Exchange database**	87%

**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.

Storage Hardware

Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI)	iSCSI SAN
Storage model and OS/firmware revision	CS460G-X2 Nimble OS 1.4
Storage cache	1 GB NVRAM Cache (Write), plus 48 GB RAM Cache, plus 4.8 TB Flash Cache (Read)
Number of storage controllers	2
Number of storage ports	(2) 10 Gbps Ethernet per Controller
Maximum bandwidth of storage connectivity to host	20 Gbps
Switch type/model/firmware revision	Extreme Summit X670-48x-FB 10 Gigabit Ethernet Switch
HBA model and firmware	Silicom PE210G2SPI9-XR v:1.4 Jumbo Frames enabled
Number of HBA's/host	2
Total number of disks tested in solution	12 SATA and 4 Flash Cache drives per array
Maximum number of spindles can be hosted in the storage	12 SATA and 4 Flash Cache drives per array

Storage Software

HBA driver	Intel 10 Gbps PCIe v. 2.1.4
HBA QueueTarget Setting	N/A
HBA QueueDepth Setting	512
Multi-Pathing	Nimble Operating System 1.4, Windows MPIO
Host OS	VMware ESXi 5.0
Guest OS	Windows 2008 R2 SP1
ESE.dll file version	14.0.639.19
Replication solution name/version	N/A for Exchange 2010 DAG Solutions

Storage Disk Configuration (Mailbox Store/Log Disks)

Disk type, speed and firmware revision	SATA 7,200 RPM, Flash Cache
Raw capacity per disk (GB)	SATA: 3 TB Flash Cache: 600 GB
Number of physical disks in test	SATA: 12 Flash Cache: 4
Total raw storage capacity (GB)	SATA: 36 TB Flash Cache: 2,400 GB Nimble native 50% compression SATA: 48 TB usable Flash Cache: 4.8 TB usable
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	12 Disks for both database and log files
Raid level	RAID 6 ***
Total formatted capacity	36 TB
Storage capacity utilization****	100% (Formatted Capacity / Raw Capacity) 75% (Formatted Capacity / Usable Capacity)
Database capacity utilization****	143% (Database & Log Size / Raw Capacity) 97.1% (Database & Log Size / Usable Capacity)

***Nimble's CASL file system uses a performance optimized RAID6 scheme that greatly reduces the rebuild time of traditional RAID6 implementations. CASL intelligently decides which drive sectors to rebuild (for example only populated data blocks are rebuilt and not empty drive space). As a result a lightly loaded array can rebuild a whole 1TB drive in about 4 hours, and a 2TB drive in about 7 hours. If the drive is not full the rebuild time would be proportionally shorter. So a half full drive 2TB drive would only take 3.5 hours to rebuild on a lightly loaded system. When the storage array is very busy, rebuild times can take longer to reduce performance latency for active applications.

**** Nimble Storage arrays use real-time compression algorithms that provide 50% capacity savings over raw disk sizes. Thus, usable capacity is 50% larger than the raw capacity of the storage and represents a more realistic expectation for sizing than raw capacity.

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 best practices on storage design, please visit http://technet.microsoft.com/en-us/library/dd346703.aspx

Nimble Storage provides an auto-tuning storage solution that leverages flash cache and greatly reduces the amount of effort to size and layout Exchange mailstores.

Core Storage

1. This solution architecture was designed to show the maximum Exchange performance capability of a Nimble Storage array. Environments smaller than 40,000 mailboxes routinely share Nimble arrays with other applications such as SQL Server and SharePoint.
2. Separate Exchange Database and Log volumes since they have different performance characteristics and will tune differently. Use the Nimble OS performance policies for Exchange Database and Exchange Logs when creating the respective volumes.
3. Nimble Storage arrays automatically provision storage from a pool of high-density disks protected by double-parity and are intelligently accelerated using NVRAM, DRAM and Flash SSD Cache. Nimble's CASL File System leverages these storage technologies natively to automatically provide scale-out performance dramatically higher than traditional tiered storage architectures. Therefore, you do not need to spend extensive time provisioning storage such that legacy storage architectures require, Nimble does this for you.
4. Please consult the "Nimble Storage Best Practices Guide for Exchange 2010" for current implementation best practices.

Backup strategy

Nimble Storage provides the first converged storage platform that combines primary storage and backup storage in the same architecture. Virtually all enterprise storage systems provide the ability to take snapshots, or point-in-time copies of data. Snapshots provide a very fast and efficient way to recover from user errors and application crashes. However, enterprises are usually limited to keeping at most a few days' worth of snapshots on disk, because storing multiple snapshots can consume large amounts of expensive and limited primary storage capacity. This mandates backing up primary data to a separate disk-based backup system frequently in order to store 60-90 days' worth of backups.

CASL enables instant, application-consistent backups on the same array with very efficient (up to 20x) backup capacity optimization. It captures compressed, incremental snapshots of the primary data on low-cost drives at pre-configured intervals. These backups are both extremely fast (on the order of milliseconds) and non-disruptive, because they do not impact application or storage performance. It therefore becomes feasible for enterprises to capture frequent backups spaced just minutes apart

(instead of the typical daily backup), enabling much finer recovery point objectives (RPOs) and improving data protection.

The Nimble approach significantly reduces costs by eliminating the need for a separate backup tier. It also eliminates the complexity and administrative overhead of managing massive data copies between tiers. Finally, it enables much faster restores, with recovery time objectives (RTOs) measured in seconds instead of hours in the event of application corruption or user error, further improving application availability.

Contact for additional information

Nimble Storage, Inc.
2740 Zanker Rd.
San Jose, CA 95134
Main: 408-432-9600
Support: 877-3NIMBLE (877-364-6253)
Fax: 408-899-5158
info@nimblestorage.com

Sales

To learn more about Nimble Storage's breakthrough technology or to schedule a briefing, please contact our technical sales team:

Sales: 877-3NIMBLE (877-364-6253)
sales@nimblestorage.com

Test Result Summary

This section provides a high level summary of the test data from ESRP and the link 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 to check Reliability tests runs for 24 hours. The goal is to verify the storage can handle high IO 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.

The following list provides an overview: (click on the underlined word will show the html report after the reliability tests run)

- There were no errors reported in any of the Event Logs collected for the reliability tests.
- 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 maximum sustainable Exchange type of IO for 2 hours. The test is to show how long it takes for the storage to respond to an IO 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 hours test duration. Each server is listed separately and the aggregate numbers across all servers is 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.

Host 1

Database I/O	
Database Disks Transfers/sec	1870.818
Database Disks Reads/sec	1,183.098
Database Disks Writes/sec	687.722
Average Database Disk Read Latency (ms)	9.1046
Average Database Disk Write Latency (ms)	7.391
Transaction Log I/O	
Log Disks Writes/sec	395.668
Average Log Disk Write Latency (ms)	2.4864

Host 2

Database I/O	
Database Disks Transfers/sec	1,980.273
Database Disks Reads/sec	1,250.345
Database Disks Writes/sec	729.929
Average Database Disk Read Latency (ms)	8.8402
Average Database Disk Write Latency (ms)	7.3298
Transaction Log I/O	
Log Disks Writes/sec	406.623
Average Log Disk Write Latency (ms)	2.4374

Host 3

Database I/O	
Database Disks Transfers/sec	1,773.79
Database Disks Reads/sec	1,124.194
Database Disks Writes/sec	649.596
Average Database Disk Read Latency (ms)	9.4152
Average Database Disk Write Latency (ms)	7.3904
Transaction Log I/O	
Log Disks Writes/sec	390.111
Average Log Disk Write Latency (ms)	2.4518

Host 4

Database I/O	
Database Disks Transfers/sec	1,942.416
Database Disks Reads/sec	1,227.406
Database Disks Writes/sec	715.011
Average Database Disk Read Latency (ms)	8.9576
Average Database Disk Write Latency (ms)	7.3338
Transaction Log I/O	
Log Disks Writes/sec	399.703
Average Log Disk Write Latency (ms)	2.4732

Aggregate Performance across all servers Metrics:

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

Database I/O	
Database Disks Transfers/sec	7,567.297
Database Disks Reads/sec	4,785.043
Database Disks Writes/sec	2,782.258
Average Database Disk Read Latency (ms)	9.0794
Average Database Disk Write Latency (ms)	7.36125
Transaction Log I/O	
Log Disks Writes/sec	1,592.105
Average Log Disk Write Latency (ms)	2.4622

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 in to 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.

Host 1

MB read/sec per database	73.55
MB read/sec total per server	367.75

Host 2

MB read/sec per database	73.578
MB read/sec total per server	367.89

Host 3

MB read/sec per database	70.138
MB read/sec total per server	350.69

Host 4

MB read/sec per database	73.904
MB read/sec total per server	369.52

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.

Host 1

Average time to play one Log file (sec)	1.116265

Host 2

Average time to play one Log file (sec)	1.19856

Host 3

Average time to play one Log file (sec)	1.207068

Host 4

Average time to play one Log file (sec)	1.227281

Conclusion

This report demonstrates Nimble Storage's ability to meet the performance needs of a 40,000 mailbox Exchange 2010 implementation at 0.15 IOPS and 1,024 MB per mailbox. Nimble's CASL file system provides true efficiency by combining high-density storage with flash SSD to enable groundbreaking performance within a small rack form factor. The Nimble storage architecture also provides additional space for highly-efficient redirect on write snapshots to provide fast full backups of Exchange using Microsoft VSS and Nimble Protection Manager.

This document is developed by Nimble Storage, 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. 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.

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

Appendix A - Test Results

This section provides representative test results from one of the four servers under test: Host 1 (ESRP). It is also representative of each server's test results, which are virtually identical.

24-Hour Stress Test Results

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1900.1	29.185	261946.778
Instance1900.2	29.243	261948.038
Instance1900.3	28.861	261962.350
Instance1900.4	29.320	261964.419
Instance1900.5	29.195	261947.527

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1900.1	1.841	232357.146
Instance1900.2	1.830	232510.462
Instance1900.3	1.830	232475.571
Instance1900.4	1.835	232357.612
Instance1900.5	1.830	232482.044

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
Instance1900.1	9.148	9.202	305.160	165.842	56825.489	34383.842	18.413	3.054	1.841	77.162	232357.146	5706.057
Instance1900.2	9.159	9.007	304.006	165.131	56983.522	34386.147	17.610	3.096	1.830	76.398	232510.462	5725.687
Instance1900.3	10.115	8.763	305.527	165.534	56452.729	34348.585	17.817	3.093	1.830	77.358	232475.571	5659.817
Instance1900.4	9.129	8.610	304.613	165.441	57060.000	34379.318	17.193	3.046	1.835	76.907	232357.612	5706.602
Instance1900.5	9.144	8.357	303.055	164.496	57017.263	34386.834	16.451	3.122	1.830	76.107	232482.044	5748.024

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	15.965	4.672	27.907
Available MBytes	13652.273	13627.000	13705.000
Free System Page Table Entries	33555606.059	33555554.000	33555612.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	65795745.942	65306624.000	66297856.000
Pool Paged Bytes	169144139.807	166268928.000	169517056.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

2-Hour Stress Test Results

Test Summary

Overall Test Result	**Pass**
Machine Name	ESRP
Test Description	
Test Start Time	7/9/2012 9:25:37 AM
Test End Time	7/9/2012 12:01:37 PM
Collection Start Time	7/9/2012 9:27:32 AM
Collection End Time	7/9/2012 11:27:24 AM
Jetstress Version	14.01.0225.017
ESE Version	14.00.0639.019
Operating System	Windows Server 2008 R2 Enterprise Service Pack 1 (6.1.7601.65536)
Performance Log	C:\Program Files\Exchange Jetstress\40,000 Mailboxes\Performance_2012_7_9_9_25_49.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	1870.818
Target Transactional I/O per Second	1500
Initial Database Size (bytes)	10737439539200
Final Database Size (bytes)	10743017963520
Database Files (Count)	5

Jetstress System Parameters

Thread Count	6 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance2832.1	Log path: C:\ESRP\log01 Database: C:\ESRP\db01\Jetstress001001.edb
Instance2832.2	Log path: C:\ESRP\log02 Database: C:\ESRP\db02\Jetstress002001.edb
Instance2832.3	Log path: C:\ESRP\log03

Database: C:\ESRP\db03\Jetstress003001.edb

Instance2832.4 Log path: C:\ESRP\log04
Database: C:\ESRP\db04\Jetstress004001.edb

Instance2832.5 Log path: C:\ESRP\log05
Database: C:\ESRP\db05\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
Instance2832.1	9.096	7.628	237.727	138.403	34338.317	36040.188	0.000	2.454	0.000	79.638	0.000	5366.380
Instance2832.2	9.058	7.479	233.681	135.737	34343.220	36014.046	0.000	2.502	0.000	78.614	0.000	5391.565
Instance2832.3	9.087	7.376	236.889	137.739	34441.047	36010.084	0.000	2.503	0.000	78.902	0.000	5354.966
Instance2832.4	9.153	7.260	237.645	137.938	34350.900	36021.805	0.000	2.455	0.000	79.437	0.000	5343.602
Instance2832.5	9.129	7.212	237.156	137.905	34294.256	36030.524	0.000	2.518	0.000	79.077	0.000	5408.994

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2832.1	30.365	261904.999
Instance2832.2	30.372	261980.119
Instance2832.3	30.376	261945.756
Instance2832.4	30.339	261975.970
Instance2832.5	30.358	261937.334

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2832.1	1.766	232539.777
Instance2832.2	1.751	232540.443
Instance2832.3	1.746	232565.237
Instance2832.4	1.753	232561.864
Instance2832.5	1.770	232550.598

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
Instance2832.1	9.096	7.628	268.093	138.403	60113.562	36040.188	10.221	2.454	1.766	79.638	232539.777	5366.380
Instance2832.2	9.058	7.479	264.052	135.737	60526.230	36014.046	10.232	2.502	1.751	78.614	232540.443	5391.565
Instance2832.3	9.087	7.376	267.265	137.739	60298.088	36010.084	8.961	2.503	1.746	78.902	232565.237	5354.966
Instance2832.4	9.153	7.260	267.983	137.938	60120.546	36021.805	8.934	2.455	1.753	79.437	232561.864	5343.602
Instance2832.5	9.129	7.212	267.514	137.905	60127.873	36030.524	9.318	2.518	1.770	79.077	232550.598	5408.994

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	13.463	6.464	19.159
Available MBytes	13761.086	13742.000	13819.000
Free System Page Table Entries	33555607.347	33555578.000	33555610.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	38016797.395	36417536.000	38662144.000
Pool Paged Bytes	90319801.453	90275840.000	90382336.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Database Backup Test Results

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance2832.1	2049060.09	07:54:23	71.99
Instance2832.2	2049052.09	07:56:21	71.69
Instance2832.3	2049052.09	07:36:14	74.85
Instance2832.4	2049068.09	07:38:40	74.46
Instance2832.5	2049068.09	07:36:46	74.76

Jetstress System Parameters

Thread Count	6 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance2832.1 Log path: C:\ESRP\log01
Database: C:\ESRP\db01\Jetstress001001.edb

Instance2832.2 Log path: C:\ESRP\log02
Database: C:\ESRP\db02\Jetstress002001.edb

Instance2832.3 Log path: C:\ESRP\log03
Database: C:\ESRP\db03\Jetstress003001.edb

Instance2832.4 Log path: C:\ESRP\log04
Database: C:\ESRP\db04\Jetstress004001.edb

Instance2832.5 Log path: C:\ESRP\log05
Database: C:\ESRP\db05\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
Instance2832.1	5.224	0.000	287.535	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2832.2	5.301	0.000	286.015	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2832.3	5.041	0.000	299.333	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2832.4	5.086	0.000	297.628	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2832.5	5.020	0.000	298.927	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	23.497	9.355	30.844
Available MBytes	15036.433	15030.000	15039.000
Free System Page Table Entries	33555584.315	33555500.000	33555612.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	50089333.636	50012160.000	50548736.000
Pool Paged Bytes	93679633.228	91615232.000	94687232.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Soft Recovery Test Results

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance1196.1	500	531.6801339
Instance1196.2	503	561.382586
Instance1196.3	515	582.8638238
Instance1196.4	512	556.8741781
Instance1196.5	500	591.3502387

Database Configuration

Instance1196.1 Log path: C:\ESRP\log01
Database: C:\ESRP\db01\Jetstress001001.edb

Instance1196.2 Log path: C:\ESRP\log02
Database: C:\ESRP\db02\Jetstress002001.edb

Instance1196.3 Log path: C:\ESRP\log03
Database: C:\ESRP\db03\Jetstress003001.edb

Instance1196.4 Log path: C:\ESRP\log04
Database: C:\ESRP\db04\Jetstress004001.edb

Instance1196.5 Log path: C:\ESRP\log05
Database: C:\ESRP\db05\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
Instance1196.1	20.647	11.040	905.016	5.597	36555.438	32768.000	5.420	0.000	8.385	0.000	232523.085	0.000
Instance1196.2	20.992	11.661	871.809	5.350	36487.379	32768.000	6.039	0.000	8.025	0.000	232414.143	0.000
Instance1196.3	19.335	10.588	878.775	5.301	36537.778	32768.000	6.190	0.000	7.941	0.000	232543.336	0.000
Instance1196.4	21.158	11.141	880.342	5.456	36616.163	32768.000	6.162	0.000	8.184	0.000	232479.492	0.000
Instance1196.5	19.810	11.335	867.089	5.044	36514.600	32768.000	6.361	0.000	7.556	0.000	232541.876	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1196.1	29.084	261954.338
Instance1196.2	29.175	261901.393
Instance1196.3	29.273	261988.699
Instance1196.4	29.119	262049.379
Instance1196.5	29.311	261953.987

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
Instance1196.1	20.647	11.040	934.100	5.597	43573.357	32768.000	5.420	0.000	8.385	0.000	232523.085	0.000
Instance1196.2	20.992	11.661	900.984	5.350	43786.666	32768.000	6.039	0.000	8.025	0.000	232414.143	0.000
Instance1196.3	19.335	10.588	908.048	5.301	43805.637	32768.000	6.190	0.000	7.941	0.000	232543.336	0.000
Instance1196.4	21.158	11.141	909.461	5.456	43834.115	32768.000	6.162	0.000	8.184	0.000	232479.492	0.000
Instance1196.5	19.810	11.335	896.400	5.044	43886.088	32768.000	6.361	0.000	7.556	0.000	232541.876	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	24.904	7.490	58.716
Available MBytes	13703.252	13685.000	14729.000
Free System Page Table Entries	33555382.020	33555050.000	33555612.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	53817622.639	50110464.000	54554624.000
Pool Paged Bytes	126198456.599	126017536.000	128983040.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Nimble Storage, Inc.
2740 Zanker Rd., San Jose, CA 95134
Tel: 408-432-9600; 877-364-6253 | www.nimblestorage.com | info@nimblestorage.com