LAB VALIDATION REPORT

Rainfinity File Virtualization Appliance

Transparent, Heterogeneous File Storage Mobility

By Tony Palmer
With Claude Bouffard

September, 2008

Table of Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC.

Introduction

The volume of unstructured data is exploding and managing these enormous datasets is becoming acutely painful for IT departments worldwide. Unstructured data is characterized by the files and objects stored outside of structured databases, most commonly on file servers or NAS devices. Rainfinity is a family of file virtualization solutions from EMC that customers utilize to optimize storage across their file sharing environment. ESG Lab validated the Rainfinity File Virtualization Appliance and its ability to non-disruptively move datasets between NAS devices and file servers in a heterogeneous, multi-protocol environment.

Background

ESG estimates that organizations will retain over 62,000 petabytes of unstructured data in digital archives over the next five years[1] in the ordinary course of doing business. To offer the reader perspective, 62,000 petabytes is equivalent to 62 million terabytes. The Library of Congress, with 29 million books in its collection at latest report, would consume only about 232 terabytes if every book were digitized with full formatting.[2] It would take more than 257,000 copies of every book in the Library to consume 62,000 petabytes.

FIGURE 1. PROJECTED GROWTH OF DATA

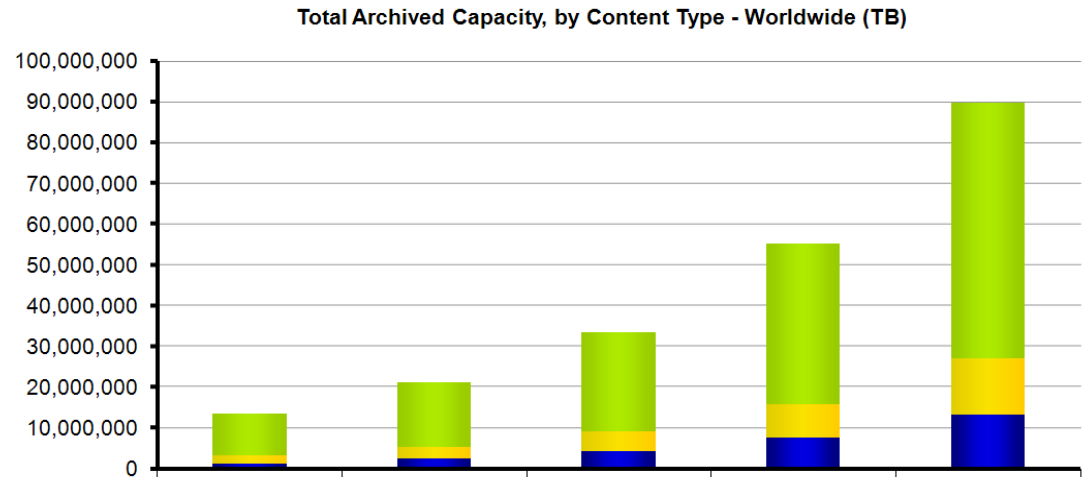

	2008	2009	2010	2011	2012
Unstructured	10,443,868	15,808,970	24,242,857	39,364,875	62,749,188
Database	1,837,780	2,991,043	4,823,578	8,110,447	13,639,302
E-mail	1,442,346	2,557,446	4,380,761	7,745,201	13,484,097

As the data volumes that organizations create and save continue to grow, stress on IT departments also increases. The tasks associated with managing newly generated and historical unstructured (file) data presents a daunting set of problems to already over-burdened IT staff. These issues are compounded by the challenges that IT is charged with, such as managing massive and growing datasets, moving data to new infrastructure and archiving aging information all while keeping data online and available.

ESG has found that data migrations are extremely common and are used to address numerous issues including file server and storage system upgrades, data center relocations, storage consolidation, end-of-lease equipment refresh and finally, but possibly most important, capacity and performance balancing. Traditional data migrations are inherently disruptive and often cause extended periods of file system unavailability due to the fact that the

[1] Source: ESG Research Report: *Digital Archiving: End-User Survey & Market Forecast 2006-2010*, January, 2006.
[2] Lyman, Peter and Hal R. Varian. "How Much Information", 2003, Retrieved from http://www.sims.berkeley.edu/how-much-info-2003 on 6/1/2008

process requires both moving data from one storage system to another and redirecting users and applications. An ESG survey[3] of over 550 storage administrators in North America and Europe indicates that the majority of migrations are conducted on weekends and application downtime is the biggest overall issue. The majority of users (75%) experienced problems with data migrations, with 58% reporting extended or unexpected downtime, and 36% experiencing application performance issues. Many users consider migrating data to be a complex process; over 72% of respondents take more than two weeks to plan an implementation and over 40% of the migrations require more than 5 people to complete.

EMC Rainfinity File Virtualization Appliance (FVA)

Rainfinity FVA is the newest addition to EMC's family of solutions that simplify, optimize, and scale the management of networked attached file systems. Purpose-built Rainfinity FVA devices provide a centralized platform for the management and on-line migration of data (including active, open files) between general Windows and UNIX file servers and/or Network Attached Storage (NAS) devices as shown in Figure 2. Integrated FVA clustering provides high availability.

While Rainfinity has always supported both Unix (NFS) and Windows (CIFS) environments, migrations of simultaneous access multi-protocol environments typically occurred in two stages. Organizations would first move data with Rainfinity using whichever protocol was dominant, and then attributes and permissions for the second protocol were applied manually. Today Rainfinity FVA provides a truly integrated multi-protocol solution for the on-line migration of unstructured file data.

FIGURE 2. TYPICAL RAINFINITY FVA IMPLEMENTATION

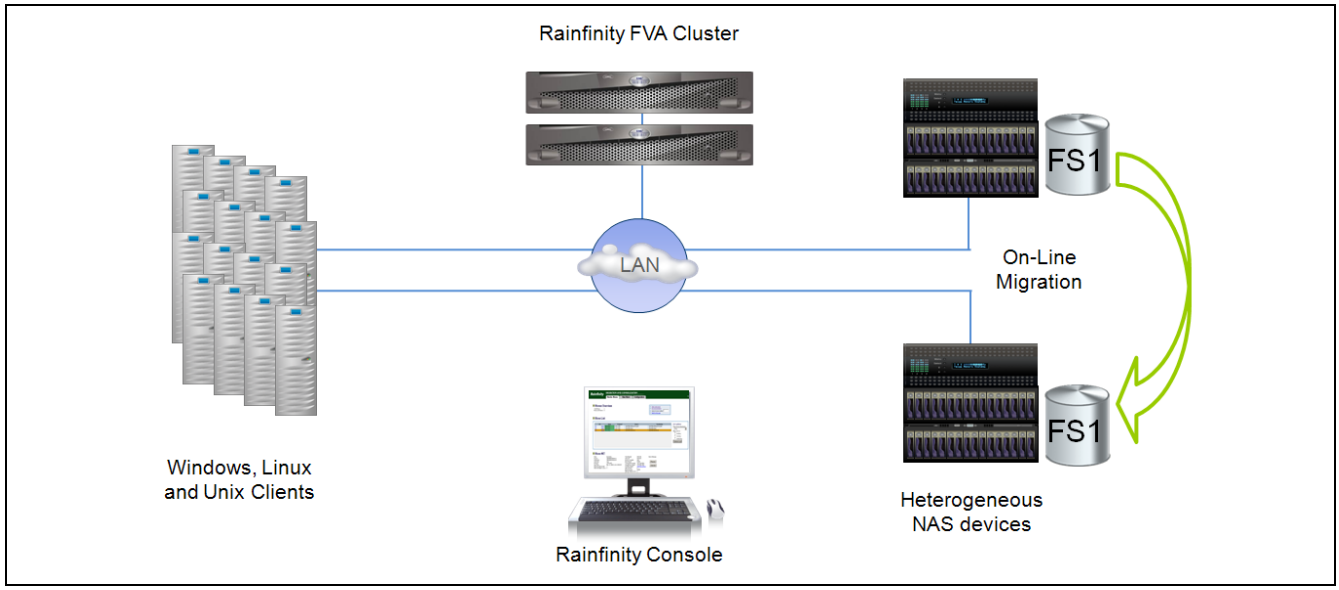

Rainfinity FVA provides a number of benefits to IT administrators and end-users including:

- Rainfinity FVA requires no agent software installation
- File systems can be moved online with no impact on end-user productivity
- End-users perceive no difference in file share location or access methods
- Data coherency is maintained when sharing multi-protocol, read-write access to the same file(s)
- Multi-protocol Windows, Unix and Linux attributes are preserved
- File-system metadata continues to reside with the native file systems
- Existing snapshot, anti-virus, backup, and advanced NAS functionality are retained

[3] Source: ESG Research Report: *Data Migration*, April 2006

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of EMC Rainfinity File Virtualization Appliance (FVA) at an EMC facility in Santa Clara, California. ESG Lab tested Rainfinity FVA's ability to perform on-line, non-disruptive data migrations of multi-protocol file systems between heterogeneous NAS systems. Also of interest were the preservation of permissions and attributes for both NFS and CIFS users and global namespace integration.

Getting Started

Rainfinity FVA provides a scalable, transparent file-protocol switching capability to migrate online datasets and leverages industry-standard global namespaces. There are many advantages to this approach: It limits risk and performance concerns by processing only files being migrated in the data path while normal data access remains outside the path, but it also leverages the continuing investments being made by customers.

Rainfinity FVA comes installed with data migration software in a purpose-built appliance. Rainfinity FVA provides complete online data migration transparency and does not require mount-point changes or the deployment of agents on clients or servers while supporting namespaces such as Microsoft DFS and Automount. Rainfinity FVA has several functions including copying and migrating files while preserving multi-protocol attributes, mapping file names and permissions while resolving conflicts, and offering the ability to synchronize simultaneous read and write access and redirect client access.

Figure 3 illustrates the four phases of multi-protocol online data migration using Rainfinity FVA. Rainfinity FVA automates these phases for the IT administrator and performs them all transparently to clients.

FIGURE 3. RAINFINITY NON-DISRUPTIVE MIGRATION PROCESS

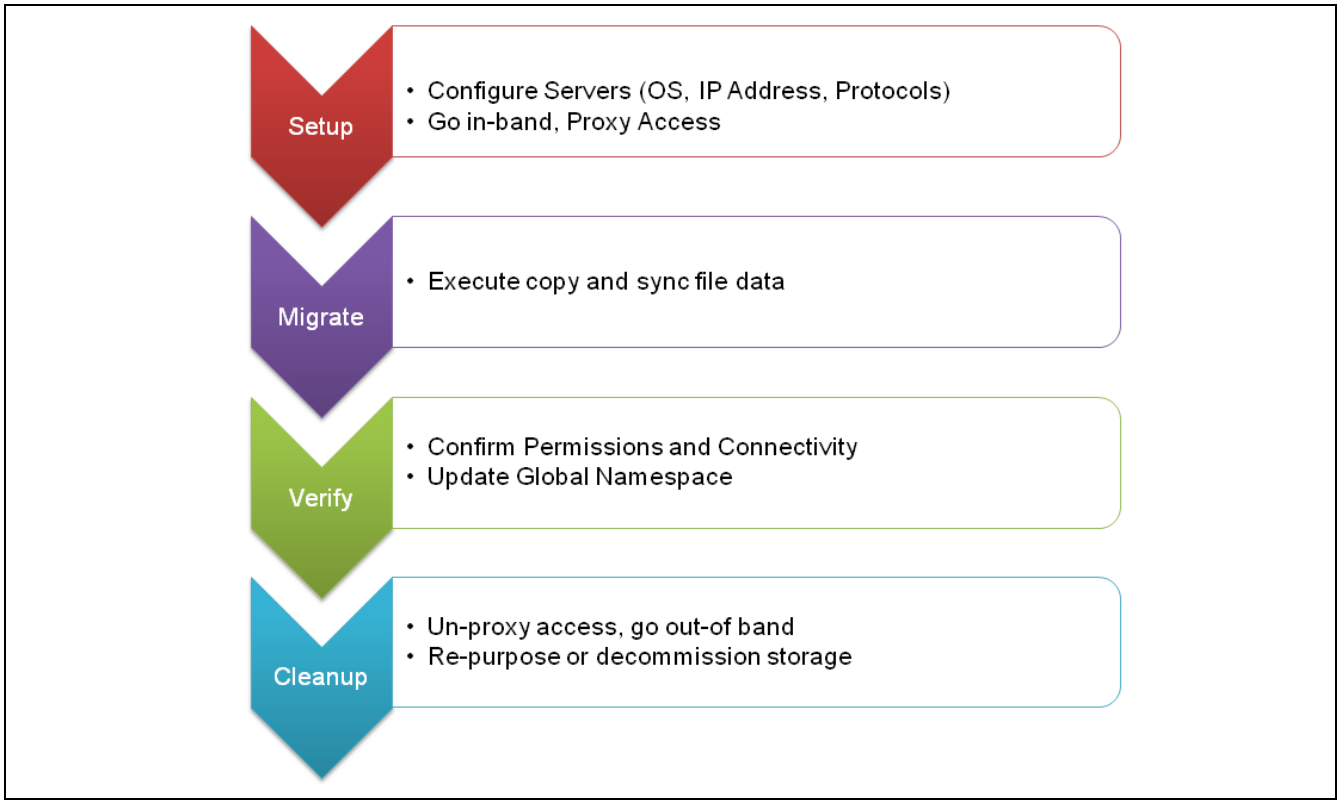

When discussing file virtualization, the terms 'in-band' and 'out-of-band' are used to describe different approaches to virtualization. A completely in-band system is installed between the client and NAS/file servers. Clients must mount the file systems through the in-band device 100% of the time and all I/O must pass through the in-band system. Out-of-band solutions manage metadata and file locking but I/O always travels directly to the NAS/file server. The advantage of in-band is control of IO and masking the back end devices, enabling data movement which is transparent to users. Out-of-band cannot offer this level of virtualization but will not become a bottleneck as no data I/O traverses the out-of-band device.

Rainfinity FVA's architecture is designed to switch between in-band and out-of-band operation as needed. The appliance is out-of-band most of the time and data flows between client systems and the NAS storage directly. Rainfinity FVA sits outside the data path until a migration is required and then switches to in-band operation. FVA moves data to the new location while the file data remains online and accessible to users. This approach is extremely scalable as it has no impact on shares, exports or folders that are not actively being migrated and is completely transparent to users and applications.

FIGURE 4. RAINFINITY OPERATION AND PERFORMANCE

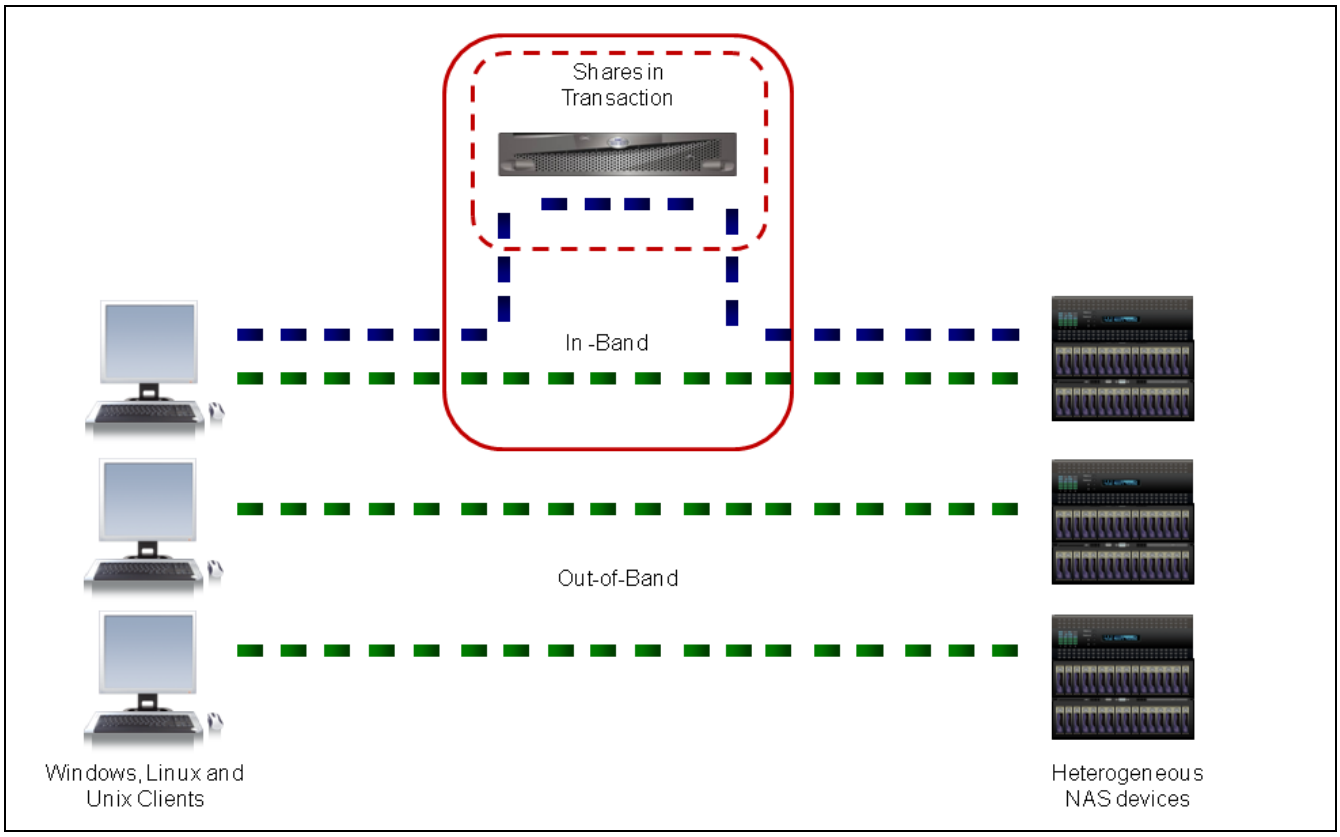

In Figure 4, a set of NAS servers are shown on the right and the traffic to and from the file systems they own and share represented by the colored blocks. Clients, on the left, read from and write to files on any of the file systems. The data streams in green represent the traffic not being migrated by Rainfinity FVA. This data moves unimpeded through the LAN as Rainfinity FVA is either completely out of band or bridging the traffic to those NAS servers. The NAS server with a file system being migrated has all of its traffic bridged by Rainfinity FVA. FVA is in-band for this server. Only the traffic belonging to the specific share being migrated or 'in transaction,' represented in blue, is processed by FVA. The rest of the traffic for that server, in green, is bridged directly to the client with no user space processing overhead. In this way, Rainfinity minimizes the impact migrations have on

an overall environment and maximizes the performance a single appliance can provide to the datasets being migrated.

Figure 5 presents the ESG Lab test bed at the EMC facility; data was migrated from a NetApp FAS840 system to an EMC Celerra NS40 using Rainfinity FVA for data movement. A complete description of the test environment is listed in the Appendix at the end of this report.

FIGURE 5. ESG LAB TEST BED

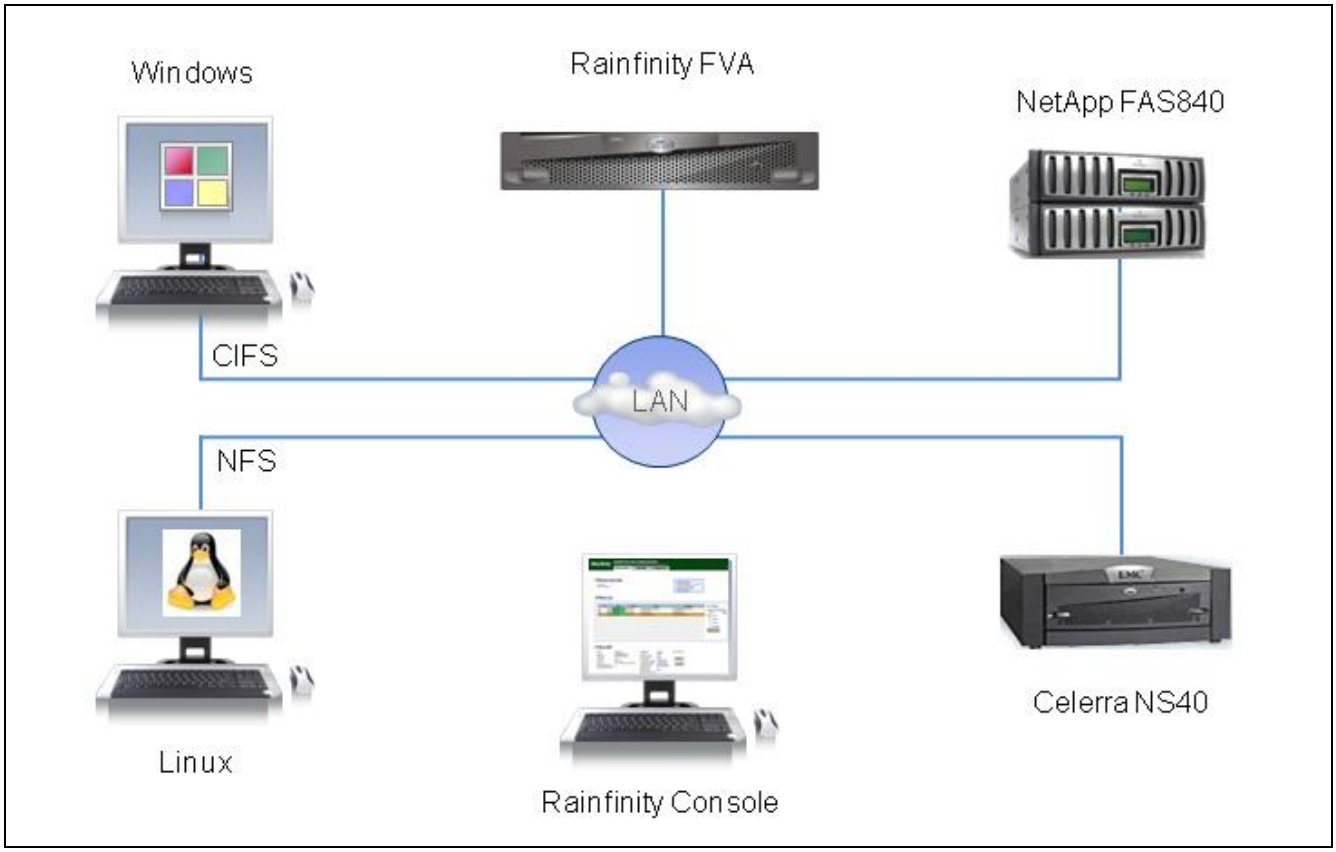

Why This Matters

Data center managers routinely move file data to a new storage system for a number of reasons including technology updates, lease refreshes, data center moves, data center consolidations, tiered storage strategies and performance/capacity remediation efforts. Despite all the improvements made to the process over the years, most data migrations are still manual, require an outage window, and are a labor intensive process.

The Rainfinity File Virtualization Appliance is an easy-to-deploy system requiring no additional client or server software and no mount point changes to existing clients. ESG Lab found that installing and configuring a purpose-built, factory-installed Rainfinity FVA is straightforward and intuitive.

Multi-Protocol, Non-Disruptive Data Migration

Multi-protocol online data migration is the main focus of Rainfinity File Virtualization Appliance. Rainfinity FVA handles each step involved with moving data transparently and automatically:

- **Checking the new destination** - Traditional migration methods require an administrator to manually check the new destination for user accessibility, file permissions, etc.
- **Transferring the data** - Manual methods like robocopy or rsync, require multiple passes with an offline pass to guarantee all data is moved to the new location before cutover.
- **Translating security IDs** - When local groups and users are migrated in the traditional manner administrators must understand Security ID (SID) translation rules and re-apply ACL's to files after data has been migrated.
- **Relocating file attributes** - Robocopy or rsync only move the attributes associated with their native protocol (CIFS and NFS respectively). A multi-protocol move requires a manual post move process to apply secondary permissions.
- **Updating the namespace** - Traditional methods require an administrator to manually edit DFS or Automount map configurations before cutover to ensure clients can access data in the new location.
- **Redirecting end-user access until all of the re-mappings are complete** - Traditional methods do not allow for this at all. Migrations require an offline window, either at the beginning or end of the migration.

Rainfinity FVA can also be used to move data between different tiers of storage. Rainfinity FVA performs data movement between heterogeneous tiers of storage transparently and non-disruptively to users or applications.

Setup

Rainfinity FVA arrives with the appliance software installed and ready to be configured. FVA can be managed using a graphical web-based interface or a command line.

ESG Lab Testing

Validation of Rainfinity FVA started with a pre-configured system attached to the network. ESG Lab configured the local NAS servers from FVA console (Local File Server Properties) as shown in Figure 6.

FIGURE 6. CONFIGURING FILE SERVER PROPERTIES

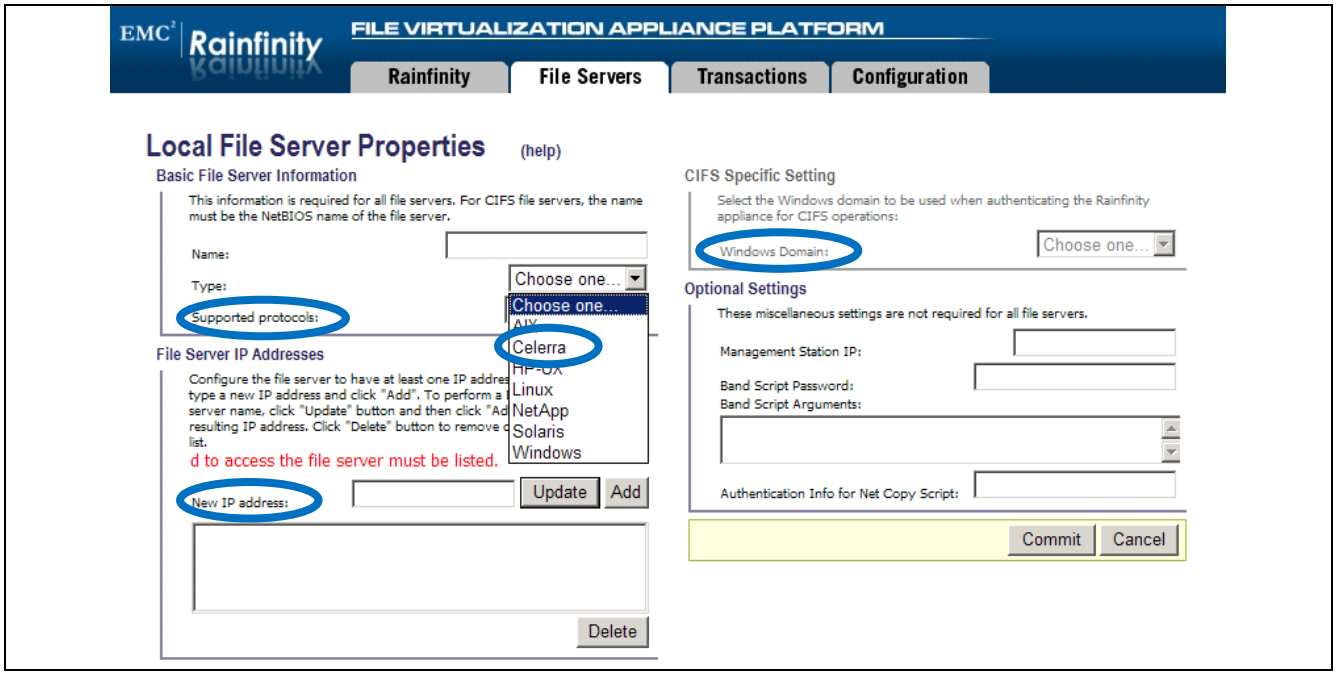

ESG Lab was provided with basic NAS server information to input. This is where an administrator specifies the type of server, protocols running on the server, IP addresses and user authentication information. Once the source and target file servers were configured, ESG Lab was able to confirm the source and target shares for the migration and review the configuration as shown in Figure 7. The system was ready for data movement less than five minutes after first logging in.

FIGURE 7. FILE SERVER OVERVIEW

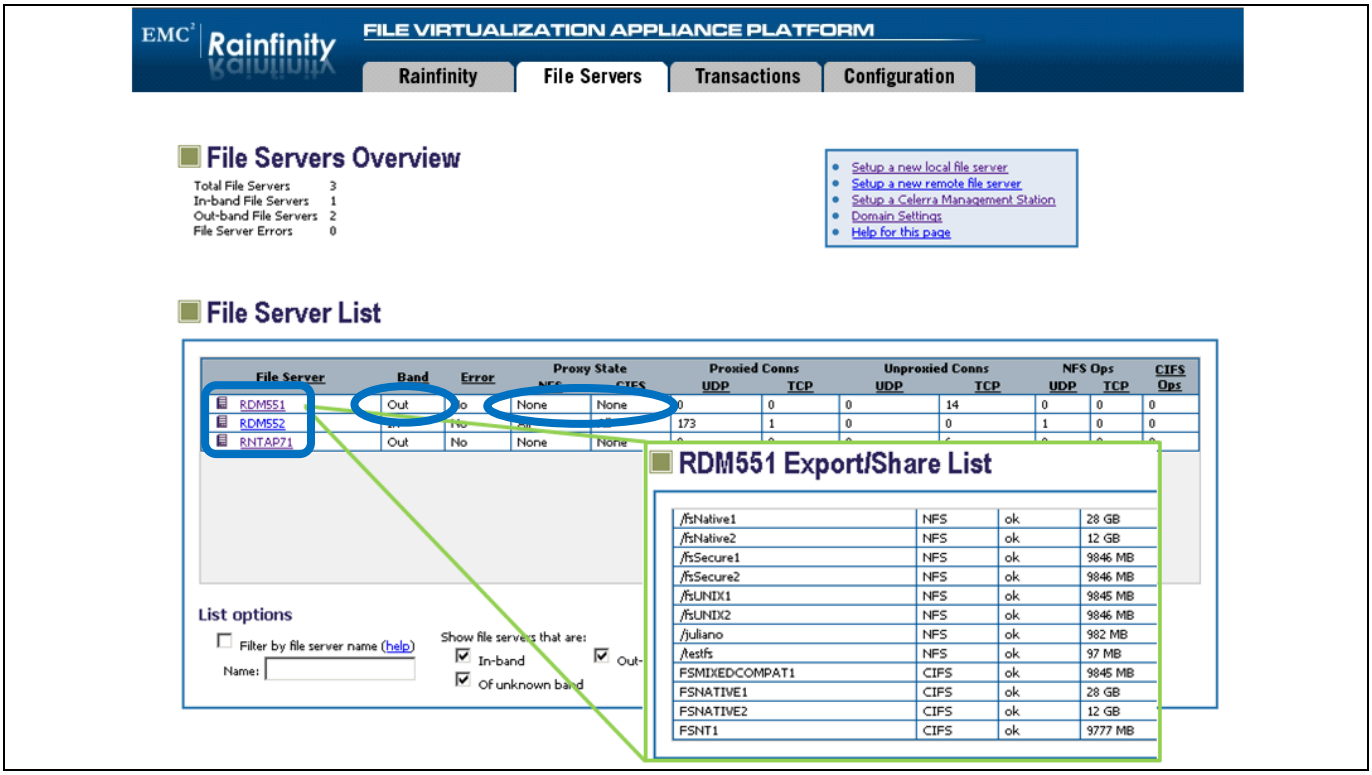

As shown in the diagram, three file servers were configured in the ESG Lab test bed. Details associated with the shares and exports owned by file server RDM551 are shown in greater detail. Also displayed on the screen are current bandedness (In-band vs. Out-of-band), proxy state of NFS exports & CIFS shares and any errors for each file server.

To prepare a file server for data movement, ESG Lab first placed Rainfinity FVA in-band for server RNTAP71. There are a number of ways to accomplish the in-band transition, but the easiest is to leverage the VLAN tagging capabilities of the NAS Server. Rainfinity can trigger a VLAN change on the file server to move it in-band. Once this change occurs, FVA becomes a transparent bridge for all traffic to the file server. Clients experienced no disruption to file access when this step was executed.

The next step was to create proxy connections to the file server before starting the data move. Without a means to gracefully proxy client connections, CIFS clients would be disconnected from their shares and have to reconnect at the start of the move. This is not an issue with NFS clients because NFS is a stateless protocol and clients will automatically reconnect. To eliminate potential disruption associated with CIFS client reconnections, an administrator would choose 'Only Proxy newly established connections.' Each time a client creates a new connection to the file server, Rainfinity will proxy the connection. Over time, all connections will be proxied—a requirement before the next phase can begin. Figure 8 displays proxy settings property screen for the CIFS protocol.

FIGURE 8. CIFS PROXY SETTINGS

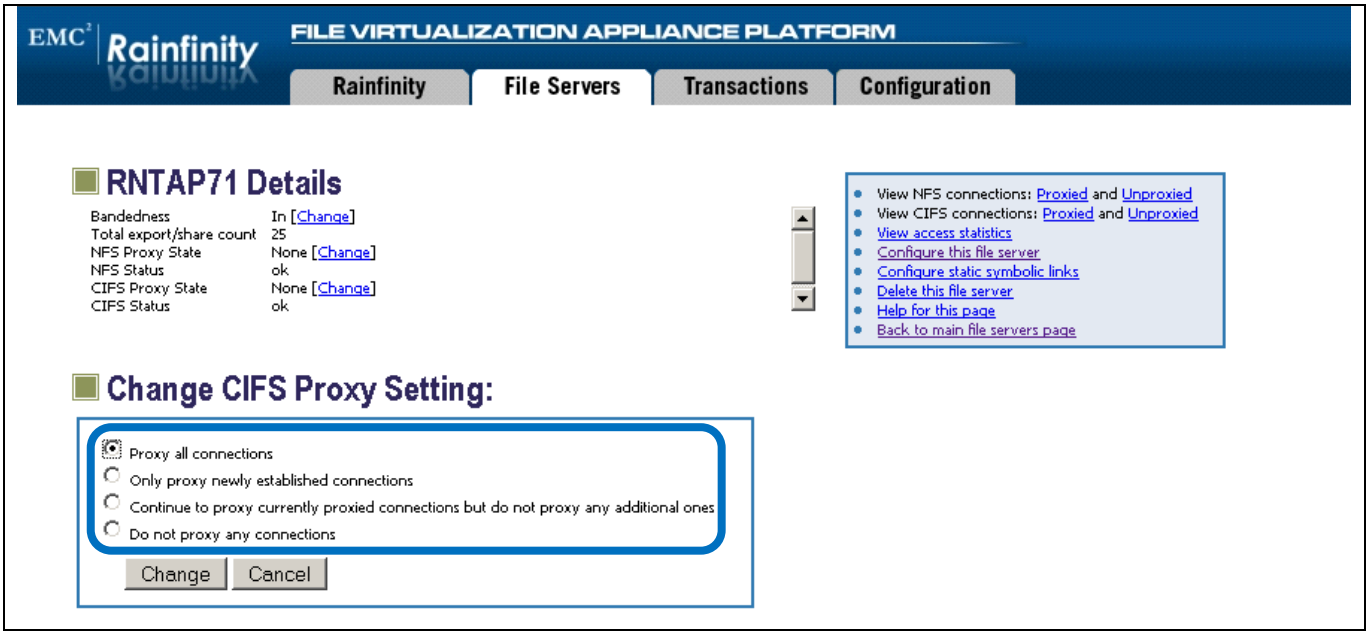

For the purposes of this test, ESG Lab selected 'Proxy all connections'. Once all connections to the source were proxied, ESG Lab started playing a video file located in the source share using both the Windows Client over CIFS and the Linux Client over NFS, as seen in Figure 9.

FIGURE 9. TESTING ONLINE MIGRATION

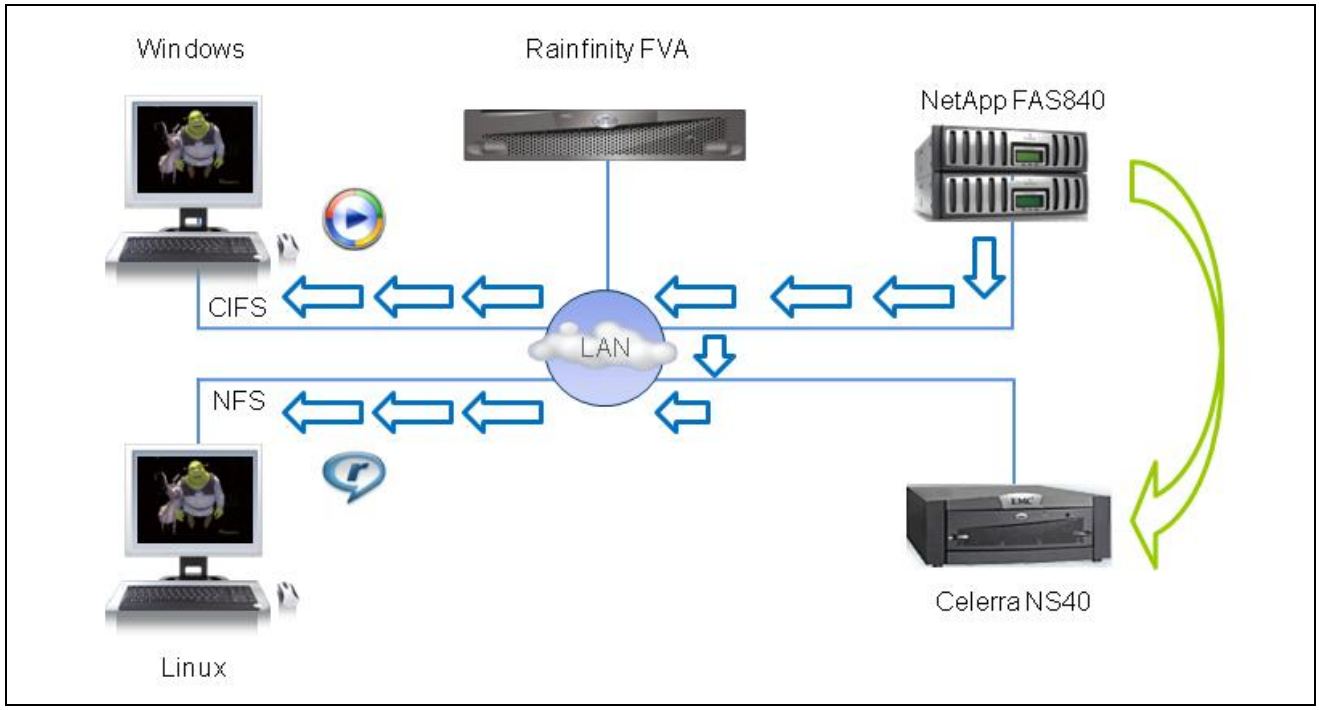

ESG Lab next proceeded to the 'Start a New Move' wizard. In this scenario, ESG Lab needed to configure a data migration to move a mixed protocol share containing the video file from the NetApp FAS840 system to an EMC Celerra NS40. There were only four steps required to configure a new move:

- Select Protocols – NFS, CIFS or both
- Select Source Server and path to data to be moved.
- Select Destination and path where data should be moved to
- Set Namespace Update options – This includes indicating whether FVA should create shares on the destination to match the source.

ESG Lab stepped through the entire process, from the initial configuration to executing the move, in less than five minutes.

Why This Matters

Traditional file system migration methods have evolved over the last several years. The simplest method, using protocol specific copy tools (rsync for NFS or robocopy for CIFS) to move data must be performed offline to ensure data consistency at the target. More advanced methods that involve modifying mount points or inserting software agents into the network have historically been complicated to set up and use. As a result, setting up a traditional file system migration can still be a complex, labor intensive, intrusive and risky affair.

ESG Lab has confirmed that Rainfinity FVA is a purpose-built plug and play appliance that is quick and easy to setup. Simply plug it in to a standard Ethernet network and its ready to go. Rainfinity FVA requires no software installation and can be used to minimize disruption to end users and applications.

Migrate

Data migration is the process of transferring data from one location to another. The new location may be within the same or between different storage systems in the same data center or to a remote location as part of a data center move or consolidation. Data migrations are used to support storage or system upgrades, to implement new NAS solutions or retire old ones. Data migrations may also be utilized to provide capacity or performance load balancing as well as consolidations to reduce the number of file servers and NAS devices to reduce power/cooling and floor space consumption. Rainfinity provides any-to-any data movement between NAS systems and/or file servers. Rainfinity FVA synchronizes the source and destination locations without affecting active data processing. Most traditionally manual steps of the migration process can be automated using Rainfinity FVA.

ESG Lab Testing
Next ESG Lab started the move and observed progress of the migration using the Monitor Moves tab in the Rainfinity console. A move automatically transitions through several states:

- **Pending –** Waiting for the transaction to start executing.
- **Preparing –** Counting files on the source and checking memory and disk space allocation.
- **Scanning** – Checking files in the source and destination for naming consistency.
- **Executing** - Baseline copy of data from the source to the destination.
- **Syncing** – Synchronizing changes between the destination and source.
- **Redirecting** – Redirecting users mounted to the source to the destination. Once a transaction is in the redirecting state, the destination directory is the authoritative copy.

FIGURE 10. EXECUTING THE MOVE

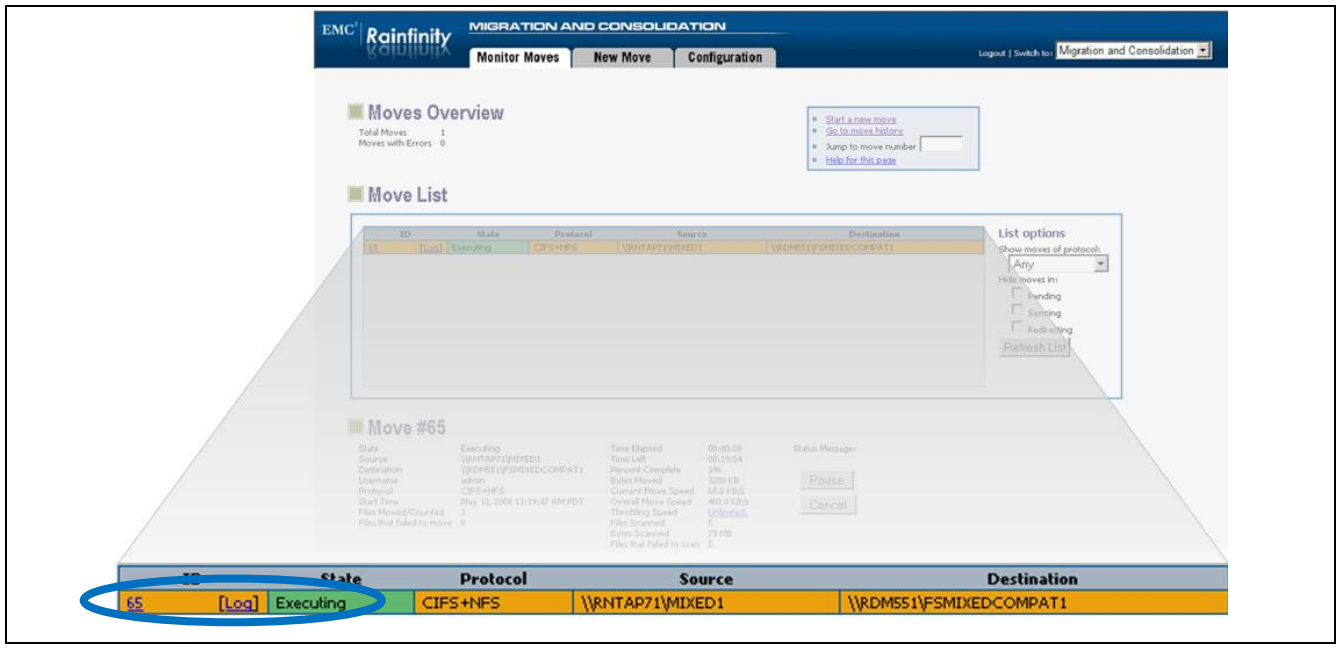

Figure 10 shows the move in the Executing phase, when data is being copied from the source to the destination. Once the baseline copy is completed, the move transitions through the Syncing phase, where the source and destination are brought 100% into sync, and then finally to Redirecting. Once in the Redirecting phase, the Destination becomes the authoritative copy and the migration can be completed. Throughout this entire process, ESG Lab watched the video on both the windows and Linux clients without impact or disruption to playback throughout the migration.

Figure 11 shows the move in the Redirecting phase. At this point any client still mounted to the original source share/export will be transparently redirected to the destination by Rainfinity FVA.

FIGURE 11. DATA MOVE STATUS

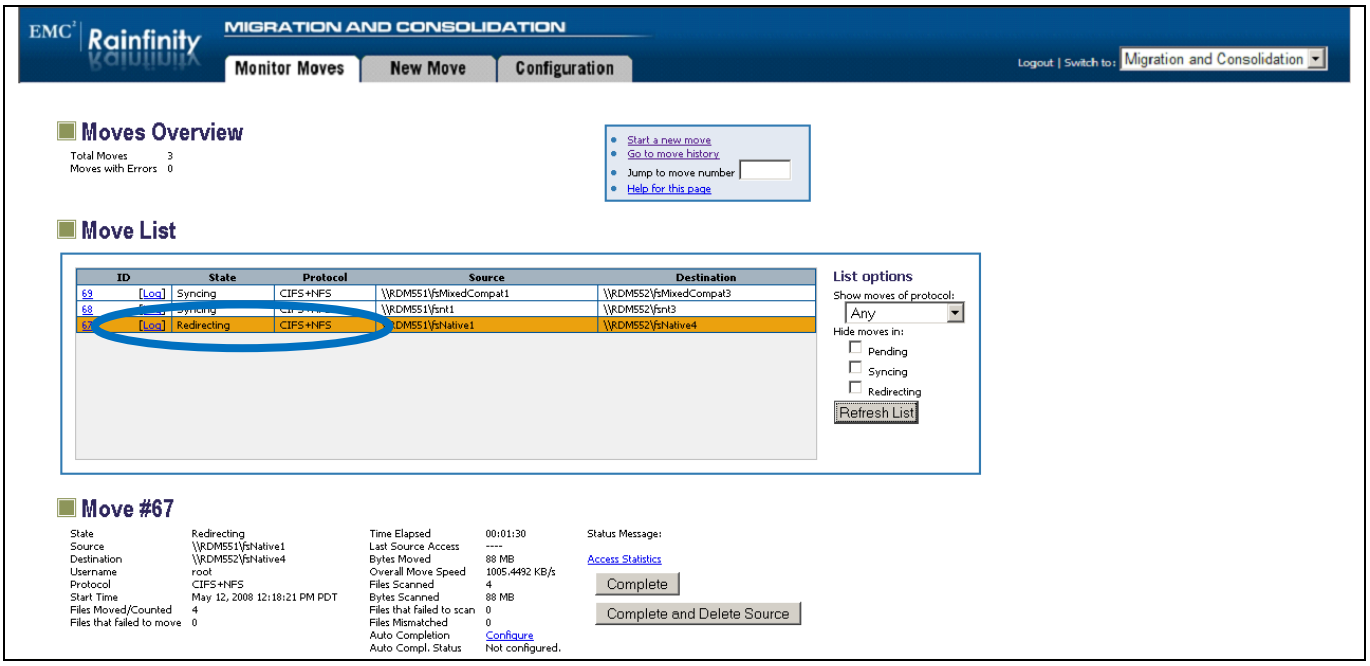

Why This Matters

Recent ESG research indicates that a majority of IT organizations (61%) can't tolerate more than four hours of downtime for their mission critical applications.[4] Five percent indicate that zero downtime has been mandated for these datasets. Given the exploding volumes of unstructured file data residing on network attached storage systems and the always-on 24x7 nature of modern organizations, traditional migration methods can't meet these strict service level agreements. ESG Lab has verified that Rainfinity FVA uses an automated, low-impact approach for on-line data migrations with no perceivable impact to end users and applications.

[4] Source: ESG Research Report, *Data Protection Survey*, January 2007, N=398

Verify

FVA facilitates both pre-migration and pre-completion safety checks throughout the data move process. Some checks are specific to either CIFS or NFS, while others apply to both protocols. As part of the verification process prior to beginning and completing the data migration, the following steps were performed by ESG Lab:

- Confirmed Source & Destination IP Addresses
- Checked Proxy State
- Verified Permissions & Connectivity before and after the move to ensure they were not altered.
- Checked Last Source & Destination Access before completing the move.

ESG Lab Testing
Figure 12 displays the access statistics for the move ESG Lab configured through Rainfinity FVA. On this screen, Source IP addresses were confirmed and the last access of both the source and destination were checked to verify that all clients were accessing the share using the destination's mount point.

FIGURE 12. FILE SYSTEM ACCESS STATISTICS

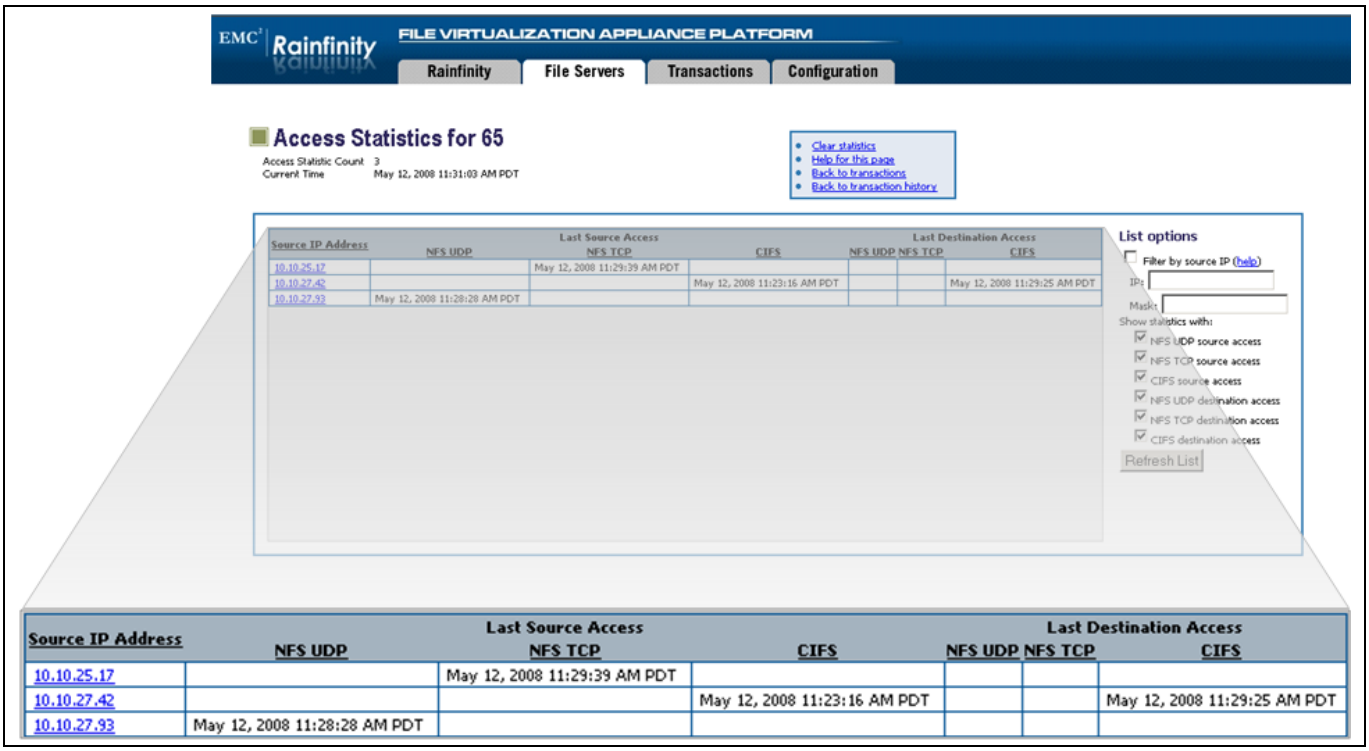

Source IP Address	Last Source Access			Last Destination Access		
	NFS UDP	NFS TCP	CIFS	NFS UDP	NFS TCP	CIFS
10.10.25.17		May 12, 2008 11:29:39 AM PDT				
10.10.27.42			May 12, 2008 11:23:16 AM PDT			May 12, 2008 11:29:25 AM PDT
10.10.27.93	May 12, 2008 11:28:28 AM PDT					

A date stamp is displayed for each NFS and CIFS file share. This screen is used to confirm that all clients have remounted or reconnected to the new destination. ESG next verified the proxy status for the file systems being moved as displayed in Figure 13.

FIGURE 13. PROXY STATUS

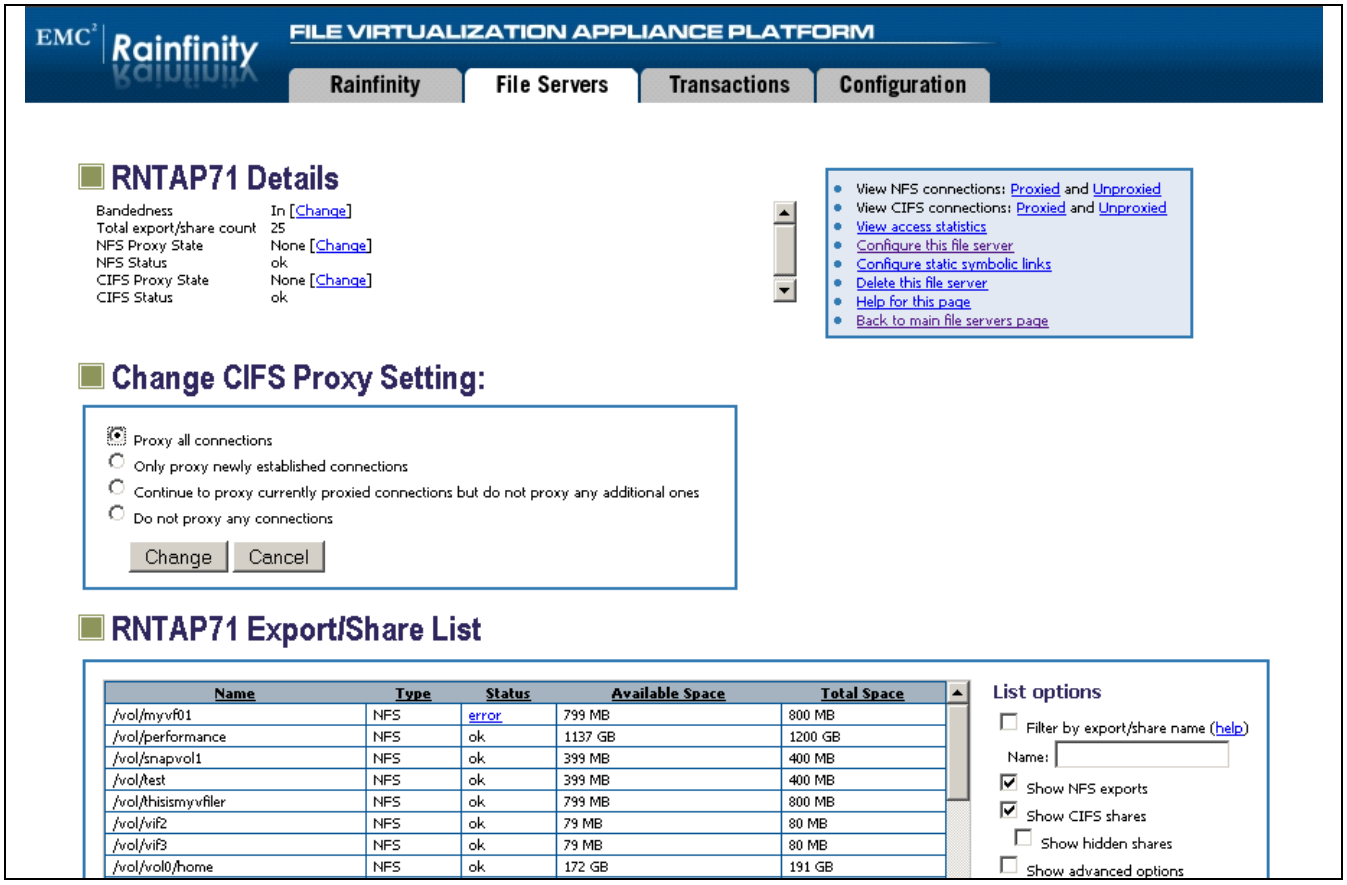

From the Rainfinity console, ESG Lab was able to check bandedness status, review the current proxy status for each export/share and change proxy settings. Once all users are confirmed to have mounted the new destination share, it can be un-proxied and the server can be taken out-of-band.

Why This Matters

Planning and verifying a data migration using traditional manual methods is complicated and error-prone. ESG research indicates that the majority of data migration processes exceed originally planned estimates.[5] Staff time is often underestimated, with 69% of respondents exceeding estimates during a migration. Additionally, 21% stated that they always exceeded planned staff time. These issues are exacerbated in customer environments where file systems are accessed by both CIFS and NFS clients because the two protocols handle things like file permissions, links, and locking very differently.

ESG Lab confirmed that Rainfinity FVA handles the entire process of transparently moving data in this type of complex environment. Rainfinity FVA provides a centralized platform which automates routine file system migration verification tasks reducing the staff time, risks and cost associated with networked file system migrations.

[5] Source: ESG Research Report, *Data Migrations*, April 2006, N=550

Cleanup

Cleanup is the set of tasks performed at the end of a data migration including the decommissioning of the old storage system. Poor cleanup not only makes it hard to figure out what file system you are using —meaning that users are not able to access files—but also can lead to inadvertent access to file systems that are no longer available resulting in data loss or data unavailable.

As part of the Rainfinity FVA cleanup phase, the following tasks are performed:

- Redirect Data Access
- Update Global Namespace
- Un-Proxy Access
- Change Bandedness (Out-of-band)

As Figure 14 illustrates, clients will be able to continue read and write access to the data set without disruption. Once the file systems have been synced at the end of a migration, Rainfinity redirects IO of clients still mounted to the source. In this diagram, the top client is still mounted to the source file system while the bottom client has already mounted the new target file system. Through Rainfinity's redirection of the top client, both systems are accessing data on the new target. This is especially useful when storage is being de-commissioned and time is critical. The old storage can be removed with zero impact to existing systems that cannot be immediately re-mounted to the new system.

FIGURE 14. POST DATA MIGRATION CLEANUP

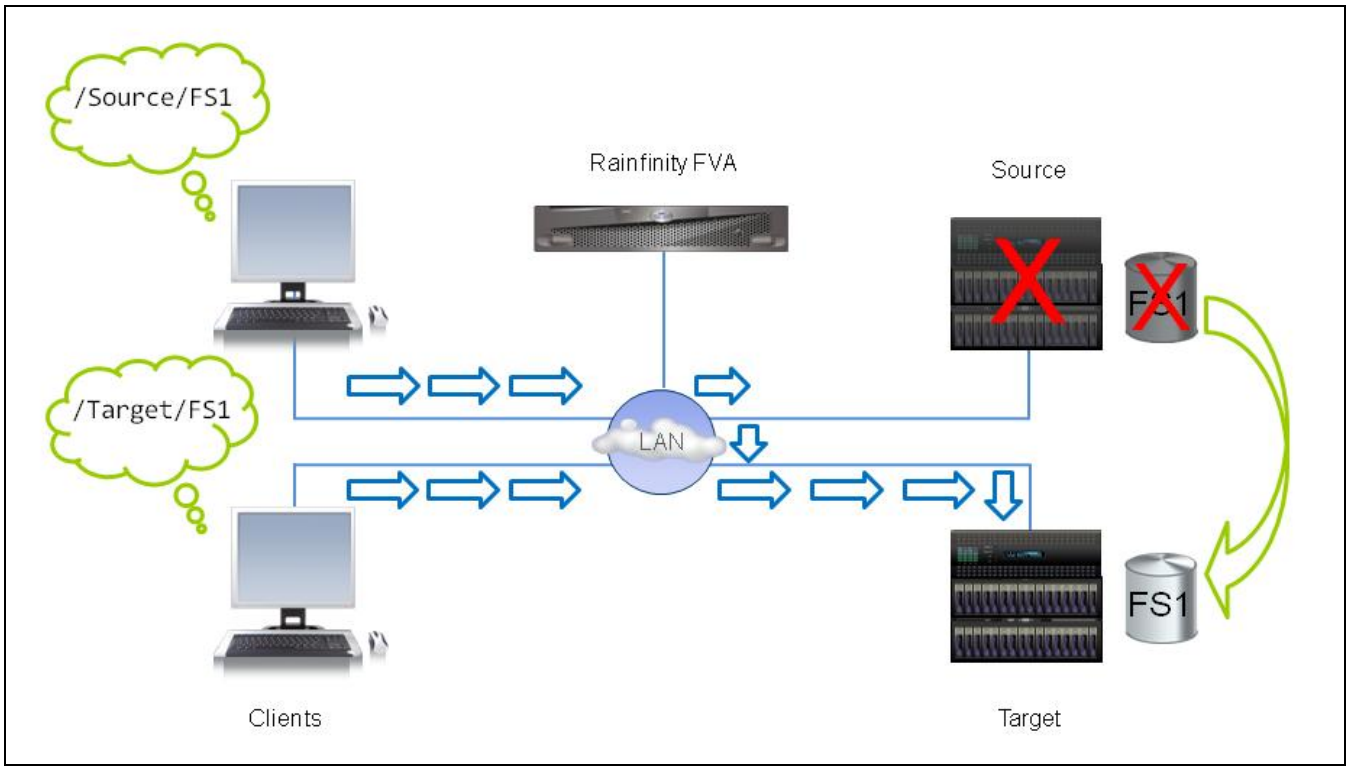

To keep the entire migration process transparent after the transaction is put into the redirect phase, Rainfinity relies on Windows and UNIX facilities (such as DFS, the automount utility and startup scripts to remount shares that are re-exported) to establish new storage locations and redirect users without interruption. The original storage can then be freed for other tasks or decommissioned.

ESG Lab Testing

While the migration was in the redirecting state, ESG Lab tested redirection and multi-protocol access. As seen in Figure 15, a file named 'redirect' was created using the Linux client, which was still mounted to the source file server (RNTAP71) over NFS, then the text 'redirected write' was added to the file.

FIGURE 15. REEDIRECTED WRITES

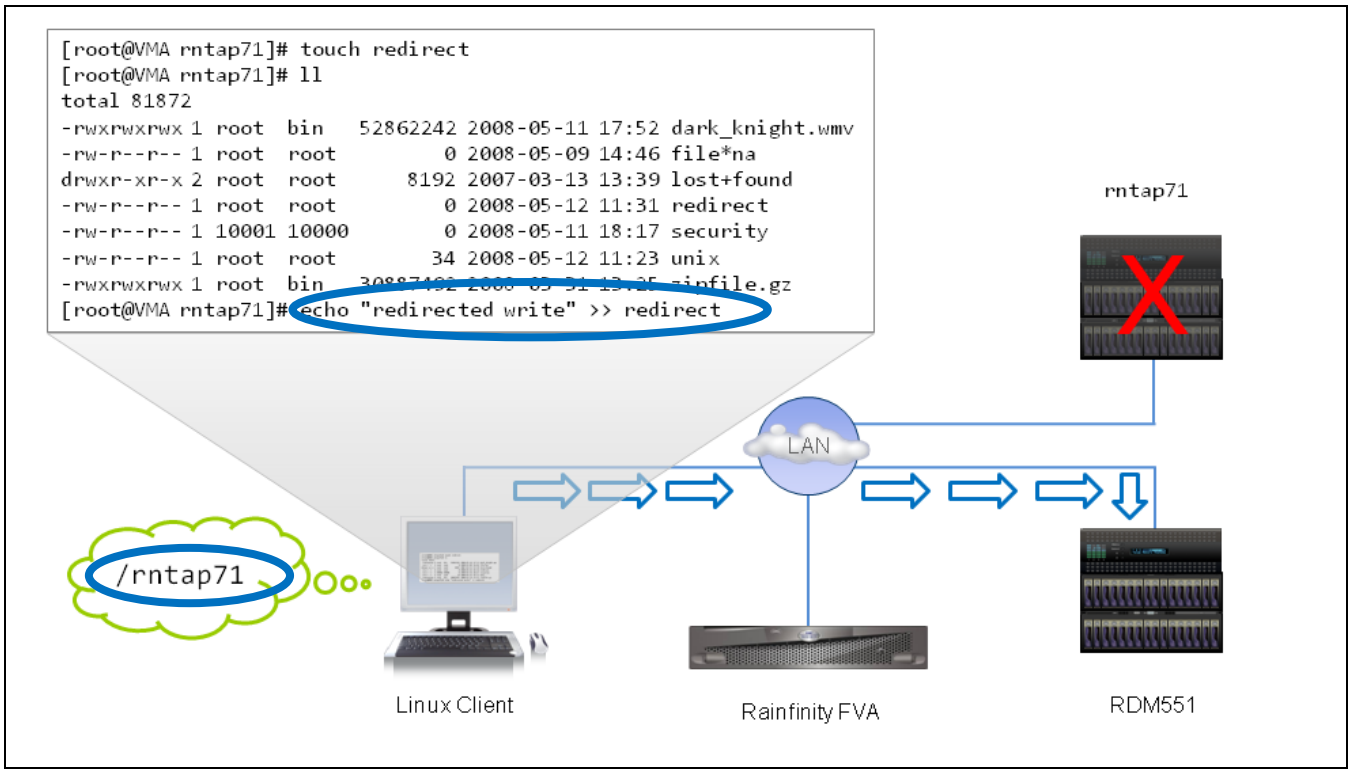

Next, ESG Lab opened the same file using the Windows client machine which had been reconnected to the new target share (on Celerra RDM551) with the same user credentials as used on the Linux client.

As seen in Figure 16, the file contained the text 'redirected write' that was inserted in the previous step. ESG Lab inserted additional text and saved the file. A quick examination of the file from the Linux client verified that the data had been written correctly and the Linux Client was able to access it while still mounted to the original source.

FIGURE 16. MULTIPROTOCOL ACCESS

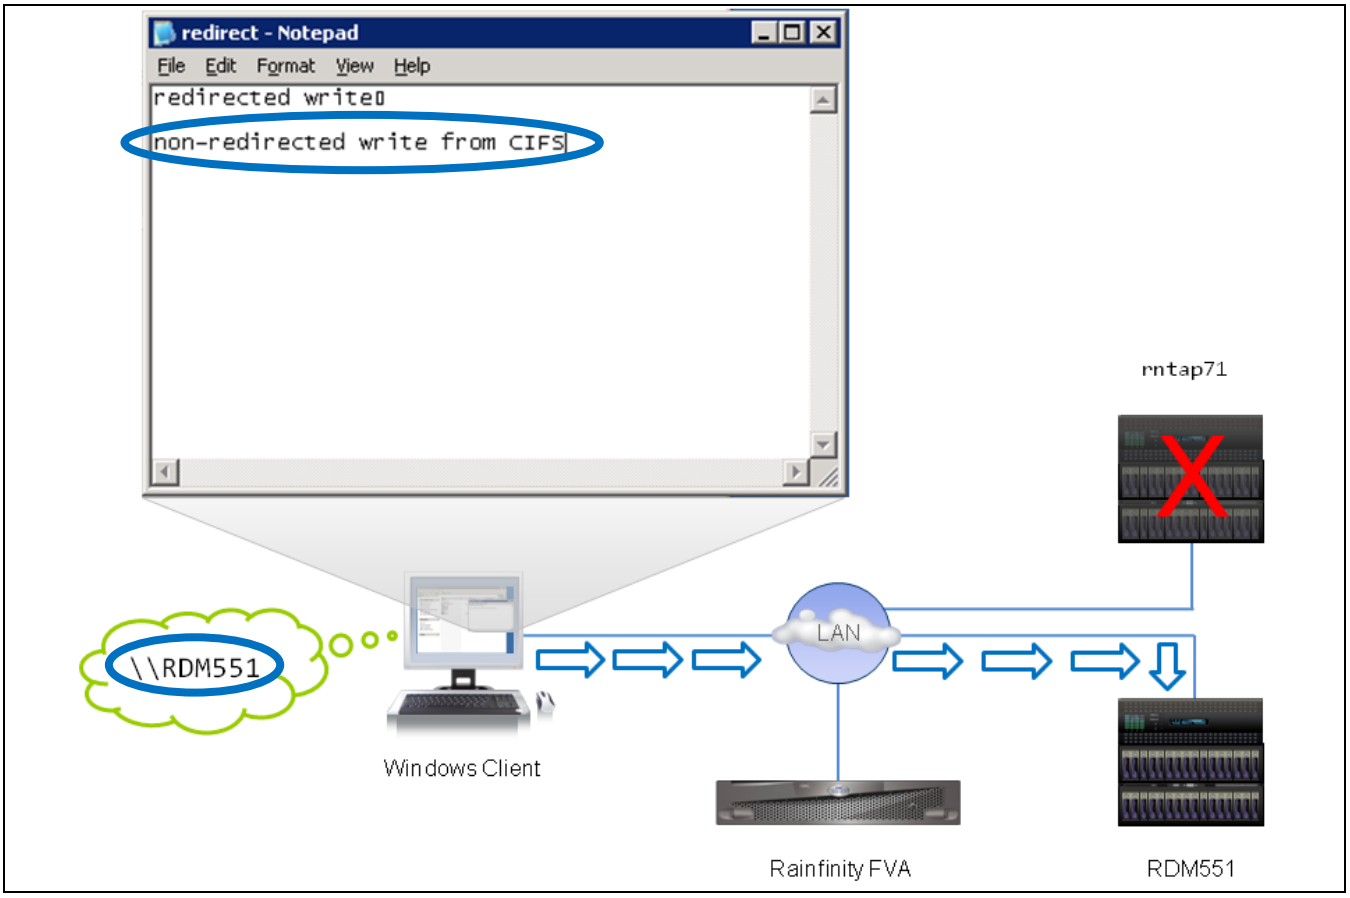

Figure 17 shows how ESG Lab examined the same file using both NFS and CIFS simultaneously. ESG Lab was logged in as the same user in both cases and confirmed that file permissions were consistent.

FIGURE 17. MULTI-PROTOCOL FILE VIEW

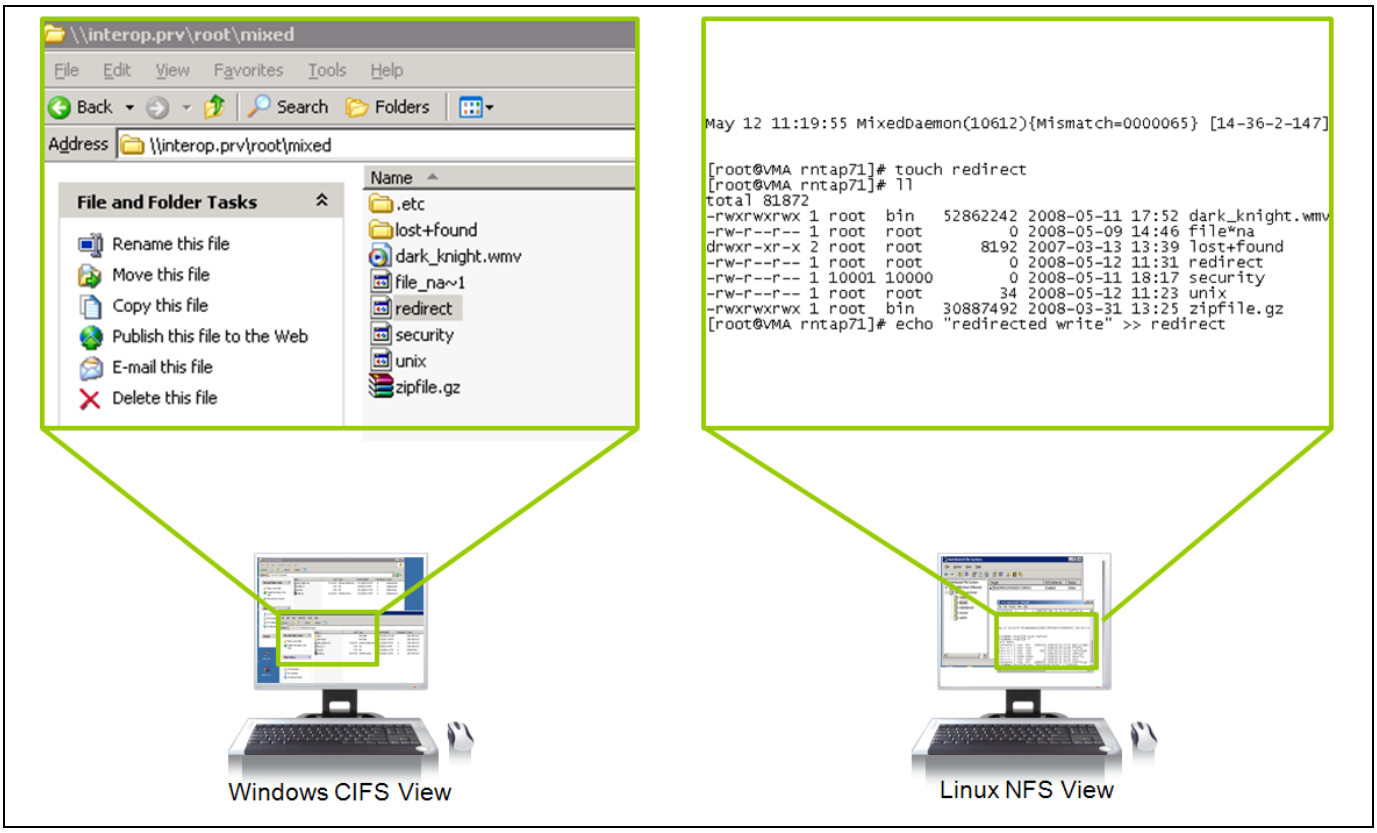

Once all clients were verified to be mounted to the new destination, ESG Lab un-proxied access and moved Rainfinity FVA back out-of-band.

Why This Matters

Sloppy cleanup after a migration can lead to a number of serious problems. Poor cleanup not only makes it hard to figure out where clients are mounting migrated filesystems from, but in the worst case it can lead to inadvertent application and end-user downtime as clients try to access decommissioned storage. Even when everything works perfectly, migrations have long been extremely disruptive, requiring either one long outage or multiple short outages to clean up the environment after a migration

ESG Lab validated that Rainfinity FVA automates the post-migration cleanup process as it automatically redirects applications and users away from decommissioned storage until all clients have been connected to the new location. This allows users and applications to stay online throughout the migration process.

Heterogeneous Global Namespace Update

Global namespace software aims to address NAS and file server proliferation by providing a single, unified namespace for file services. In the world of NAS, each file server or NAS device contains its own namespace, meaning it manages its own shares or exports independently and clients must mount each share or export individually.

A global namespace aggregates the shares and presents clients with what looks like one huge file system with all of the shares and/or exports in one place. A global namespace should not be confused with a clustered filesystem. A clustered file system is designed to address capacity and/or performance scalability for a single dataset. A global namespace simply presents a single logical view of multiple shares or exports from multiple servers/clusters.

When a client mounts a global namespace, the user sees a root directory with all of the aggregated shares arrayed below it. Figure 18 shows the logical name and UNC path of a share as seen inside Microsoft DFS, the global namespace application that Microsoft provides for Windows Environments. DFS integrates with Active Directory for security and file permissions.

FIGURE 18. GLOBAL NAMESPACE MANAGEMENT

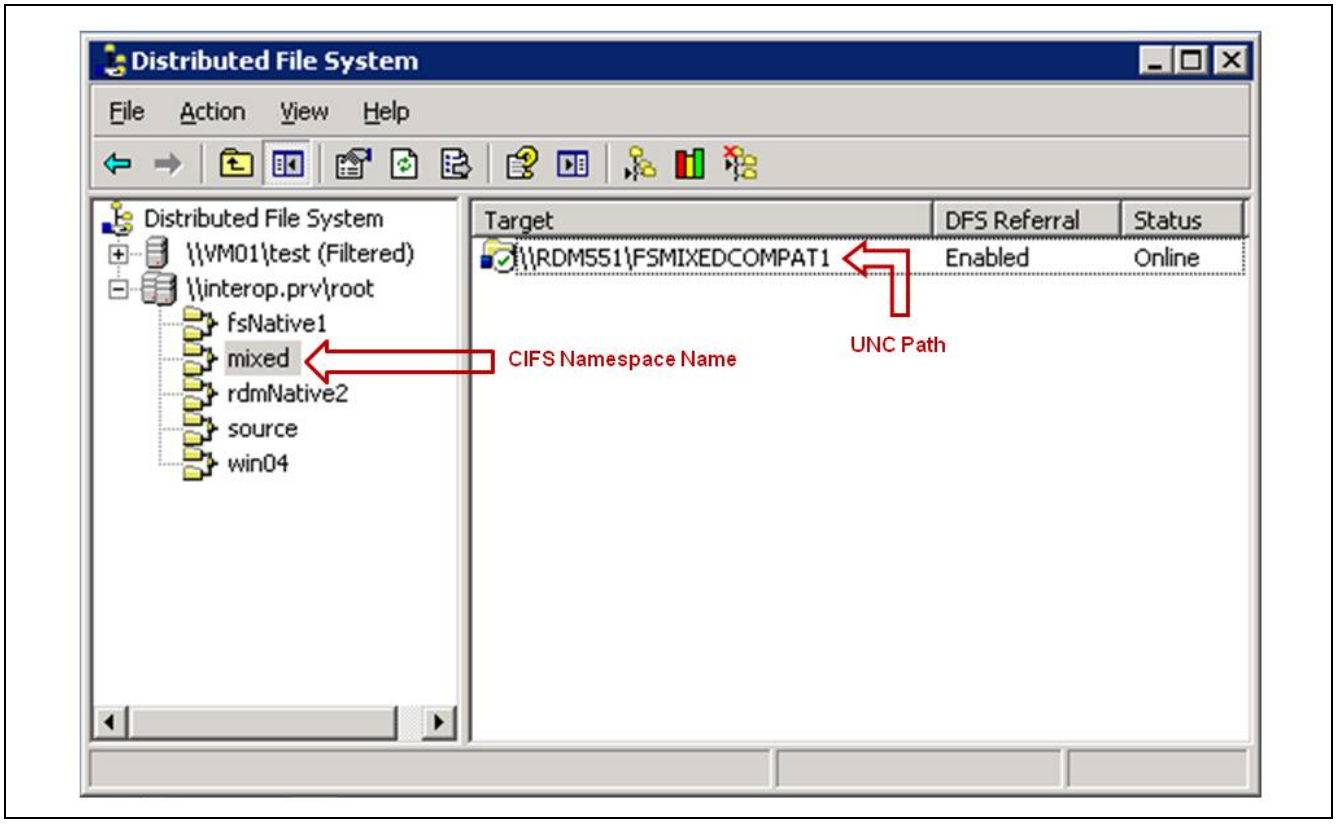

UNIX/Linux provides global namespace services using automount maps managed by NIS and OpenLDAP servers. This has the same net effect for NFS that DFS has for CIFS, providing a browseable namespace made up of many aggregated NFS exports.

Rainfinity FVA's Global Namespace Management application automatically updates and maintains proper mapping of shares and exports to heterogeneous NAS environments (using Microsoft DFS or Automount maps).

The Global Namespace Management application manages distributed global namespaces by subscribing and publishing namespace schemas stored in DFS, NIS, LDAP, or Rainfinity's own namespace server. When migrations are performed, both physical locations' namespaces are automatically updated.

In this way, Rainfinity FVA provides administrators with a common method of updating multiprotocol namespaces served by independent CIFS and NFS servers, ensures continuous read-write access during data migrations, and provides multiprotocol namespace synchronization that eliminates manually maintaining heterogeneous namespaces.

The DFS global namespace allows CIFS users to access shares from a single logical mount point that resides on separate and distinct file and NAS servers.

ESG Lab used the Rainfinity FVA Global Namespace Management application to update the underlying standards-based global namespaces and confirmed that Rainfinity can simplify and enhance management of Microsoft Windows DFS roots and UNIX/Linux Automount maps managed by NIS and OpenLDAP servers.

Rainfinity enabled ESG Lab to maintain a virtual view of shares and mount points across heterogeneous NAS servers.

Why This Matters

In the world of NAS, each file server or NAS device contains its own isolated namespace, meaning that the shares or exports owned by the server are managed independently and clients must mount shares individually. Providing global access to files for heterogeneous users is increasingly important as the volume of data under management grows. A global namespace aggregates networked file storage and presents clients with what looks like one huge filesystem with all of the data a client has access to in one logical place. In multi-protocol environments, administrators need to manage each set of users (Windows and UNIX) as separate entities, maintaining both namespaces separately. Data migrations present additional risk as administrators must manually update both namespaces to ensure data access.

Rainfinity FVA reduces time and effort by providing a common interface to unify and update both global namespace environments. Clients access their files as they always have, without any proprietary software or additional network devices in the way.

ESG Lab Validation Highlights

- ☑ The Rainfinity File Virtualization Appliance was integrated into an existing NAS environment quickly and easily with no software agents.
- ☑ ESG Lab used FVA to move a live file system hot and online with no disruption to multiprotocol clients actively reading from and writing to files.
- ☑ Rainfinity was able to preserve both UNIX attributes and Windows ACL's in a single move operation.
- ☑ The redirect function provided continuous access for clients without having to remount or reconnect.
- ☑ Integration with Microsoft DFS and NIS Automount for global namespace functionality was transparent to the administrator and end-users.

Issues to Consider

- ☑ When ESG tested Rainfinity FVA for this report, Rainfinity integrated with industry standard global namespace products for both Windows and UNIX. ESG did not test Rainfinity FVA with EMC's own global namespace product (Rainfinity Global Namespace Appliance). Microsoft DFS and AutoMount for NFS are both complex and difficult to use. ESG lab believes implementing a Global Namespace solution on the Rainfinity platform should be an excellent way to leverage Rainfinity's technology and is a logical extension to the Rainfinity family.
- ☑ Capacity and utilization reporting, a useful feature of the Rainfinity Global File Virtualization product, was absent from the version of Rainfinity FVA tested by ESG Lab. This functionality enables administrators to examine their environment and see capacity and performance hotspots at a glance and would be a useful enhancement to the Rainfinity FVA.

ESG Lab's View

IT departments are managing larger-than-ever pools of unstructured file data that are growing at an accelerating rate. As the scale of managed data increases by orders of magnitude, new management approaches are required. Also, as datasets scale, data migrations are playing a more prominent role in the data center to address consolidation, capacity and performance demands.

Potential solutions must address the issues of exploding capacity and continuous access without injecting risk into the environment. Data integrity, performance, high availability and disaster recovery are essential requirements. Traditional and highly manual migration methods are no longer acceptable. There needs to be a better, faster and less painful way to complete these migrations while keeping the datasets online.

Automating global namespace updates makes life easier for both administrators and users; administrators don't have to manually update namespace after the data is migrated and users don't need to do anything to find their data in its new location. This allows administrators and end-users to be more productive and more likely to store their data on shared devices versus their own local drives. This enables better corporate compliance and improved ability to archive aging data, reducing the capacity of data that needs to be backed up.

Rainfinity FVA drops directly into existing NAS environments with no special configuration or software required. Clients access their NAS storage as they always have, with no modification and no interruption during data movement. Rainfinity is changing the way one looks at file services and brings a comfort level to file migration projects by automating the tedious and error prone aspects of file data migrations. Rainfinity proves it is possible to stay ahead of massive capacity growth and move at the speed of business.

ESG Lab was impressed with the flexibility and power of Rainfinity FVA. Rather than managing NAS resources one system at a time, IT needs the ability to manage data more logically at the dataset level, with control over the properties and attributes of those datasets and the ability to act on them. ESG Lab validated that EMC Rainfinity FVA provides organizations with the power to simplify multiprotocol data mobility, increase utilization, optimize performance and reduce TCO for their NAS environment.

Appendix

Table 1. TEST CONFIGURATION

Hardware	Software
EMC Rainfinity FVA 32 GB memory; 12 GigE Ports	EMC Rainfinity Software Version 1.2
EMC Celerra NS40 Dart 5.5	Windows 2003 Server (SP2)
NetApp FAS840 OnTap 7.1	Linux Red Hat FC7 (Intel)
HP-DL380 (G3) Windows Servers	Namespace – DFS R2 – Win; NIS - INUX
SuperMicro x64 6GB Ram Linux Servers	
Cisco 2950 Gigabit Ethernet Switch	

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com