NetApp®

Technical Case Study:
Red Hat Deploys Network Appliance™ Storage to Support a Rapidly Expanding All-Linux® Environment

Network Appliance | May 2005 | TR-3398

TECHNICAL CASE STUDY

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

redhat.

Executive Summary

Red Hat, Inc., the world's leading Linux and open source provider, is headquartered in Raleigh, North Carolina, with satellite offices spanning the globe. Red Hat is the leading Linux and open source solutions provider to the enterprise. Red Hat offers operating system solutions as well as middleware, applications, and management solutions and services.

Over the past seven years, Red Hat has deployed a wide variety of NetApp solutions to address its storage needs throughout the company. Network Appliance and solutions integrator Datalink (*www.datalink.com*) have teamed to provide Red Hat with an optimal combination of products and services for its highly demanding, high-growth environments:

- **Red Hat Network**. The Red Hat systems management platform, Red Hat Network (*www.redhat.com/software/rhn/*), utilizes clustered NetApp storage in three locations. NetApp SnapMirror® software is used to distribute software to each Red Hat Network location and to mirror copies of Red Hat Network content to multiple sites as part of a robust disaster recovery strategy.

- **Corporate IT.** Red Hat's mission-critical business applications—including Oracle® E-Business Suite—rely on clustered NetApp storage. Red Hat leverages SnapMirror to mirror all data from headquarters to a standby site, which can be operational within 20 minutes if necessary.

- **Engineering.** Red Hat's engineering organization utilizes a variety of NetApp primary and nearline storage systems to provide highly available, well-protected storage for build trees and source repositories.

The combination of highly available NetApp storage; NetApp Snapshot™ technology; and NetApp SnapRestore®, SnapMirror, and SnapVault® software throughout the company allows Red Hat to ensure data availability, increase staff productivity, and simplify data management.

Introduction

Network Appliance and its dedicated team of channel partners are committed to the delivery of storage solutions that help enterprise customers simplify storage management, increase flexibility, improve data protection, and decrease total cost of ownership.

This case study describes the selection and global deployment of NetApp technology by Red Hat, Inc., the world's leading provider of Linux and open source solutions. Working with NetApp STAR partner Datalink, Red Hat has leveraged NetApp solutions to meet critical storage, data management, and disaster recovery requirements. A strong commitment to Linux ensures that NetApp solutions work seamlessly with Red Hat's all-Linux infrastructure.

Customer Challenges

- Nearly 40% annual growth
- Core revenue-generating application requires 24x7 availability
- >40 hour full backups

NetApp Value Proposition

- Flexible, unified storage
- Commitment to Linux
- Comprehensive data protection
- Replication and disaster recovery
- Reduced TCO

Business Benefits

- Utility-like availability for an all-Linux environment
- Simplified mirroring between sites
- Full disaster recovery for critical Oracle ERP software
- Improved protection for critical engineering data
- Centralized backup of non-NetApp storage in remote offices

Background

Founded in 1993, Red Hat rapidly became a leading enterprise software and support provider. Red Hat also contributes directly to the growth of open source solutions through sponsorship of community-supported projects such as the Fedora Project. For fiscal year 2004, Red Hat reported a 39% increase in revenues over the previous year. Accelerating growth creates significant challenges for the company's IT infrastructure.

Red Hat relies on a complex environment that includes three main data centers in the United States, five development centers in three countries, and dozens of engineers worldwide working from their homes. Red Hat uses Linux exclusively in all aspects of its operations.

NetApp Deployment at Red Hat

When Red Hat first began using NetApp storage over seven years ago, it relied on direct-attached SCSI storage arrays and software RAID. At that time, most of the widely used Linux file systems lacked journaling. If a large server went down, it could take hours for the file systems to be checked and the system to come back online. Moving to reliable, networked storage from NetApp was therefore an immediate win.

In the ensuing years, Red Hat has expanded its strategic relationship with Network Appliance and deployed highly flexible NetApp storage to provide a centrally managed storage architecture to support key business areas: Red Hat Network, Corporate IT, and Engineering.

Business Area	NetApp Systems	Capacity
Red Hat Network	FAS960 (4)	10TB
Corporate IT	FAS960 (2)	10TB
Engineering	FAS940, F840, F760, NearStore® R200	26TB

Table 1) Red Hat NetApp deployment.

Red Hat Network

Red Hat Network is the cornerstone of Red Hat's subscription-based business model. This complete systems management platform for Linux enables customers around to globe to instantly access software updates, manage Linux deployments, and take advantage of sophisticated provisioning capabilities. "If Red Hat Network is down, we're not providing the automated services our customers depend on to help run their Linux-based infrastructure," explains Jay Madison, director of IS Operations.

To provide 24x7 uptime, Red Hat deployed NetApp storage at its primary Red Hat Network data center in Arizona, as well as at corporate headquarters and a secondary site for downloads in Florida. Each FAS960 cluster consists of two systems in an active-active configuration. RAID-DP™ functionality provides data protection from two-disk failures on a single system. Should one system in a cluster fail for any reason, the other automatically takes over, and data access continues unimpeded without any noticeable impact on the end user.

"Our NetApp storage delivers exceptional availability for Red Hat Network," says Madison. "We are now able to treat storage as if it were a utility. The fact that our data is always available and always protected gives us great peace of mind."

All the various Red Hat software distributions and updates available for download are stored on the Arizona cluster. Red Hat uses NetApp SnapMirror software to centrally manage the distribution of its many software packages from its corporate headquarters in Raleigh, North Carolina, to both the Arizona and Florida locations.

Red Hat also uses SnapMirror to replicate data throughout the day between these three sites. Should services from the primary Arizona site become disrupted, both the Florida and North Carolina sites are able to take over all Red Hat Network capabilities within minutes.

Red Hat Network content is currently archived directly from the NetApp systems to tape using NDMP, but the company is in the process of transitioning to an online backup solution based on NetApp SnapVault software and NetApp NearStore nearline storage. In this model, backups can be economically retained on disk for as long as necessary and archived to tape as needed.

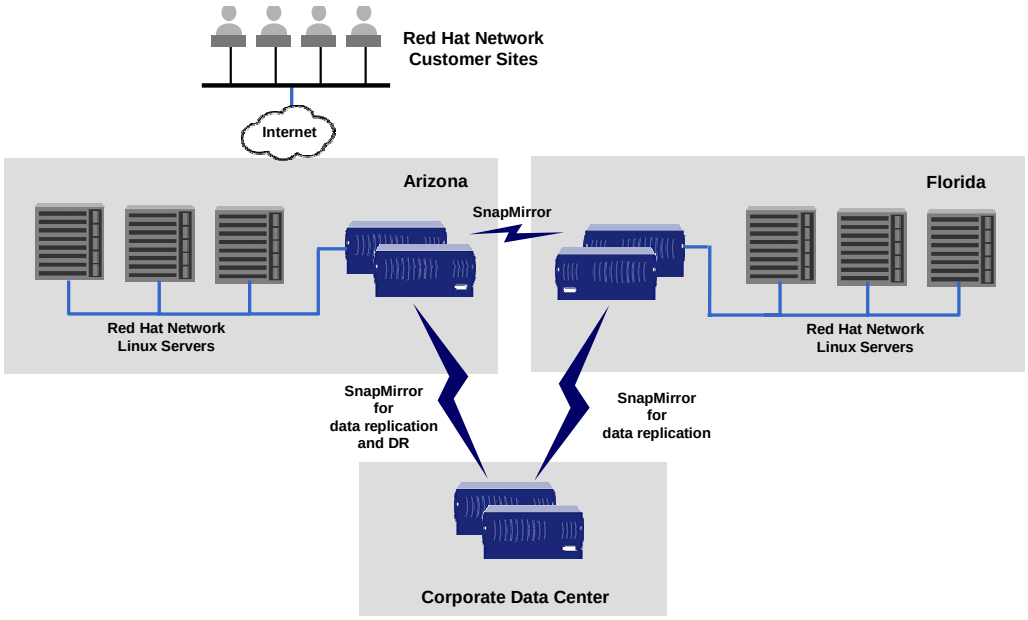

Figure 1. Red Hat Network infrastructure.

Corporate IT

Red Hat has implemented Oracle E-Business Suite to manage all of its business and financial transactions. To maximize availability for its Oracle infrastructure, Red Hat deployed FAS960 systems in its corporate data center in North Carolina.

The databases supporting Red Hat's implementation of Oracle E-Business Suite are backed up using a tight coupling of NetApp Snapshot copies and the Oracle hot backup utility. NetApp Snapshot technology enables the creation of incremental point-in-time copies on a regularly scheduled basis. If an application error were to corrupt the database, Red Hat can almost instantly recover from any of more than 250 saved Snapshot copies. Once the copy is restored, it would be merely a matter of playing back the database redo logs, and the database would again be operational. This process can be accomplished in minutes rather than the hours or days that would be required to recover from traditional tape backups.

Red Hat also maintains a standby environment in its Florida data center. To ensure that data is protected in the event of a sitewide disaster, data on NetApp storage at the corporate data center is continuously mirrored to the Florida site using NetApp SnapMirror software.

If a disaster were to shut down the Raleigh data center, the results of disaster recovery tests give Red Hat confidence that it could resume business operations in Florida in 20 to 30 minutes. Business data is also archived to tape on a daily basis to provide backups and to ensure compliance with regulatory requirements.

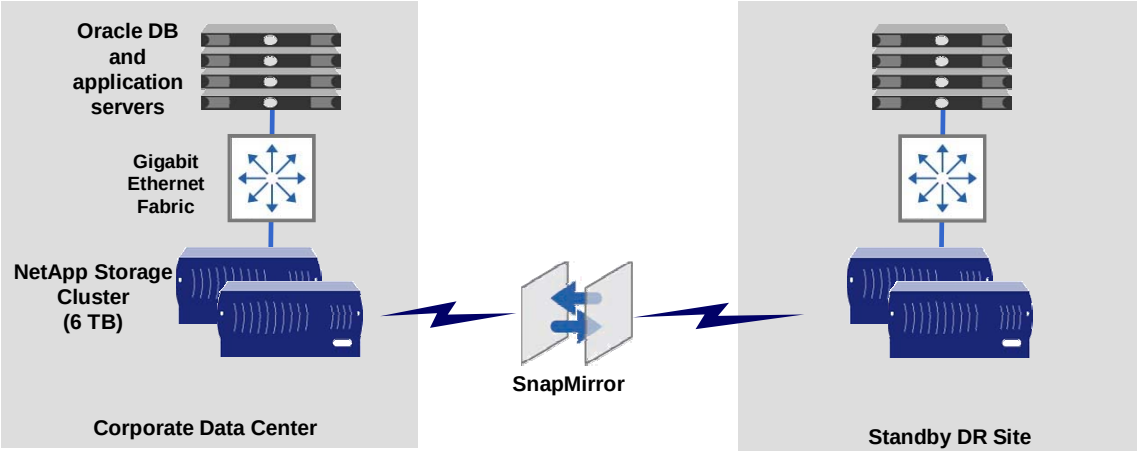

Figure 2. Red Hat infrastructure for Oracle E-Business Suite.

Red Hat Engineering

Initially, Red Hat engineering was responsible for managing its own data storage environment. Engineers would use the data on their storage servers as test beds, sometimes with unexpected results to data integrity. Since full backups of several terabytes of archived engineering data required up to 40 hours to complete, tape archiving did not occur on as frequent a basis as needed. Data restorations from tape frequently took even longer and were not 100% reliable.

Today, the primary repository of Red Hat engineering data is a standalone FAS960 archiving 2TB of the most critical engineering data at headquarters in Raleigh. The data repository is protected using NetApp SnapVault software to back up the data to a NetApp NearStore system, which is then backed up to tape weekly to create space for the next week's backups.

In addition to the headquarters location, Red Hat has development centers in Massachusetts, Alabama, and California. NetApp storage appliances serve much of the engineering data, including build trees and historical repositories.

The Power of Partnerships

Red Hat trusts Datalink, a Network Appliance channel partner, to supply all of its NetApp technology needs. Datalink is an information storage architect specializing in the analysis, design, implementation, and support of information storage infrastructures and a member of the exclusive NetApp Storage Authorized Reseller (STAR) program. The company is also a NetApp Registered Service Provider, enabling it to provide a very high level of service to enterprise customers.

NetApp depends on highly skilled partners such as Datalink to extend its reach and help ensure that all purchasers of NetApp solutions get the personal touch that will enable them to succeed. Red Hat's Madison cites a recent example: "A few years ago Red Hat was shipping a preconfigured NetApp cluster from Raleigh to another data center for immediate deployment. When the cluster was damaged after it fell off the shipping truck, Datalink worked tirelessly with NetApp to get us an expedited replacement within 48 hours. Datalink and NetApp continue to work closely together to meet our needs."

Overall Impact

Network Appliance solutions offer substantial benefits for fast-growing companies such as Red Hat, especially companies with a strong commitment to Linux. "The snap technologies—Snapshot, SnapRestore, SnapMirror, and SnapVault—provide the biggest benefit to me on a daily basis," points out Red Hat senior system administrator Nathanial Golnik. "Once I configure the software, data transfers and backups occur automatically without any problems. My life has gotten a lot easier since we began using these tools."

NetApp contributes directly to the productivity of Red Hat's end users and IT staff. "NetApp allows us to deliver a higher SLA for both Red Hat Network and our internal end users," continues Madison. "Data is always available, and users can usually recover any files they accidentally delete from online Snapshot copies without administrator intervention. This not only improves end-user productivity but also saves my IT staff a lot of time, so they are more productive."

"Network Appliance has made a significant effort to provide storage solutions that are optimized to work with Linux," sums up Red Hat executive vice president, Worldwide Operations Joanne Rohde. "The great synergy between Red Hat Enterprise Linux and NetApp technology ensures that our storage infrastructure delivers the utmost in performance and availability for our busy servers and desktops."

Conclusion

With NetApp technology, Red Hat has experienced greater availability, improved productivity, and reduced administrative burden. As important, NetApp people and partnerships with companies like Datalink also make a tremendous contribution. "We rarely know what's coming at the end of the day," concludes Madison. "Unusual requests and large requests for storage are par for the course. Because of Datalink and NetApp, we are always able to deliver. We enjoy a mutually beneficial relationship between solution provider, vendor, and customer. Everything is tightly integrated and coordinated, and that is helping us succeed."

NetApp®

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com