EMC AND CISCO: GLOBAL FILE VIRTUALIZATION INTEGRATING RAINFINITY WITH A CISCO NETWORK INFRASTRUCTURE

TABLE OF CONTENTS

1. INTRODUCTION

The movement toward the virtualization of services is a key trend affecting the enterprise and data center networks. Increasingly intelligent IP network infrastructure allows for virtualization services to be incorporated directly into the network and delivered within a single framework.

This white paper elaborates how EMC® Rainfinity® Global File Virtualization™ and Cisco® networking technologies interoperate by discussing common deployment scenarios in customer environments.

Rainfinity's split-path virtualization architecture leverages the existing IP network infrastructure and interfaces with Cisco Catalyst® LAN switches. Rainfinity communicates with the Catalyst switch software, and acts as a layer-2 device, like Cisco PIX® firewall's transparent mode, when performing virtualization functions in-band.

EMC and Cisco Systems® share a common goal to develop joint networked storage solutions that are based on open architectures and industry standards. This jointly produced white paper examines Cisco compatibility and best practices to ensure the combined solutions meet network availability, scalability, security, and manageability requirements.

2. RAINFINITY OVERVIEW

EMC Rainfinity Global File Virtualization is an infrastructure management platform that simplifies management of file services by creating a virtualized pool of storage across heterogeneous NAS and file servers. Rainfinity abstracts file access and increases data mobility, so that networked storage resources can be managed independently of users and applications that utilize them. Rainfinity benefits include lowering capital expenditures and reducing the total costs of ownership, by increasing storage utilization, optimizing performance, deploying storage tiers, accelerating consolidation, and improving disaster recovery.

Diagram 1. EMC Rainfinity Global File Virtualization

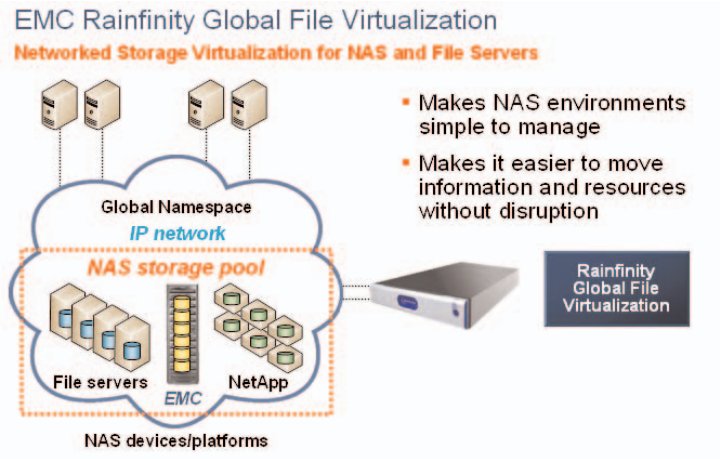

Rainfinity is built on a unique split-path virtualization architecture that integrates standards-based out-of-band global namespaces with selective in-band file protocol processing capabilities. In-band processing enables Rainfinity to synchronize and redirect file access over the network, which delivers the effect of expanding the storage and file serving capacity beyond the physical implementation. The transparent file protocol switching technology understands CIFS and NFS protocols, and provides continuous read/write access while relocating files. This protocol switching technology filters and forwards packets. Most processing occurs below the TCP/IP stack at the kernel-mode driver level.

Rainfinity is configured as a transparent layer-2 bridge when performing virtualization functions in-band, so that it can inspect packets without requiring changes to existing network routes and IP addresses. Transparent bridging is a common technique used by network security devices such as firewalls, VPN concentrators, and intrusion detection devices, for its benefits and strengths. Transparent bridging reduces configuration and deployment time, and requires less processing overhead than routing. Transparent bridges do not appear as an IP route hop to connected devices as they function at the Ethernet layer. Rainfinity operates similarly to Cisco PIX firewall's transparent mode, by creating a virtual interface that bridges a pair of VLAN ports on the same network segment.

Rainfinity is connected to the Catalyst LAN switch where NAS heads and file servers are connected. In the following simple diagram, the source and destination NAS heads are on the same IP subnet. A management IP address is required to access Rainfinity. There are minimum changes required to the network topology. Before deploying Rainfinity, a new private VLAN, "VLAN2", is created with no IP address or routing interface. Rainfinity is connected to VLAN1 and VLAN2 on that switch, and is configured to create a bridge between these VLANs. Deploying Rainfinity as a transparent bridge has no effect on the traffic path between clients and NAS heads, unless specific NAS heads are virtualized by changing the VLAN associations for the ports connected to them.

Diagram 2. Before Rainfinity Virtualization

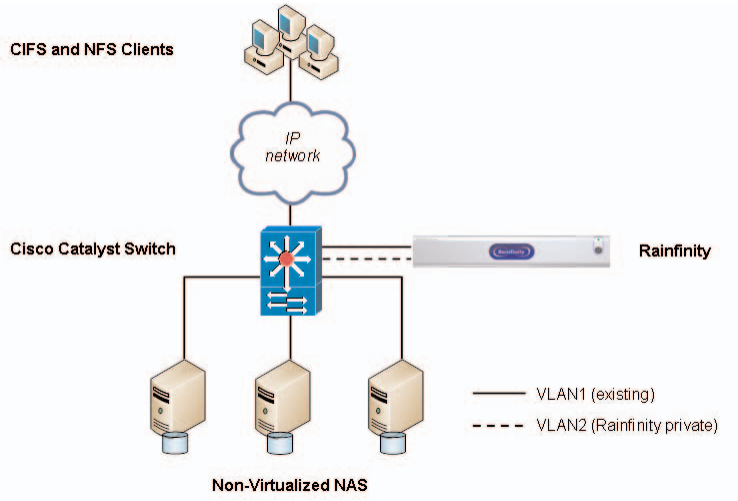

When the ports connected to the NAS heads are reassigned from VLAN1 to VLAN2 on the Catalyst switch, all network traffic to and from these NAS heads flows through Rainfinity. Packets from clients are switched to the VLAN1 port connected to Rainfinity, and Rainfinity forwards them to the source and destination NAS heads using the VLAN2 interface of its bridge. Network traffic to and from the non-virtualized NAS heads completely bypasses Rainfinity and flows as it did before.

Diagram 3. After Rainfinity Virtualization

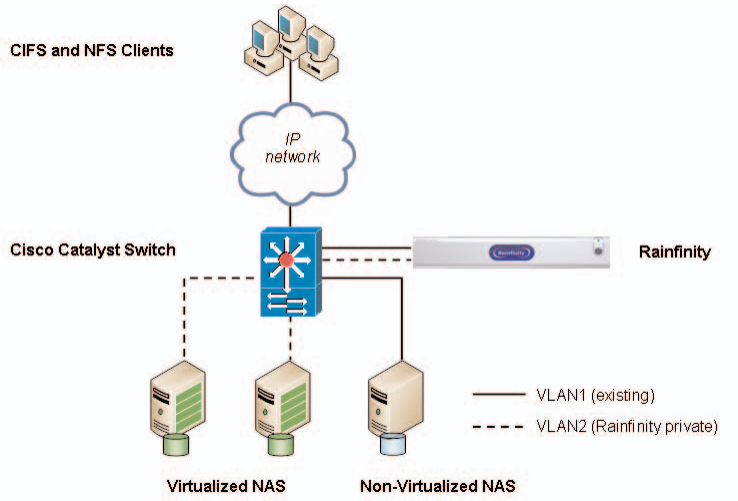

Rainfinity is flexible to support different network topologies where the source and destination NAS heads are on different IP networks, and when they are multi-homed with multiple different IP addresses on different networks. In order to integrate with complex network topologies, Rainfinity can be configured with multiple bridges, and multiple Rainfinity units can be used in a distributed deployment configuration where NAS heads are on different networks and in geographically diverse locations. To determine the configurations that best support your specific network topologies, consult with the EMC's Rainfinity Global Practice group.

When the source and destination are on different IP networks, there are two configuration options. One option is to deploy a Rainfinity unit at the source network, configure it with two bridges, one each connected to the VLAN of the source and destination, and span VLANs across the networks. Another option is to deploy two Rainfinity units, one at the source network and another at the destination network, and configure each unit with one bridge. The local and remote Rainfinity units communicate with each other through an IP-based tunnel, so that packets from both units are routed through the existing LAN and WAN infrastructure. To accommodate for a multi-homed NAS head, Rainfinity can be configured with one bridge for each different network from which the NAS head receives connections.

3. LAYER-2 INTEROPERABILITY

Rainfinity filters and forwards packets by MAC addresses as it operates as part of the Ethernet network. Rainfinity's transparent bridge implementation complies with the IEEE 802.1D Spanning Tree Protocol (STP) standard, and it is compatible with the IEEE standard BPDU format. The use of MAC addresses, MTU, frame status, and multicast address ranges follows the standard.

Rainfinity interoperates with other STP-capable devices in the network to recover from link failures and guard against undesirable loops. Because Rainfinity acts as a switch and not as a server, ports connected to the Catalyst switch should disable PortFast so that it can provide path redundancy and interoperate with BPDU Guard. The Spanning Tree bridge priority should be appropriately set relative to other STP-capable devices, in order not to have Rainfinity as the root bridge and minimize STP convergence.

A properly functioning layer-2 network has one active path between two end points. STP forces redundant paths into the blocking state, and changes a redundant path to the forwarding state when a network segment becomes unavailable from an active link.

When Rainfinity is integrated into a highly available redundant network, it is connected to both the active and standby Catalyst switches. The active-standby mode prevents multiple bridges from connecting to the same network segment so that bridge loops do not occur. Rainfinity transparently passes layer-3 traffic from routers, allowing interoperability with high availability IP services such as HSRP and VRRP.

Diagram 4. Rainfinity in a highly available redundant network

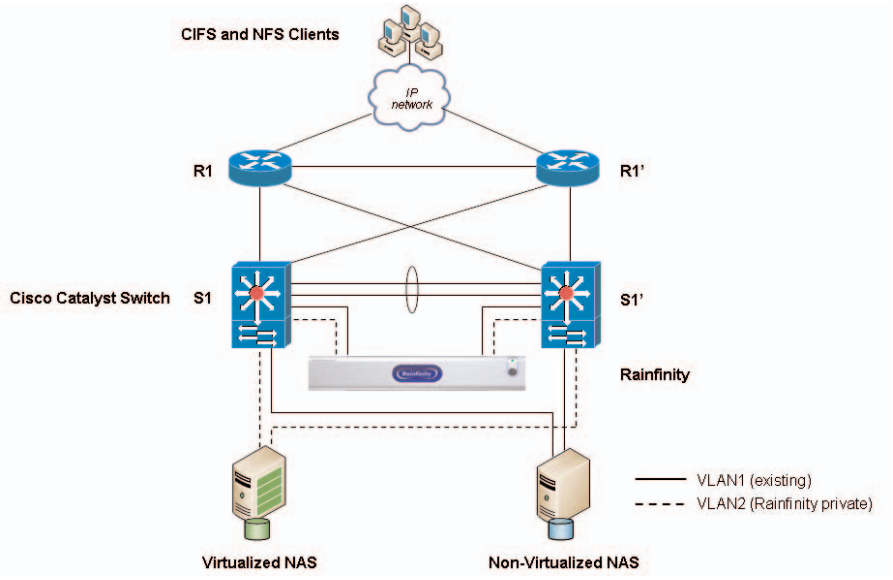

Rainfinity is a highly available hardware configuration with dual hot-pluggable power supplies, mirrored SCSI disk drives, and clustering capabilities that provide enhanced fault tolerance. Rainfinity utilizes a redundant heartbeat network to support failovers when deployed as a cluster. The cluster operates in the active-standby mode where the secondary unit automatically becomes active when the primary unit fails. Upon failover, the active-standby designation is switched. Administrators can also perform a manual failover. Both units are connected to the Catalyst switch, and Rainfinity manages multiple bridges between units to prevent bridge loops, similar to the previous example of integrating Rainfinity into a highly available redundant network.

Diagram 5. Rainfinity cluster in a highly available redundant network

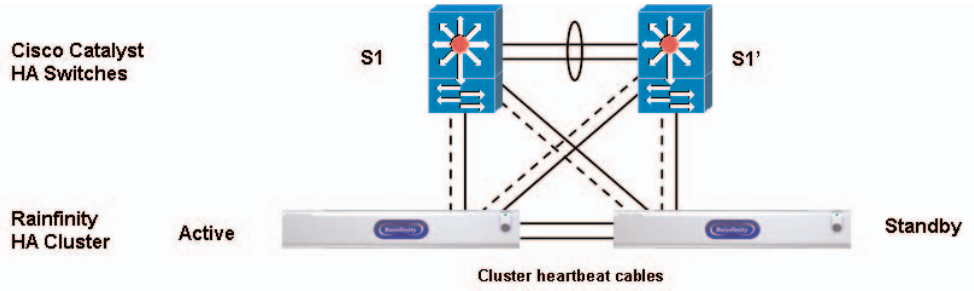

Rainfinity comes with 8 Gigabit Ethernet ports and a virtual bridged interface can aggregate multiple physical ports to create a logical link. Transparent bridging operates at wire speed. This port aggregation feature supports the Cisco EtherChannel and IEEE 802.3ad Link Aggregation standards. STP sees aggregated ports as a single logical link.

4. INTERFACING WITH CISCO CATALYST SWITCHES

Rainfinity's transparent and scalable architecture allows it to move into or out of the data path for a particular NAS head on the fly. To automate this process, Rainfinity interfaces with the Catalyst switch running IOS® or CatOS software, and sends commands to reassign specific switch ports between the original VLAN to the newly created Rainfinity's private VLAN. To reduce layer-2 disruptions and accelerate STP convergence, Rainfinity also sends commands to clear the MAC address of the NAS head from the CAM table when a VLAN change occurs. In addition, Rainfinity clears its own bridge table. Clearing the layer-2 forwarding table of the switch and Rainfinity ensures continuous CIFS and NFS access.

Rainfinity logs into the Catalyst switch remotely by running an Expect script that establishes a Telnet session. The script is configured to prompt for a login user and password or to store them in an encrypted format. Rainfinity executes this script when an administrator requests a NAS head to become virtualized or non-virtualized. The appropriate switch commands, with the specific port numbers and VLAN IDs associated with a particular NAS head, are sent to the switch software. The process of dynamically changing the split-path mode is seamless and transparent. Alternatively if a secure method is preferred, this process can be integrated into the site specific change control procedure by automatically sending a notification message to the network administrators responsible for making configuration changes to the Catalyst switches.

Rainfinity is accessible for management from the web browser using HTTPS (port 443) or HTTP (port 80), or from the command line using SSH (port 22). Rainfinity is a secure system that uses only necessary inbound ports. Rainfinity is a purpose-built hardware appliance with an integrated operating system kernel that is modified and hardened. Rainfinity uses a local authentication system that uses a password hashing algorithm. Administrative and operator consoles can be defined by configuring read/write and read-only access control.

5. HOW RAINFINITY MEETS NETWORK AVAILABILITY, SCALABILITY, SECURITY, AND MANAGEABILITY REQUIREMENTS

Cisco's intelligent switching and routing infrastructure allows for EMC's Rainfinity Global File Virtualization services to be delivered within a single network framework. A properly implemented solution increases the effectiveness of virtualization and achieves a greater business impact. This section summarizes how Rainfinity meets the availability, scalability, security, and manageability design criteria when it is integrated into the network infrastructure.

Availability

- Rainfinity is not a routing device and it operates at the layer-2 in Ethernet networks like the Cisco PIX firewall transparent mode.
- Transparent bridging complies with the IEEE 802.1D STP and it is compatible with the IEEE BPDU format. The use of MAC addresses, MTU, frame status, and multicast address ranges follows the standard.
- Rainfinity integrates with redundant switched networks and participates in STP convergence. Rainfinity also seamlessly interoperates with highly available IP services such as HSRP and VRRP.
- Rainfinity provides a high availability cluster that automatically fails over and converges to a to a new network topology with STP. This fault tolerant design is integrated with the transaction-based system at the application level to ensure data integrity.

Scalability

- Transparent bridging operates at Gigabit wire speed. It is a low overhead and simple device that is implemented below the TCP/IP stack at the kernel-mode driver level.
- Gigabit port aggregation supports the Cisco EtherChannel® and IEEE 802.3 ad Link Aggregation standards.
- Gigabit ports support jumbo frames with up to 9,000 byte MTU.
- Selective in-band file protocol processing filters ports connected to NAS and file servers at the Catalyst switch.
- Global namespace services are provided out-of-band.
- Distributed deployment allows for multiple Rainfinity physical units to work cooperatively as a single logical unit.

Security

- Transparent bridging is a common technique used by network security devices. The stealth mode makes its invisible from the outside world.
- Management console access is limited to using a few inbound ports: HTTPS and HTTP for GUI, and SSH for CLI access. Rainfinity can be protected using firewall and intrusion detection security devices to control access.
- Local authentication uses password hashing algorithm. Passwords are stored in an encrypted format when interfacing directly with Catalyst switches.
- Role-based security provides administrator and operator access control.
- Integrated operating system kernel is modified and hardened for security.

Manageability

- Transparent bridging simplifies deployment, as existing routes and IP addresses do not need reconfiguration.
- Rainfinity is a file virtualization device that is typically used by storage administrators managing NAS and file servers. The management console provides access to all the features including availability and health monitoring capabilities.
- Rainfinity as a network device can be monitored and managed using the existing network management framework.
- Rainfinity provides logging and trouble-shooting utilities.

6. CONCLUSION

EMC Rainfinity Global File Virtualization is designed for Cisco interoperability and seamlessly integrates with the existing Cisco network infrastructure.

Transparent bridging is a well-established concept that is used by firewalls and other types of network security devices including Cisco PIX Security Appliances for its benefits: Ease of deployment, performance, and security. Rainfinity uses the same technique to be able to filter and forward packets.

Rainfinity's layer-2 bridging interoperates with the IEEE 802.1D Spanning Tree Protocol standards. When deploying Rainfinity in a cluster configuration or integrating it into a highly available network with redundant Ethernet switches, the hot-standby topology prevents bridge loops by keeping only one active path open.

Rainfinity adapts to various network topologies. Multiple physical ports may be aggregated to create a logical interface to increase bandwidth. Complex network topologies are supported by creating multiple bridge groups and deploying multiple Rainfinity units in the distributed deployment configuration.

Rainfinity uses a standard method to communicate with the Cisco Catalyst LAN switches, and supports both IOS and CatOS software. Rainfinity sends commands to these switches that reassign VLAN associations for the ports connected to NAS and file servers. Changing the VLAN associations establishes a transparent bridge and enables file protocol switching. This process can be fully automated or integrated into a site-specific change control procedure by sending a message to the network administration group.

The combined EMC and Cisco solution allows for file virtualization services to be incorporated into the intelligent switching and routing infrastructure, and delivered within a single network framework. The EMC Rainfinity Global File Virtualization solution meets the availability, scalability, security, and manageability design criteria for enterprise and data center networks.