Technical Report

200,000 Exchange Server 2013 Mailboxes on NetApp FAS8060

An Overview of Performance and Scalability

Wei Liu, NetApp
February 2014 | TR-4268

Abstract

This tech report describes the testing of 200,000 Exchange Server 2013 mailboxes in a virtualized environment on a 6-node NetApp® FAS8060 storage cluster running the clustered Data ONTAP® operating system. The workload was generated by using Jetstress 2013.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

This report describes the testing of 200,000 Exchange Server 2013 mailboxes in a virtualized environment by using a configuration consisting of 2 database copies and 3 database availability groups (DAGs) on 6 Fujitsu RX300S7 servers and a NetApp FAS8060 storage system running the clustered Data ONTAP operating system. The report also presents the test results and analysis.

1.1 NetApp FAS8060

Powered by Data ONTAP and optimized for scaling out, the FAS8000 series unifies your storage infrastructure and has the flexibility to keep up with changing business needs while delivering on core IT requirements for uptime, scalability, and cost efficiency.

The FAS8000 features a multiprocessor Intel® chip set and leverages high-performance memory modules, NVRAM to accelerate and optimize writes, and an I/O-tuned PCIe Gen3 architecture that maximizes application throughput. Building on a decade of multicore optimization, Data ONTAP drives the latest cores and increased core counts to keep up with continuous growth in storage demands. The result is a flexible, efficient I/O design capable of supporting large numbers of high-speed network connections and massive capacity scaling.

By delivering a wide range of on-board ports to support drive, cluster, and host connectivity, the FAS8000 offers exceptional flexibility and expandability in an extremely dense package. In addition to providing 1GbE, 10GbE, and 6Gb SAS on-board SAS ports, the FAS8060 also offers integrated unified target adapter (UTA) ports support for 16Gb Fibre Channel, 10GbE, and FCoE, so your storage is ready on day one for whatever choices you make in the future.

Customers can easily expand data management capabilities as their needs grow by adding optional software features that provide enterprise-proven capabilities for automated provisioning and restoration; simple, efficient disaster recovery; automated application; and virtual-machine-aware backup, recovery, and cloning.

1.2 Jetstress 2013

Jetstress 2013 was released in March 2013. Like its predecessor, Jetstress 2013 is a tool provided by Microsoft for evaluating an Exchange 2013 storage subsystem's performance, scalability, and reliability in a lab environment. Jetstress 2013 works with the Microsoft® Exchange Server 2013 database engine and simulates the Exchange Server's I/O workload. Jetstress is typically run prior to real-world deployment to validate the theoretical design of the Exchange storage solutions. In this project, the Jetstress 2013 tool was used to generate the 200,000-mailbox workload.

2 Test Environment

The following sections describe the test topology, the server and storage configurations, and the targeted customer profile.

2.1 Test Topology

Figure 1 depicts the high-level topology of the 200,000-mailbox virtualized test environment. At the physical level, the testbed is composed of 6 Fujitsu RX300S7 servers connected to a 6-node FAS8060 storage cluster through a pair of Brocade 6510 Fibre Channel (FC) switches and dual redundant 8Gbps FC paths. Windows Server® 2012 with Hyper-V® was used to create the virtualized environment in which 48 virtualized Exchange 2013 mailbox servers were organized in 3 DAGs. Each DAG was built on a pair of the Fujitsu servers and a pair of the FAS8060 controllers or nodes. And each DAG has 16 mailbox servers (8 active and 8 passive).

Figure 1) High-level topology and DAG architecture of the 200,000-mailbox test environment.

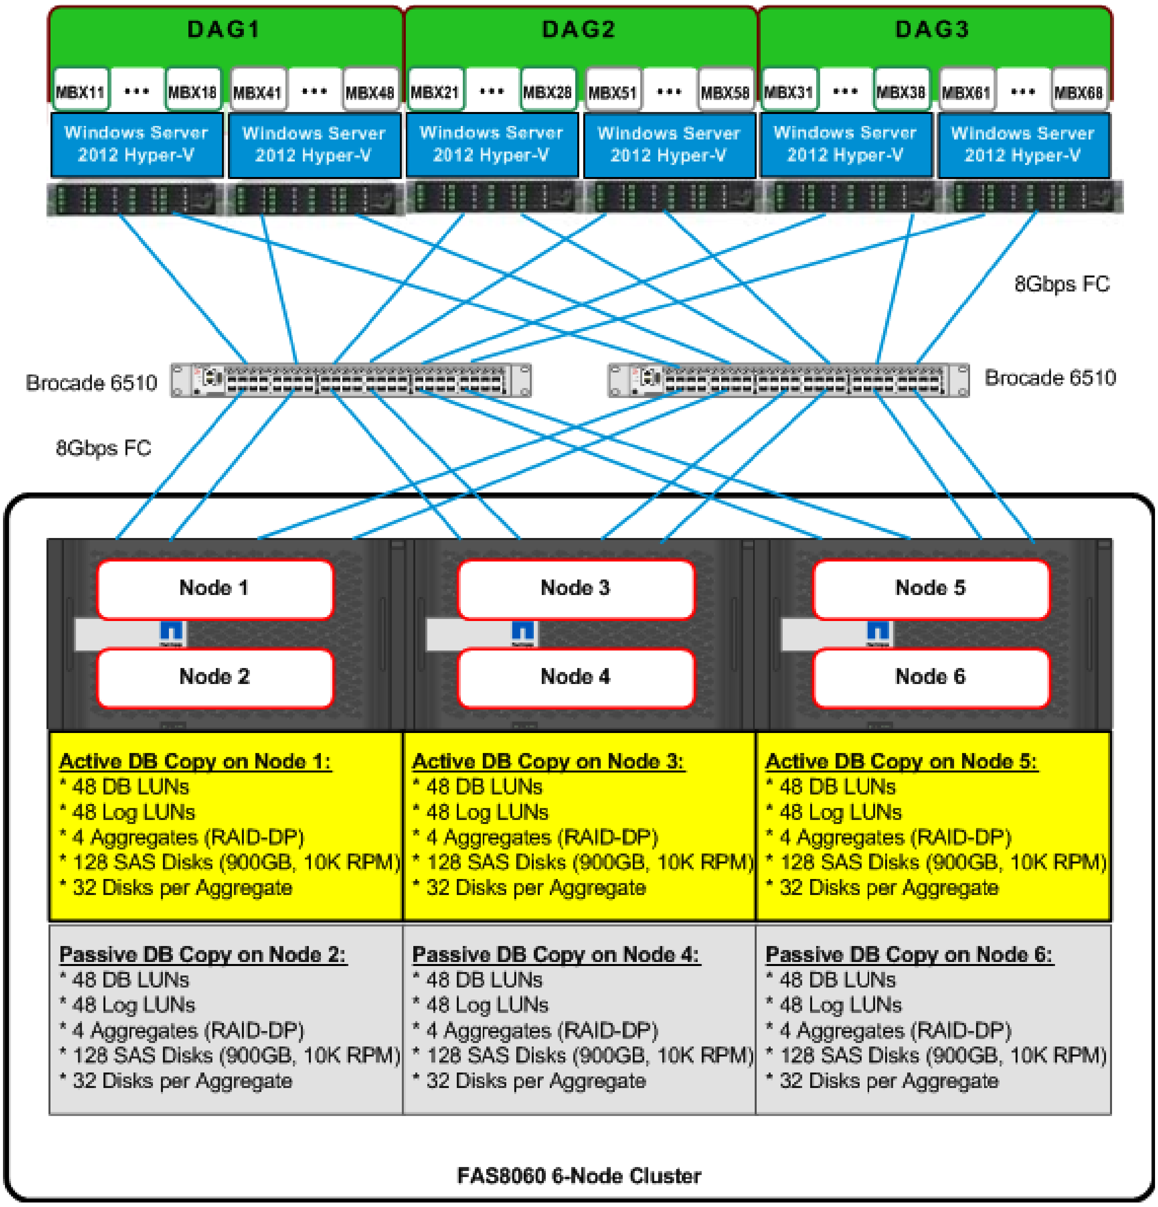

2.2 Server and Storage Configurations

Table 1 shows the Fujitsu RX300S7 server configuration details, together with information about virtualized Exchange 2013 mailbox servers.

Table 1) Server configuration details.

Entity	Description
Processor	12-core Intel Xeon® CPU E5-2630 @2.30 GHz
RAM	128GB
HBA model and firmware	QLogic QLE2562, v2.2.3
HBA driver	STOR miniport v9.1.11.20
Multipathing software	Data ONTAP DSM 4.0.1 for Windows® MPIO
Hardware virtualization software	Windows Server 2012 with Hyper-V

Entity	Description
Exchange virtual machine guest OS	Windows Server 2012 Datacenter (6.2.9200.0)
Number of virtual machines per host	8
Jetstress 2013 version	15.00.0658.004
ESE.dll version	15.00.0516.026

Table 2 shows the FAS8060 node configuration details.

Table 2) FAS8060 node configuration details.

Entity	Description
Data ONTAP version	Clustered Data ONTAP 8.2.1
Storage cache	576GB (64GB, plus 512GB Flash Cache™ intelligent caching)
Number of storage ports tested	2
Disk type	SAS 900GB 10K RPM
Number of disks tested	128
Number of aggregates	4
Number of storage virtual machines	4
Number and size of LUNs	96 (48x 1.5TB DB LUNs and 48x 20GB log LUNs)

The architecture of the test environment was designed to take advantage of the scale-out capability of clustered Data ONTAP. For instance, each pair of FAS8060 nodes (together with a pair of the Fujitsu servers) form a building block on which a DAG can be created. To scale a solution, you simply add more nodes and create more DAGs.

Clustered Data ONTAP 8.2.1 allows two or more controllers (or nodes) to operate as one shared resource pool or storage cluster. The storage cluster can be expanded, contracted, and subdivided nondisruptively into one or more secure partitions, or storage virtual machines (SVMs). An SVM is a logical storage container that includes allocated storage resources from within the cluster as well as security parameters such as rights and permissions.

For example, the SVMs (VS11, VS12, and so on) in Figure 2 behave like virtual storage controllers, providing hosts with data access to the storage resources inside the cluster. Figure 2 shows the configuration details of DAG1 in the test environment. For example, eight active mailbox servers (MBX11–MBX18) were configured on server1 and four active SVMs (VS11–VS14) were created on node-1 for servicing MBX11 to MBX18. DAG1 has two Exchange database copies. The primary copy (Active DB Copy) is on node-1. The secondary copy (Passive DB Copy) is on node-2. The primary copy storage and the secondary copy storage are configured identically (from the host on down to the spindle, including brand, model, firmware, drivers, and so on).

In addition, Figure 2 shows that node-1 was configured with four aggregates (aggr11–aggr14). Each aggregate contains 12 database LUNs (DB11–DB16, DB21–DB26) and 12 transactional log LUNs (not shown). Each aggregate provides storage, while its corresponding SVM provides data access for two of the eight active mailbox servers. Therefore, node-1 has 48 DB LUNs and 48 log LUNs for DAG1's primary copy storage. The DAG1 configuration is as follows:

- 8 active mailbox servers: MBX11–MBX18
- 1,390 mailboxes per database
- 8,340 mailboxes per server (6 databases per server)
- A total of 48 active databases for 66,720 active mailboxes:
  - DB11, DB12, …, and DB16 hosted on MBX11

- DB21, DB22, …, and DB26 hosted on MBX12
- DB31, DB32, …, and DB36 hosted on MBX13
- DB41, DB42, …, and DB46 hosted on MBX14
- DB51, DB52, …, and DB56 hosted on MBX15
- DB61, DB62, …, and DB66 hosted on MBX16
- DB71, DB72, …, and DB76 hosted on MBX17
- DB81, DB82, …, and DB86 hosted on MBX18

- A total of 48 passive databases: hosted by 8 additional servers (MBX41–MBX48); the passive database copy is on separate aggregates and on the separate but identical storage controller: node-2
- 2 copies of databases

Figure 2) DAG1 configuration details.

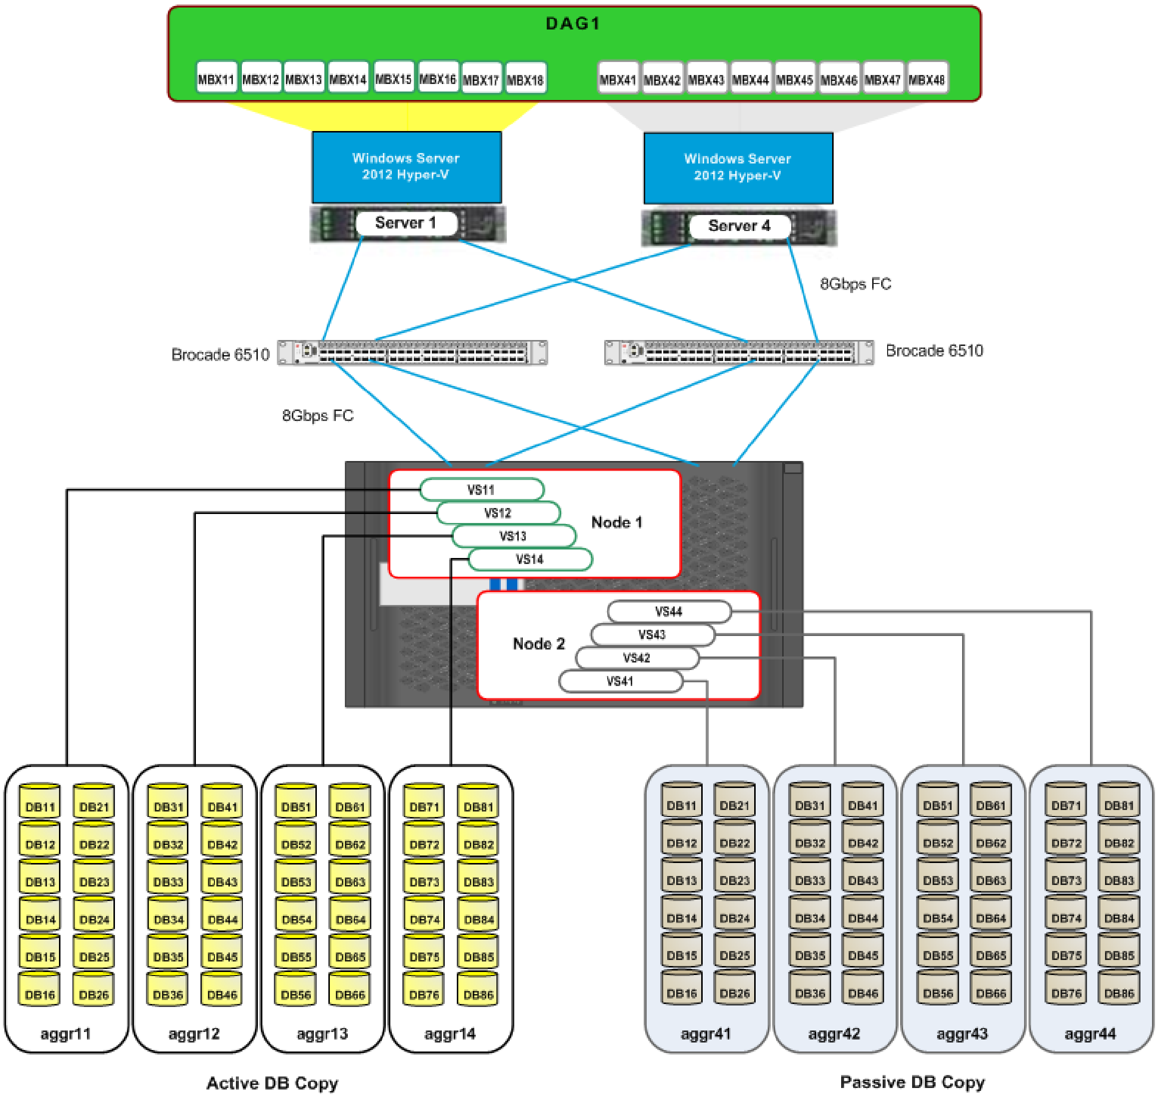

DAG2 and DAG3 were configured using the same method as DAG1. Therefore, their configuration details are omitted for brevity. The entire 200,000-mailbox test environment consists of the following:

- 3 DAGs
- 24 active mailbox servers
- 1,390 mailboxes per database
- 8,340 mailboxes per server (6 databases per server)

- A total of 144 active databases for 200,160 active mailboxes
- A total of 144 passive databases, hosted by 24 additional servers; the passive databases are on separate aggregates and on the separate but identical storage controllers
- 2 copies of databases

2.3 Targeted Customer Profile

The Jetstress 2013 targeted customer profile is as follows:

- 200,160 mailboxes
- 48 Exchange servers (24 tested)
- 0.05 IOPS (0.06 IOPS tested for additional 20% headroom)
- 1.0GB per mailbox
- Maintaining ample controller headroom available for storage infrastructure activities such as data replication and disk reconstructs
- 24/7 background database maintenance
- Mailbox resiliency (2-copy)

3 Results and Analysis

This section shows the Jetstress 2013 performance test results.

3.1 Performance

The Jetstress 2013 performance test was run for two hours to validate that the performance of the storage subsystem is adequate for Exchange Server 2013's mailbox server role.

Figure 3 summarizes the test results including the actual achieved IOPS (bars), the target IOPS (blue line), the average database read latency (green line with boxes), and the average database write latency (purple line with dots). The target IOPS of 0.06 per mailbox (0.05 plus 20% headroom) resulted in 500.4 IOPS per virtual Exchange Server. The achieved IOPS exceeded the target IOPS substantially, by 36% on average. This means that the storage subsystem can easily handle the I/O load of 200,000 mailboxes, with ample headroom for potential I/O expansion. The average database read latency was 10.66ms, below the required 20ms threshold. The average database write latency was 1.69ms, far below the required 20ms threshold.

Note: For latencies, lower is better.

Figure 3) Jetstress 2013 performance test results.

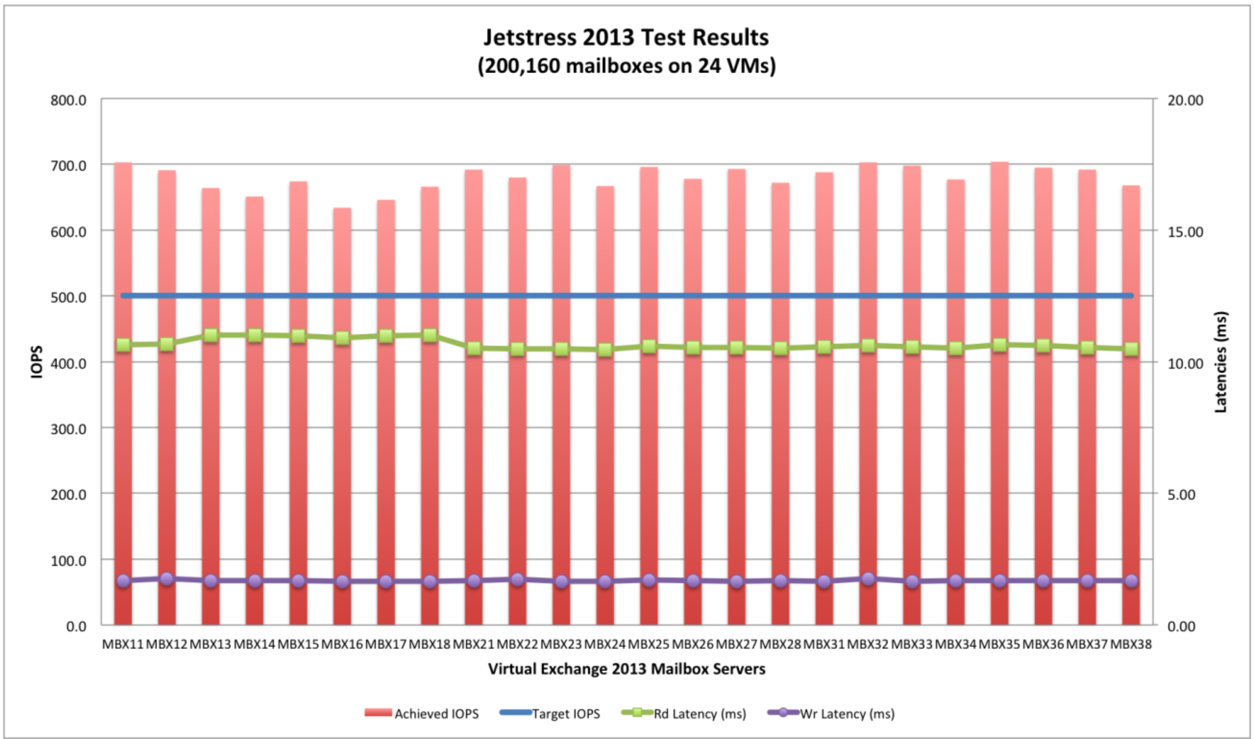

3.2 Scalability

One of the benefits of clustered Data ONTAP is scale-out; that is, by scaling the number of nodes, one can scale the workload handled by the storage subsystem.

Figure 4 highlights the scale-out capability. It shows that by increasing the number of nodes inside the storage cluster from 2 to 4 and 6, the number of mailboxes supported also increased linearly, from 66,720 to 133,440 and 200,160.

Figure 4) The scaling of Exchange 2013 mailboxes.

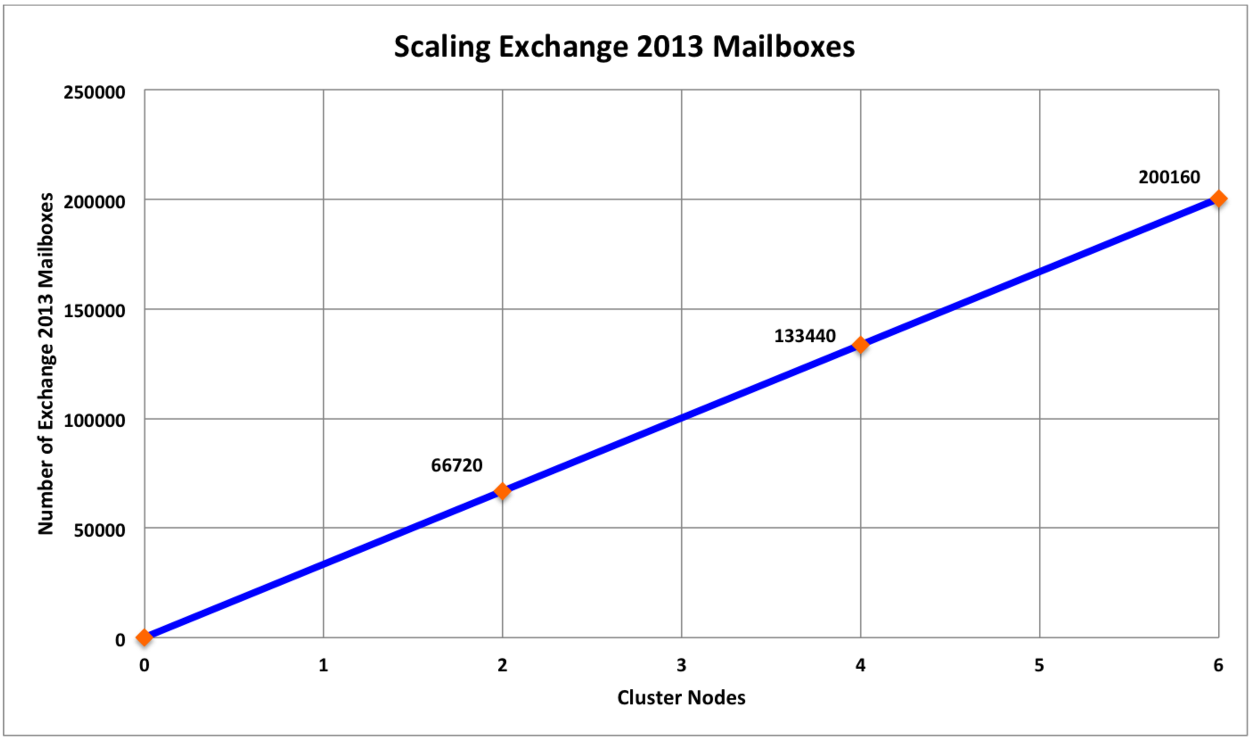

4 Conclusion

Jetstress 2013 was used to simulate the workload of 200,000 Exchange Server 2013 mailboxes on FAS8060 running clustered Data ONTAP 8.2.1. This work demonstrated the performance and scalability of FAS8060:

- **Performance**: Average database read and write latencies are 10.66ms and 1.69ms, respectively. Both latencies are well below the 20ms threshold.
- **Scalability**: By doubling or tripling the cluster nodes, the workload can also be doubled or tripled.

Acknowledgements

The author wishes to thank the following colleagues for their contributions during various stages of the project: Dave Tanis, Reena Gupta, Keith Griffin, Duaine Fioritto, Lee Dorrier, Chris Lueth, John Elliott, A.J. Mahajan, Joshey Lazer, Dhanesh Kumar, Rahul Menon, Scott Lane, Carol Whiteley, Steve Young, and Nathan Moffitt.

Version History

Version	Date	Document Version History
1.0	February 2014	Initial release

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.