



Red Hat Global File System for scale-out web services

by Subbu Krishnamurthy (Based on the projects by ATIX, Munich, Germany)

Red Hat leads the way in delivering open source storage management for Linux with the integration and support of LVM, ext3, NFS, autofs, volume mirroring, I/O multipathing, and more. Capabilities found in technologies like LVM, LVM2 and ext3--the default now for most enterprise Linux implementations--are, in great part, due to the effort of Red Hat-sponsored community engineering efforts.

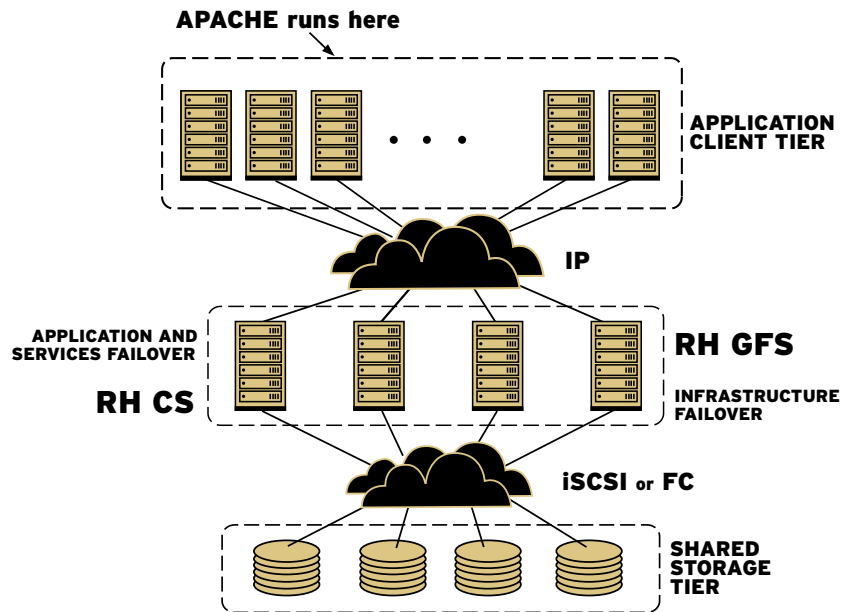
The delivery of an even greater set of storage management capabilities like those found in Red Hat® Global File System (GFS), Red Hat Cluster Suite (CS) and Distributed Lock Manager (DLM) continued with the 2005 introduction of Red Hat Enterprise Linux® 4. GFS and CS enable high availability and data sharing for key applications like Oracle® Real Application Clusters (RAC), NFS, and Apache. The highly successful 2.6-kernel-based LVM2 follows the 2.4 kernel-based LVM with a cluster-aware version named CLVM. CLVM is included in the Red Hat Enterprise Linux distribution as a key part of scale out storage management.

Device Mapper and its Red Hat Enterprise Linux 4 support of multipath IO, iSCSI, and InfiniBand complements the maturing fibre channel protocol (FCP). The promise of virtualization through Xen brings all of these technologies together into a highly advanced storage management platform with the scheduled release with Red Hat Enterprise Linux 5.

Open source storage management technologies are beginning to offer an attractive alternative to expensive proprietary solutions. Storage management capabilities found in Red Hat Enterprise Linux help customers reduce costs and simplify complex mission-critical IT environments.

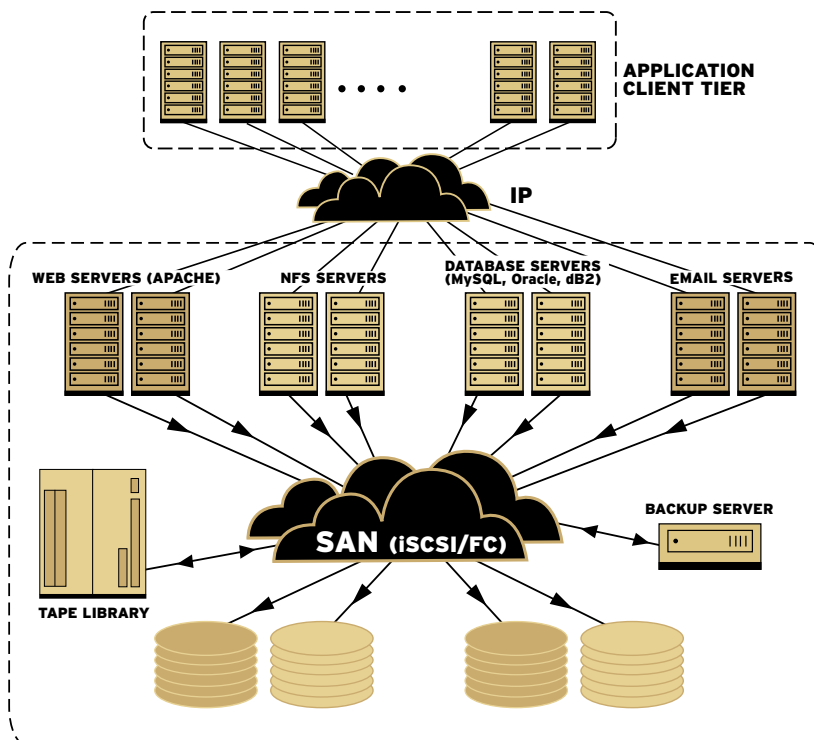
Below are two case studies of customers who have benefited from the deployment of open source storage management from Red Hat. IP-Tech and Messe München have deployed a combination of GFS, CS, and Red Hat Enterprise Linux as part of the mission-critical IT infrastructure that delivers web services to their constituents.

WEB SERVICES



RED HAT GFS & CS FOR WEB SERVICES

Customer: IP TECH and Messe München



CASE STUDY 1: IP TECH

IP Tech is one of the largest Internet hosting and service providers in Switzerland. The success of their business requires 100% availability of their web serving infrastructure.

Before deploying Red Hat technologies, IP Tech used NFS to meet data sharing requirements. They encountered significant problems maintaining their hosting environment as it became unstable under heavy load. IP Tech could not afford instability—they support more than 100,000 email users and expect 5-7 million daily web hits, 5 million daily email connections, and 4 million daily MySQL connections. With NFS, file services and mounts would come and go without warning, stopping operations at critical times and during high load periods. This caused downtime with maximum negative impact.

In 2002, IP Tech deployed 14 nodes of Red Hat GFS to support 1,500 MySQL database applications, 10,000 gmail domains, and 28,000 Apache web domains in special configurations. Most recently, IP Tech migrated to a centralized blade-based infrastructure consisting of 22 diskless blade servers for GFS nodes and two terabytes of redundant shared root storage.

By using a Red Hat Enterprise Linux cluster with Red Hat GFS, IP Tech achieved both high availability and performance. If any of the servers crashed or if an application (like http or gmail) hung, the server could be re-booted quickly and brought back into the cluster infrastructure without disrupting the services being provided by the other servers.

Hardware repair times are minimal for blade servers, as they are simply replaced when they fail, and the replacement boots off the shared root boot image. Server and storage scalability can be achieved during ongoing operation. Each night the data is replicated onto a second storage system via a LAN-free backup using GFS and the storage area network (SAN).

IP Tech also uses the shared GFS root disk feature that simplifies the process of updating software and performing static application service load-balancing in the cluster. For example, when IP Tech has a spam attack they can quickly change some web servers to act as mail servers, keeping mail service up while countering the attack. In this scenario, all cluster services continue to operate uninterrupted.

IP Tech also uses Red Hat GFS to take snapshots of existing file system and database volumes on the storage hardware. These volumes are then mounted read-only and backed-up in parallel with regular file system operation.

In summary, the key benefits IP Tech found when using Red Hat GFS were:

1. High performance and scalability at a level that was unattainable with NFS.
2. Reduced complexity via data sharing and shared root disk images.
3. On-the-fly load-balancing of services across the cluster to provide on-demand computing and application performance.
4. No downtime impact from backups.

By using Red Hat Enterprise Linux and Red Hat GFS, IP Tech has achieved excellent performance, availability, and scalability that was unattainable with NFS. Data sharing via Red Hat GFS has allowed IP Tech to reduce management complexity and scale performance to meet customer demand—all the while using a low-cost, Linux-based hardware and software infrastructure. In contrast to their NFS-based storage infrastructure, IP Tech found that the cluster running Red Hat GFS “ran by itself.”

CASE STUDY 2: MESSE MÜNCHEN INTERNATIONAL

Messe München International (Munich International Trade Fair – MMI), one of the top organizers of trade fairs in the world, organizes about 40 trade fairs for capital goods, consumer goods, and new technologies. Every year, more than 30,000 exhibitors and two million visitors from around the world take part in the events in Munich. MMI also organizes trade fairs in Asia, North America, and South America. With five subsidiary companies and 75 foreign branches, MMI operates in 89 countries worldwide.

MMI is clearly set for further growth and as the business expands, so must the IT infrastructure. By the middle of 2005, it was apparent that the existing infrastructure for providing web services--a Debian® Linux cluster using NFS--was no longer able to cope with the increased requirements. Any further scaling of the current system would have involved significant incremental expense.

“We have always had a very good experience with Linux,” says Martina Ritzer, head of the web services department at MMI. “With the new solution, we wanted to continue to benefit from the flexibility and vendor independence of Linux while deploying a scalable overall solution which was professionally supported and properly certified by the main hardware and software providers.”

MMI managers had heard about the capabilities offered by Red Hat GFS and were interested in the clustering and data sharing capabilities of the system for possible use in supporting MMI web services.

Red Hat engaged Munich-based ATIX, an Advanced Partner for clustering and storage, to assist in the project. ATIX specializes in deploying highly scalable IT platforms for use in data centers, and has extensive experience with Red Hat GFS and Red Hat Enterprise Linux.

To meet all of MMI’s web services requirements, ATIX helped build a solution that includes remote monitoring and administration, as well as a “diskless shared root cluster”atechnology implementation developed by ATIX.

MMI’s new cluster has 16 HP® ProLiant servers, each with two Intel® Xeon 2.8-3.2 GHz processors andthe storage subsystem is HP EVA. Red Hat Cluster Suite provides high availability ensuring services failover from one server to another and for load balancing between the cluster nodes. The volume management layer (CLVM/LVM2) virtualizes the storage devices and provides easy-to-manage logical volumes and allows “striping” . Any changes to the storage pool are visible to all clustered servers and the file system can be expanded dynamically online.

Red Hat GFS also coordinates the distributed accesses called “locking service,” thus guaranteeing data consistency and integrity within the file system. Standard tools of the HP storage system are used for backup. By using the snapshot function of the HP EVA 5000, the whole system can be backed up without service interruption.

The configuration of the cluster at MMI provides high performance and flexibility. Booting the server nodes directly off the storage system provides for easy scalability on a “plug and play” basis. Since the operating system is also installed on the shared storage system, maintaining and updating is much easier because there is only one image to update. The separation of cluster nodes and storage means that all information about the structure and content of the cluster is consolidated in the shared storage system. If one server fails, no information is affected or has to be reinstated. This minimizes the recovery time for a clustered node because the server hardware only has to be restarted or replaced in order to return the node to its normal status.

“With the new system, we have an extremely high-performance solution offering us maximum scalability for the future,” says 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 architecture using the Red Hat Global File System. We are also well-equipped for future growth. Resource scaling using the plug & play principle is a whole new experience for us.”

By using Red Hat Enterprise Linux with GFS in combination with HP servers and storage , MMI has achieved

excellent performance and scalability. The infrastructure change reduced complexity while improving availability to a level unattainable with NFS. Red Hat GFS has allowed MMI to scale performance to meet growing customer demand.

The value of the implemented solution is summed up by Ritzer: "Our requirements for the new clustering system are very demanding. Our experience with Red Hat Enterprise Linux and Red Hat Global File System proved conclusively that it is possible to succeed with critical high-end projects using open-source software."