PolyServe™ A company of hp

Case Study

Music India OnLine

The Challenge

- Migrate the local content stores of 12 web servers to a centralized repository serviced by a scalable NAS gateway solution powered by PolyServe
- Ensure high scalability and availability, now and in the future
- Improve site performance
- Simplify and reduce cost of system management

The Solution

- Two-node cluster of Intel Architecture-based servers servicing NFS requests
- Twelve Web servers running Real Networks Helix and Apache
- PolyServe Matrix Server software
- Red Hat Enterprise Linux
- Two Brocade Silkworm switches
- EMC Clariion CX200 Storage Area Network array

Results

- Saved more than 50% of the CAPEX costs compared to a traditional NAS solution
- Superior NFS scalability and system availability at a lower cost
- Simplified system management by creating a single repository of audio files for multiple Web servers
- Enabled 24x7 client-access to digital media
- Created operational flexibility for executing file serving, back-up, and scripting tasks from the cluster
- Facilitated non-disruptive expansion of server and storage capacity

MusicIndiaOnline.com Migrates to a High Performance, Affordable NAS Platform Powered by PolyServe

Entertaining nearly seven million fans of Indian music, MusicIndiaOnline.com is the world's leading online source for Indian music. In fact, the Web site currently offers a selection of more than 130,000 Indian music titles dating back to the 1950s. MusicIndiaOnline.com is one of several online businesses, including greetings.com and Artineed.com, operating under the ownership of ASV Cyber Solutions.

The purpose behind MusicIndiaOnline.com is to provide fans of Indian music an online outlet to search, access and listen to the songs of their choice. Rather than provide a site for music downloads, the site streams audio files, allowing users to listen to their favorite Indian music, on demand, through a Real Audio or Window Media player. The Web site has been up and running for nearly six years. Truly a pioneer in this type of online audio streaming service, MusicIndiaOnline.com (MIO) primarily markets to the Indian population in the U.S. Today, the site is enjoying very aggressive growth, remaining the market leader in its musical category.

THE CHALLENGE

As a result of MIO's growing popularity, the site has consistently experienced a doubling of its traffic every three months. Such quantifiable growth figures offered solid proof that MIO's existing storage methodology would simply not work. MIO used a combination of direct-attached storage in each Web server, and replication software to synchronize content among the servers. As a result, MIO's architecture had many data silos that required significant management to ensure the data was synchronized and there was enough storage capacity to accommodate growth. As the content repository continued to grow, storage capacity of each machine, which was limited to between 330GB to 400GB, became a major headache.

As the number of servers required to support the system grew, matters of replication took their toll on the site's performance. At one point, MIO was waiting up to nine hours to replicate the site's content. At the same time, no interim upgrades could occur during the replication. These types of performance deficiencies had customer complaints on the rise, as users were having difficultly getting on the site and accessing audio files in a consistent manner. Amidst their sea of servers, MIO was having trouble pinpointing the actual source of particular problems. In the end, they realized the source of the problem was the architecture itself. The storage backbone had hit a wall.

To properly address their concerns, MIO decided they needed a central repository of data to enable an efficient, non-disruptive method for scaling out their content, now and in the future. Given the growth pattern in site traffic, it was critical to have a storage platform that could grow with the increased demands, as well as an ever-increasing library of audio content.

Additionally, MIO needed to ensure high reliability and back-ups within a new architecture. The search for a single, highly available, scalable and redundant storage and delivery platform began.

Keeping cost to a minimum, while making significant gains in scalability, availability and performance, was the goal. The following criteria were used in evaluating the viability of Network Attached Storage (NAS) offerings:

- Minimize cost of a new architecture
- Industry leading price-performance ($ per MB/s)
- Eliminate any single point of failure
- Ensure high scalability in terms of storage capacity and performance (MB/s)
- Maximize availability
- Ensure redundancy
- Simplify system management

In addition to PolyServe Matrix Server cluster file system, MIO evaluated a multi-head NAS gateway from a major storage vendor and a clustered single-head NAS solution, as well as an additional cluster file system.

> *"With the PolyServe solution, system and storage management is dramatically simplified. All my content is in one place – no more issues with replication or running out of storage capacity."*
>
> ***Bharath Aiyer***
> ***Vice President of Operations***
> ***MusicIndiaOnline.com***

The multi-head NAS gateway solution was eliminated quickly from any practical consideration due to its high price point. It was more than three times the

cost of a PolyServe powered NAS offering. MIO found that it had no additional features that could justify the significantly higher cost.

The clustered single-head NAS solution also came with a high price tag — double the cost of Matrix Server. Along with the high price tag, MIO discovered this NAS solution suffered from substantially less aggregate I/O throughput, had limited expansion capabilities and proved a bit difficult and rigid to manage from an operational standpoint.

In addition to PolyServe Matrix Server, MIO evaluated another clustered file system (CFS). However, they quickly eliminated it from consideration because the product had a single point of failure and proved difficult to setup and configure.

After extensive evaluation, MIO concluded that the superior option that provided even better performance and a true cost advantage was PolyServe Matrix Server shared data clustering software. Neither the other CFS nor any of the other NAS offerings could provide what PolyServe Matrix Server was delivering:

- No single point of failure
- Easy scale out of storage capacity to accommodate projected content growth
- Easy scale out of performance by simply adding more low-cost industry-standard Linux servers
- Multi-purpose Linux servers that could be deployed concurrently as NAS gateways and for other tasks such as executing scripts, data back-up and Web serving

THE SOLUTION

The scalable NAS gateway solution from PolyServe was unique in its ability to combine the benefits of NAS with the advantages of SAN. What's more, PolyServe Matrix Server cost well over 50 percent less than other NAS products considered. With Matrix Server, MIO was able to deploy a cluster of affordable Linux-based, industry-standard servers in an environment where each server in the cluster can simultaneously read and write to volumes of shared data on their SAN. Because their Linux server cluster could be treated as a single NFS server, there was no need to divide or replicate data among separate servers and no need to assign clients manually to individual servers.

And with Matrix Sever, MIO's data can be shared among all servers. Storage provisioning and back-up procedures were greatly simplified. As a result, MIO could deliver the desired high availability for audio streams through their Real Networks Helix servers with full protection against server, network and storage failures. With Matrix Server, the cluster detects and takes prescriptive action to failures at the network, storage, server, and software layers within the cluster. PolyServe Matrix Server also provides transparent fail-over of NFS sessions and improves system uptime and service quality for the entire cluster, plus offers support for Multi-Path I/O (MPIO).

Matrix Server's symmetrical cluster file system also gives MIO seamless, simple, on-the-fly scalability to meet growing workloads, essential to maintaining and growing the site's audio content and performance with no degradation in service. In the Matrix Server cluster, each individual server contributes incremental performance to the cluster. No individual server can impose a bottleneck on system performance. As MIO's needs grow, they can simply add additional servers to the cluster, or storage to the SAN, without service disruption to the NFS cluster.

"I don't understand why anyone would buy another traditional NAS appliance. The PolyServe NAS solution provides better performance, at half the cost of a competing NAS solution. Moreover, I get critical features like high availability and pay-as-you-grow expansion for free."

Bharath Aiyer
Vice President of Operations
MusicIndiaOnline.com

RESULTS

With the new Matrix-Server-enabled NAS solution, MIO users are now accessing audio files via any of the servers in the cluster. Also, multiple users who have mounted the audio files via different servers are using the same files simultaneously, receiving the same results as if they were using a single server.

MIO expects to see the PolyServe NAS gateway's performance grow proportionately with its site traffic — doubling every three months. The site architecture, in total, is handling nearly 1.2TB, yet has the capacity to scale out to 2.5TB.

"With the Matrix Server shared data cluster solution, I can pay as I grow. I just plug in additional servers and storage capacity as my operations need them. That is critical when my user traffic is doubling every three months!"

Bharath Aiyer
Vice President of Operations
MusicIndiaOnline.com

A cost sensitive project from the outset, the Matrix Server scalable NAS solution saved MIO more than 50 percent in capital expenditures, compared to competing NAS solutions. To address increasing demands on their site, MIO has the available capital to easily scale out their architecture by adding low-cost, industry-standard servers to the cluster, enhancing the site's available bandwidth. The highly scalable and available Matrix Server solution has also introduced tremendous advantages in terms of manageability. Data replication is no longer necessary, there is no single point of failure and cluster, and file system configurations can be changed from a central point of control.

CONFIGURATION

The two-node PolyServe Matrix Server cluster comprises two Linux-based Dell 1750s, providing a highly scalable and available environment to support 12 Web servers running Real Networks Helix and Apache software. The two-node cluster is connected to an EMC Clariion CX200 SAN via two Brocade Silkworm switches, creating a scalable NAS platform capable of accessing data concurrently at high speeds and with full data integrity.

"Bottom line — the PolyServe Matrix Server gives me flexibility. Unlike, the software that goes into the expensive, 'high-end' NAS appliances, Matrix Server is configurable, doing what I want it to do, and doing it very well."

Bharath Aiyer
Vice President of Operations
MusicIndiaOnline.com

PolyServe, Inc.
20400 NW Amberwood Drive, Suite 150
Beaverton, OR 97006
Toll Free: 877-765-7378
Tel: 503-617-7574
Fax: 503-617-7592

MusicIndiaOnline_CaseStudy_042407.doc