PolyServe™ A company of hp

Case Study

CASPUR

The Challenge

- Find the highest performance cluster file system for new Linux email application - Refero
- Support more than 800,000 concurrent users
- Integrate easily with open source software
- Ensure high scalability and availability, now and in the future
- Keep acquisition and operational costs to a minimum

The Solution

- Eight-node cluster of dual-processor Dell PowerEdge 2850 Intel® Xeon® processor-based servers for initial setup
- PolyServe Matrix Server™ clustering software
- EMC Clariion cx300 Storage Array
- Brocade fibre channel switches

Results

- Highest performing Linux cluster file system when compared to 4 competitors
- 13.6 times faster than Red Hat Global File System in transactions per second
- 6.8 times faster than Veritas™ Storage Foundation™ CFS in Read and Write KBps
- On-the-fly scalability requiring no system downtimes for system upgrades
- Simplified administration by managing servers, storage and network from a central control point
- Delivered a highly reliable system to support central high-volume email application facilitated non-disruptive expansion of server and storage capacity

PolyServe Matrix Server™ Blows Away the Competition in CASPUR's Benchmark Comparison for eMail

CASPUR, Consorzio Interuniversitario per le Applicazioni di Supercalcolo Per Università e Ricerca, is a computing consortium for nine state universities in Italy and the Italian Ministry of Education, Universities and Research (MIUR). Founded in 1992, it is a non-profit organization, financed by MIUR and by the associated Universities. CASPUR was established to build and manage a High Performance Computing (HPC) center for the universities' computing research and intensive computing needs. Today it is also responsible for the many other computing needs of its university community.

Its mission is to promote the use of the most advanced information processing systems and to develop research programs aimed at a more effective and innovative usage of information and communication technology.

CASPUR's charter includes management of the High Performance Computing (HPC) center, support of the research programs it helps develop and consultation with various university groups on new computing services being considered. CASPUR is also responsible for the development and management of

"Basic Services" for universities and MIUR, including web services and email. To this end, and consistent with its mission to promote the most advanced information technology, CASPUR is deploying a new Linux-based email application.

THE CHALLENGE

CASPUR's new email application, Refero, must support the I/O demands of over 800,000 concurrent users and provide file storage for each user's mailbox. To meet these very high performance and large data requirements, CASPUR needed the support of a commercial-grade cluster file system (CFS).

CASPUR's primary selection criterion was I/O performance, but the solution had to be cost-effective, open, scalable and reliable as well. Refero is the email application CASPUR will offer as a service to its universities and to the Italian public administration as part of an e-government program. With this sort of responsibility, reliability is of utmost importance. Application downtime could bring the work of 800,000 users to a halt.

As with most IT organizations around the world, CASPUR is cost sensitive. It therefore chose to build an email solution largely on open source software and the highest price-performance hardware.

Furthermore, the IT professionals at CASPUR wanted a CFS that was designed specifically for Linux, not one that had been ported to Linux from another operating system. They had seen performance and compatibility issues with file systems not designed with Linux in mind and did not want to fall into that trap.

THE BENCHMARK STUDY

To ensure that CASPUR chose the fastest and most robust cluster file system (CFS), the IT professionals at CASPUR undertook a comprehensive evaluation of the available Linux CFSs. They ran a tightly controlled benchmark test of five cluster file systems:

- Veritas Storage Foundation Cluster File System™
- Red Hat Global File System
- IBM TotalStorage® SAN File System
- PolyServe Matrix Server™
- Standard Linux ext3

ADIC StorNext® File System was also considered, but eventually rejected due to its limited support of large number of files. CASPUR used Linux ext3, a standard single node file system, as a control point to show the overhead associated with a clustered approach.

CASPUR evaluated these file systems on two separate benchmarks and in multiple scenarios. The first benchmark was Postmark 1.5, an industry standard benchmark that simulates file usage of email and Internet applications. The second benchmark was a test CASPUR developed to simulate the IMAP (Internet message access protocol) environment. (Full evaluation details and results can be found at http://www.caspur.it/en/activitiesservices/basicservices/documentation.html.)

The performance test conducted by CASPUR showed one clear winner and an obvious choice for highest performance. PolyServe Matrix Server™ vastly outperformed the competition and CASPUR selected it as the only commercial software package in its open source solution stack.

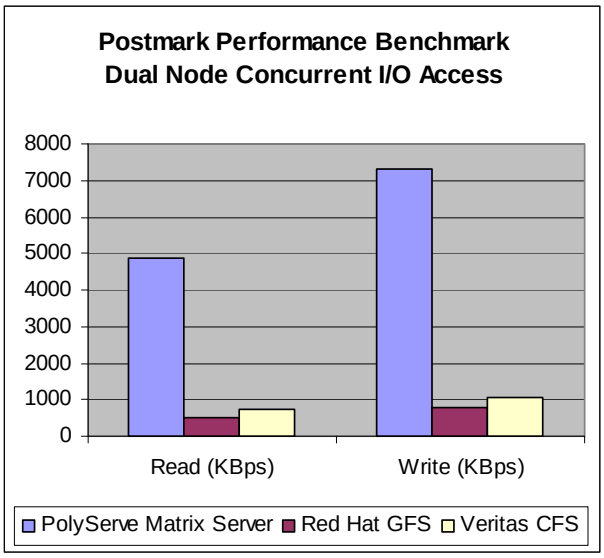

Throughout the performance evaluation, PolyServe Matrix Server outperformed the others on nearly every benchmark test. On over 70% of the tests, Matrix Server showed better than 2x the performance of any other competitor.

"We're impressed by the performance of PolyServe Matrix Server. No need for software tuning or tweaking, it performed right out of the box."

Leonardo Valcamonici
Network Applications and Services Group Leader
CASPUR

In a single node, single application instance configuration, PolyServe consistently performed among the best (The single node ext3, the control point, often led in single node performance). But in the multi-node / multiple process configuration, PolyServe stood in a class all by itself.

The following table highlights the two-node performance results of PolyServe in each test compared to the competition. The numbers indicate the relative performance benefit of PolyServe Matrix Server over the respective competitive CFS, in terms of percent improvement:

Test	Red Hat GFS	Veritas CFS
Transactions per sec.	1260%	509%
Creates per sec.	331%	851%
Reads per sec.	1260%	509%
Deletes per sec.	6815%	1521%
Reads (KBps)	843%	582%
Writes (KBps)	842%	581%

For a complete discussion of the study and full details of the results read CASPUR's Performance Evaluation of Linux Cluster File System for Large Scale Email Service whitepaper, which can be found at http://www.caspur.it/en/activitiesservices/basicservices/documentation.html

THE SOLUTION

To then assess the scalability and reliability of PolyServe Matrix Server using the Refero application, CASPUR then deployed PolyServe in a 5-node pilot cluster.

Again CASPUR was impressed with the PolyServe solution. Due to its symmetric architecture—a key differentiator from the other cluster file systems evaluated—Matrix Server's performance scales linearly with the number of nodes deployed in the cluster. Matrix Server is designed for on-demand scalability; additional nodes can be added to a live environment easily, quickly and without halting the

application. This enables CASPUR to react to the changing demands of the e-mail users' environment without interrupting the application Service Level Agreements (SLA) of its customers.

In terms of reliability, CASPUR found PolyServe to be more reliable than other cluster file systems it has used. Indeed, PolyServe Matrix Server's failover capabilities and cluster management tools were essential to making the Linux and open source stack an enterprise-class, reliable system.

"PolyServe provides management capabilities that really make managing a large cluster easy."

Leonardo Valcamonici
Network Applications and Services Group Leader
CASPUR

CASPUR found a solution with PolyServe that met its needs across the board. Matrix Server allowed CASPUR to support a large enterprise-class application reliably on Linux and Intel architecture-based servers. And the performance of PolyServe Matrix Server was so much better than the other cluster file systems, CASPUR could save money by using fewer resources and still provide higher levels of performance for the end user.

CONFIGURATION

For its production system, CASPUR will initially deploy an eight node cluster consisting of dual-processor Dell PowerEdge 2850's. These will be connected on a SAN to an EMC Clariion cx300 storage array starting with 10 TB of data. The cluster will run Red Hat Enterprise Linux 3.0.

CASPUR plans to gradually phase in its user population in 200,000 increments. To meet the eventual storage needs of 800,000 or more users, the storage cluster will require roughly 30 TB of data.

To support this, CASPUR intends to upgrade to an EMC Clariion cx700. If further increases in performance are required, CASPUR can simply increase node count at any time, without disrupting service.

Pleased with Matrix Server's capabilities as a file server, CASPUR has also implemented a three node cluster, based on PolyServe Matrix Server, for a Web server farm.

"As a public agency, we are always very cost sensitive. PolyServe Matrix Server made it possible for us to use Linux and an open source email application for a customer base of over 800,000 users."

Leonardo Valcamonici
Network Applications and Services Group Leader
CASPUR

PolyServe enabled CASPUR to build very high performance enterprise systems on cost-effective industry-standard servers and storage. PolyServe brought performance levels to CASPUR that no one else in the industry could deliver with Linux. CASPUR now considers PolyServe Matrix Server a key component for standardizing its file serving applications on Linux.

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