

Global File Systems

Is Red Hat's GFS worthy of the name?

Douglas McClendon

Department of Electrical Engineering & Computer Science
The University of Kansas

dmc@ku.edu

<http://douglas.mcclendon.org/school/ku/eecs645/termpaper>

Global File Systems

abstract

A simple to use, simple to deploy, high performance Global File System (**GFS**) is something of a holy grail for computer network users and administrators. This is no doubt why Red Hat Inc. chose the name "Red Hat Global File System" (**RH-GFS**) for one of its key enterprise products. This paper presents an introduction to the broad topic of Global File Systems implemented via the Internet, and a brief evaluation of Red Hat's product bearing the same name.

GFS Overview

- Definition: A Global System for Sharing Files
- Key Components
 - A data storage medium
 - A means of data communication/transportation
 - A file addressing/identification mechanism
 - A protocol to tie it all together
- A Ubiquitous Example: The Web as a GFS
 - Has the 4 key components
 - Breakthrough simplicity, at least for client users
 - * Limitation: does not use traditional OS filesystem interface

Traditional Filesystems

- Traditional Filesystem Interface
 - “Mounted” Filesystems. E.g. Mount /dev/disk /mnt/data
 - “Mapped” Drives, E.g. Map N:\ to \\alphaserver\sharedvol
 - Standard OS Interface that most apps already use
- Traditional Networked Filesystems
 - Single server, perhaps many global clients via WAN
 - Sun's NFS, Microsoft's SMB
 - * Limitation: filesystem resides on a single server

Filesystem Advances: RAID

- The Problems
 - Resiliancy, fault tolerance
 - Bottlenecked data bandwidth channels
 - Individual drive storage limits
- The Solutions
 - Distribute data across multiple drives
 - Put drives on separate buses
- The corresponding issues with networked filesystem
 - Distribute data across multiple servers
 - Put servers on separate networks
 - * Let's now look more closely at Red Hat's RH-GFS

Red Hat's Global File System

Really a Filesystem for Tightly Coupled Clusters

- Strength
 - Does utilize RAID-like distribution across servers
- Weaknesses
 - Designed with high speed LAN/SAN in mind
 - No network security: i.e. authentication, encryption
 - Maximum 256 server nodes
- The WAN workaround
 - Use IPsec or VPN to make WAN seem like LAN
 - This also provides network security features
- * More alternatives on the horizon...

The Global File System Future

Other Products, Projects, and the Future

- Many Current Options
 - pvfs2 (nasa, nsf, doe), lustre (hp, cray) dfs (ms), googlefs
 - peerfs (radiant), ocfs (oracle), teragrid gfs-wan (nca)
 - FUSE hacks: gmailfs, wikipediafs, bittorrentfs
- What I'd like to see
 - A standard filesystem interface for p2p filesharing networks
 - `mount -t truly_global_fs internet:videoarchive /mnt/videos`
 - `ls /mnt/videos/pbs/2005/November/*/*Nova*`
 - Just as easy to upload/write files as it is to download/read