

SOFTWARE PLATFORM

Aim/Requirement	Solution
Highly scalable cluster system for web services, professionally supported and able to be managed efficiently	Red Hat Enterprise Linux®, Red Hat Cluster Suite and Red Hat Global File System as the basis for a diskless shared root cluster.
High-performance Linux	Red Hat Enterprise Linux is the result of intensive cooperation with the community and integrates high-performance technologies for use in business.
Vendor independence	Red Hat uses no code that is not supported by the community. Red Hat Enterprise Linux is the distribution with by far the fewest patches. Continued maintenance of the software by the community is therefore guaranteed and ensures that Red Hat's customers enjoy top product quality and vendor-independence.
Wide choice of certified hardware and software	Red Hat Enterprise Linux is the Linux operating system platform with the largest range of certified partner solutions.
Professional support	Red Hat and Advanced Partner ATIX.

HARDWARE PLATFORM

Aim/Requirement	Solution
Good value and high-performance standard hardware for cluster nodes	HP ProLiant servers, each with two Intel Xeon processors
High-performance network with tried and tested technology	HP SAN switches and Cisco network switches
Virtualizable storage solution	HP EVA 5000 Storage System

RED HAT SALES AND GENERAL INQUIRIES

USA/Canada
1-888-REDHAT1
1-866-273-3428 x44555
sales@redhat.com

Worldwide offices
+1-919-754-3700
visit www.redhat.com/about/corporate/wwoffices

CASE STUDY

AT THE MUNICH TRADE FAIR: RED HAT® GLOBAL FILE SYSTEM (GFS) PROVIDES A MODERN CLUSTER SOLUTION WITH VIRTUALIZED STORAGE

Messe München International (Munich International Trade Fair – MMI), which organizes about 40 trade fairs for capital goods, consumer goods, and new technologies is one of the world's leading trade fair companies. Over 30,000 exhibitors representing more than 90 countries and two million visitors from about 180 countries take part in the events in Munich every year. With five subsidiary companies abroad, global trade fairs and 75 foreign branches dealing with 89 countries, MMI has a very strong worldwide network.

MMI is clearly set for further growth, and, as the business expands, so must the IT infrastructure. This need became apparent by the middle of 2005; the existing infrastructure for providing web services was no longer fully able to cope with the increased requirements. Further scaling of the system would have involved great expense. At that time, the web services were based on a Linux cluster using the NFS file system; the operating system used on the cluster nodes was Debian.

"We had previously always had very good experience of Linux," explains Martina Ritzer, Head, Web Services of MMI. "With the new solution, we really wanted to continue to benefit from the flexibility and vendor-independence of Linux, and have a scalable overall solution that was also professionally supported and properly certified by the main hardware and software providers."

MMI managers had already heard about the opportunities offered by the Red Hat Global File

System (GFS). Red Hat had acquired the GFS technology when it took over the Sistina company, released it under GPL license, and developed it in partnership with the community. MMI turned to Red Hat to investigate the possibility of using a dynamic cluster system for its web services. Red Hat brought ATIX, its Advanced Partner for clustering and storage, on board. ATIX specializes entirely in highly scalable IT platforms for use in data centers and had extensive experience with Red Hat GFS, going back to the time when the preceding technology was developed and marketed under the ownership of Sistina. In this case, the service provider developed the "diskless shared root cluster," an overall solution based on Red Hat GFS meeting all of MMI's requirements.

There are various ways of building a high-availability cluster. In the most rudimentary method, IT services run simultaneously on several servers and user queries are allocated by a load distributor. For static services, this approach is



“Our requirements for the new cluster system were very demanding. Our initial experience with the new solution based on Red Hat Enterprise Linux® and Red Hat Global File System proves conclusively that it is possible to realize even high-end projects successfully using open source software.”

perfectly practical, but for dynamic services it soon reveals its limitations. Here you have to separate the application servers from the data storage layer and give them shared access to a storage system. The traditional approach used for shared data storage is a file server, as was used by MMI before the installation of the new solution. This server often became a bottleneck and a single point of failure for the whole system.

In order to overcome these problems inherent in the traditional approach, every server in the cluster would have to have direct access to the storage devices and be both able to read them and write to them concurrently. A cluster file system connects application servers to a linked Storage Area Network (SAN), allowing parallel file system access from all cluster nodes to a central storage system. The diskless shared root cluster, developed by ATIX and in use at MMI, takes this approach to its logical conclusion and avoids having any hard disks in the cluster nodes.

The starting configuration for the new cluster at MMI has 16 nodes: HP ProLiant servers each with two Intel® Xeon™ 2.8-3.2GHz processors. Red Hat Cluster Suite is responsible for ensuring that the

services of one server are taken over by another if it fails and, for distributing the load between the cluster nodes. At the heart of the system is Red Hat GFS, which allows parallel file system access by all cluster nodes to a central storage system. The server cluster is linked via the network to an HP Eva storage system. The GFS pool layer virtualizes the storage devices, making them available and non-hardware-dependent. This means that several devices can be combined in a pool using “striping” or linear concatenation. Any changes to the pool configuration are visible to all cluster servers and, a volume manager ensures that the file system can be expanded online. Since, in a GFS storage cluster, many servers are accessing the same physical data blocks, a mechanism exists for coordinating the distributed accesses—the so-called “locking service.” This feature guarantees data consistency within the file system.

Already, the configuration of the cluster at MMI provides a high-performance and flexible solution using the latest technology. The innovative next

“With the new cluster system, we have an extremely high-performance solution which offers us maximum scalability for the future.”

step is to do away with hard disks in the cluster servers and to boot them directly from the storage system. This configuration is very easily scalable new resources in the form of new server hardware can simply be added on the “plug & play” principle because the operating system is also installed

centrally on the storage system. In addition, it makes maintaining the operating system much easier, because there is only one version of the operating system to update.

The complete separation of cluster nodes and central disk storage means that all information about the structure and content of the cluster is consolidated

“All the systems are performing brilliantly with the new operating system and the new architecture using the Red Hat Global File System.”

in the central storage system. This means that if one server fails, no information has to be reinstated. This reduces the restart time for a cluster node to a minimum because only the server hardware has to be replaced in order to return the system to its normal status, increasing the overall availability of the cluster. The standard tools of the HP storage system are used for backup. By using the Snapshot function of the HP Eva 5000, tape copies of the whole system can be produced while the system is running. The cluster can also be remotely monitored and administered by using the ATIX solution com.oonics GrayHead.

With Red Hat solutions, migration to the new cluster took place quickly and—because of ATIX’s experience—also went very smoothly: “The customer wanted a brisk schedule,” recalls Mark Hlawatschek, director of ATIX. “In the end, the process went smoothly and without problems. In this respect,

too, the combination of standard hardware and open source software proved to be highly effective and reliable.”

“With the new cluster system, we have an extremely high-performance solution which offers us maximum scalability for the future,” says Martina Ritzer. “We run MySQL, Tomcat PHP and email services in the cluster as well as FTP, CVS and our staging front-end software. All the systems are performing brilliantly with the new operating system and the new architecture using the Red Hat Global File System. We are also really well equipped for future growth. Resource scaling using the plug & play principle is a whole new experience for us.”