SELinux Support over GlusterFS

Jiffin Tony Thottan

Software Engineer, Red Hat

Thank you for contribution

- Brain Foster
- Niels De Vos
- Manikandan Selvaganesh

Agenda

- GlusterFS
- SELinux with GlusterFS
- Challenges
- Clients
- How it is going?

GlusterFS

- An open source, scale-out distributed file system(posix like)
- Software Only and operates in user-space
- Aggregates Storage into a single unified namespace
- No metadata server architecture
- Provides a modular, stackable design
- Runs on commodity hardware

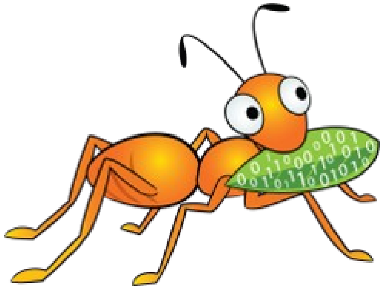

GlusterFS Terminologies

- Data is stored on disk using native formats (e.g. ext4, XFS)
- Has following components
  - Servers known as storage bricks (glusterfsd daemon), export local filesystem as volume
  - Clients (glusterfs process), creates composite virtual volumes from multiple remote servers using stackable translators
  - Management service (glusterd daemon) manages volumes and cluster membership
  - Gluster cli tool

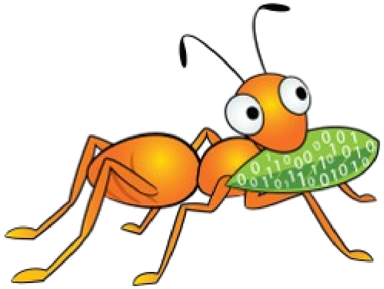

GlusterFS Architecture

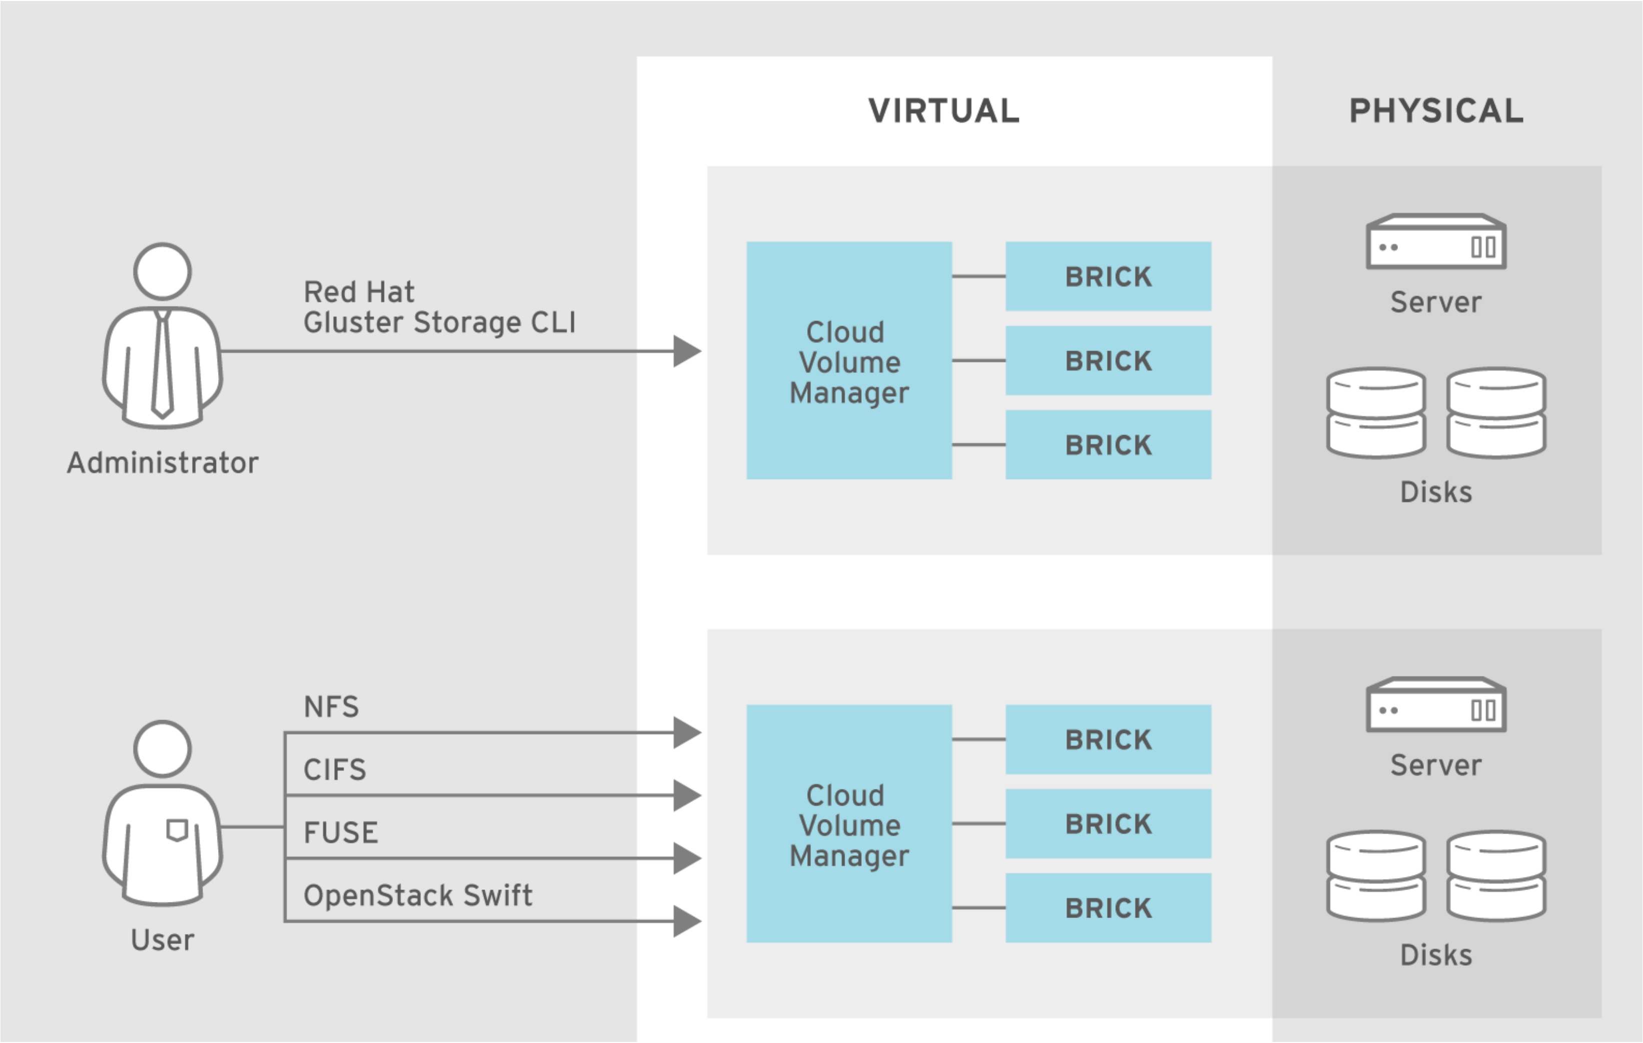

GlusterFS internals : Translators

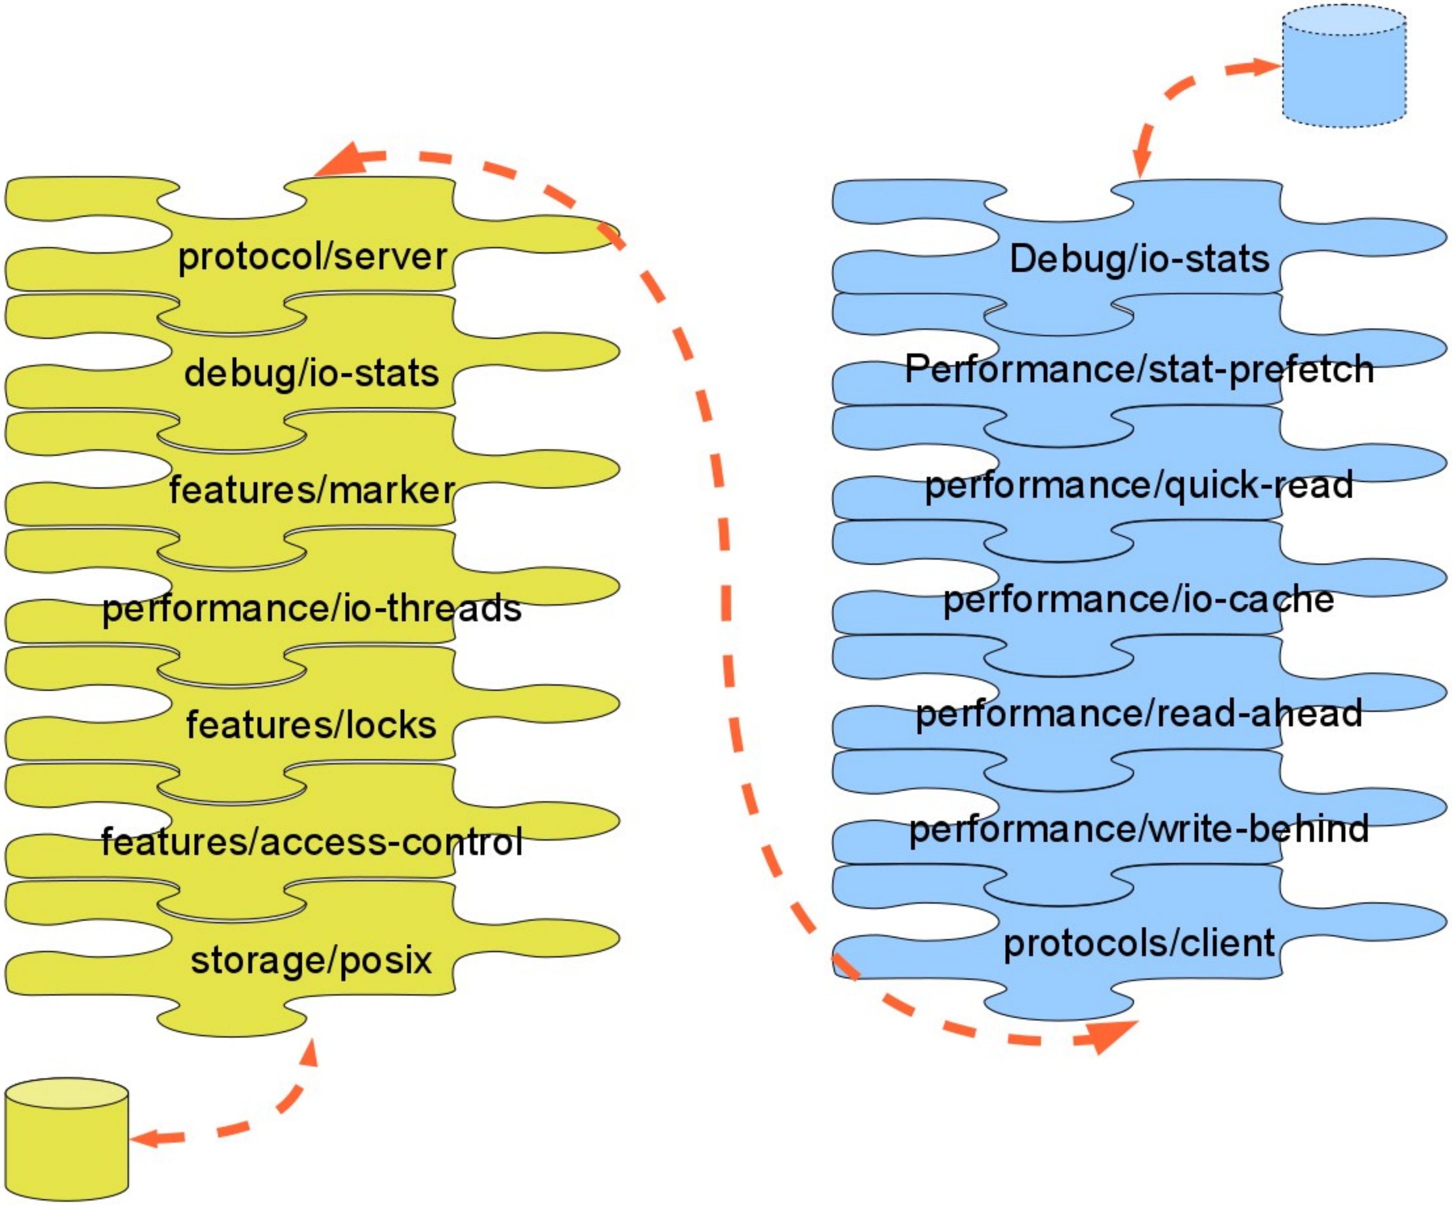

protocol/server

debug/io-stats

features/marker

performance/io-threads

features/locks

features/access-control

storage/posix

Debug/io-stats

Performance/stat-prefetch

performance/quick-read

performance/io-cache

performance/read-ahead

performance/write-behind

protocols/client

Brief Intro: SELinux aka Security Enhanced Linux

- Implementation of a mandatory access control
- SELinux can enforce rules on files and processes based on policies
- Processes and files are labeled with an SELinux context

  ls -Z file1

  -rwxrw-r-- user1 group1 unconfined_u:object_r:user_home_t:s0   file1

  SELinux contexts follow the user:role:type:level syntax.
- At backend it stored as extended attribute with key "security.selinux"

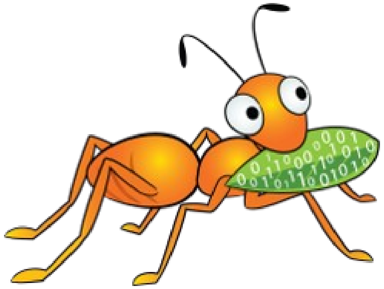

GlusterFS with SELinux

- GlusterFS is an application which works very well with SELinux : system_u:system_r:glusterd_t:s0
- SELinux context on files accessed by gluster processes
  - /var/log/glusterfs
    - system_u:object_r:glusterd_log_t:s0
  - /var/run/gluster
    - system_u:object_r:glusterd_var_run_t:s0
  - /var/lib/gluster
    - system_u:object_r:glusterd_var_lib_t:s0
  - /etc/glusterfs
    - system_u:object_r:glusterd_conf_t:s0
  - Bricks
    - system_u:object_r:glusterd_brick_t:s0

Applications which uses GlusterFS

- Depending on the application context may vary
- For example
  - Fuse clients (gluster native client)
    - system_u:object_r:fusefs_t:s0
  - NFS clients
    - system_u:object_r:nfs_t:s0

Do this good enough ???

- Nope
- The applications cannot save context for their users
- Security being one of key aspects of File System
- And SELinux was one of trending one
- Being posix compliant file system, GlusterFS is missing this feature

Why it is not working ???

- Bricks has its own context
- Application cannot overwrite these context
- If overwrites everything will go into chaos

How it can be done ???

- Introducing translator , of course known as selinux at server side
- It does the following :
  - Stores SELinux context as “trusted.glusterfs.selinux”
  - Does the mapping for server and client
- It interrupts following system calls (aka fops) :
  - setattr, getattr, create, mkdir, mknod
- This translator loaded by default in server graph

- Default SELinux context for a entry in a volume “system_u:object_r:glusterd_brick_t:s0”
- Internal operations such as self-heal, rebalance should be ignored
- Enforcement should be done at client side

SERVER

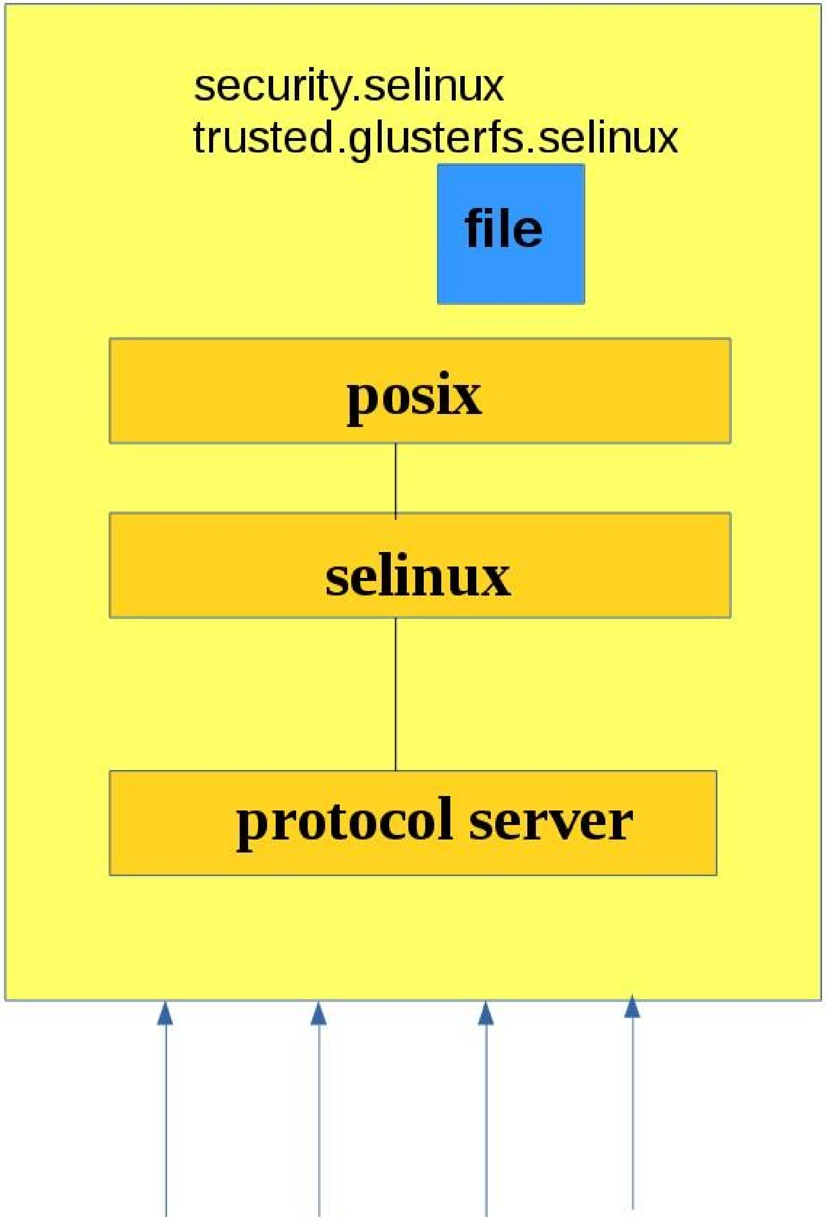

Request from clients sends
semanage, restorcon etc

Clients

- fuse clients
  - Bug : https://bugzilla.redhat.com/1272868
  - Patch : http://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/id=102aefdda4d8275ce7d7100bc16c88c7
272b260
- NFS clients
  - Labelled NFS

Where are we now ???

- Planned it for 3.10, but didn't make it
- Two patches posted upstream
  - Implementation of SELinux translator
  - SELinux brick file context management scripts
- Two patches yet to be started
  - Provide SELinux context from parent
  - Provide gfapis for managing SELinux context

References

Mailing lists:

gluster-users@gluster.org

gluster-devel@gluster.org

IRC:

#gluster and #gluster-dev on freenode

Feature page : https://github.com/gluster/glusterfs-specs/blob/master/accepted/SELinux-client-support.md

Links (Home Page):

http://www.gluster.org

Q & A

Thank You