



NETAPP TECHNICAL REPORT

Performance Report: Vhayu Velocity Compatibility and Performance on NetApp Storage Systems Using FC SAN

Providing FC SAN Storage for Vhayu Velocity

Chris Lemmons, NetApp and David Tobin, Vhayu Technologies Corporation
February 2009 | TR-3745

ABSTRACT

NetApp and Vhayu engineers recently completed an extensive project to demonstrate the suitability of NetApp® storage for Vhayu® Velocity™ implementations. This technical report documents the successful outcome of that project.

TABLE OF CONTENTS

- 1 INTRODUCTION 3**
- 2 COMPATIBILITY AND PERFORMANCE TEST RESULTS 3**
 - 2.1 FC SAN COMPATIBILITY TESTING 3
 - 2.2 FC SAN PERFORMANCE TESTING 4
- 3 VHAYU AND NETAPP BACKGROUND 4**
 - 3.1 VHAYU VELOCITY SOFTWARE 4
 - 3.2 NETAPP STORAGE TERMINOLOGY 5
- 4 TESTED CONFIGURATION INFORMATION 6**
 - 4.1 FAS3170 CONFIGURATION 7
 - 4.2 VHAYU ANALYTICS SERVER CONFIGURATIONS 7
- 5 CONFIGURING THE VHAYU ANALYTICS ENGINES TO USE NETAPP STORAGE 7**
 - 5.1 CREATE THE HISTORICAL ARCHIVE LUNS ON THE FAS3170 8
 - 5.2 PREPARE THE ANALYTICS ENGINE SERVERS FOR FIBRE CHANNEL 8
 - 5.3 BUILD THE HISTORICAL ARCHIVE 8
 - 5.4 CREATE SPACE-EFFICIENT COPIES OF THE HISTORICAL ARCHIVE 9
- 6 RESULTS OF FC SAN COMPATIBILITY AND PERFORMANCE TESTING 12**
 - 6.1 FC SAN COMPATIBILITY TESTING RESULTS 13
 - 6.2 FC SAN PERFORMANCE TEST RESULTS 14
- 7 BEST PRACTICES AND SIZING RECOMMENDATIONS 17**
- 8 APPENDIX A: CREATE AND DISTRIBUTE THE HISTORICAL ARCHIVE TO ALL VHAYU AE SYSTEMS 18**
 - 8.1 PROVISIONING THE HISTORICAL ARCHIVE STORAGE ON THE FAS3170 19
 - 8.2 PRESENTING THE HISTORICAL ARCHIVE STORAGE TO THE AE SYSTEM 19
 - 8.3 CREATING SPACE-EFFICIENT CLONED COPIES OF THE HISTORICAL ARCHIVE 21
- 9 APPENDIX B: MANAGE NIGHTLY UPDATE OF THE HISTORICAL ARCHIVE 25**
 - 9.1 REMOVE HISTORICAL ARCHIVE LUNS FROM AE SERVERS 25
 - 9.2 DELETE EXISTING HISTORICAL ARCHIVE DATA ON FAS3170 26
 - 9.3 CREATE AND PRESENT THE UPDATED HISTORICAL ARCHIVE 26
- 10 APPENDIX C. VHAYU AND NETAPP TEST CONFIGURATION NETWORK DIAGRAM 30**

1 INTRODUCTION

Vhayu's Velocity suite of products contains high-performance market data solutions for the capture and high-speed analysis of massive amounts of streaming and historical data. Existing client implementations include applications for pretrade analysis, algorithmic trading, customized VWAP calculations, transaction cost analysis, quantitative analysis, and regulatory compliance for MiFID and Reg NMS. Velocity for Equities gives users the ability to crunch massive amounts of real-time and historical data in milliseconds for concise analysis that provides brokers and hedge funds with a clear competitive trading advantage.

A critical piece of Vhayu's solution is the ability to use fast and reliable shared storage to access historical archives of trading data. NetApp provides a unified storage solution that is second to none in allowing storage access by using multiple protocols—including FC, iSCSI, NFS, and CIFS—in a single architecture.

Recently, Vhayu and NetApp engineers collaborated to define how to optimally configure NetApp storage solutions into a Velocity environment. This technical report describes the details of that venture, including:

- Best practices for laying out the storage for use as a Velocity historical archive using NetApp storage
- How to use NetApp Snapshot™ and clone technologies to create extremely space-efficient copies of the Velocity historical archive for presentation to Velocity analytic servers
- How to use these NetApp technologies listed above to update the historical archives from the daily feed logs
- Results of performance tests relating to random querying of the historical archive and recovery of analytical engines through feed log processing
- Results of compatibility tests designed to make sure that the Vhayu Applications and tools work correctly with NetApp storage.
- How to use simple Windows® batch files to automate the presentation and updating of the historical archives

2 COMPATIBILITY AND PERFORMANCE TEST RESULTS

At the Vhayu facility in Los Gatos, California, Vhayu and NetApp engineers conducted a series of tests to gauge the compatibility of the NetApp FAS3170 storage controller when configured to work with the Vhayu applications. In addition to the compatibility tests, Vhayu and NetApp conducted a series of performance tests to determine that NetApp storage met or exceeded the performance requirements for optimal operation of the Vhayu applications.

2.1 FC SAN COMPATIBILITY TESTING

Once configured to use the NetApp storage system, these tests consisted primarily of starting each of the normal Vhayu services and using those services to process data as they would under normal operation. The goal of the tests was to make sure that the files residing on the NetApp storage system used by the Vhayu applications could be created and accessed normally.

Successful completion of these tests required that the Vhayu application software work correctly with the configured NetApp storage. **We found that all of the compatibility tests completed successfully and that the NetApp storage was fully compatible with the Vhayu application stack.**

2.2 FC SAN PERFORMANCE TESTING

The performance tests were designed primarily to gauge how well the NetApp FAS3170 supported the Vhayu applications by providing FC storage for the following basic tasks required for optimal performance of the Vhayu software:

- Random queries of the Vhayu historical archive database using two and four Vhayu Analytics Engine (AE) systems
- Recovery of Vhayu AE systems from a simulated system failure by processing the Vhayu feed log file from one, two, and four AE systems

By using multiple Vhayu analytic engine systems, these tests placed higher loads on the FAS3170 as the number of AE systems was increased.

Overall, we found that the FAS3170 provided excellent performance when configured with the Vhayu applications. Also worth noting, we found that when storing the Vhayu historical archive on the FAS3170, the Vhayu applications exhibited linear performance scaling when multiple analytic engines were accessing the archive. Additionally, the Vhayu applications were able to process nearly 900,000 records per second when simulating the recovery of 4 analytic engines.

3 VHAYU AND NETAPP BACKGROUND

Before discussing the specifics of the tests and results, this section describes the underlying technologies of the Vhayu Velocity software and the NetApp storage products.

3.1 VHAYU VELOCITY SOFTWARE

Vhayu Velocity™ offers the fastest market data solutions for the capture and high-speed analysis of massive amounts of streaming and historical data, to brokers, hedge funds, alternative trading systems, and exchanges. When discussing the results of the compatibility and performance tests later in this document, we use the following terms to describe the different aspects of the Vhayu Velocity software stack.

- **Analytics engine (AE):** One or more of these systems provides access to both current and historical market data. These systems are queried by multiple clients and perform a variety of analyses.
- **Broadcast server (BS):** The broadcast server handles incoming real-time feeds of market data and transmits the data to the analytics engines in Vhayu's data format.
- **Feed logger (FL):** Records the data stream processed by the broadcast server. The logs generated by the feed logger are used to update the historical archive database and to recover an analytics engine after a failure.
- **Historical archive:** A read only set of data that provides historical ticker information for use by the analytics engines. This data is generally updated on a daily basis using the current day's data maintained by the feed logger. Data in the historical archive is a prime candidate for placement on shared storage like the NetApp FAS3170, where space efficient copies can be distributed to the analytics engines.
- **Real-time archive:** Current day's data maintained on the local storage of the analytics engines.

3.2 NETAPP STORAGE TERMINOLOGY

This section presents an overview of the terminology used to configure and describe NetApp storage systems. For more information, see the Data ONTAP 7.3 Storage Management Guide.

In their most basic form, NetApp storage systems combine large numbers of individual disk drives to provide flexible, high-performance storage to a variety of applications by using multiple protocols—including Fibre Channel (FC), iSCSI, NFS, and CIFS—in a single architecture. This section defines the storage building blocks used by NetApp to provide storage. The terms defined here are used throughout the rest of this technical report.

AGGREGATES AND RAID GROUPS

Data ONTAP® allows users to configure individual disk drives into large blocks of storage called *aggregates*. The individual disk drives contained in an aggregate are configured in RAID groups. RAID groups can contain differing numbers of disks, some of which are used to provide parity protection if a disk failure in a specific RAID group requires the data in the RAID group to be restored.

NetApp recommends using RAID-DP® parity protection when creating RAID groups and aggregates. RAID-DP uses two parity disks and can provide protection from the simultaneous loss of two data disks in a RAID group. RAID-DP provides significantly better data protection compared to more traditional single parity RAID schemes (RAID 5), while providing significantly better storage utilization compared to RAID schemes that offer more than single disk parity protection through mirroring (RAID 1).

Aggregates are a collection of one or more RAID groups configured into a single block of storage. When you create an aggregate, Data ONTAP assigns data disks and parity disks to RAID groups, depending on the options you choose, such as the size of the RAID group (based on the number of disks to be assigned to it) or the level of RAID protection. You can increase the usable space in an aggregate by adding disks to existing RAID groups or by adding new RAID groups to the aggregate.

FLEXVOL VOLUMES

A FlexVol® volume holds user data that is accessible through one or more of the access protocols supported by Data ONTAP, including Network File System (NFS), Common Internet File System (CIFS), HyperText Transfer Protocol (HTTP), Web-based Distributed Authoring and Versioning (WebDAV), Fibre Channel Protocol (FC), and Internet SCSI (iSCSI). A volume can include files (which are the smallest units of data storage that hold user- and system-generated data) and, optionally, directories and qtrees in a network-attached storage (NAS) environment, and also LUNs in a storage area network (SAN) environment.

FlexVol volumes are created in an existing aggregate. A FlexVol volume can share its containing aggregate with other FlexVol volumes. Thus a single aggregate is the shared source of all the storage used by the FlexVol volumes it contains. Flexible volumes are therefore no longer bound by the limitations of the disks on which they reside. A FlexVol volume is simply a "pool" of storage that can be sized based on how much data you want to store in it, rather than on what the size of your disk dictates.

A FlexVol volume can be shrunk or increased on the fly without any downtime. Flexible volumes have all the spindles in the aggregate available to them at all times regardless of the size of the volume. As a result, even the smallest FlexVol volume enjoys the performance benefits of all the disks in the containing aggregate.

FlexVol volumes are a breakthrough technology in which volumes are logical data containers that can be managed and moved independently from the underlying physical storage. Because FlexVol volumes have access to all available physical storage in the NetApp system, FlexVol

technology provides substantial improvements in storage provisioning, utilization, and volume sizing. Data volumes can be sized and resized quickly and dynamically as application requirements change, and sizing can be automated based on space management policies.

LUNS

In SAN environments, NetApp storage systems are targets that have storage target devices, which are referred to as LUNs (logical unit numbers). With Data ONTAP, you configure the storage systems by creating aggregates to contain FlexVol volumes to store LUNs. LUNs created on NetApp storage systems in a SAN environment are used as targets for external storage that is accessible from initiators, or hosts. You use these LUNs to store files and directories that are accessible through a UNIX® or Windows host via FC or iSCSI. Multiple LUNs can be created in a single NetApp FlexVol volume, limited only by the amount of available space in the containing FlexVol volume and aggregate.

SNAPSHOT COPIES

A Snapshot copy is an online read-only logical copy of a FlexVol volume. Typically a Snapshot copy takes only a few seconds to create, regardless of the size of the volume or the level of activity on the NetApp storage system. After a Snapshot copy is created, changes to data objects are reflected in updates to the current version of the objects, as if the Snapshot copy did not exist. Meanwhile, the Snapshot version of the data remains completely stable. The presence of one or more Snapshot copies in a volume does not affect the performance of that volume.

FLEXCLONE VOLUMES

NetApp FlexClone® technology uses Snapshot copies of FlexVol volumes to enable instant cloning of data volumes and data sets without requiring additional storage space at the time of creation. Each cloned volume is a transparent, virtual copy that you can use for essential enterprise operations, such as testing and bug fixing, platform and upgrade checks, multiple simulations against large data sets, remote office testing and staging, and market-specific product variations.

FlexClone is a powerful and unique feature introduced in Data ONTAP 7G that adds a new level of agility and efficiency to storage operations by allowing a user to create an instant clone of a flexible volume (FlexVol volume). A FlexClone volume is a writable point-in-time image of a FlexVol volume or the FlexVol volume itself. With FlexClone, it takes only a few seconds to create a clone of a FlexVol volume without interrupting access to the parent volume the clone is based on.

The clone volume uses space very efficiently, allowing the original FlexVol volume and the cloned volume to share common data, storing only the data that changes between the original volume and the clone. This provides a huge potential saving in storage space, resources, and cost. In addition, a FlexClone volume has all the features and capabilities of a regular FlexVol volume, including the ability to grow or shrink and the ability to be the source of another FlexClone volume.

4 TESTED CONFIGURATION INFORMATION

This section describes the basic hardware and software configuration of the NetApp FAS3170 and the Vhayu Velocity application. For a network diagram of the configured testbed, see Appendix C.

4.1 FAS3170 CONFIGURATION

The NetApp FAS3100 series of storage systems offers customers a single unified storage architecture to meet their needs and a common suite of application-aware software and management tools that are part of the Data ONTAP operating system. In addition, with technology such as thin provisioning, deduplication, RAID-DP, and FlexClone, customers can boost storage efficiency and performance without breaking the bank.

For these tests, we configured the FAS3170 as follows:

- Six shelves of 300GB/15K RPM FC disks, providing approximately 13TB of usable storage
- Standard active-active controller configuration, as shown in the network diagram in Appendix C
- Running Data ONTAP 7.3P3
- FC connectivity provided by using 2 @ Brocade Fibre Channel switches
- Standard Gigabit Ethernet connection for remote management

4.2 VHAYU ANALYTICS SERVER CONFIGURATIONS

The Vhayu AE systems used in these tests consisted of four servers containing a minimum of eight processors and 16GB of RAM configured as follows:

- 64-bit versions of Microsoft Windows Server 2003 Enterprise Edition and SP2
- Qlogic 2642 dual port Fibre Channel host bus adapter
- FC connectivity provided by using two Brocade Fibre Channel switches
- Connected for remote configuration by using standard Gigabit Ethernet

5 CONFIGURING THE VHAYU ANALYTICS ENGINES TO USE NETAPP STORAGE

We found the process of connecting the Vhayu AE servers to the NetApp FAS3170 to be very straightforward, consisting of the following general steps:

1. On the FAS3170 storage controllers, create a set of LUNs to hold the Vhayu historical archive database.
2. Following NetApp best practices, install a set of host utilities, FC HBA drivers, and multipathing software on each of the Vhayu AE servers.
3. Present the LUNs created on the FAS3170 to the AE server designated to have read/write access to the historical archive by mapping the host side FC initiators to the same set of LUNs.
4. Using the standard Windows disk management tools, scan for additional disks.
5. When the LUNs appear under Windows, use the standard Windows disk management tools to initialize and format the disks by using NTFS and use them as you would any other Windows disk.
6. Using Vhayu tools, create the historical archive by loading the archive files across the LUNs created to hold the archive.
7. Using NetApp Snapshot and FlexClone technology, create space-efficient copies of the historical archive and present them to the remaining Vhayu AE systems so that all AE systems have the same copy of the historical archive.

The rest of this section describes these steps in detail. Although we used the NetApp GUI to perform the initial configuration, Appendix A provides the full set of command line items required to initially configure the NetApp FAS3170 for use with the Vhayu Velocity software. These commands can be used as the basis for creating scripts to perform the initial configurations.

5.1 CREATE THE HISTORICAL ARCHIVE LUNS ON THE FAS3170

As shown in Appendix C, the FAS3170 used for this exercise contained two storage controllers, each configured with a total of forty-two 300GB/15K RPM Fiber Channel disk drives. To maximize the usable storage as well as performance, we created a single aggregate on each storage processor that contained a total of 36 disks configured as 3 RAID-DP groups of 10 data disks and 2 parity disks per group (10+2). This left 3 spare disks per controller.

The result was approximately 13TB of usable storage across both of the FAS3170 storage controllers. We then created a single FlexVol volume on each storage controller that provided 6TB of storage and included a 5% Snapshot reserve. Additionally, we set the fractional reserve for this FlexVol volume to zero. In each of the FlexVol volumes we then created a total of four 1.5 TB LUNs to present to the analytics engine servers for use in deploying the historical archive and feed log files.

We then presented all eight of the 1.5TB LUNs to the AE system, which required both read and write access to the historical archive in order to perform the nightly updates. Presenting the LUNs to Windows is a simple matter of determining the WWPNs of the Fibre Channel initiators installed into the AE systems and using the Data ONTAP management GUI to map those specific initiators to each of the eight 1.5 TB LUNs.

5.2 PREPARE THE ANALYTICS ENGINE SERVERS FOR FIBRE CHANNEL

Each of the four AE server systems was configured to access the FAS3170 using NetApp best practices for using the Fibre Channel Protocol (FCP). This includes the following items:

- Install the NetApp FC Host Utilities Kit
- Install the latest NetApp certified driver for the Qlogic 2462 FC HBAs installed in each server
- Install the Data ONTAP DSM 3.2R1 for the Windows MPIO multipathing module to manage all the paths to the LUNs that make up the historical archive

5.3 BUILD THE HISTORICAL ARCHIVE

We used Vhayu-specific tools to create an archive for the testing that contained two years' worth of historical data that was spread across the eight disks roughly by month. After mapping, we used the standard Windows disk management utilities to rescan the disks; the LUNs appeared to Windows like any other basic disk, and they can be formatted by using NTFS. Figure 1 shows how the LUNs appear on the AE system that has both read and write access to the historical archive.

We found it helpful to assign a meaningful volume name to each of the LUNs under Windows to make it easier to remember which disk was associated with which LUN on the FAS3170.

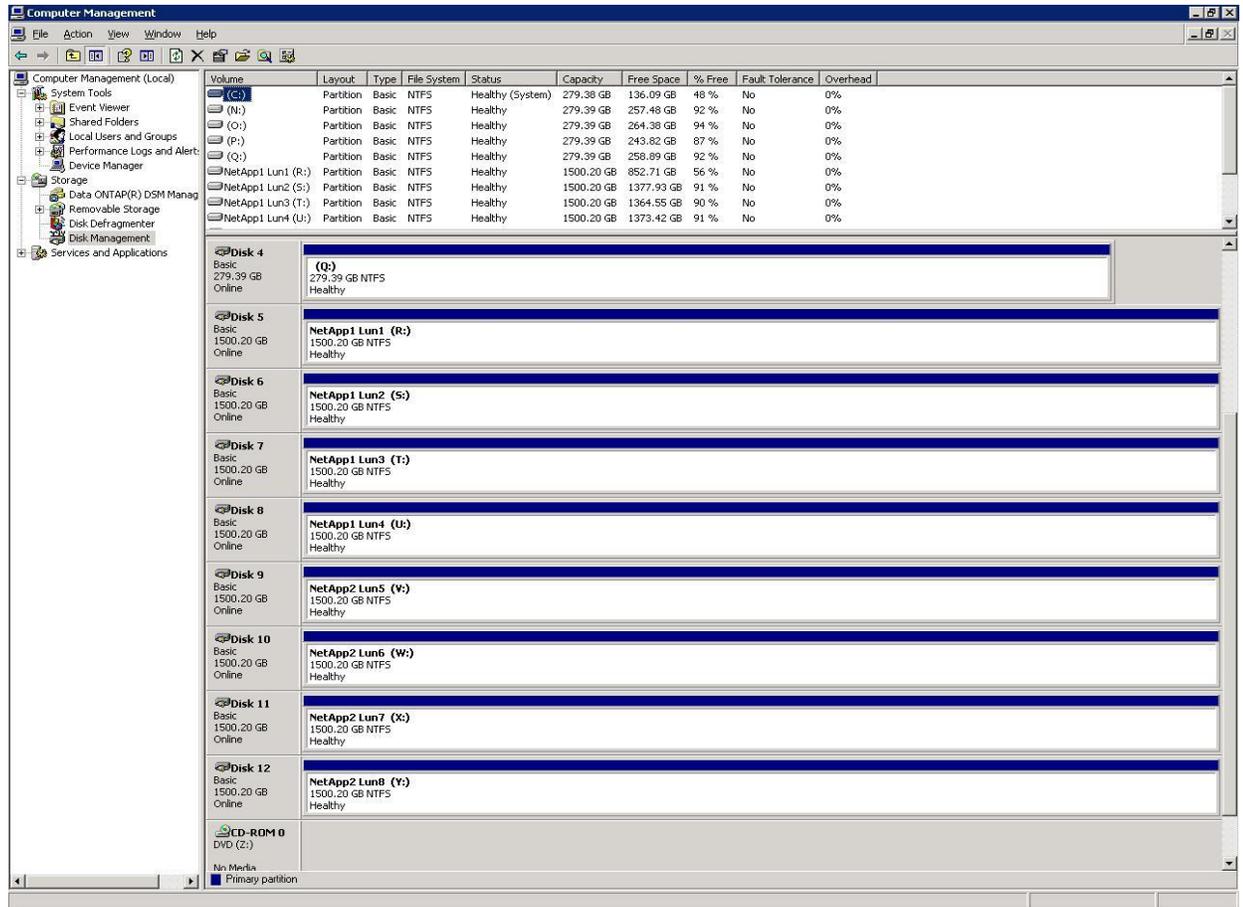


Figure 1) Historical archive LUNs mapped to Vhayu AE system.

5.4 CREATE SPACE-EFFICIENT COPIES OF THE HISTORICAL ARCHIVE

Once the historical archive was created, we used the NetApp Snapshot and FlexClone technologies described earlier to create a separate, space efficient clone of the FlexVol volume containing the historical archive for each of the Vhayu AE servers. Using FlexClone volumes allowed us to create extremely space-efficient copies of the historical archive for each analytics engine that initially consumed no additional space above that allocated for the original FlexVol volume.

Creating a FlexClone volume begins by creating a standard Snapshot copy of the FlexVol volume on which we created the historical archive on each of the FAS3170 controllers. These Snapshots contained copies of all of the original eight 1.5 TBLUNs that make up the historical archive. The FlexClone volumes use these Snapshot copies as the basis for presenting these LUNs to the remaining analytic engines systems. For our tests, we created a single Snapshot copy on each of the FAS3170 controllers by using the Data ONTAP command line interface.

Figure 2 shows the three FlexClone volumes created on one of the FAS3170 controllers for use by each of the three AE systems requiring read-only access to the historical archive. In this case, we used a date and time format to name the Snapshot copies on which the FlexClone volumes are based. This is shown in the Parent Snapshot column in Figure 2. A similar set of FlexClone volumes was created on the second FAS3170 controller as well.

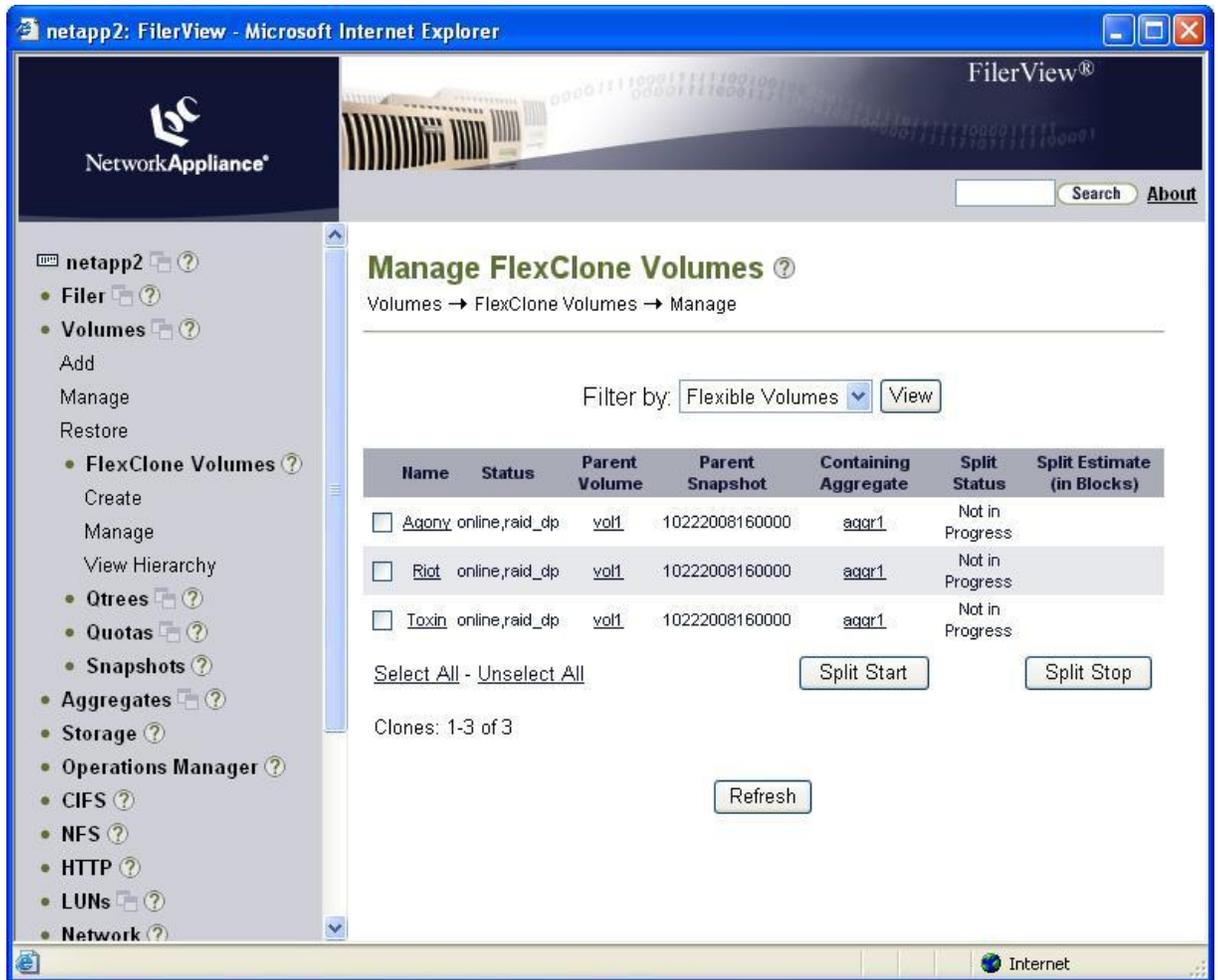


Figure 2. FlexClone volumes for read-only historical archive access.

Although these FlexClone volumes are copies of the original FlexVol volume used for the historical archive, they are managed from with Data ONTAP as normal FlexVol volumes. Figure 3 shows the original FlexVol volume (vol1) used by the AE system, which requires read and write access to the historical archive, and the three FlexClone volumes, which provide a copy of the historical archive to the remaining AE systems. Each of the four FlexVol volumes provides a total of 6TB of storage for historical archives (24TB of total space) while actually consuming only the initial 6TB of space required to initially provision the historical archive. Again, we created a similar set of FlexClone volumes on the second FAS3170 controller.



Figure 3) All historical archive FlexVol volumes on FAS3170.

When the FlexClone volumes have been created for the remaining analytics engine servers, the process to present the LUNS contained in each of these clones to the AE systems is very similar to that described earlier for initially presenting the FlexVol volume. The following general steps describe the process and can be done by using either the Data ONTAP command line or the GUI interface:

1. Set each LUN in each FlexClone volume to an online state.
2. Gather the WWPN names of the FC initiators on the remaining AE systems and map them to the specific LUNs.
3. On each AE system, use the standard Windows disk management tools to rescan the disks. All 8 of the LUNs should appear on each system. Because they are cloned from the original set of LUNs, they are presented as NTFS formatted disks.
4. Assign the same set of drive letters to each of the disks in the same fashion as when the historical archive was created.

Figure 4 shows all 16 LUNs (4 LUNs per AE system) on one of the FAS3170 storage controllers after completing the steps just outlined. This shows the state of each LUN as online; the name of the AE system to which each LUN is mapped; and the LUN ID assigned to each LUN.

Note: Each AE system that requires read-only access to the historical archive uses the same set of LUN IDs. A similar set of LUNs exists on the second FAS3170 controller, making a total of eight 1.5 TB LUNs presented to each AE system that is using the historical archive.

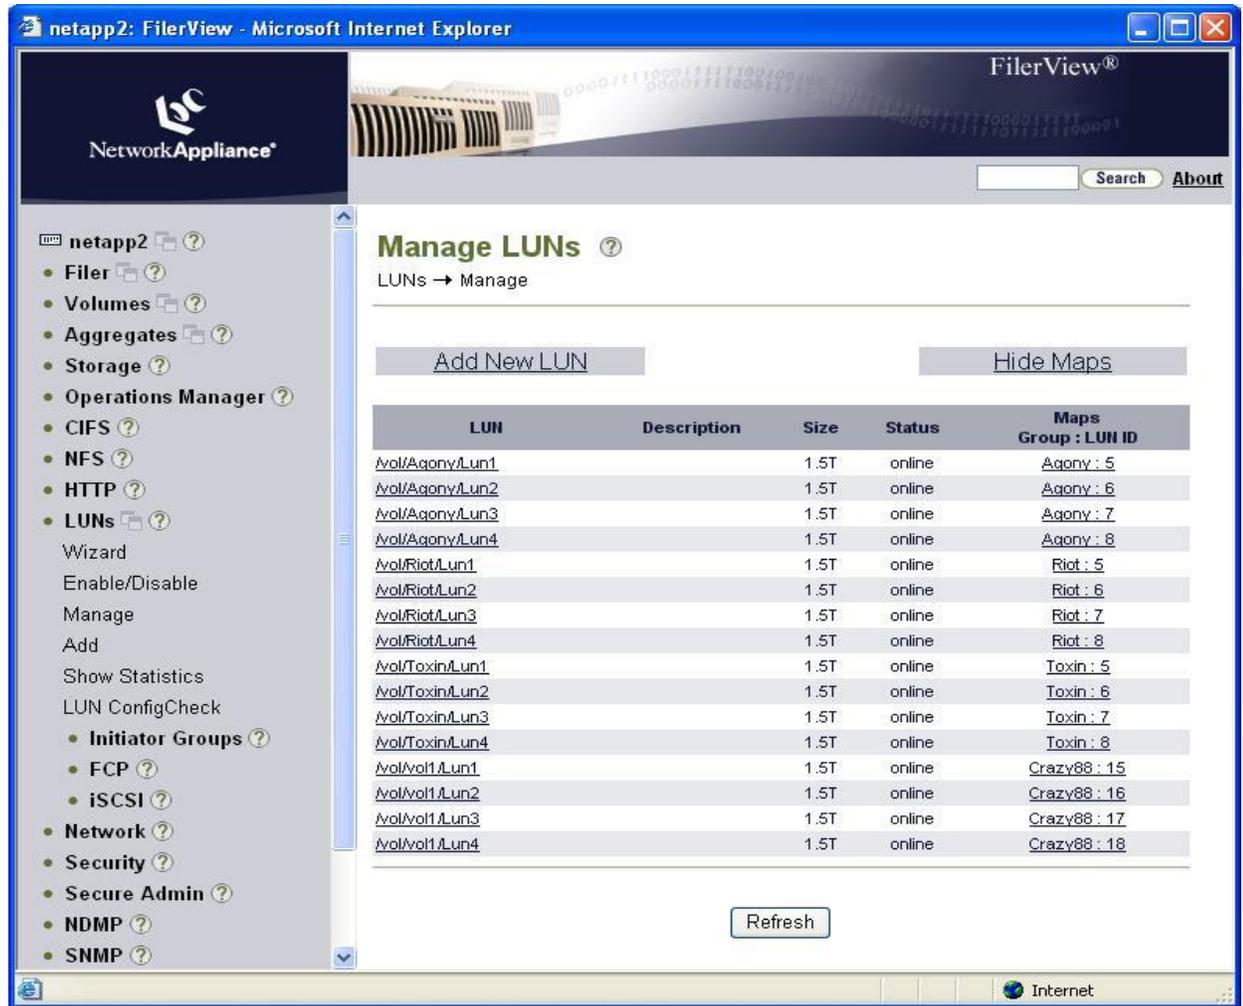


Figure 4) FAS3170 LUNs provide historical archives for AE systems.

6 RESULTS OF FC SAN COMPATIBILITY AND PERFORMANCE TESTING

In the past, Vhayu has observed compatibility, performance, and sharing issues when implementing the Vhayu applications on SAN storage systems. In addition to verifying overall compatibility, NetApp and Vhayu engineers were also interested in demonstrating the ability of the NetApp storage systems to accept the high volume of continuous data writes, especially writing multiple files from multiple systems, as well as the ability to support an equally high volume of continuous random access to shared files from multiple systems simultaneously.

At the Vhayu facility in Los Gatos, California, Vhayu and NetApp engineers jointly conducted a series of tests to gauge the compatibility of the NetApp storage when configured to work with the Vhayu applications. In addition to the compatibility tests, Vhayu and NetApp also conducted a series of performance tests to determine whether NetApp storage met or exceeded the performance requirements for optimal operation of the Vhayu applications.

The following sections describe the tests conducted and summarize the test results.

6.1 FC SAN COMPATIBILITY TESTING RESULTS

After configuring the Vhayu applications to use the NetApp storage system, these tests consisted primarily of starting each of the normal Vhayu services and beginning to process data independently, as in normal operation. The goal of the tests was to determine whether the files used by the Vhayu applications can be created and accessed normally.

Successful completion of these tests simply required that the Vhayu application software work correctly with the configured NetApp storage. All of the compatibility tests completed successfully, and the NetApp storage was fully compatible with the Vhayu application stack.

The following paragraphs describe the compatibility tests that we conducted.

AE RECOVERY TESTING

This test case simulated the recovery of a crashed AE server. We started an AE server and recovered it by using an existing Vhayu formatted feed file located on the NetApp FAS3170 storage system.

BROADCAST SERVER COMPATIBILITY TESTING

This test case simulated using a BS to create a Vhayu formatted feed file on the NetApp FAS3170 storage system. For this test we used Vhayu's internal Feed4Test tool.

FEED LOGGER COMPATIBILITY TESTING

This test case simulated the continuous updating of a Vhayu formatted feed log file residing on the FAS3170 storage system. We configured a BS to generate data on a common multicast address and used the feed logger to create and populate the feed log file.

ARCHIVE TOOL COMPATIBILITY TESTING

These tests built and manipulated historical archives where the source and destination files both existed on the NetApp storage system. We conducted the following tests:

- Used ArcBuild to build a new archive on the FC SAN from a Vhayu feed file on the FC SAN
- Used ArcBuild to update an existing archive on the FC SAN from a Vhayu feed file on the FC SAN
- Used ArcMaint to copy an existing archive on the FC SAN to another archive on the FC SAN
- Used ArcMaint to split an existing archive on the FC SAN into multiple archives by record type

UNBUFFERED I/O TESTING

Vhayu Velocity software uses unbuffered I/O from within Windows. These tests were designed to make sure that unbuffered I/O to and from the FAS3170 worked correctly with the Vhayu software. We used Vhayu's internal CopyNoBuffer application and the publicly available PassMark load generator to verify compatibility by reading and writing files to and from the FAS3170.

To make sure that 512 byte sector compatibility exists, we configured the internal buffer size of the CopyNoBuffer program to 512 bytes for all tests.

MULTIPLE COMPONENT TESTING

After verifying the compatibility of the individual Vhayu software components by using the tests just described, we conducted this test to determine the compatibility of the NetApp FAS3170 with all the components of the Vhayu software running simultaneously, as would be the case in a real deployment.

This test case deployed multiple AE servers, created a historical archive on the NetApp storage system from a SAN-based feed log, and deployed a BS system to send data at an average rate of 150,000 records per second. Additionally, we used a series of scripts to simulate each of the AE servers retrieving data simultaneously from the historical archive.

6.2 FC SAN PERFORMANCE TEST RESULTS

Basic performance tests were conducted to compare the performance of typical tasks by using the Vhayu applications with NetApp storage and local disks. The following tests were performed:

- Randomly querying the historical archive database and measuring the throughput achieved by using the FAS3170 and multiple AE servers
- Sequentially reading the feed logger file to simulate what happens when one or more of the analytical engine systems fails and has to have the current information (information that has come in after the crash) updated onto the local disks

Overall performance measurement was composed of the real-time data storage tests just described and the historical data retrieval part of the performance tests. This simulated a realistic environment as well as evaluating the performance of simultaneous reading and writing of the NetApp storage system.

HISTORICAL ARCHIVE PERFORMANCE TESTING

As more AE servers access the same historical archive, the performance demands on the NetApp storage system increase significantly. The goal of this test was to measure how well the performance of the FAS3170 held up as the number of AE servers accessing the historical archive increased from two to four.

We evaluated the historical data access performance by starting up to four AE servers, all accessing a single Vhayu historical archive stored on the FAS3170. We then ran a series of scripts to have each AE server access the historical archive by using a variety of random access patterns. We measured the throughput observed by the AE servers by having each AE system use Windows Performance Monitor to capture the throughput observed on each of the individual disks that made up the historical archive.

The results of this test showed linear performance scaling for all test cases. Figure 5 compares the throughput observed using four AE systems relative to that observed using two AE systems. In this case, the throughput generated using two AE systems is represented by a value of 100. The other column on the graph represents the throughput observed when using four AE systems. For example, a value of 200 on this graph would indicate a throughput value twice that generated using two AE systems.

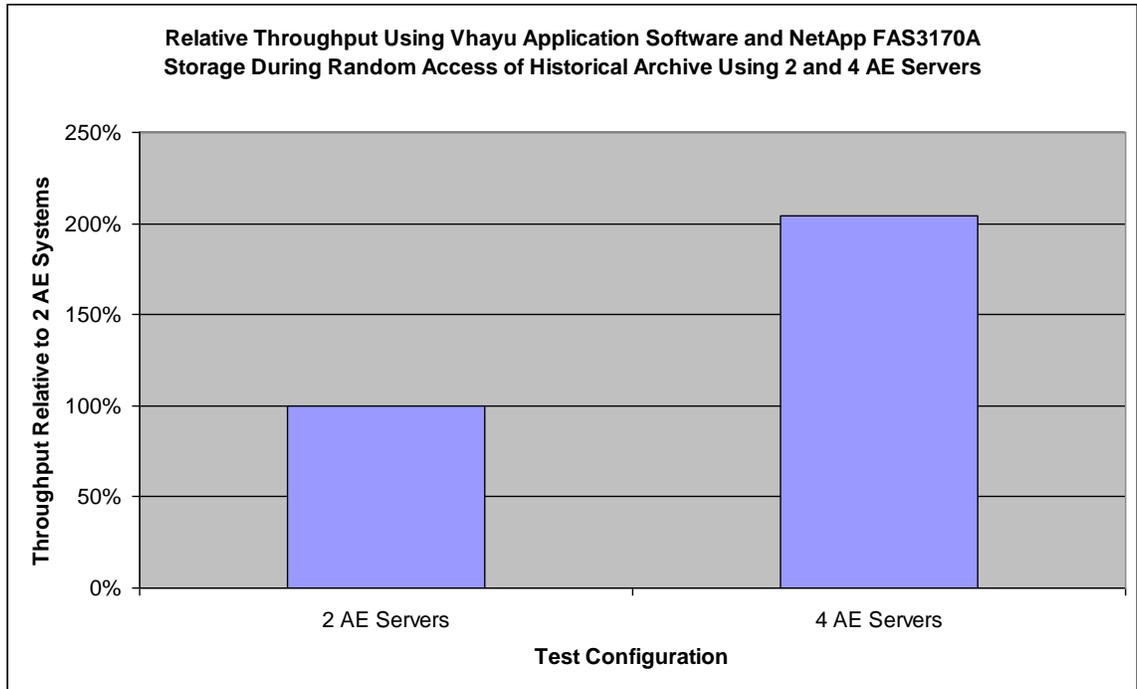


Figure 5) Relative throughput performance accessing the historical archive using two and four AE servers.

Figure 6 presents the CPU and disk utilization observed on the FAS3170 during each of the performance test cases using the historical archive. This figure shows that there was still plenty of headroom left on the FAS3170 in terms of both CPU and disk, indicating that there is capacity to handle additional AE systems accessing the historical archive if need arises.

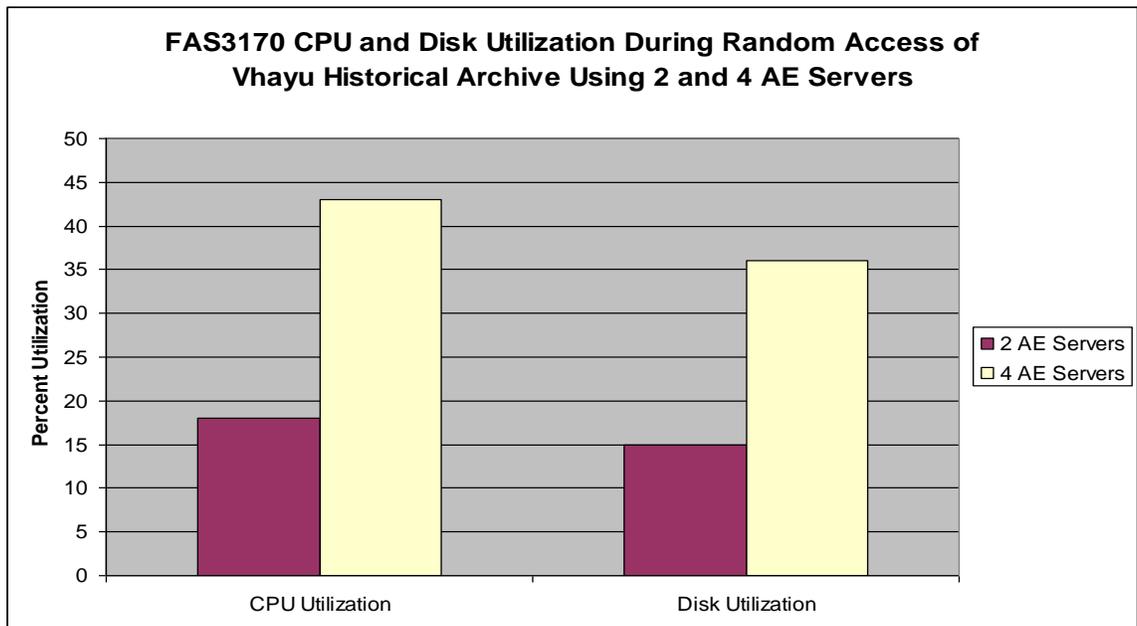


Figure 6) FAS3170 CPU and disk utilization during performance testing using the historical archive.

We evaluated the feed log recovery performance by starting up one, two, and four AE servers in recovery mode to force the reading of data from a Vhayu feed log stored on the FAS3170 storage system to update the real-time data stored on the AE's local disks. The goal of the test was to observe the number of records per second processed when recovering from one, two, and four simultaneous AE failures.

Figure 7 shows the results of this test. To recover from a failure, each AE has to process the entire feed log file. This is basically a sequential read through all of the records in the feed log file. These results show that the number of records processed per second increased to almost 900,000 when all 4 AEs were recovering. For these tests, the average record size was 139 bytes.

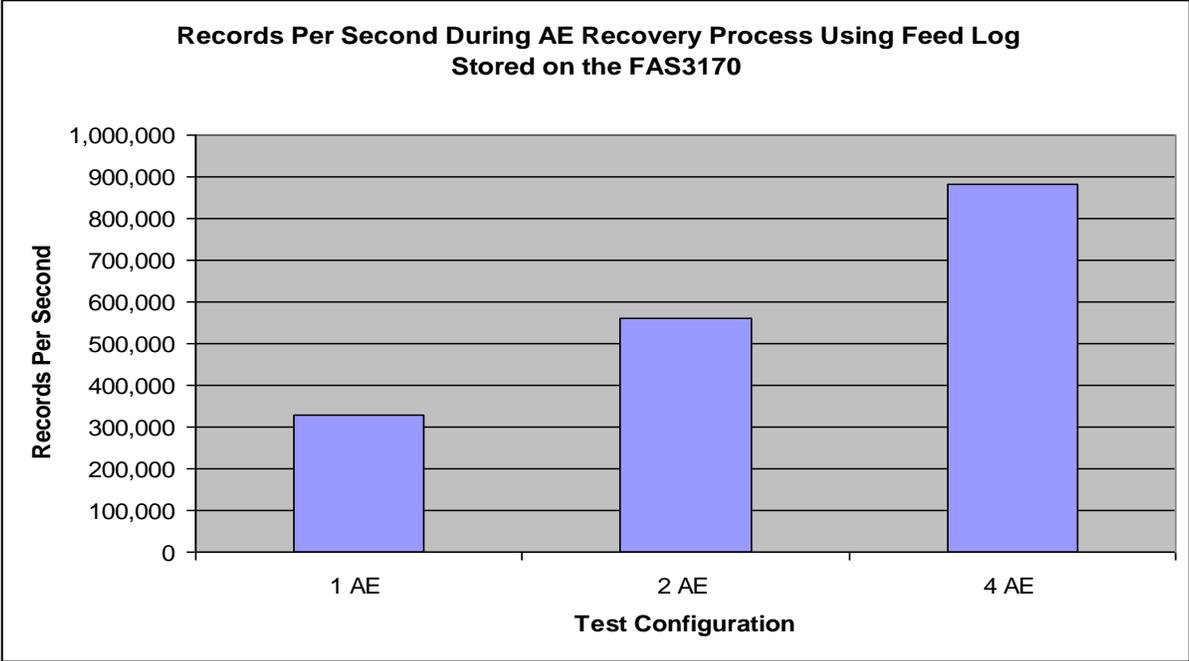


Figure 7. Records per second processed when recovering one, two, and four AEs from a feed log located on the FAS3170.

7 BEST PRACTICES AND SIZING RECOMMENDATIONS

We followed the standard NetApp best practices for FC SAN when configuring the FAS3170 for use with the Vhayu applications. The remainder of this section discusses specific recommendations for use when deploying Vhayu Velocity software with NetApp storage.

- The Vhayu Velocity software can generate 50 to 100GB of new data each day that must be incorporated into the historical archive that night. Using a smaller number of larger LUNs on the FAS3170 makes it easier to manage the nightly rollover process.
- If possible, spread the individual files that make up the historical archive as evenly as possible across all of the LUNs for better I/O concurrency and overall storage performance.
- Disable any automatic Snapshot creation schedules for all aggregates and FlexVol volumes. It is not necessary to have more than one Snapshot copy of the historical archive volumes active at any time.
- Because there is virtually no overwriting of the current historical archive data held in the active Snapshot copy, NetApp recommends configuring the FlexVol volume used for the historical archive LUNs with a 5% Snapshot reserve. For the same reason, NetApp recommends setting the fractional reserve to zero for the FlexVol volume that contains the initial set of LUNs to make up the historical archive.
- When creating your FlexVol volume for the historical archive, make sure that the size is not larger than 95% of the size of the containing aggregate. This leaves enough free space in the aggregate for Data ONTAP to function correctly.
- When initially creating the LUNs in the FlexVol volume, make sure that the total size of the LUNs does not exceed 95% of the total space in the FlexVol volume. This leaves enough free space in the volume for Data ONTAP to function correctly.
- Make sure that the LUN IDs you use in the original historical archive and clones are used each time you present a new set of historical archive LUNs through the FlexClone process.

8 APPENDIX A: CREATE AND DISTRIBUTE THE HISTORICAL ARCHIVE TO ALL VHAYU AE SYSTEMS

Each Vhayu AE system in production is required to access the same copy of the historical archive. This section contains the details and set of commands required to create the initial historical archive and distribute it to each of the four Vhayu AE systems we used for the tests.

The commands to manage the NetApp storage assume the use of remote shell (`rsh`) to remotely execute commands on the FAS3170 and the use of the `psexec` utility to remotely execute Windows storage configuration commands on the Vhayu AE systems. Although the steps in this appendix are performed manually, it is a straightforward process to use either Windows batch files or PERL to script these commands for remote administration. Also, the size values used in the commands to create aggregates, FlexVol volumes, and LUNs are specific to our test configuration; they will change depending on the actual configuration.

In reading this appendix, assume the following:

- The four Vhayu AE systems are named `agony`, `riot`, `toxin`, and `crazy88`.
- None of the Vhayu AE systems are currently using an existing historical archive.
- `Crazy88` is the Vhayu AE system with read and write access to the historical archive.
- The Vhayu AE systems named `agony`, `toxin`, and `riot` require read-only access to the historical archive.

The following general steps describe how to create copies of the historical archive for all of the AE systems:

1. Create an aggregate on each of the FAS3170 storage controllers consisting of 36 disks per aggregate.
2. Create a 6 TB FlexVol volume inside each of the aggregates.
3. Create four 1.5TB LUNs in each of the FlexVol volumes to present to the Vhayu AE system that has read and write access to the historical archive files (`crazy88`). This is a total of eight 1.5TB LUNs.
4. Present all of the LUNs to the AE system that will have read and write access to the historical archive.
5. Using the Windows disk management tools, scan for the new disks. When they are visible under Windows, use NTFS to initialize and format each of the eight 1.5TB disks.
6. Using Vhayu tools, create the historical archive using the eight 1.5TB disks.
7. After the archive is created, create a set of FlexClone volumes on the FAS3170 to provide space-efficient, read-only copies of the historical archive to the remaining AE systems/

The rest of this appendix describes the commands required to perform these steps. To configure and manage the storage, we used the FAS3170 command line interface. In all cases, commands were executed on both FAS3170 storage controllers, named `netapp1` and `netapp2`.

8.1 PROVISIONING THE HISTORICAL ARCHIVE STORAGE ON THE FAS3170

The following commands create an aggregate on both FAS3170 storage controllers, with each aggregate containing a total of thirty-six 300GB/15K RPM disk drives:

- `rsh netapp1 aggr create aggr1 -r 12 -t raid_dp 36`
- `rsh netapp2 aggr create aggr1 -r 12 -t raid_dp 36`

After creating the aggregates, use the following commands to create a FlexVol volume in each aggregate with 6 TB of available space:

- `rsh netapp1 vol create vol1 -s volume aggr1 6t`
- `rsh netapp2 vol create vol1 -s volume aggr1 6t`

Set the snapshot reserve on the volumes just created to 5 percent:

- `rsh netapp1 snap reserve vol1 5`
- `rsh netapp2 snap reserve vol1 5`

Set the fractional reserve on the volumes just created to zero:

- `rsh netapp1 vol options vol1 fractional_reserve 0`
- `rsh netapp2 vol options vol1 fractional_reserve 0`

After creating the FlexVol volumes, use the following commands to create four 1.5TB LUNs on each of the FAS3170 storage controllers:

- `rsh netapp1 lun create -s 1500g -t windows /vol/vol1/Lun1`
- `rsh netapp1 lun create -s 1500g -t windows /vol/vol1/Lun2`
- `rsh netapp1 lun create -s 1500g -t windows /vol/vol1/Lun3`
- `rsh netapp1 lun create -s 1500g -t windows /vol/vol1/Lun4`
- `rsh netapp2 lun create -s 1500g -t windows /vol/vol1/Lun1`
- `rsh netapp2 lun create -s 1500g -t windows /vol/vol1/Lun2`
- `rsh netapp2 lun create -s 1500g -t windows /vol/vol1/Lun3`
- `rsh netapp2 lun create -s 1500g -t windows /vol/vol1/Lun4`

8.2 PRESENTING THE HISTORICAL ARCHIVE STORAGE TO THE AE SYSTEM

After the historical archive LUNs have been created, they must be presented to the Vhayu AE system that requires both read and write access to the archive. NetApp uses initiator groups (igroups) to map between specific storage LUNs and the FC HBAs residing on the Vhayu AE systems. This mapping allows the LUNs to be managed on the Windows hosts as local basic disks.

Each igroup contains the two WWPN names for the two FC initiators on each of the four AE systems. We obtained these WWPN names from the Qlogic management software installed on each AE system. Because all four of the AE systems will eventually need to access LUNs on the FAS3170, the following commands create a set of four igroups, one for each AE system.

Note: The following commands do not show the actual WWPNs; they merely provide placeholders.

- `rsh netapp1 igroup create -f -t windows crazy88_igrp <crazy88_wwpn1 crazy88_wwpn2>`
- `rsh netapp1 igroup create -f -t windows agony_igrp <agony_wwpn1 agony88_wwpn2>`
- `rsh netapp1 igroup create -f -t windows toxin_igrp <toxin_wwpn1 toxin_wwpn2>`
- `rsh netapp1 igroup create -f -t windows riot_igrp <riot_wwpn1 riot_wwpn2>`
- `rsh netapp2 igroup create -f -t windows crazy88_igrp <crazy88_wwpn1 crazy88_wwpn2>`
- `rsh netapp2 igroup create -f -t windows agony_igrp <agony_wwpn1 agony88_wwpn2>`
- `rsh netapp2 igroup create -f -t windows toxin_igrp <toxin_wwpn1 toxin_wwpn2>`
- `rsh netapp2 igroup create -f -t windows riot_igrp <riot_wwpn1 riot_wwpn2>`

After creating the igroups, map the LUNS for the AE named crazy88 that requires read and write access to the historical archive:

- `rsh netapp1 lun map /vol/vol1/lun1 crazy88_igrp 11`
- `rsh netapp1 lun map /vol/vol1/lun2 crazy88_igrp 12`
- `rsh netapp1 lun map /vol/vol1/lun3 crazy88_igrp 13`
- `rsh netapp1 lun map /vol/vol1/lun4 crazy88_igrp 14`
- `rsh netapp2 lun map /vol/vol1/lun1 crazy88_igrp 15`
- `rsh netapp2 lun map /vol/vol1/lun2 crazy88_igrp 16`
- `rsh netapp2 lun map /vol/vol1/lun3 crazy88_igrp 17`
- `rsh netapp2 lun map /vol/vol1/lun4 crazy88_igrp 18`

Once the LUNs are mapped, scan the disks under Windows on Crazy88. Use the `diskpart` utility on each of the AE systems to rescan the disks under Windows. For this we created a file on each AE system called `rescan.txt` that simply contained the line `rescan`. When used in conjunction with the `-s` switch, `diskpart` opens the specified file and reads the set of commands to execute from the file. In this case, we're telling `diskpart` to conduct a rescan of all the disks. This is the same operation as performing a rescan of the disks using the Windows-based disk management tools.

- `psexec \\crazy88 diskpart -s rescan.txt`

Once the eight 1.5TB LUNs appear under Windows on the AE system, you need to initialize and format them using NTFS and assign each one a volume name and drive letter. After formatting the disks, we used the Vhayu ArchBuild tools to create a historical archive using two years of data that was spread across all eight of the disks. By spreading the data across the disks, we ensured that the I/O generated during random access of the historical archive would be spread evenly across the disks and across the FAS3170 storage controllers, resulting in more efficient storage.

8.3 CREATING SPACE-EFFICIENT CLONED COPIES OF THE HISTORICAL ARCHIVE

Once the historical archive was created, we used NetApp Snapshot and FlexClone technology to create space efficient copies of the historical archive to present to the other three Vhayu AE servers used in these tests. Here are the general steps for this process:

1. Create a Snapshot copy of the FlexVol volume on both FAS3170 storage controllers that contains the LUNs used by the historical archive.
2. Using this Snapshot copy, create a separate FlexClone volume for each of the remaining three AE servers.
3. Present the LUNs from each of the FlexClone volumes to the respective AE server so that they each have a copy of the historical archive that is identical to the historical archive created initially.

Use the following commands to create the Snapshot copy of the historical archive FlexVol volume. A Snapshot copy is created on both FAS3170 storage controllers:

- `rsh netapp1 snap create voll archive_snap`
- `rsh netapp2 snap create voll archive_snap`

Create a new set of FlexClone volumes by using the Snapshot copy just created as a base. Create one for each AE system that requires read-only access to the archive.

Note: All of the FlexClone volumes are created by using the same Snapshot copy of the updated historical archive. This ensures that all of the AE systems are using the same historical archive.

- `rsh netapp1 vol clone create agony_clone -s none -b voll archive_snap`
- `rsh netapp1 vol clone create toxin_clone -s none -b voll archive_snap`
- `rsh netapp1 vol clone create riot_clone -s none -b voll archive_snap`
- `rsh netapp2 vol clone create agony_clone -s none -b voll archive_snap`
- `rsh netapp2 vol clone create toxin_clone -s none -b voll archive_snap`
- `rsh netapp2 vol clone create riot_clone -s none -b voll archive_snap`

Although the newly created FlexClone volumes contain the LUNs on which the historical archive are stored, the LUNs are not initially in an online state and are not initially mapped to any of the FC initiators on the remaining AE systems. Each FlexClone volume contains four different LUNs. Each AE system is presented with a total of eight LUNs, so we set the LUNs for all AE systems to the online state.

For the AE named agony:

- `rsh netapp1 lun online /vol/agony_clone/lun1`
- `rsh netapp1 lun online /vol/agony_clone/lun2`
- `rsh netapp1 lun online /vol/agony_clone/lun3`

- rsh netapp1 lun online /vol/agonny_clone/lun4
- rsh netapp2 lun online /vol/agonny_clone/lun1
- rsh netapp2 lun online /vol/agonny_clone/lun2
- rsh netapp2 lun online /vol/agonny_clone/lun3
- rsh netapp2 lun online /vol/agonny_clone/lun4

For the AE named toxin:

- rsh netapp1 lun online /vol/toxin_clone/lun1
- rsh netapp1 lun online /vol/toxin_clone/lun2
- rsh netapp1 lun online /vol/toxin_clone/lun3
- rsh netapp1 lun online /vol/toxin_clone/lun4
- rsh netapp2 lun online /vol/toxin_clone/lun1
- rsh netapp2 lun online /vol/toxin_clone/lun2
- rsh netapp2 lun online /vol/toxin_clone/lun3
- rsh netapp2 lun online /vol/toxin_clone/lun4

For the AE named riot:

- rsh netapp1 lun online /vol/riot_clone/lun1
- rsh netapp1 lun online /vol/riot_clone/lun2
- rsh netapp1 lun online /vol/riot_clone/lun3
- rsh netapp1 lun online /vol/riot_clone/lun4
- rsh netapp2 lun online /vol/riot_clone/lun1
- rsh netapp2 lun online /vol/riot_clone/lun2
- rsh netapp2 lun online /vol/riot_clone/lun3
- rsh netapp2 lun online /vol/riot_clone/lun4

Once the LUN's are set to the online state, they can be mapped to FC initiators on the different AE systems. We used the same set of igroups and LUN IDs that we initially used to set up the historical archive.

For the AE named agony:

- rsh netapp1 lun map /vol/agonny_clone/lun1 agony_group 1
- rsh netapp1 lun map /vol/agonny_clone/lun2 agony_group 2
- rsh netapp1 lun map /vol/agonny_clone/lun3 agony_group 3
- rsh netapp1 lun map /vol/agonny_clone/lun4 agony_group 4

- `rsh netapp2 lun map /vol/agony_clone/lun1 agony_group 5`
- `rsh netapp2 lun map /vol/agony_clone/lun2 agony_group 6`
- `rsh netapp2 lun map /vol/agony_clone/lun3 agony_group 7`
- `rsh netapp2 lun map /vol/agony_clone/lun4 agony_group 8`

For the AE named toxin:

- `rsh netapp1 lun map /vol/toxin_clone/lun1 toxin_group 1`
- `rsh netapp1 lun map /vol/toxin_clone/lun2 toxin_group 2`
- `rsh netapp1 lun map /vol/toxin_clone/lun3 toxin_group 3`
- `rsh netapp1 lun map /vol/toxin_clone/lun4 toxin_group 4`
- `rsh netapp2 lun map /vol/toxin_clone/lun1 toxin_group 5`
- `rsh netapp2 lun map /vol/toxin_clone/lun2 toxin_group 6`
- `rsh netapp2 lun map /vol/toxin_clone/lun3 toxin_group 7`
- `rsh netapp2 lun map /vol/toxin_clone/lun4 toxin_group 8`

For the AE named riot:

- `rsh netapp1 lun map /vol/riot_clone/lun1 riot_group 1`
- `rsh netapp1 lun map /vol/riot_clone/lun2 riot_group 2`
- `rsh netapp1 lun map /vol/riot_clone/lun3 riot_group 3`
- `rsh netapp1 lun map /vol/riot_clone/lun4 riot_group 4`
- `rsh netapp2 lun map /vol/riot_clone/lun1 riot_group 5`
- `rsh netapp2 lun map /vol/riot_clone/lun2 riot_group 6`
- `rsh netapp2 lun map /vol/riot_clone/lun3 riot_group 7`
- `rsh netapp2 lun map /vol/riot_clone/lun4 riot_group 8`

After mapping the LUNs to the remaining AE systems, use the `diskpart` utility on each AE system to rescan the disks under Windows. For this we again used a file on each AE system called `rescan.txt` that simply contained the line `rescan`. When used in conjunction with the `-s` switch, `diskpart` opens the specified file and reads the set of commands to execute from the file. In this case, we're telling `diskpart` to conduct a rescan of all the disks:

- `psexec \\agony diskpart -s rescan.txt`
- `psexec \\toxin diskpart -s rescan.txt`
- `psexec \\riot diskpart -s rescan.txt`

After rescanning all of the disks on each of the remaining AE servers, all eight of the 1.5TB LUNs become available using the standard Windows disk management tools. Because the Snapshot

copy was taken after the disks were formatted, all that is necessary to access the historical archive is to assign a volume name and drive letter to each disk. In order for all four AE servers to function to see the same files, we assigned the same drive letters to each disk as were used when the historical archive was created.

9 APPENDIX B: MANAGE NIGHTLY UPDATE OF THE HISTORICAL ARCHIVE

Every night the Vhayu application updates the historical archive by adding information saved about the current day's events and presents the updated archive to each of the AE systems in use. This means that all of the AE systems are always using the same historical archive. Although there are a number of ways to accomplish this, this appendix contains a set of commands that could be the basis of a script created to perform this nightly maintenance task.

The commands to manage the NetApp storage assume the use of remote shell (`rsh`) to remotely execute commands on the FAS3170 and the use of the `psexec` utility to remotely execute Windows storage configuration commands on the Vhayu AE systems. Although the steps in this appendix are performed manually, it is a straightforward process to use either Windows batch files or PERL to script these commands for remote administration.

A number of the steps here are virtually identical to those used to create the initial historical archive and present copies of it to the AE servers described in Appendix A.

In reading this appendix, assume the following:

- All of the Vhayu AE systems are currently using an existing historical archive.
- The Vhayu AE systems are named `agony`, `riot`, `toxin`, and `crazy88`.
- `Crazy88` is the Vhayu AE system with read and write access to the historical archive.

The following general steps describe how to update and present an updated copy of the historical archive:

1. Set the FlexClone volumes containing the read-only copies of the historical archives to an offline state.
2. Rescan the disks under Windows on the AE systems with read-only access to the historical archive so that those LUNs disappear.
3. Delete the FlexClone volumes that provide the backing for the read-only historical archives.
4. Delete the Snapshot copy that provides the backing for the read-only historical archives.
5. Using the AE system with read and write access to the historical archive, update the historical archive with the current day's activity.
6. After updating, create a new Snapshot copy of the historical archive.
7. Create a new set of FlexClone volumes by using the new Snapshot copy and present them back out to the AE systems that have read-only access to the historical archive.
8. Rescan the disks on those AE systems that have read-only access to the historical archive.

The rest of this appendix describes the commands required to perform these steps. To configure and manage the storage, we used the FAS3170 command line interface. In all cases, commands were executed on both FAS3170 storage controllers, named `netapp1` and `netapp2`.

9.1 REMOVE HISTORICAL ARCHIVE LUNS FROM AE SERVERS

The first step in the process is to remove the LUNs that make up the existing historical archive from those AE systems that have read-only access. This is done by setting the FlexClone volumes providing the LUNs to these systems to offline by using the following commands:

- `rsh netapp1 vol offline /vol/agony_clone`

- `rsh netapp1 vol offline /vol/toxin_clone`
- `rsh netapp1 vol offline /vol/riot_clone`
- `rsh netapp2 vol offline /vol/agonny_clone`
- `rsh netapp2 vol offline /vol/toxin_clone`
- `rsh netapp2 vol offline /vol/riot_clone`

Once the FlexClone volumes are offline, use the `diskpart` utility on each of the AE systems to rescan the disks under Windows. Because the FlexClone volumes are offline on the FAS3170, all of the disks supported by these volumes disappear from Windows after the rescan is complete. For this we created file on each AE system called `rescan.txt` that simply contained the line `rescan`. When used in conjunction with the `-s` switch, `diskpart` opens the specified file and reads the set of commands to execute from the file. In this case, we're telling `diskpart` to conduct a rescan of all the disks:

- `psexec \\agonny diskpart -s rescan.txt`
- `psexec \\toxin diskpart -s rescan.txt`
- `psexec \\riot diskpart -s rescan.txt`

9.2 DELETE EXISTING HISTORICAL ARCHIVE DATA ON FAS3170

Next, delete the FlexClone volumes on the FAS3170 that provide the historical archive data to those AE systems with read-only access to the historical archive. These commands also delete all of the LUNs associated with the FlexClone volumes:

- `rsh netapp1 vol destroy /vol/agonny_clone -f`
- `rsh netapp1 vol destroy /vol/toxin_clone -f`
- `rsh netapp1 vol destroy /vol/riot_clone -f`
- `rsh netapp2 vol destroy /vol/agonny_clone -f`
- `rsh netapp2 vol destroy /vol/toxin_clone -f`
- `rsh netapp2 vol destroy /vol/riot_clone -f`

Because it is no longer needed, delete the Snapshot copy of the historical archive that provided the basis for the read-only copies of the historical archives:

- `rsh netapp1 snap delete vol1 archive_snap`
- `rsh netapp2 snap delete vol1 archive_snap`

9.3 CREATE AND PRESENT THE UPDATED HISTORICAL ARCHIVE

When the Snapshot copy is deleted, update the historical archive by adding the current day's data into the archive mounted read/write on the AE system `crazy88`. Then create a new set of FlexClone volumes for each of the remaining AE systems and present the copies of the archive to the AE systems.

You start this process by creating a new Snapshot copy of the newly updated FlexVol volume containing the updated historical archive on each of the FAS3170 storage controllers:

- `rsh netapp1 snap create voll updated_archive_snap`
- `rsh netapp2 snap create voll updated_archive_snap`

After creating the Snapshot copies, create a new set of FlexClone volumes using the Snapshot copy as the base. You need to create one for each AE system that requires read-only access to the archive.

Note: All of the FlexClone volumes are created by using the same Snapshot copy of the updated historical archive. This ensures that all of the AE systems are using the same historical archive.

- `rsh netapp1 vol clone create agony_clone -s none -b voll updated_archive_snap`
- `rsh netapp1 vol clone create toxin_clone -s none -b voll updated_archive_snap`
- `rsh netapp1 vol clone create riot_clone -s none -b voll updated_archive_snap`
- `rsh netapp2 vol clone create agony_clone -s none -b voll updated_archive_snap`
- `rsh netapp2 vol clone create toxin_clone -s none -b voll updated_archive_snap`
- `rsh netapp2 vol clone create riot_clone -s none -b voll updated_archive_snap`

Although the newly created FlexClone volumes contain the LUNs on which the historical archive are stored, the LUNs are not initially in an online state and are not initially mapped to any of the FC initiators on the AE systems. Each FlexClone volume contains four different LUNs. Each AE system is presented with a total of eight LUNs, so we set the LUNs for all AE systems to the online state.

For the AE named agony:

- `rsh netapp1 lun online /vol/agony_clone/lun1`
- `rsh netapp1 lun online /vol/agony_clone/lun2`
- `rsh netapp1 lun online /vol/agony_clone/lun3`
- `rsh netapp1 lun online /vol/agony_clone/lun4`
- `rsh netapp2 lun online /vol/agony_clone/lun1`
- `rsh netapp2 lun online /vol/agony_clone/lun2`
- `rsh netapp2 lun online /vol/agony_clone/lun3`
- `rsh netapp2 lun online /vol/agony_clone/lun4`

For the AE named toxin:

- `rsh netapp1 lun online /vol/toxin_clone/lun1`

- rsh netapp1 lun online /vol/toxin_clone/lun2
- rsh netapp1 lun online /vol/toxin_clone/lun3
- rsh netapp1 lun online /vol/toxin_clone/lun4
- rsh netapp2 lun online /vol/toxin_clone/lun1
- rsh netapp2 lun online /vol/toxin_clone/lun2
- rsh netapp2 lun online /vol/toxin_clone/lun3
- rsh netapp2 lun online /vol/toxin_clone/lun4

For the AE named riot:

- rsh netapp1 lun online /vol/riot_clone/lun1
- rsh netapp1 lun online /vol/riot_clone/lun2
- rsh netapp1 lun online /vol/riot_clone/lun3
- rsh netapp1 lun online /vol/riot_clone/lun4
- rsh netapp2 lun online /vol/riot_clone/lun1
- rsh netapp2 lun online /vol/riot_clone/lun2
- rsh netapp2 lun online /vol/riot_clone/lun3
- rsh netapp2 lun online /vol/riot_clone/lun4

Once the LUNs are set to the online state they can be mapped to FC initiators on the different AE systems. The following commands use the same set of igroups and LUN IDs that are used to initially set up the historical archive for the AE systems requiring read-only access.

For the AE named agony:

- rsh netapp1 lun map /vol/agony_clone/lun1 agony_group 1
- rsh netapp1 lun map /vol/agony_clone/lun2 agony_group 2
- rsh netapp1 lun map /vol/agony_clone/lun3 agony_group 3
- rsh netapp1 lun map /vol/agony_clone/lun4 agony_group 4
- rsh netapp2 lun map /vol/agony_clone/lun1 agony_group 5
- rsh netapp2 lun map /vol/agony_clone/lun2 agony_group 6
- rsh netapp2 lun map /vol/agony_clone/lun3 agony_group 7
- rsh netapp2 lun map /vol/agony_clone/lun4 agony_group 8

For the AE named toxin:

- rsh netapp1 lun map /vol/toxin_clone/lun1 toxin_group 1

- rsh netapp1 lun map /vol/toxin_clone/lun2 toxin_group 2
- rsh netapp1 lun map /vol/toxin_clone/lun3 toxin_group 3
- rsh netapp1 lun map /vol/toxin_clone/lun4 toxin_group 4
- rsh netapp2 lun map /vol/toxin_clone/lun1 toxin_group 5
- rsh netapp2 lun map /vol/toxin_clone/lun2 toxin_group 6
- rsh netapp2 lun map /vol/toxin_clone/lun3 toxin_group 7
- rsh netapp2 lun map /vol/toxin_clone/lun4 toxin_group 8

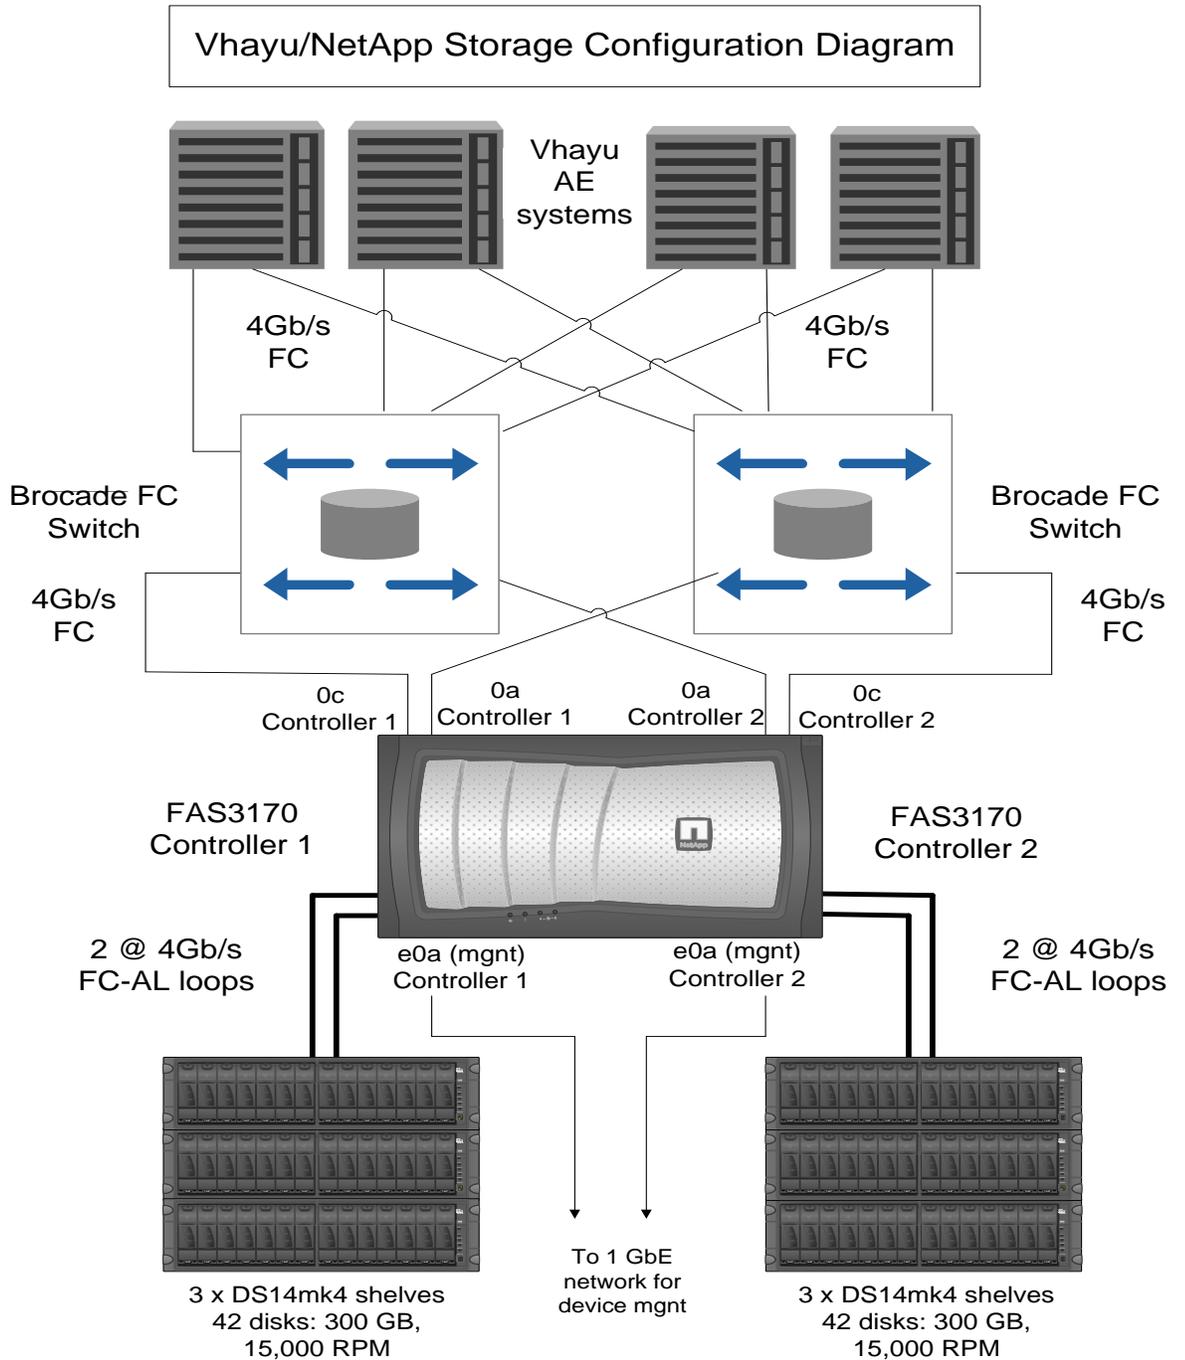
For the AE named riot:

- rsh netapp1 lun map /vol/riot_clone/lun1 riot_group 1
- rsh netapp1 lun map /vol/riot_clone/lun2 riot_group 2
- rsh netapp1 lun map /vol/riot_clone/lun3 riot_group 3
- rsh netapp1 lun map /vol/riot_clone/lun4 riot_group 4
- rsh netapp2 lun map /vol/riot_clone/lun1 riot_group 5
- rsh netapp2 lun map /vol/riot_clone/lun2 riot_group 6
- rsh netapp2 lun map /vol/riot_clone/lun3 riot_group 7
- rsh netapp2 lun map /vol/riot_clone/lun4 riot_group 8

After mapping the LUNs back to the correct AE systems, simply rescan the disks on the Windows hosts. All of the disks appear again with the updated historical archive information.

- psexec \\agony diskpart -s rescan.txt
- psexec \\toxin diskpart -s rescan.txt
- psexec \\riot diskpart -s rescan.txt

10 APPENDIX C. VHAYU AND NETAPP TEST CONFIGURATION NETWORK DIAGRAM



© 2009 NetApp and Vhayu Technologies Corporation. All rights reserved. Specifications subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, FlexVol, RAID-DP, and Snapshot are trademarks or registered trademarks of NetApp, Inc. Vhayu is a registered trademark and Vhayu Velocity is a trademark of Vhayu Technologies Corporation. Windows is a registered trademark of Microsoft Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

