glusterFS

Cluster File System Manager

Features and Benefits:

• Scales to Petabytes using commodity servers and storage

• Flexible volume and performance management via stacked modular configuration

• No single point of failure, no centralized MDS. Files can be recovered without glusterFS

• Linear performance increases to several times that of conventional SANs

• User friendly – basic cluster can be up and running in thirty minutes

• Runs over standard TCP/IP connections or high performance IB channels

• Sits on top of existing POSIX file systems – ext2/3, xfs, reiser, etc.

• Provides unified VFS

• Provides multiple load balancing options, Automatic File Replication (AFR) and file striping for high performance

• Self-healing – no FSCK required

• Tunable I/O accelerators on both server and client sides

Typical clustered file systems work to aggregate storage and provide a unified view to clients, but in doing so, most place unreasonable demands on the systems administrator and severe restrictions on supported OS's and hardware. In contrast to this scenario, glusterFS is simple to deploy and easy to maintain.

A system administrator with basic networking skills can have a glusterFS cluster up and running in less than an hour. GlusterFS enhances existing storage operations by simplifying the scaling of capacity and I/O using commodity servers and storage. It provides a highly extensible framework that spans across the majority of the storage spectrum – from 10's of TBs to PBs --- and can address the needs of SMBs, large corporate data centers and HPC clusters. The key to this flexibility is glusterFS's modular, stackable design and full POSIX compliance. GlusterFS's design incorporates pluggable modules (called translators) which fine tune the file systems to your application. GlusterFS translators are available for both the server and the client.

GlusterFS was designed without the need for centralized meta-data servers which can rob performance and compromise scalability. The directory structure is present on all data servers which allows user data to be uniquely recovered even without glusterFS. With glusterFS sitting on top of your existing POSIX compliant file system, the learning curve needed to deploy and administer glusterFS is minimal. This allows accelerated testing and rapid roll-out to your production systems.

GlusterFS is compact and non-kernel intrusive. It is supported on all of the major linux packages and has been ported to FreeBSD, OS X and Open Solaris. With glusterFS running outside the kernel, new features are easily incorporated and rolled into your cluster. Most updates only require a simple "make install" followed by a restart of the glusterFS process.

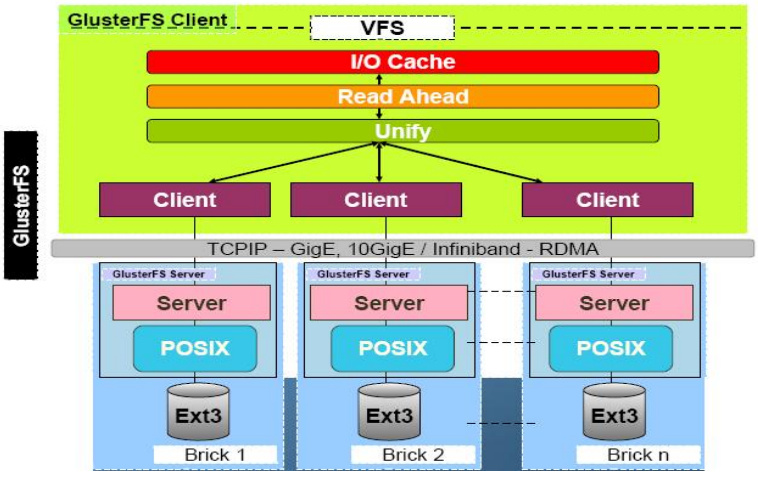

To learn more about the glusterFS offering or to see our full range of storage solutions, visit our website at: www.dasdi.com.

GlusterFS Capabilities

Operating Systems Supported	Linux, FreeBSD (other POSIX compliant OSs, website for current support)
Cluster Transport	TCP/IP, IB-VERBS, IB-SDP
Storage Interface Support	Any interface supported by the server brick
Number Server Bricks	No hard limit, tested with 64 server bricks
Number of Client Nodes	No hard limit, tested with 700 clients
Protocol Translators	Server (exports volumes over the network), Client (attaches to remote glusterFS volumes)
Cluster Translators	AFR (Automatic File Replicator), Stripe (stripes a single file across multiple server bricks), Unify (aggregates multiple storage bricks into a single super server)
Cluster I/O Schedulers	ALU (distributes files across server bricks based on user defined criteria), NUFA (in an HPC environment gives local storage higher priority), Random (scatters files over all available server bricks), Round-robin (rotates file creation through all server bricks)
Performance Translators	Read-ahead (predictive pre-fetch), Write-behind (delayed write and write-gathering), Threads (adds AIO capability to better utilize CPU idle time), IO-cache (reduces server loading for reads) , Stat-prefetch (fetches stat info for all files in a folder in a single operation)
Extra Feature Translators	filter (advanced filtering on files names/attributes), posix-locks (storage independent POSIX record locking support), trash (a recycle bin), fixed-id (all calls through this translator will be from a single UID and GID)
Storage Translators	POSIX (binds glusterFS server to the underlying file system)

Block Diagram of a Complete glusterFS Cluster