PolyServe™ A company of hp

Case Study

HESS

# Amerada Hess' Data Center Boosts Performance and Achieves Huge Savings with a PolyServe NAS Cluster

## The Challenge

- Build a high-performance NAS solution leveraging inexpensive, industry-standard server and storage building blocks
- Deliver 1.2 GB per second of throughput and support 1,300 HPC clients by 2005
- Eliminate any single point of failure
- Ensure high scalability and availability, now and in the future
- Simplify and reduce cost of system management

## The Solution

- Nine-node NAS Cluster of dual-processor Dell PowerEdge 2650 servers with 2 GB RAM each
- 1,000-node seismic processing HPC cluster
- PolyServe Matrix Server software
- Red Hat Enterprise Linux 2.1
- Brocade Fibre Channel switch
- Six 10 TB Nexsan InfiniSAN ATAbeast storage subsystems

## Results

- Saved more than 80% in hardware and storage costs compared to a traditional NAS solution
- Scaling I/O throughput from peak of 625 MB per second to average of 1.2 GB per second by CY05
- Supporting 1,000 HPC clients growing to 1,300 clients by 2005
- Simplified administration by managing servers, storage and network from a central control point
- No single point of failure, system wide
- Scale capacity and performance as demand dictates

Amerada Hess Corporation (common stock ticker symbol AHC, New York Stock Exchange) is a global, integrated energy company engaged in the exploration for, and the production of, crude oil and natural gas, the refining of crude oil and the sale of refined products, natural gas and electricity.

At the heart of this story are the ongoing, high-impact exploration programs being conducted by Amerada Hess across the globe, including the deep-water exploration in the Gulf of Mexico, West Africa and Southeast Asia. At its E&P offices in Houston, Texas, Amerada Hess maintains a critical data center servicing the seismic processing needs of its global exploration and production operation. Here, the E&P division must process huge amounts of raw seismic data, collected worldwide, to help determine the presence and location of oil deposits.

## THE CHALLENGE

At the Houston data center, the E&P division maintains a high-performance computing (HPC) cluster of approximately 1,000 Linux client servers.

The 1000-node HPC cluster performs seismic processing of raw data collected in the field using several mature, in-house applications. Much like the processing requirements found in biosciences, video production and government, the data access and I/O throughput demands by the 1,000-node seismic processing cluster are immense. At any given point in time, the E&P division estimates it handles 50 to 70 terabytes (TB) of active data. With only a few single file servers in place, performance was suffering, replication was taxing, failover was a problem and capacity was all but tapped out. Furthermore, the E&P division knew their data center would have to keep accommodating annual growth of 25% or more. In fact, the E&P estimates that by 2005, its HPC cluster will grow to 1,300 nodes and need access to 150 TB of data online. The E&P division also had a performance goal of delivering average throughput of 1.2 Gigabytes (GB) per second by 2005.

*"PolyServe software enables us to build a high-performance NAS system with industry-standard building blocks at a price point dramatically below a traditional NAS approach."*

***Vic Forsyth***
***Manager, Exploration and Technical Systems***
***Amerada Hess Corporation***

Amerada Hess is a large user of NAS to support seismic interpretation activities. Naturally, a NAS appliance-based solution had to be considered as a possible solution for its seismic processing requirements. But just as quickly, Amerada Hess' E&P Division dismissed traditional NAS offerings. The division was concerned the NAS appliances would come with a large price tag, would not be able to inexpensively handle E&P's throughput requirements and would drastically increase its costs for high availability and scalability.

Already a large Linux shop, the E&P Division was determined to bring together inexpensive, industry-standard components to create a more robust platform for its heavy and growing data-processing demands. Jeff Davis, Amerada Hess' Technical Lead of Global Technical Computing Infrastructure, was already familiar with the concept of the PolyServe Matrix Server™ shared data software for clusters of Linux- and Intel architecture-based servers. Enamored with the concept, the E&P Division wasted no time putting the PolyServe NAS Cluster solution – powered by Matrix Server – to the test. In the fall of 2003, the E&P Division simulated its desired environment with the PolyServe NAS Cluster solution. Not only did it work; it worked as promised.

Unlike other NAS offerings, the PolyServe NAS Cluster solution delivered:

- A global file system namespace across a cluster of servers and storage
- The ability to aggregate I/O throughput across the cluster for more performance vs. traditional NAS
- The ability to allow all nodes to see and share the same data concurrently with no replication
- A single storage pool for simplified storage management and backup
- A reduced number of single points of failure
- Easy scale out of storage capacity to accommodate projected data growth
- Easy scale out of performance by simply adding more low-cost industry-standard Linux servers

For its 1000-node, HPC seismic processing environment, Amerada Hess' E&P division selected the PolyServe NAS Cluster solution as the enabling software that would couple together its industry standard components to deliver superior NAS performance at a drastically reduced price.

*"The PolyServe NAS Cluster solution, coupled with the outstanding technical support from PolyServe, virtually guaranteed our success. We got the NAS cluster up and running quickly and the performance and simplicity of administration has exceeded my expectations."*

***Jeff Davis***
***Technical Lead***
***Amerada Hess Corporation***

## THE SOLUTION

Amerada Hess has deployed a five-node PolyServe NAS Cluster. For an initial investment of approximately $100,000 including servers, storage and software, Amerada Hess was able to deploy a 3-node PolyServe NAS Cluster solution with 20 TB of redundant storage. They have subsequently added another six nodes to the NAS cluster based on their initial findings and success.

*"To match the performance and capacity of the 60-terabyte, fully redundant PolyServe NAS Cluster, we would have had to purchase a traditional NAS solution costing seven times the price of the PolyServe solution and incurred a significant administration headache."*

***Vic Forsyth***
***Manager, Exploration and Technical Systems***
***Amerada Hess Corporation***

For the initial $100,000 investment, other NAS offerings that Amerada Hess had evaluated would have only supplied a single server, 4 TB solution. In effect, Amerada Hess would have had to buy five of these NAS appliances to match the capacity of the PolyServe NAS Cluster solution. To enable server redundancy, the costs would have even been more significant — 7 times the cost of the PolyServe-enabled cluster. Furthermore, five clustered NAS filers would have meant individual systems to manage, further adding to the administrative complexity.

The PolyServe NAS Cluster solution integrates Network File System (NFS) protocol functionality with a true symmetric cluster file system, high-availability services and cluster and storage management capabilities. The product aggregates up to 16 low-cost Linux- and Intel Architecture-based servers for high-performance, fault-tolerant file serving across a SAN. Shared data and management capabilities ease IT administration by enabling servers and storage to be managed at a single central control point.

The PolyServe NAS Cluster solution acts as a single NAS filer with a single pool of storage and single set of exported files systems. Administration of the system is greatly eased because there is no requirement to create redundant copies of data, or to divide clients manually among different servers. Additionally, the PolyServe NAS Cluster solution enables servers in the cluster to automatically fail over each other's client connections immediately in the event of server failure.

To meet the demands of Amerada Hess' aggressive growth plans, the PolyServe NAS Cluster solution has also introduced unprecedented levels of scalability and performance. With traditional NAS offerings, performance is capped and increased performance means investing in a new, expensive NAS appliance and partitioning data between them. With the PolyServe NAS Cluster solution, Amerada Hess need only invest in more affordable industry-standard servers when they need more performance.

*"The PolyServe software makes it possible for us to purchase and easily add servers and storage when we need it – not before. I'm not forced to buy a whole new NAS system. PolyServe's pay-as-you-grow model is a much smarter NAS deployment option."*

***Jeff Davis***
***Technical Lead***
***Amerada Hess Corporation***

## RESULTS

In initial performance assessments, Amerada Hess saw its three-node NAS cluster service 270 simultaneous clients and deliver I/O throughput performance of 375 MB/second from multiple filesystems. In E&P's environment, each individual PolyServe NAS server can deliver 125 MB/second while servicing approximately 90 simultaneous clients. That said, it would only take a cluster of two nodes running PolyServe NAS Cluster software to match the 250 MB/second peak performance of the most expensive NAS appliances. The unprecedented levels of NAS performance from the PolyServe-enabled cluster are only rivaled by its low cost, which no other competitive offering can match.

Because the PolyServe NAS Cluster can aggregate the performance of up to 16 servers, Amerada Hess will easily support the anticipated growth of its HPC cluster from 1,000 nodes to as many as 1,300 in 2005. To keep pace with the E&P Division's seismic processing demands, the PolyServe cluster is expected to expand to as many as 16 nodes with storage capacity growing from 20 TB to as much as 150 TB.

*"With the PolyServe-enabled NAS cluster in production, performance and capacity expansion on demand has been terrific. We are well on our way to meeting our goal of 1.2 Gigabytes per second of throughput and servicing 1,300 HPC clients from the PolyServe NAS cluster."*

***Vic Forsyth***
***Manager, Exploration and Technical Systems***
***Amerada Hess Corporation***

What's most impressive is that with the relatively low investment Amerada Hess has made in the PolyServe NAS Cluster solution, the E&P Division has the available capital to easily scale out their architecture by adding low-cost, industry-standard servers to their cluster. They saved more than 80% in hardware and storage costs compared to a traditional NAS solution and will continue to enjoy superior NAS performance. Plus, these additional servers will still remain under one central point of control for easy administration.

## CONFIGURATION

The data center recently upgraded the PolyServe NAS Cluster solution to be the storage platform supporting the entire 1,000-node HPC cluster. The file-serving cluster now consists of PolyServe NAS Cluster software, Red Hat Enterprise Linux 2.1, nine dual-processor Dell PowerEdge 2650 servers with 2 GB of memory each, connected to a storage area network (SAN) consisting of six 10 TB Nexsan infiniSAN ATAbeast storage subsystems (60 TB storage capacity total).

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