Technical Report

Microsoft Exchange Server 2010 on NetApp Data ONTAP Operating in Cluster-Mode on FAS2240 Storage Systems Solution Guide

Wei Liu, NetApp
January 2012 | TR-4020

ABSTRACT

This report describes how to set up an entry-level configuration for Microsoft® Exchange Server 2010 mailbox server storage using NetApp® Data ONTAP® 8.1 architecture operating in Cluster-Mode on a FAS2240 storage system.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 EXECUTIVE SUMMARY

This NetApp solution guide provides guidelines for configuring Microsoft Exchange 2010 mailbox server storage on a FAS2240 storage system running NetApp Data ONTAP 8.1 operating in Cluster-Mode. The solution was tested and validated using Microsoft Exchange Server Jetstress 2010, and it meets all of the requirements outlined in the Microsoft Exchange Solution Reviewed Program (ESRP).

1.1 DOCUMENT PURPOSE

This document describes the design, implementation, and validation of a small configuration of Exchange 2010 on FAS2240 using Cluster-Mode. The document can be used for both presales and postsales engagements to provide examples of how the solution can be designed, implemented, tested, and validated.

Note: The configuration described in this document is not the only way the solution can be built.

1.2 OBJECTIVE

The objective of this document is to provide one of the three tested example solutions for Exchange 2010 on storage systems that use Cluster-Mode. Cluster-Mode enables nondisruptive operations to maintain availability of storage capacity during routine maintenance, upgrades, or even volume movement within a Vserver through Data ONTAP operating in Cluster-Mode. This document covers the small (up to 1,600 active users) entry-level configurations on a FAS2240 two-node cluster.

Note: This document does not provide information about how to scale up your Exchange 2010 environment, because it focuses only on small entry-level configurations. Such information will be covered in implementations of solutions on medium and large configurations in the future.

1.3 SCOPE

The scope of this document is limited to the performance and reliability aspects of the storage solution. Other advanced features of Cluster-Mode, such as online scale-up and scale-out, data migration, and consolidation, are not covered here, but will be described in separate technical reports or white papers in the future.

1.4 AUDIENCE

The target audience for this document includes NetApp field engineers, professional service engineers, sales engineers, and customers who plan to deploy Exchange 2010 mailbox servers on NetApp storage systems operating in Cluster-Mode.

1.5 PREREQUISITES

In addition to having relevant access to the storage system, readers of this document should have a basic understanding of the following:

- NetApp storage administrative skills, such as using the CLI
- Data ONTAP clustered configuration
- Microsoft Exchange 2010
- Jetstress 2010

For more information on these prerequisites, refer to the References section later in this document.

2 SOLUTION DESIGN

This section describes the design of an Exchange 2010 mailbox resiliency storage solution with 1,600 active mailboxes that uses a NetApp FAS2240 running Data ONTAP 8.1 operating in Cluster-Mode. This solution is designed for small to midsize enterprise Exchange environments. The target customer is typically one looking for a reliable, efficient, high-performance, and highly available Exchange storage solution.

2.1 TOPOLOGY

Figure 1 illustrates the topology of a solution that simulates a two-member Exchange 2010 database availability group (DAG). The simulated DAG configuration includes the following components:

- 2 mailbox servers (3650-1 and 3650-2)
- 800 mailboxes per database
- 2 active databases (DB1 and DB2 hosted on 3650-1)
- 2 passive databases (DB1 and DB2 hosted on 3650-2)

  Note: Passive databases are on a separate aggregate and on a separate but identical storage controller (FAS2240-2).

- 2 copies of the database

Figure 1) Topology of the Exchange 2010 storage solution on FAS2240.

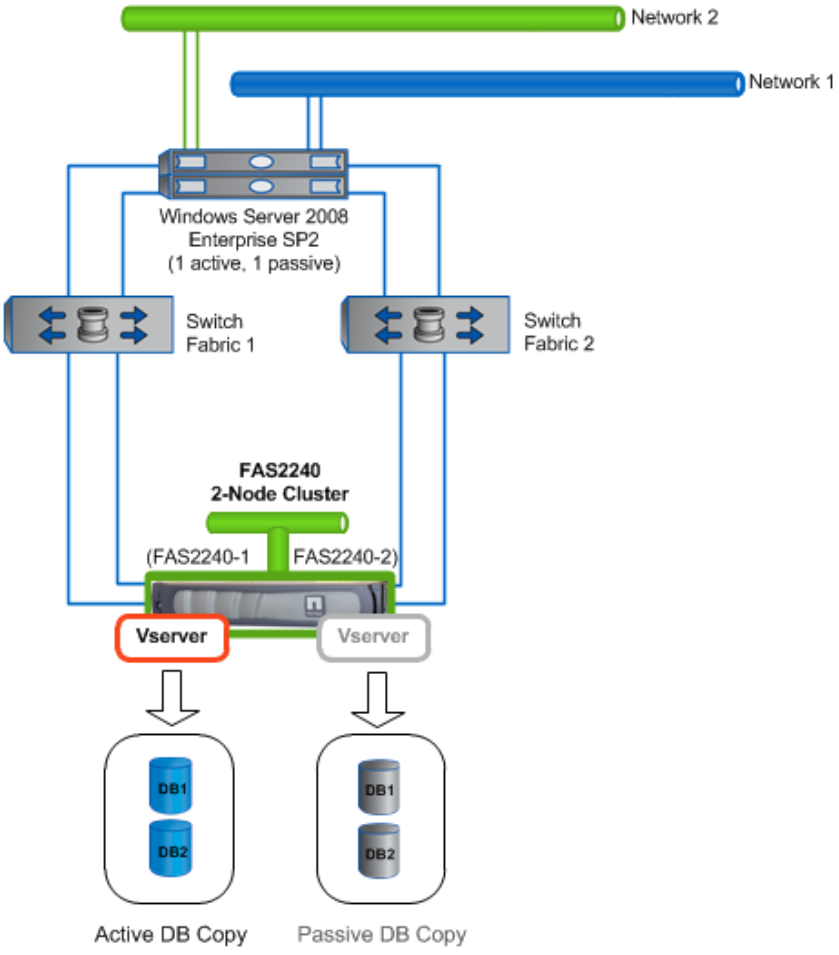

2.2 SERVER

The IBM x3650 server tested in this solution was equipped with 4 Intel® Xeon® processors and 18GB of RAM. Table 1 provides the server configuration details.

Note: Two identical servers are included in this solution, but only the active server was tested.

Table 1) IBM x3650 server configuration.

Component	Description
Processor	4 Intel Xeon CPU X5355s (2.66GHz)
Memory	18GB
Network interface card (NIC)	2 Broadcom BCM5708C NetXtreme II 1 Gigabit Ethernet (GbE) network adapters
Operating system	Microsoft Windows Server® 2008 Enterprise SP2
Jetstress 2010	Version: 14.01.0225.017
ESE.dll	Version: 14.01.0322.000
Multipathing software	Microsoft iSCSI Initiator and multipath input/output (MPIO)

JETSTRESS TARGETED CUSTOMER PROFILE

The solution tested is designed for:

- 1,600 mailboxes
- 2 servers
- 0.100 IOPS (0.120 tested for additional 20% head room)
- 2GB mailboxes
- Mailbox resiliency (2-copy)
- Background database maintenance 24/7

2.3 NETWORK SWITCHES

The network between the two FAS2240 nodes was established using two 10GbE links through a pair of Cisco Nexus 5020 switches.

ISCSI STORAGE AREA NETWORK

The data network between the IBM x3650 servers and the FAS2240s was established using dual redundant 1GbE links through a pair of Cisco 4948 switches. The iSCSI storage protocol used the Microsoft iSCSI Initiator and Microsoft MPIO software.

2.4 STORAGE

The FAS2240 is a member of NetApp's new generation FAS2000 series, which doubles the performance of prior-generation systems. FAS2240 is ideal for midsize businesses or distributed sites of larger organizations that require high performance or capacity needs and the latest storage technology. FAS2240 offers enterprise-class functionality at an affordable price. With its new and more powerful system manager interface, FAS2240 is easy to use and installs in minutes.

In a 2U form factor, FAS2240 comes with 24 internal SAS drives, 6GB of cache per controller, dual independent hot swappable controllers, and dual hot swappable power supplies. Additional expansion options include FC, SAS, or SATA shelves, which are connected through SAS or Fibre Channel interfaces that are built into the two controllers, for a total maximum capacity of 147 drives and total raw capacity of 374TB.

The FAS2240 tested used standard NetApp features such as NetApp Snapshot™ technology, NetApp RAID-DP® technology, NetApp FlexVol® flexible volumes, and NetApp SnapDrive® software. The tested solution included disk space allocated for backing up and restoring Exchange databases and logs using Snapshot technology, which provides in-place logical backups of an entire Exchange environment, usually in seconds, regardless of whether the data backed up is several gigabytes or hundreds of terabytes. Additionally, NetApp Snapshot technology causes virtually no loss of performance following Snapshot copy creation and use going forward. The performance results listed in this report were produced with one Snapshot copy in place on the tested storage systems.

RAID-DP is a high-performance implementation of RAID 6 that provides double parity across the disk subsystem and protects against the failure of up to two disks per RAID group. Calculations have shown double-parity RAID offers over 160 times the protection against data loss than RAID 10, and almost 4,000 times the protection against data loss than RAID 5. The performance results listed in this report were produced with RAID-DP enabled.

FlexVol volumes enable the creation of logical storage volumes across a large pool of disk drives known as an aggregate. FlexVol volumes provide three core advantages for workloads such as Microsoft Exchange:

- Using FlexVol volumes, the capacity and performance bandwidth of a large collection of fast drives can be made available to all volumes. Even very small FlexVol volumes have the performance benefit of a large number of drives.
- A large number of volumes can be created, all with independent Snapshot copy schedules, mirroring events, and so on.
- All FlexVol volumes can be managed independently while receiving the maximum I/O benefit of a much larger pool of drives.

SnapDrive for Microsoft Windows® offers a rich set of capabilities that virtualizes and enhances storage management for Windows environments. SnapDrive is tightly integrated with NTFS to provide a layer of abstraction between application data and physical storage associated with that data.

Table 2 provides details of the FAS2240 configuration used in the solution.

Table 2) NetApp FAS2240 configuration in this solution.

Component	Description
FAS2240 controller	2 controllers
Data ONTAP	Version 8.1.1 operating in Cluster-Mode
Controller cache	6GB per controller
Number of network ports	2 10GbE and 5 1GbE ports per controller
Disk drives	16 1TB 7.2K RPM SATA drives
Vserver	VVESPA4 (for Jetstress testing)
Aggregate	AGGR5 (for Jetstress testing)
FlexVol volumes	2 for databases, 2 for logs
Logical unit numbers (LUNs)	2 for databases, 2 for logs

Table 3 provides storage layout information, including aggregate, FlexVol volume, and LUN configuration details.

Table 3) Storage layout with aggregate, volume, and LUN configuration details.

Controller Name (Vserver)	Aggregate	RAID Type	Number of Data + Parity Drives	FlexVol Volume	LUN Name	LUN Size	LUN Mapping
FAS2240-1 (vvespa4)	aggr5	RAID-DP	14 + 2	db_fv	DB1	2000GB	E:\
				db2_fv	DB2	2000GB	F:\
				log_fv	LOG1	20GB	G:\
				log2_fv	LOG2	20GB	H:\

The FAS2240-1 storage controller was configured using RAID-DP with one aggregate for databases and log files. Four FlexVol volumes—two for the two database LUNs and the remaining two for two log LUNs—were created for the active server. Only the active (or primary) copy was tested for this guide. However, the passive (or secondary) copy should be configured the same way as the active copy, from the host to the spindle, including brand, model, firmware, and drivers.

3 SOLUTION IMPLEMENTATION

This section discusses storage solution implementation details.

3.1 ADD LICENSES

The following licenses are required prior to configuring storage on FAS2240 controllers using Cluster-Mode:

- Base
- NFS
- iSCSI

Use the CLI to verify that the required licenses are present on the system. If the licenses are not present on the system, you must add them.

```
 vespa-exch::*> license show
  (system license show)
Feature          Cluster SN   Limit     Description
---------------  -----------  -------  --------------
Base             1-80-000011  666       Base License w/cluster size limit (nodes)
NFS              1-80-000011  666       NFS License
iSCSI            1-80-000011  666       iSCSI License
3 entries were displayed.
```

3.2 ESTABLISH THE TWO-NODE CLUSTER

Use Cluster Wizard to create a cluster on the first node, and then join the second node to the cluster. For more information, refer to the Data ONTAP Cluster-Mode Software Setup Guide.

After you create the two-node cluster, enter the following command to verify that the cluster was properly established:

```
vespa-exch::*> cluster show
Node                  Health  Eligibility   Epsilon
--------------------- ------- ------------  ------------
vespa-exch-01         true    true          true
vespa-exch-02         true    true          false
2 entries were displayed.
```

3.3 OPTIMIZE STORAGE LAYOUT

Enter the following command to optimize the storage layout. For more details, refer to TR-3647: Configuring and Tuning NetApp Storage Systems for High-Performance Random-Access Workloads:

```
vespa-exch::*> node run -node vespa-exch-01 options wafl.optimize_write_once off

// verify the setting as below:
vespa-exch::*> node run -node vespa-exch-01 options wafl.optimize_write_once

wafl.optimize_write_once     off        (value might be overwritten in takeover)
```

Note: The `setflag` command described in TR-3647: "Configuring and Tuning NetApp Storage Systems for High-Performance Random-Access Workloads" has become `options` in Data ONTAP 8.1 RC1.

3.4 CREATE THE AGGREGATE

Enter the following command to create a 16-disk, 32-bit aggregate on the first node, `vespa-exch-01`:

```
vespa-exch::*> aggr create -aggregate aggr5 -node vespa-exch-01 -diskcount 16 -t
raid_dp -s 16 -B 32-bit -volume-style flex
```

3.5 CREATE THE VSERVER

Use the following command to create a Vserver using the aggregate that you created in section 3.4:

```
vespa-exch::*> vserver create -vserver vvespa4 -rootvolume root_vol -rootvolume-
security-style ntfs -aggregate aggr5 -ns-switch file -language en_US
```

Enter the following command to view the details of the Vserver:

```
vespa-exch::*> vserver show vvespa4

                                    Vserver: vvespa4
                                Vserver Type: cluster
                                Vserver UUID: 4cbc03c9-80cc-11de-b01e-123478563412
                                 Root Volume: root_vol
                                   Aggregate: aggr5
                         Name Service Switch: file
                         Name Mapping Switch: file
                                    NIS Domain: -
                  Root Volume Security Style: ntfs
                                 LDAP Client: -
                                    Language: en_US
                               Snapshot Policy: default
                                     Comment:
                 Anti-Virus On-Access Policy: default
                                  Quota Policy: default
                 List of Aggregates Assigned: -
 Limit on Maximum Number of Volumes allowed: unlimited
                       Vserver Admin State: running
```

```
                        Allowed Protocols: nfs, cifs, fcp, iscsi
                     Disallowed Protocols: -
            Protocol Services use Data LIFs: true
```

Note: Vserver UUID is used to construct a unique target name for the iSCSI service. Make sure `Allowed Protocols` includes `iscsi`. If `iscsi` is not listed in `Allowed Protocols`, use `vserver modify…` to make the necessary change.

3.6 CREATE ISCSI SERVICE

Enter the following command to create the iSCSI service associated with the Vserver:

```
vespa-exch::vserver*> iscsi create

// to verify
vespa-exch::vserver*> iscsi show
           Target                           Target                         Status
Vserver    Name                             Alias                          Admin
---------- -------------------------------- ------------------------------ ------
vvespa4    iqn.1992-08.com.netapp:sn.4cbc03c980cc11deb01e123478563412:vs.6
                                            vvespa4                        up
```

3.7 CREATE LOGICAL INTERFACES

Enter the following command to create two logical interfaces (LIFs) for the iSCSI service on the Vserver:

```
vespa-exch::network interface*> create -vserver vvespa4 -lif lif1 -role data -data-
protocol iscsi -home-node vespa-exch-01 -home-port e0c -address 192.168.101.20 -
netmask 255.255.255.0

vespa-exch::network interface*> create -vserver vvespa4 -lif lif2 -role data -data-
protocol iscsi -home-node vespa-exch-01 -home-port e1b -address 192.168.101.21 -
netmask 255.255.255.0

// to verify
vespa-exch::network interface*> show
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
vespa-exch
            cluster_mgmt up/up    192.168.20.222/24   vespa-exch-02 e1a     false
vespa-exch-01
            clus1        up/up    192.168.10.10/24   vespa-exch-01 e1a     true
            clus2        up/up    192.168.10.11/24   vespa-exch-01 e1b     true
            mgmt1        up/up    192.168.20.220/24   vespa-exch-01 e0M     true
vespa-exch-02
            clus1        up/up    192.168.10.12/24   vespa-exch-02 e1a     true
            clus2        up/up    192.168.10.13/24   vespa-exch-02 e1b     true
            mgmt1        up/up    192.168.20.221/24   vespa-exch-02 e0M     true
vvespa4
            lif1         up/up    192.168.101.20/24   vespa-exch-01 e0c     true
            lif2         up/up    192.168.101.21/24   vespa-exch-01 e0d     true
9 entries were displayed.
```

Note: The IP addresses shown are examples only. You must use the actual IP addresses that are assigned to ports e0c and e0d.

3.8 CREATE AN IGROUP

To create an iGroup, you must know the iSCSI initiator name. Obtain the iSCSI initiator name from the Microsoft iSCSI Initiator tool. Typically, the format of the iSCSI initiator name is `iqn.1991-05.com.microsoft:<server name>`. In this case, the iSCSI initiator name is `iqn.1991-05.com.microsoft:3650-1`.

```
vespa-exch::lun igroup*> igroup create -vserver vvespa4 -igroup igrp_3650-1 -protocol
iscsi -t windows -initiator  iqn.1991-05.com.microsoft:3650-1

// to verify
vespa-exch::lun igroup*> show
Vserver   Igroup       Protocol OS Type  Initiators
--------- ------------ -------- -------- ------------------------------------
vvespa4   igrp_3650-1  iscsi    windows  iqn.1991-05.com.microsoft:3650-1
```

3.9 CREATE FLEXVOL VOLUMES

Four FlexVol volumes are required for the Jetstress test. Enter the following commands to create the volumes:

```
vespa-exch::*> vol create -vserver vvespa4 -volume db_fv -aggregate aggr5 -size 3050GB
-type RW -max-write-alloc-blocks 512

vespa-exch::*> vol create -vserver vvespa4 -volume db2_fv -aggregate aggr5 -size
3050GB -type RW -max-write-alloc-blocks 512

vespa-exch::*> vol create -vserver vvespa4 -volume log_fv -aggregate aggr5 -size 20GB
-type RW

vespa-exch::*> vol create -vserver vvespa4 -volume log2_fv -aggregate aggr5 -size 20GB
-type RW
```

Note: For db_fv and db2_fv, 1000GB of 3050GB is reserved for Snapshot copies.

3.10 IMPROVE READ PERFORMANCE

Enter the following command on database volumes db_fv and db2_fv to improve read performance:

```
vespa-exch::*> vol modify -vserver vvespa4 -volume db_fv -read-realloc on
vespa-exch::*> vol modify -vserver vvespa4 -volume db2_fv -read-realloc on
```

For more information on the `read-realloc` command and improving read performance, refer to the Data ONTAP 8.1 Cluster-Mode System Administration Guide.

3.11 CREATE THE LUNS

Four LUNs are required for the Jetstress tests. Create the LUNs by entering the following commands:

```
vespa-exch::lun*> lun create -vserver vvespa4 -volume db_fv  -lun  db1_fv.lun -s 1.9TB
-t windows_2008 -space-reserve disable

vespa-exch::lun*> lun create -vserver vvespa4 -volume db2_fv  -lun  db2_fv.lun -s
1.9TB -t windows_2008 -space-reserve disable

vespa-exch::lun*> lun create -vserver vvespa4  -volume log_fv -lun log1_fv.lun -s 19GB
-t windows_2008 -space-reserve disable
```

```
vespa-exch::lun*> lun create -vserver vvespa4  -volume log2_fv -lun log2_fv.lun -s
19GB -t windows_2008 -space-reserve disable
```

3.12 MAP THE LUNS

Map the four LUNs to the server by entering the following commands:

```
vespa-exch::lun*> lun map -vserver vvespa4 -volume db_fv -lun db1_fv.lun -igroup
igrp_3650-1

vespa-exch::lun*> lun map -vserver vvespa4 -volume db2_fv -lun db2_fv.lun -igroup
igrp_3650-1

vespa-exch::lun*> lun map -vserver vvespa4 -volume log_fv -lun log1_fv.lun -igroup
igrp_3650-1

vespa-exch::lun*> lun map -vserver vvespa4 -volume log2_fv -lun log2_fv.lun -igroup
igrp_3650-1

// to verify
vespa-exch::lun*> lun show -m
Vserver    Volume        Qtree         LUN           Igroup   LUN-ID  Protocol
---------  ------------  ------------  ------------  -------  ------  --------
vvespa4    db2_fv        ""            db2_fv.lun    igrp_3650-1   1  iscsi
vvespa4    db_fv         ""            db1_fv.lun    igrp_3650-1   0  iscsi
vvespa4    log2_fv       ""            log2_fv.lun   igrp_3650-1   3  iscsi
vvespa4    log_fv        ""            log1_fv.lun   igrp_3650-1   2  iscsi
4 entries were displayed.
```

Now the LUNs are presented to the server. The host-side configuration for Jetstress is virtually the same for both Cluster-Mode and 7-Mode and is not described in this document.

4 SOLUTION VALIDATION

ESRP methodology was followed to test and validate the Exchange 2010 storage solution on a FAS2240 running Data ONTAP operating in Cluster-Mode.

4.1 PERFORMANCE TEST

The performance test validates the performance of the Exchange storage solution by running the Jetstress test for two hours. The overall test result should be a Pass, as shown in Figure 2. The 1,600-mailbox Jetstress performance test completed successfully.

Figure 2) Jetstress performance test results example.

Microsoft Exchange Jetstress 2010

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	3650-1
Test Description	
Test Start Time	10/29/2011 7:10:45 AM
Test End Time	10/29/2011 9:15:19 AM
Collection Start Time	10/29/2011 7:15:17 AM
Collection End Time	10/29/2011 9:15:10 AM
Jetstress Version	14.01.0225.017
ESE Version	14.01.0322.000
Operating System	Windows Server (R) 2008 Enterprise Service Pack 2 (6.0.6002.131072)
Performance Log	C:\Program Files\Exchange Jetstress\results\Performance_2011_10_29_7_10_50.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	249.024
Target Transactional I/O per Second	192
Initial Database Size (bytes)	3435997036544
Final Database Size (bytes)	3436793954304
Database Files (Count)	2

Table 4 shows the database transactional input/output (I/O) performance results, including latencies and input/output operations per second (IOPS).

Table 4) Jetstress performance test results.

Metrics	Results
Database disk transfers/sec (IOPS)	249.0
Database disk reads/sec	157.4
Database disks writes/sec	91.6
Average database disk read latency (ms)	17.20
Average database disk write latency (ms)	1.02

4.2 STRESS TEST

The stress test validates the reliability of the Exchange storage solution by running the Jetstress test for 24 hours; the overall test result should be a Pass. The 1,600-mailbox Jetstress stress test passed successfully. Table 5 shows the database transactional I/O performance results, including latencies and IOPS.

Table 5) Jetstress stress test results.

Metrics	Results
Database disk transfers/sec (IOPS)	217.3
Database disk reads/sec	136.4
Database disks writes/sec	80.9
Average database disk read latency (ms)	19.65
Average database disk write latency (ms)	0.96

4.3 DATABASE BACKUP TEST

The database backup test measures the sequential read performance of the storage system. Table 6 shows the results of the database backup test.

Table 6) Database backup test results.

Metrics	Results
MB read per second per database	52.44
MB read per second total per server	104.88

4.4 SOFT RECOVERY TEST

The soft recovery test measures log replay performance. Table 7 shows the results of the soft recovery test.

Table 7) Soft recovery test results.

Metrics	Results
Average time to play one log file (sec)	2.20

5 RECOMMENDATIONS

This section provides some recommendations based on experiences gained from this project as well as past knowledge of building Exchange storage solutions in 7-Mode.

- Turn `wafl.write_once` off before aggregate creation.
- Turn `read_realloc` on for database volumes.
- Use 32-bit aggregates.

  Note: This recommendation is specific to FAS2240 only.
- Use the same aggregate for both database and transaction logs.
- Separate database volumes from transaction log volumes.
- Separate database LUNs from log LUNs, one LUN per volume.
- Use fewer large databases rather than many small databases.
- Use GPT disk type to get correct disk sector alignment when formatting disks on the host.
- Turn on background database maintenance 24/7.

- Tune the thread count per database manually rather than using autotune.
- Place copies of the same database in separate aggregates.

6 CONCLUSION

This document describes the design, implementation, and validation of an Exchange 2010 storage solution (small configuration) using Data ONTAP operating in Cluster-Mode on a FAS2240 storage system. It demonstrates how the solution can be designed and implemented, and how to validate the solution using standard ESRP methodology.

7 REFERENCES

This document uses the following references.

7.1 MICROSOFT REFERENCES

- Exchange 2010 Solution Reviewed Program (ESRP)—Storage v3
  http://technet.microsoft.com/en-us/exchange/ff182054
- Microsoft Exchange Server Jetstress 2010
  http://technet.microsoft.com/en-us/library/ff706601.aspx

7.2 NETAPP REFERENCES

- Configuring and Tuning NetApp Storage Systems for High-Performance Random-Access Workloads
  http://media.netapp.com/documents/tr-3647.pdf
- Data ONTAP Cluster-Mode Software Setup Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel81crc2/pdfs/ontap/csetup.pdf
- Data ONTAP 8.1 Cluster-Mode System Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel81crc2/pdfs/ontap/csysadmin.pdf
- Fibre Channel and iSCSI Configuration Guide for the Data ONTAP 8.1 Cluster-Mode Release Family
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel81crc2/pdfs/ontap/fc_config.pdf

8 ACKNOWLEDGEMENTS

- Kamal Charan, Tony Gaddis, Tomislav Grcanac, Phil Larson, and Manpreet Singh, Engineering
- Lee Gates, John Parker, and Robert Quimbey, MSBU
- Lisa Fail and A.J. Mahajan, Product Management
- Rick Jooss, SANiSAN
- Keith Griffin and Chris Lemmons, Workload Engineering

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com