



Mass Transit and the Evolution of File Virtualization

A Guide to Understanding EMC Rainfinity

EMC Proven™ Professional Knowledge Sharing 2008

Craig T. Kensey
EMC Corporation
Technology Consultant
kensey_craig@emc.com

Table of Contents

Introduction.....	3
History	4
Basics of Global File Virtualization	5
GFV Essentials	7
Basically Networking	8
Out-of-band Features	11
GFV GUI – Adding File Servers	12
Advanced Networking	12
Pre In-band procedures	14
Moving a file server in-band.....	16
Proxy State.....	16
The In-band Process.....	17
Global Namespace or not.....	19
GFV Wrap-up	21
FMA Essentials.....	22
Accolades.....	25
The Future Direction of Rainfinity	25

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

Introduction

This article provides you with a high level understanding of EMC Rainfinity Global File Virtualization (GFV) concepts and its numerous capabilities. It is meant to educate current and potential EMC customers, sales professionals and technology teams on how Rainfinity’s innovative approach can be used to centrally manage and maintain active and inactive data in Network Attached Storage (NAS) file serving infrastructures. My role at EMC is to provide pre-sales support for Rainfinity opportunities. Therefore, this article combines sales, technology and marketing of the GFV solution.

The first question is, “What does Mass Transit have to do with managing NAS?” It’s really just a metaphor. My aim throughout this article is to periodically use the analogy of a simple commuter train ride to present the Rainfinity solution.

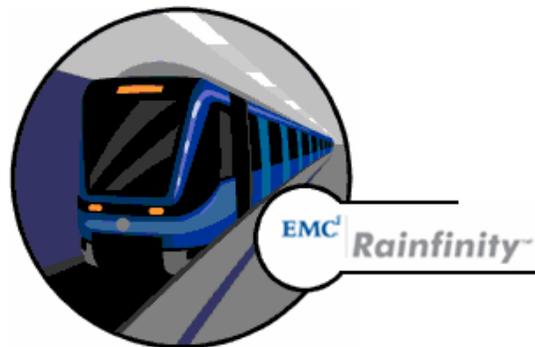


Figure 1 – Author’s fictitious analogy of GFV

“The only way to be sure of catching a train is to miss the one before it.” G. K. Chesterton

History

EMC purchased Rainfinity in 2005 to bolster our “Virtualizing Information Infrastructure” arsenal that currently includes server and data protection virtualization, block storage virtualization and file virtualization, allowing EMC customers to:

- Reduce total Infrastructure costs
- Simplify management
- Increase service levels
- Provide a flexible infrastructure
- Increase energy efficiencies

Originally coined Rainstorage, life began as a research project at the California Institute of Technology (Caltech), working with NASA's (the U.S. National Aeronautics and Space Administration's) Jet Propulsion Laboratory and the Defense Advanced Research Projects Agency. The original research sought to identify software building blocks for developing distributed applications and was called RAIN (Reliable Array of Independent Nodes). Rainfinity spun off from Caltech in 1998.

The original RainStorage software virtualized Windows, UNIX, and Linux file systems across heterogeneous network-attached storage systems and file servers, making those devices appear as a single unit. The technology simplified management and made it easier for businesses to carry out data migrations. Today, EMC leverages those initial developments and has transformed Rainfinity into the fully functional virtualization tool that we deliver today.



Figure 2 – 2U Rainfinity Appliance

Basics of Global File Virtualization

Enterprise network file serving environments have become major management headaches. Companies are being forced to re-evaluate how they handle this large, critical chunk of storage infrastructure in response to complicated and time consuming file migrations, capacity and performance balancing, and heightened compliance concerns. Rainfinity GFV is an appliance based solution (hardware and software) that recaptures stranded storage capacity on network file servers that support the CIFS (Common Internet Files System) protocol used mainly by Microsoft clients and the NFS (Network File System) protocol used by UNIX and Linux clients to perform file migrations. GFV operates at the share/export level; FMA operates at the file level.

GFV's major benefits:

- Transparent Data Migrations
 - *Example:* Migrate data off an older NetApp filer to an EMC Celerra during normal business hours without impacting end users.
- Providing an industry standard global namespace
 - *Example:* Using Microsoft DFS to automate and redirect clients upon completion of a migration.
- Centralizing views to manage the entire NAS infrastructure
 - *Example:* Reporting and trending analysis allows administrators to identify, analyze and resolve capacity related issues.
- Tier storage infrastructures
 - *Example:* Migrate data from expensive Fibre channel drives to lower cost SATA (Serial Advanced Technology Attachment) drives.
- Archive data to meet compliance regulations
 - *Example:* Financial data will be archived for a minimum of 7 years to an EMC Centera® to meet strict regulatory requirements.

Rainfinity provides seven purpose-built applications to simplify common tasks and assist storage administrators to optimize storage decisions. Each of the following Rainfinity applications focuses on a particular objective:

- Capacity Management
- Performance Management
- Tiered Storage Management
- Global Namespace Management
- Migration and Consolidation
- Synchronous Replication
- Archiving

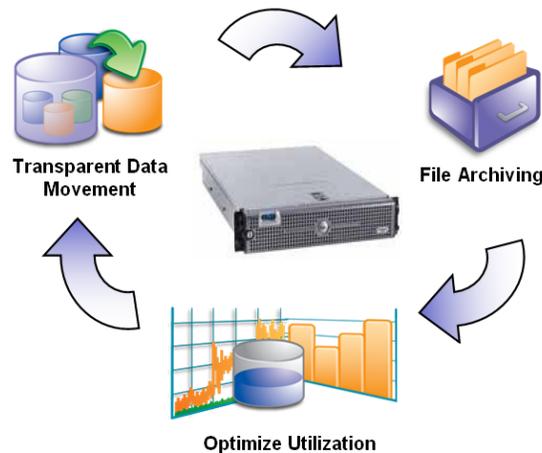
Virtualization is also a strategic part of a long term infrastructure. With many of the initial use cases mentioned above, where do you start? What are the criteria you need to consider when evaluating file virtualization solutions? Typically, before virtualization is first used in an organization, storage management costs are high, data migrations are difficult, and utilization rates are relatively low. Migration or consolidation is the first phase for file virtualization in these environments.

Migrations and Consolidations can be time consuming, disruptive and extremely costly. With file virtualization, the storage administrators' major obstacle is end-user disruption. This is eliminated and storage consolidations are dramatically different. Non-disruptive, transparent data migration preserves customer service levels, takes a fraction of the time when compared with standard consolidations, saves significant administrative time, and avoids the cost of over-allocation.

Now that storage is consolidated, real-time **optimization** or capacity balancing saves even more money and time as data can be continually moved across storage devices in response to environment and business changes. Virtualization balances capacity, performance, and leverages tiered storage saving money, resources, and unnecessary challenges by improving resource utilization and allowing intelligent allocation.

Industry surveys reveal that, on average, 70% of stored data hasn't been accessed in 90 days or longer. **Archiving** inactive data is an additional benefit of file virtualization. Archiving presents another compelling solution to further optimize primary NAS storage by automatically moving inactive files to less expensive secondary storage based on policies. Files that are moved appear as they are on primary storage to users and applications. File archiving dramatically improves storage efficiency and backup/restore times, while supporting additional business requirements such as adhering to corporate mandated compliance and retention.

Figure 3 –
Compelling
repetitive use
cases for GFV



GFV Essentials

Networking, in-band/out-of-band, and global namespace are important terms and technologies that need to be acknowledged to successfully deploy, comprehend and manage a Rainfinity solution. Networking is the first and easily the most important technology that we need to explore; then we need to examine how GFV integrates within a corporate network topology.

Basically Networking

For readers that are new to Rainfinity and perhaps new to networking, I will now draw upon my analogy by using a mass transit map.

Figure 4 – NJ Path Rail System Map



Many commuters travel from Penn Station, Newark to the 33rd street station in Manhattan on a daily basis. Once the train makes its usual stops at Journal Square or Grove Street, passengers destined for 33rd street or even Hoboken have an option to exit the train and continue their route on another line. All along this journey, and once aboard the yellow line, a vast network of tracks and communication switches guide our train on its merry way, changing the landscape from 4 tracks to 3 to 2 and finally to the 1 track that takes us to our final destination.

In many respects, networked environments behave similarly to our mass transit scenario. They leverage devices such as switches and cables; they offer different paths for data. Networks are usually referred to as a LAN (Local Area Network – in the same building) or WAN (Wide Area Network – separated by geography). Network switches make quick decisions by inspecting data and determining the source and destination device that it will travel on; cables are responsible for carrying the data.

By using this logic, let us assume the following from Figure 4 and the information above:

- **WAN** = the entire rail system that ties NJ and NY together
- **Switches** = Stations and switching
- **Cables** = Train tracks
- **Passengers** = Data

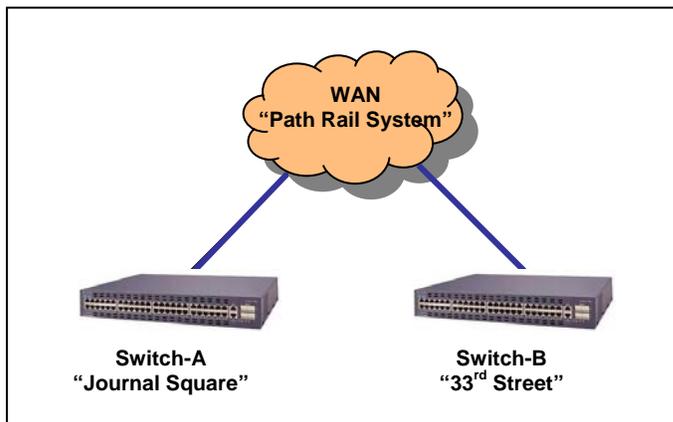
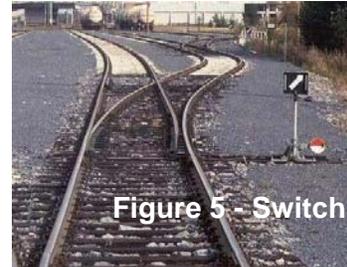


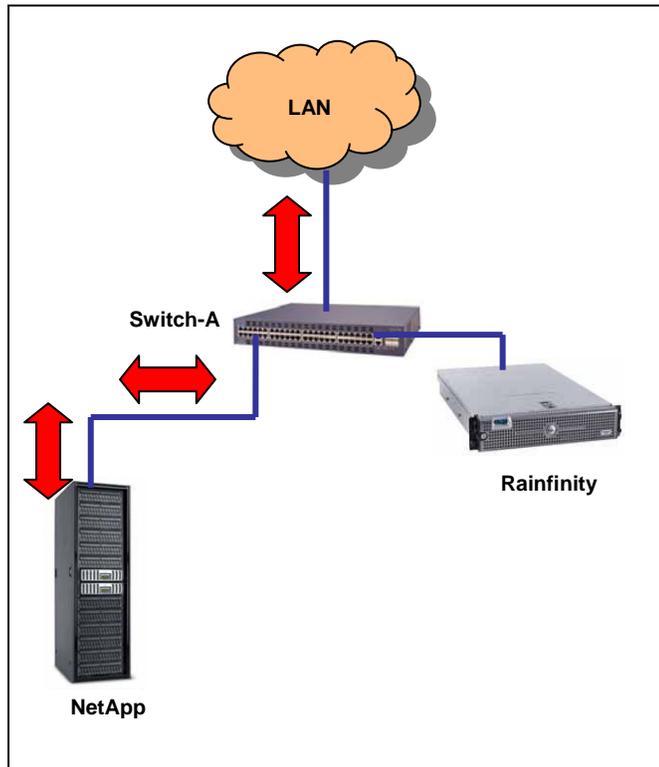
Figure 6 – Simple network layout

This simple network diagram illustrates a WAN connected to 2 switches. Each is attached by a cable. Data traverses the WAN and is routed to the appropriate switch similar to the way passengers leave from one station and are switched to arrive at another.

Figure 7 – Red arrows indicates data flow from end users to the NetApp

Our network now includes three devices connected to the LAN: Rainfinity, NetApp and the switch. By plugging each device into the switch and with some additional configuration, these devices are now accessible from anywhere on our LAN.

Users will be configured to map a drive to the NetApp filer, to a home directory or corporate share from their workstations to store important data. Just like our passengers, they will take daily train trips to each respective stop.



Out-of-band Features

Out-of-band means “out of the path of data.” It refers to how Rainfinity is currently configured in *Figure 7*; no data (traffic) is flowing through the GFV appliance. Rainfinity is currently configured to gather valuable statistics at this point such as Capacity, Performance and Tiered Storage Management from the NetApp.



Figure 8 – Capacity Management graph – GFV

This type of information is valuable to any storage administrator since an entire data center of file servers can be monitored for capacity and performance trending. This allows for a well balanced, top performing file services infrastructure.

The Performance Management module monitors CPU utilization on file servers and reports on issues called “hotspots” that can potentially cause bottlenecks in the environment. The Tiered Storage module identifies frequent or infrequent access to data. This allows storage administrators to pool different levels of storage and to assign data appropriately within the environment. The graphical reports for each module all share the same look and feel.

Out-of-band is also how Rainfinity FMA is architected and will be explained in full detail later in this article.

GFV GUI – Adding File Servers

Adding file servers into the Rainfinity GUI is a simple process that requires a name, an IP address and, for a CIFS server, a domain user account. Once this process is completed, all file servers that have been added to the GUI can be monitored out-of-band.

■ File Server List

File Server	Band	Error	Proxy State		Proxied Conns		Unproxied Conns		NFS Ops		CIFS Ops
			NFS	CIFS	UDP	TCP	UDP	TCP	UDP	TCP	Ops
 Celerra76	Out	No		None	0	0	0	0			0
 Celerra77	Out	No		None	0	0	0	0			0
 Celerra78	Out	No		None	0	0	0	0			0

List options

Filter by file server name ([help](#))
Name:

Show file servers that are:

In-band
 Out-of-band
 Of unknown band

Show remote file servers

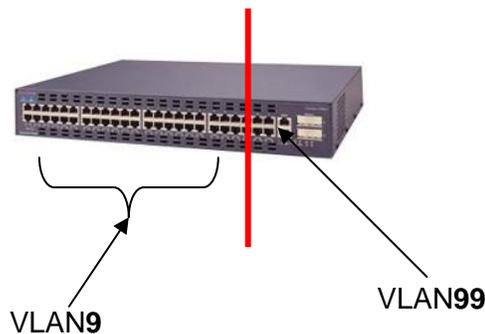
Figure 9 – Configured file server list in the Rainfinity GUI, note the **Band** state (Out).

Advanced Networking

Now we will explore some networking terms that will allow GFV to perform its most fundamental use, migrations. Since networking plays such an important role with migrations, very specific switch configurations are required to provide data to be migrated. The following section contains industry standard networking terms and does not require any proprietary configurations, a very important aspect of Rainfinity.

VLAN (Virtual Local Area Network) allows network administrators to logically re-segment their networks without physically rearranging the devices or network connections. As an example, you might create and segment a VLAN for the payroll department and one for guest users thus making the network more secure and easy to manage. Creating a new VLAN is the first step to integrating Rainfinity into the network as illustrated in *Figure 10*.

Figure 10 – Switch showing 2 VLANS 9 & 99



The switch has 48 ports (switchports) so it can handle many different devices and VLANs. For GFV, we are configuring just 1 port on this switch to be part of VLAN 99 and the rest of the ports are in our Public VLAN. It is a very easy task for a network administrator to create a new VLAN.

Bridges carry trains over bodies of water; bridges also carry frames of data from source to destination. **Transparent bridging** is the methodology by which Rainfinity will be able to bridge the VLAN 9 to VLAN 99. By bridging these VLANs, we come up with the phrase **in-band** as data now flows through the GFV appliance. Here is how it works.

Pre In-band procedures

Now that you have created this new private VLAN, we must prepare GFV to transparently bridge the 2 VLANs.

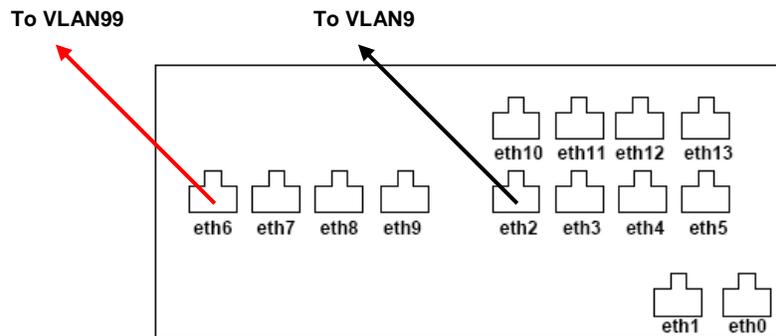


Figure 11 – Rear view of the GFV appliance and cabling to the switch

GFV ships with up to 12 Ethernet ports so you can create up to 6 bridges. In the figure above, we see that eth6 is plugged into the switch designated as the newly created “Rainfinity VLAN99” and eth2 is plugged into the switch that accesses the public VLAN9. This is referred to as creating a bridge in Rainfinity; additional Command Line Interface (CLI) configurations are required. Now that the bridge is created, data still flows as suggested in *Figure 7*, we’ve just prepared the network and Rainfinity. Our next step ties this whole configuration together, powering up this transparent bridge and beginning the in-band process.

There are several different methods to move file servers in-band. VLAN tagging and Switchport changes are the two preferred methods.

VLAN tagging (Industry Standard 802.1Q) classifies ethernet packets into specific LANs by attaching additional data to each packet. The additional data stored in the Ethernet header contains an ID number corresponding to a particular VLAN (VLAN99 in our case). Both Celerra® and NetApp support VLAN tagging, and this method has the potential to be fully transparent when automated.

When a file server needs to be inband, it is reconfigured to use the VLAN tags of the isolated VLANs. All traffic sent out of the file server will arrive on the private side of the GFV Bridge. To move a file server out of band, we reconfigure the network interfaces to use VLAN tags of the existing VLANs (move back to VLAN9). Configuring VLAN tagging on the switchports leading to both the Celerra and NetApp is an important, additional step. *Figure 12* ties all this together.

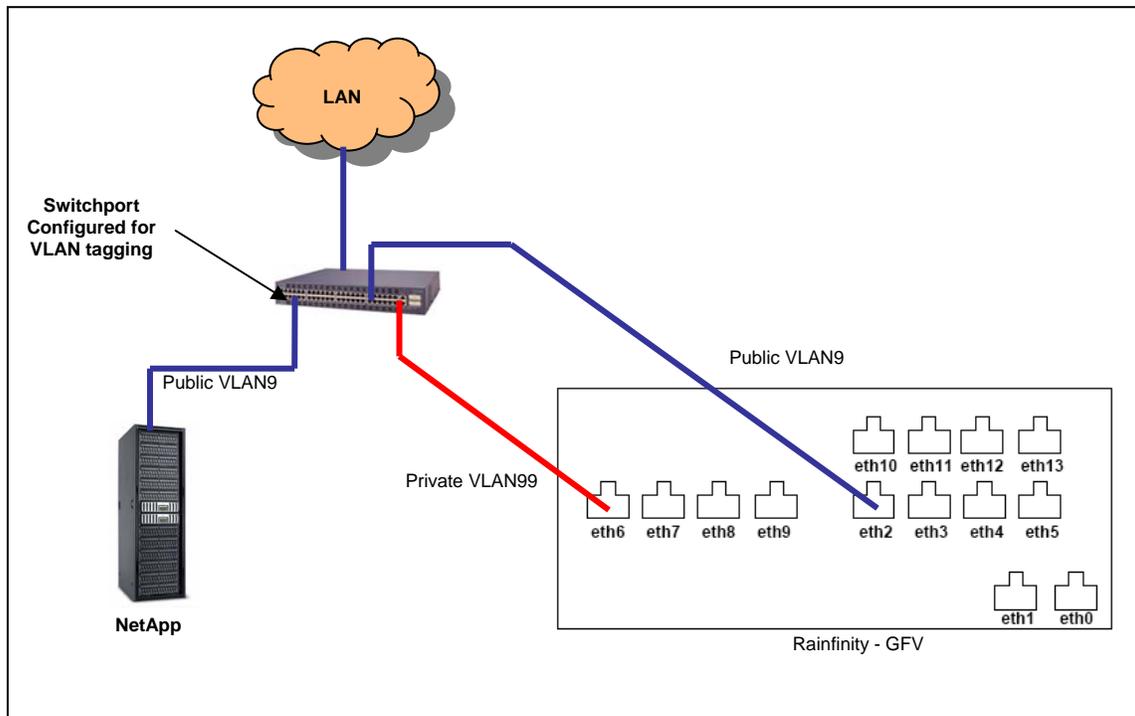


Figure 12 – VLAN configuration, the switchport connected to the NetApp is configured for VLAN tagging and GFV's port eth6 is plugged into the private VLAN, this is still considered an out-of-band configuration.

The **switchport method** of moving file servers in-band is a manual process performed by a network administrator. A simple command to change the switchport designation from VLAN 9 to VLAN 99 is all it takes to reconfigure a specific switchport. The decision to use either method relies on the customer's network team and how the infrastructure is currently set up. Rainfinity ships with scripts that help to automate the process of moving file servers that may also be utilized in band.

Moving a file server in-band

Our NetApp filer will be coming off lease next month and we decide to purchase an EMC Celerra[®]. I've configured my network environment for Rainfinity including configurations to the NetApp and Celerra file servers. Our goal is to migrate all data as soon as possible. The data size is about 10TB, all CIFS and 1000 determined users will not accept explanations for not having access to their data. This is a very common theme and one that our Rainfinity solution is well adept at solving.

Traditional utilities such as ROBOCOPY and EMCopy fall short of meeting our customer's goals of not impacting end users and applications. The number one obstacle that our customers face while planning for a migration is to provide unimpeded access to data by end users; Rainfinity solves this major business challenge.

It is very important to understand the type of data to be migrated. In our scenario, it is all Microsoft Office documentation and can take a network change easily when we move the NetApp in-band. A short, planned outage would be required if NetApp was hosting an Oracle application.

Proxy State

For this migration, we will be using the **switchport** change to bring the NetApp and the newly installed Celerra in-band. Once configured to the in-band state for the migration, Rainfinity must proxy all connections to track and synchronize client access during a migration. If the underlying connection of a CIFS session is ended or broken, the CIFS session is also terminated. When this occurs, CIFS will not attempt to reconnect.

Applications that attempt to utilize a terminated session will receive errors when performing CIFS operations. The handling of such errors is application specific. For instance, a user editing a Microsoft Word document stored on a CIFS share would not necessarily notice if their CIFS session was terminated because the application is designed to automatically reconnect transparently.

The In-band Process

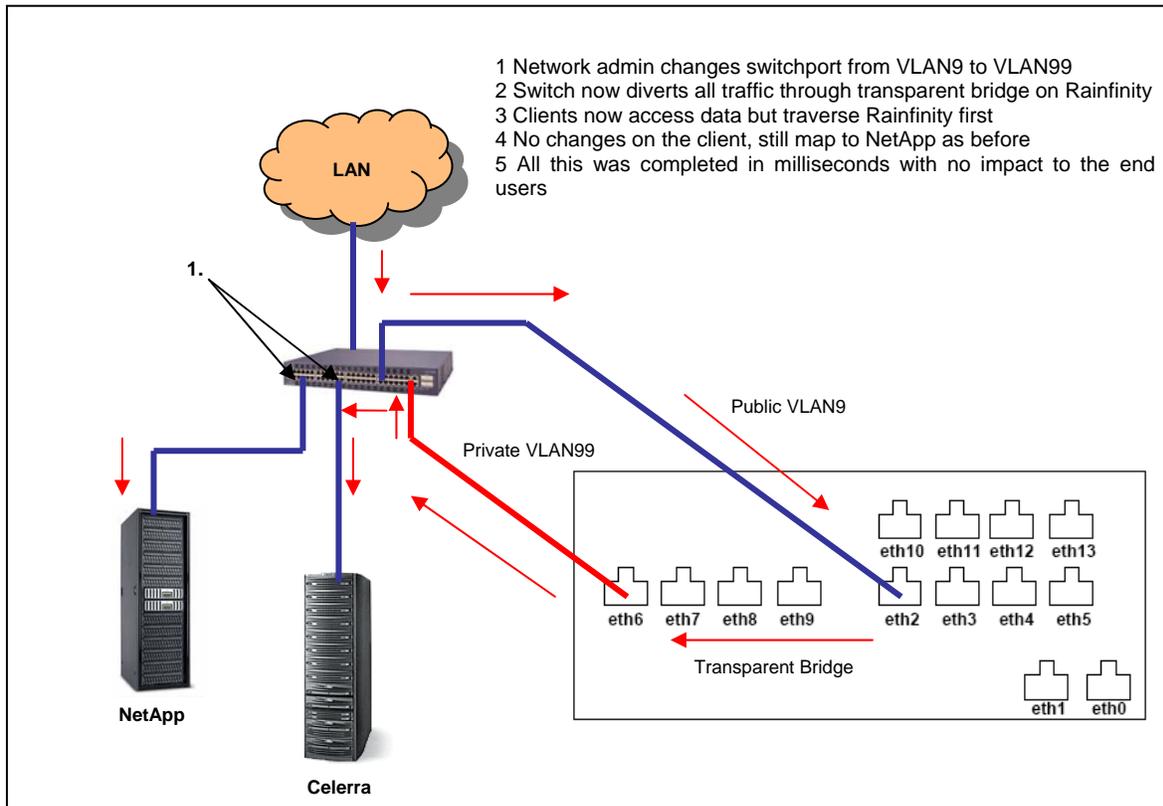


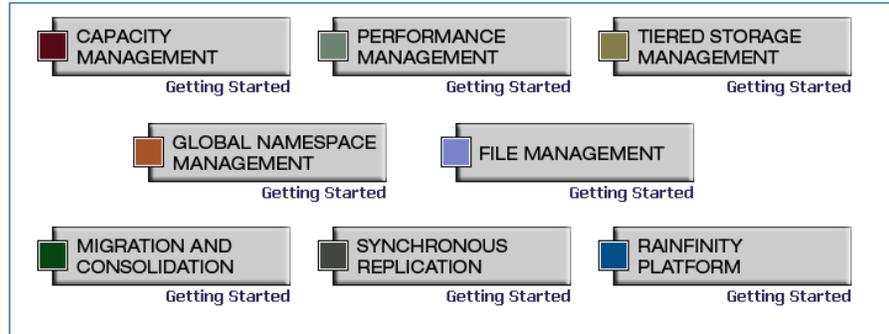
Figure 13 – Switchport change method to make data flow through Rainfinity

Migrating Data

Before jumping directly into a production migration, many if not all customers want to see GFV in action. We draw up thoroughly documented test and acceptance plans for this purpose. We are ready for the migration now that our file servers are in-band and the GUI shows those 2 servers as in-band.

Navigating from the GUI's home screen, we select the Migration and Consolidation module. From there we select "New Move" and we proceed to follow the Wizard that guides us the rest of the way.

Figure 14 –
Rainfinity
GUI



The following are the high level steps to initiate the migration that require user intervention:

- I. NFS or CIFS
 - a. the current version of 7.2 can only migrate 1 protocol at a time
- II. When to start the move, now or at a later date
- III. Preserve stubs?
 - a. If the environment contains CFA (Centera File Archiver) or FMA, stubs or pointers to the data are present (more details to come), GFV has the ability to migrate and preserve these stubs
- IV. What is the source and path to move?
- V. What is the destination file server and path?
 - a. GFV also can create this on the destination
- VI. Use Rainfinity to copy data
 - a. The move can also be done with files that have already been copied to the destination; at this point Rainfinity will sync the changes, this is referred to as a delta move.
- VII. The final step is to click Finish and the migration process begins.

Rainfinity's unique transaction based feature protects data throughout the migration and allows for complete data integrity in the event of failure. During the migration, users have full control of their data. Rainfinity also honors locking features and maintains current security styles. The migration can be throttled "on the fly" to consume the bandwidth necessary to allow for optimal network utilization.

The following are the high level states in this CIFS migration:

- I. Rainfinity checks to make sure the file servers are in-band.
- II. Rainfinity proxy's all connections.
- III. Executing
 - a. This is the state when GFV copies data from the source to the destination. GFV copies data and does not move data, this is an important point.
- IV. Upon completion of the copy, GFV will synchronize (syncing state) any changes
 - a. This can be referred to as an active mirror; both the source and destination are identical.
 - b. Users are still mapped to the original source file server; the source is still the authoritative file server.
- V. The next phase is the most important and is referred to as two-way-syncing. In this state, both the source and the destination are identical but now the destination becomes authoritative. This is when users will be transitioned to the new destination server. This will be the focal point of the next section.
- VI. Complete the transaction.

Global Namespace or not

As stated in step V above, completing the transaction and providing end users with access to the new destination filer is the most critical stage of the migration. Rainfinity supports industry standard global namespaces such as DFS (Distributed File System) for Windows clients, and Automount for UNIX/Linux clients. The process of cutting over those users is automatic if an environment has it fully integrated. Without a namespace, Rainfinity contains an important monitoring tool called "access statistics" that report on each users' access to data and allows the administrator full control on cutting over client access. Once the administrator is assured that all end users have been moved to the new destination, the transaction is complete and the original filer can be repurposed or removed from production.

Let's use our mass transit analogy to describe a namespace and its importance. Let's say that you are waiting for the train. Communication is important if the train's departure time has changed or if the train has been delayed. Commuters can access this information by looking at the large arrival and departure information board. Everyone has the same information about their travel plans.



Figure 15 –
Effective global communication ensures that every commuter knows where and when to go to catch the next train.

Global namespace is key to efficient management of distributed file storage. A namespace allows clients to access files without knowing their location (just as they access websites without knowing their IP addresses).

Transactions progress from syncing to the two-way syncing state during a CIFS migration. The Global Namespace Management application automatically commits the changes to the namespace schema and automatically updates the physical location thus providing for complete client cutover transparency. Generally speaking, customers today rarely fully deploy a unified global namespace; however, the **Client Access Statistic** link is another valuable tool that helps customers manage that critical phase of the transaction. Rainfinity tracks every users' access during a migration and allows for simplified management during the cutover. CIFS clients need to be manually pointed to the new destination by editing current logon scripts. Logging off and then logging back on completes the cutover.

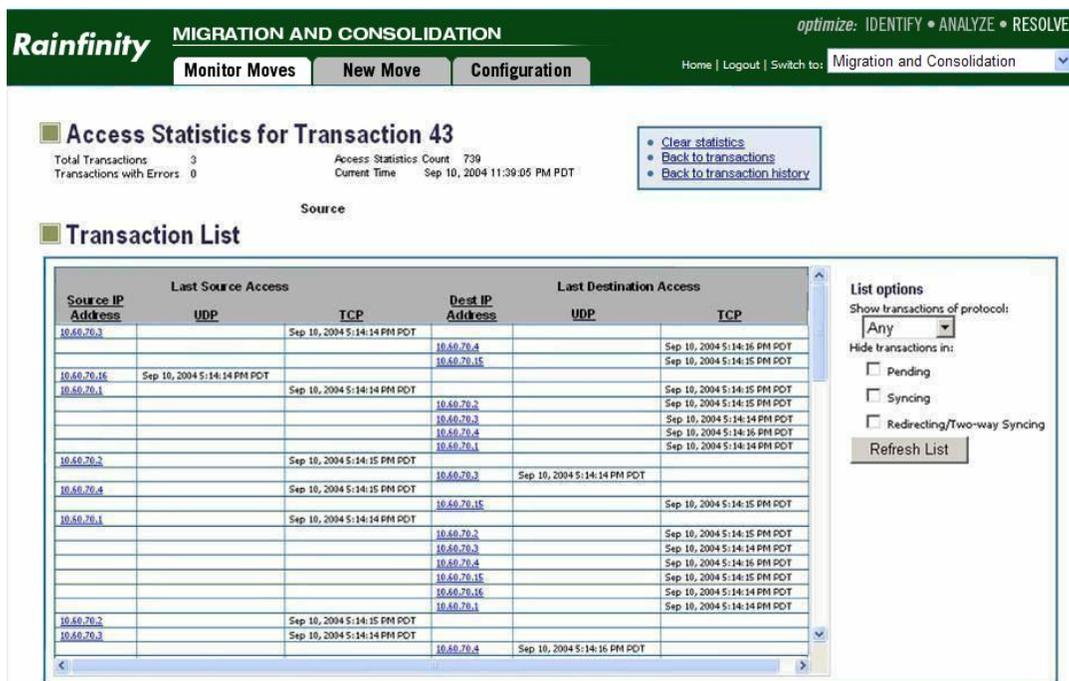


Figure 16 – Access Statistics used to manage client cutovers

GFV Wrap-up

To review the important features of GFV, it's important to note that there are critical network configurations that are required to integrate Rainfinity into a customer's networking environment. Consider multiple face-to-face meetings and white board discussions with the networking group to facilitate and solidify understanding. Once this is done, the criterion of goal oriented migration procedures need to be individually created for each customer based on their unique environment. The primary goal is not to impact end users during a migration. GFV has the capabilities to meet this requirement but certain processes and procedures must be fully implemented and followed to ensure success.

FMA Essentials

Archiving with FMA is considered an entry point for GFV as these solutions can be sold separately. FMA becomes more attractive as customers begin to appreciate how Rainfinity can unify and manage active and inactive data within a single solution for their file serving environment. The biggest statistical eye opener is that on average, 70% of NAS data is static. Customers benefit greatly by moving this large chunk of data to lower cost storage. They can reduce the amount of data to be backed up, improve capacity utilization, transition to a customized ILM (Information Lifecycle Management) strategy, and fulfill goals for meeting new compliance standards. This article will describe how Rainfinity FMA archives static and under-utilized data.

FMA optimizes primary NAS storage by automatically moving inactive files based on policies to less expensive secondary storage (either NAS or CAS). Moved files appear as they are on primary storage to users and applications. FMA operates out-of-band as opposed to in-band like GFV (for a migration) and does not require a persistent database. FMA creates stub files that contain all the necessary information required for users to recall data. Maintaining critical databases only adds to the complexities of managing a separate solution. FMA also provides the capability to recover stub files, identify and address orphan files, and track versions. Orphaned files exist when the stub file linking the files is deleted by an end user or application.

FMA leverages a policy engine to define which files should be archived. Users can combine and evaluate multiple rules in a single policy. FMA includes the following rule types for archiving and data collection:

- Last Accessed Time – the last time this specific file has been accessed or read
- Last Modified Time – the last time this file was edited
- File Size – size in terms of the space it consumes; i.e. 11 MB
- File Name – extensions such as .doc, .jpeg and .xls

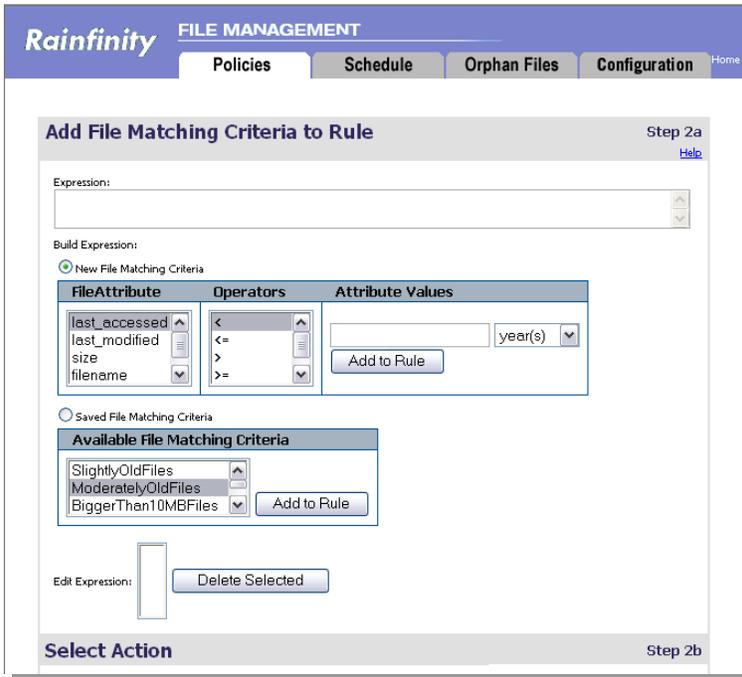


Figure 17 – FMA GUI showing rules to create a policy.

Building a policy is one thing but how do you know, once in production, how much that policy may yield in terms of archived data size? Before accepting a bottle of fine wine, a simple taste test is performed to ensure you are pleased.

Similarly, FMA has an optional solution called a “what if” analysis to preview your new policy prior to moving it into production. Once the policy meets the required goals, schedule the policy to automatically run at a pre determined date and time.

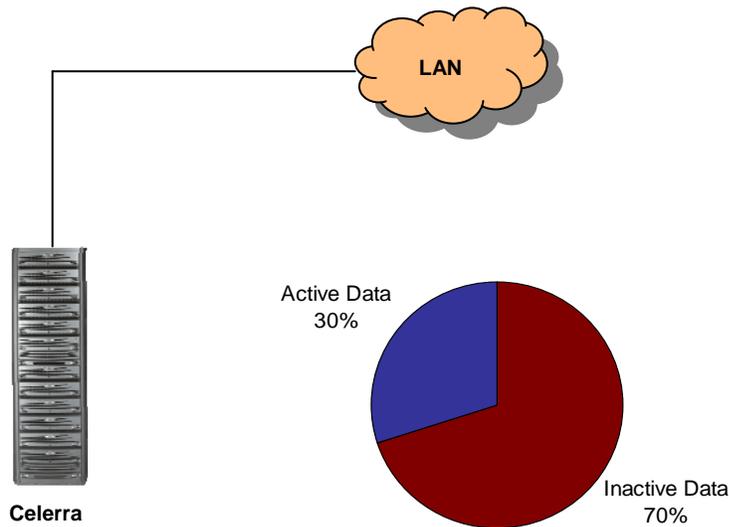


Figure 18 – Typical ratio of active to inactive data.

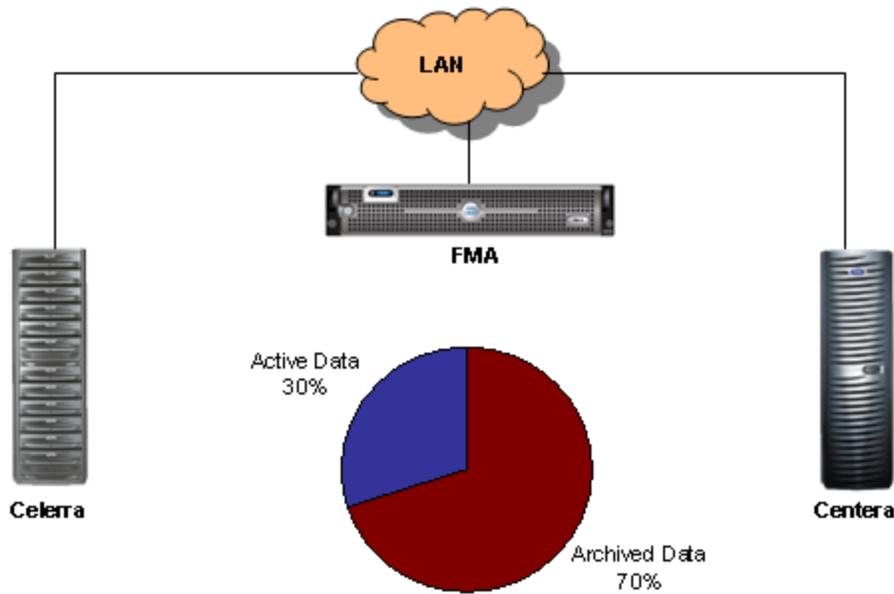


Figure 19 – Potential archive yields utilizing FMA & Centera on initial scheduled archive freeing up valuable space on the Celerra.

FMA supports archiving from Celerra or NetApp to a Celerra, NetApp or Centera. Using a Centera solution includes additional benefits such as single instance storing, assured authenticity and self healing. Stub files created by FMA are 8KB in size so if an original document was 10MB in size and 1000 people each had the same exact document, only 1 of those documents would be archived to our Centera. In effect, this 8KB (Kilobyte) stub references 10GB (Gigabyte - 10GB = ten thousand megabytes) of data!

Rainfinity FMA and GFV form a complete file virtualization solution. It simplifies NAS management, making it easy for storage administrators to manage capacity and performance utilization, reduce storage TCO, and transparently migrates open files without end user disruption.

Accolades

"The EMC Rainfinity Global File Virtualization solution has helped our IT organization achieve dramatic efficiency gains and a higher level of control over our enterprise file systems. Our IT organization has been impressed with the Rainfinity solution's ease of use across our mixed-vendor environment, and we're thrilled it now provides native file archiving support."

- Sr. Systems Administrator, Applied Physics Laboratory of Major U.S. University

"EMC's Rainfinity continues to be the #1 in use file virtualization technology, with NetApp trailing right behind".

Source: TheInfoPro

The Future Direction of Rainfinity

EMC developers continue to drive innovative new solutions to make Rainfinity GFV the sustainable leader in this space. Recent improvements to bolster FMA's archiving credibility allow for a complete conversion from CFA to FMA. Other additions include enhanced reporting, SNMP support and Celerra VDM (Virtual Data Mover) support to name just a few. Recent GFV enhancements include SNMP support to ease monitoring and management, support for new operating systems, and Celerra VDM support.

The adoption rate to implement a global unified namespace (DFS or Automount) has been rather low for enterprise customers since deploying this type of technology requires a very long and tedious rollout. To address this obstacle, the development of a Global Unified Namespace appliance is underway. This will ease the burden and to allow all the benefits of deploying a namespace where end users will be managed from a single appliance.

File serving environments with both CIFS and NFS may combine these two protocols to allow simultaneous access; this is referred to as multi-protocol. The upcoming release will allow GFV to migrate both protocols simultaneously. This is a very compelling advantage as GFV is only capable of migrating this type of data off-line today.

Additional Information

EMC PowerLink and Rainfinity Central contain a wide variety of information pertaining to Rainfinity GFV and FMA for marketing, sales and technical information. In fact, much of the information contained in this paper is referenced in some of these documents. By reading and understanding what Rainfinity has to offer, I hope to aide you in your sales efforts or provide your organization with a more robust tool to manage your existing NAS infrastructure.

May I suggest a next step? Please contact your local Rainfinity Technology Consultant, or IP Representative. We are a world wide, world class organization and I am sure that any of my colleagues would gladly discuss Rainfinity Global File Virtualization in more detail if offered the opportunity.