


Gluster roadmap: Recent improvements and upcoming features




Niels de Vos
GlusterFS co-maintainer

ndevos@redhat.com
ndevos on IRC
@nixpanic on Twitter


Agenda

- Introduction
- Release Schedule
- History of feature additions
- Plans for upcoming features




What is GlusterFS?

- Scalable, general-purpose storage platform
 - POSIX-y Distributed File System
 - Object storage (swift)
 - Distributed block storage (QEMU)
 - Flexible storage (libgfapi)
- No Metadata Server
- Heterogeneous Commodity Hardware
- Flexible and Agile Scaling
 - Capacity – Petabytes and beyond
 - Performance – Thousands of Clients




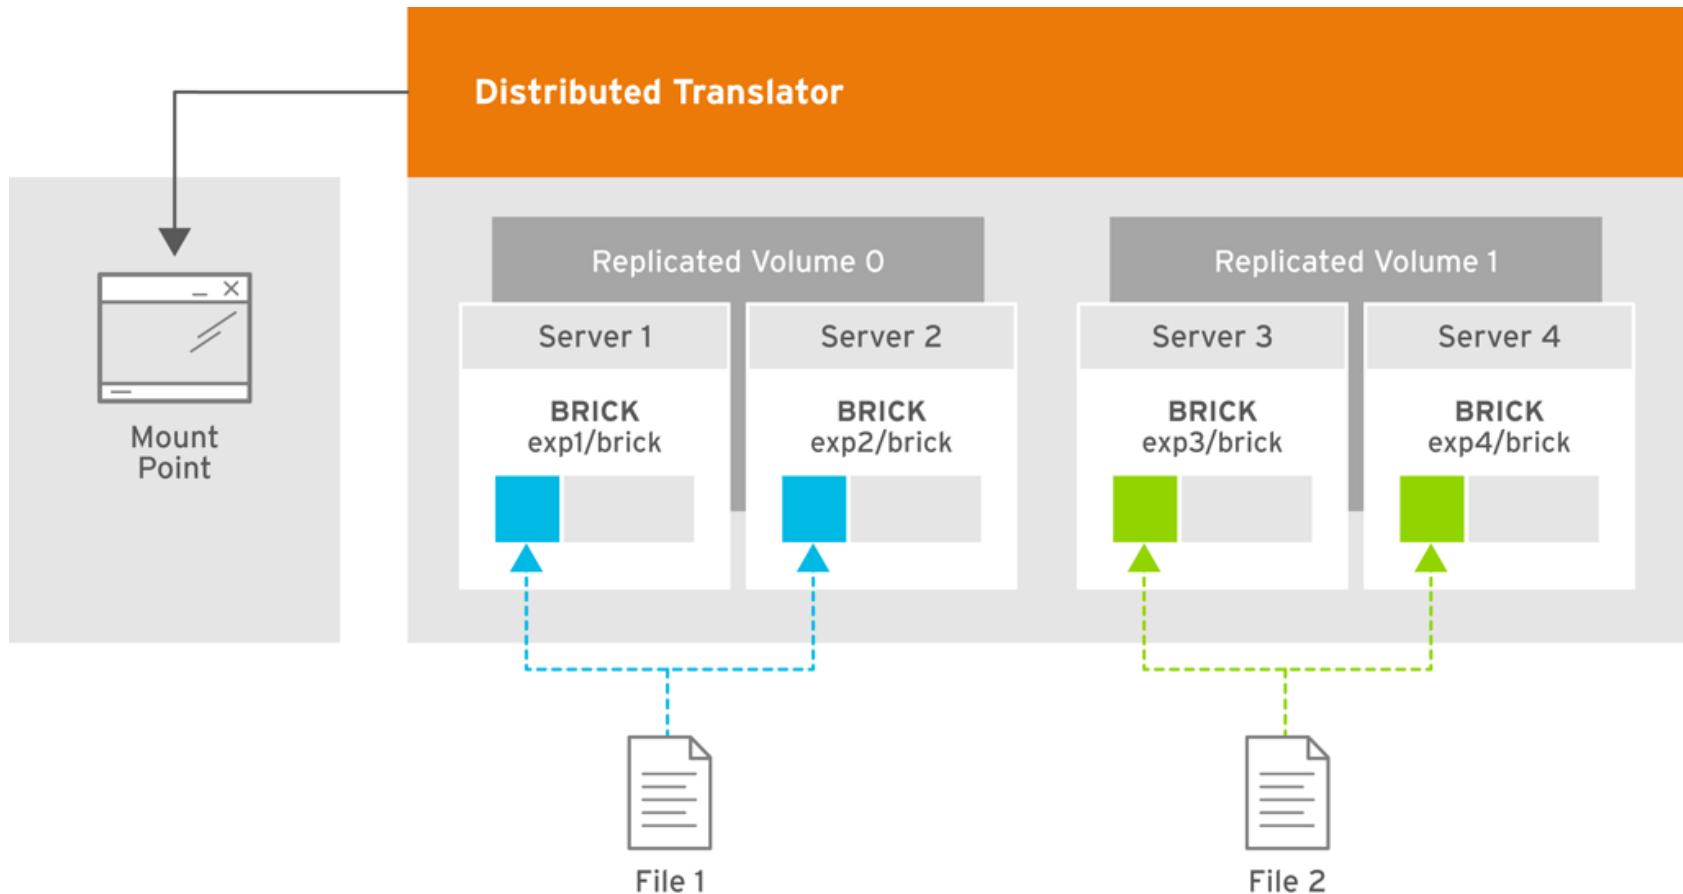
Data Access Overview

- GlusterFS Native Client
 - Filesystem in Userspace (FUSE)
- NFS
 - Built-in Service, NFS-Ganesha with libgfapi
- SMB/CIFS
 - Samba server required (libgfapi based module)
- Gluster For OpenStack (Swift-on-file)
 - Object-based access via Swift
- libgfapi flexible abstracted storage
 - Integrated with QEMU, Bareos and others




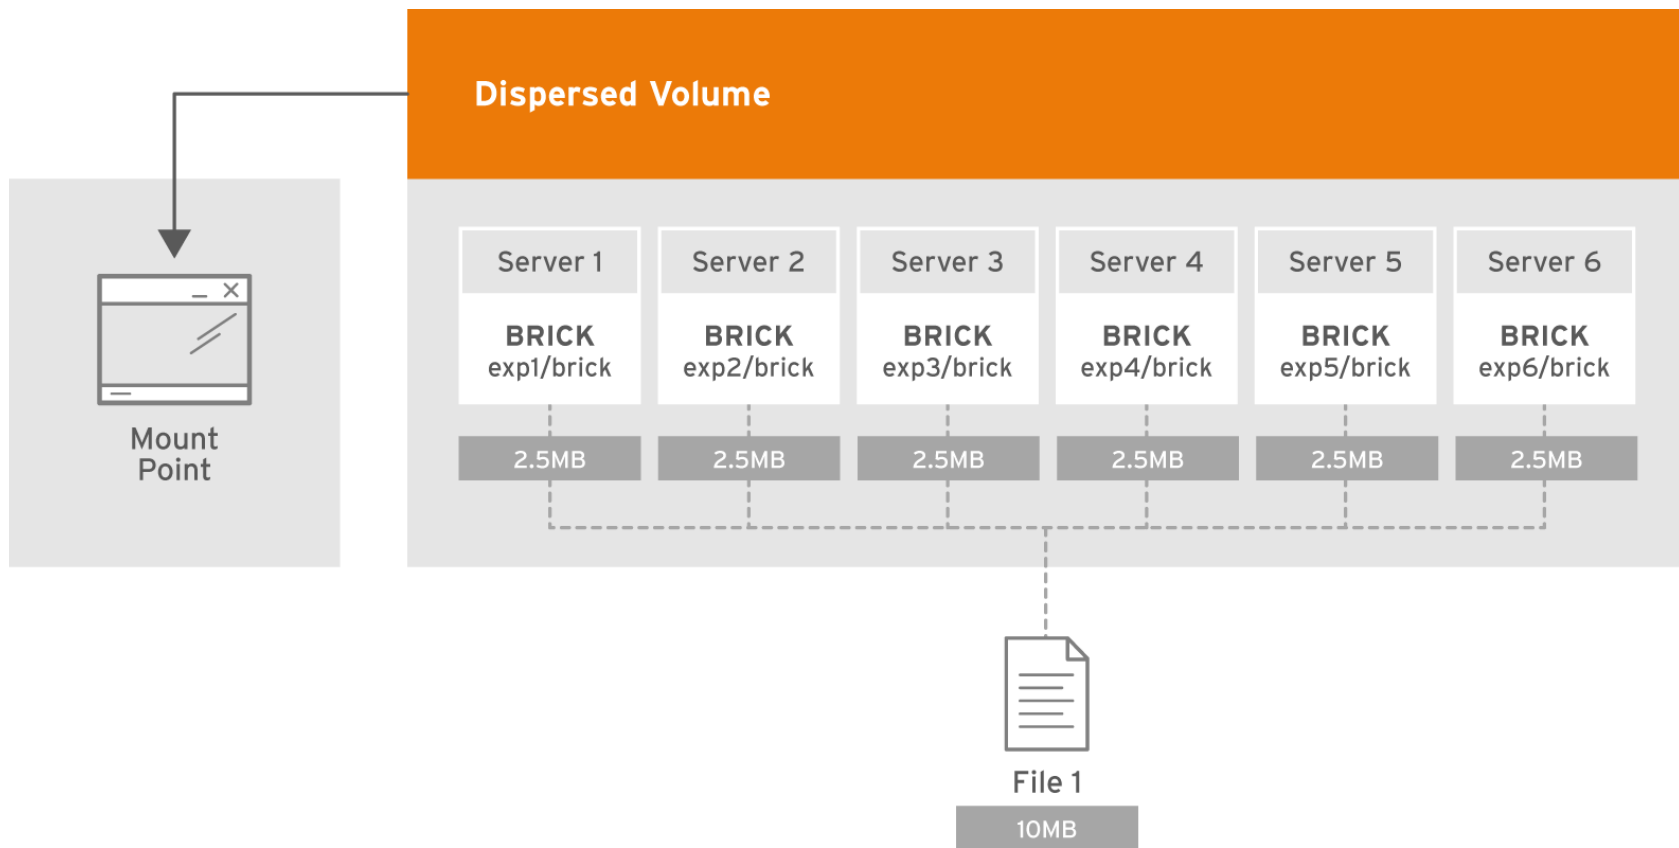
Distributed Replicated Volume

- Distributes files across replicated bricks



Disperse / Erasure Coding

- *Similar* to RAID 5/6 over the network
- Encoded fragments of files




Distribution Integration

- Available in Fedora, Debian, NetBSD and others
- CentOS Storage SIG packages and add-ons
- Community packages in multiple versions for different distributions on <http://download.gluster.org/>
- Quick Start guides on <http://gluster.org> and CentOS wiki

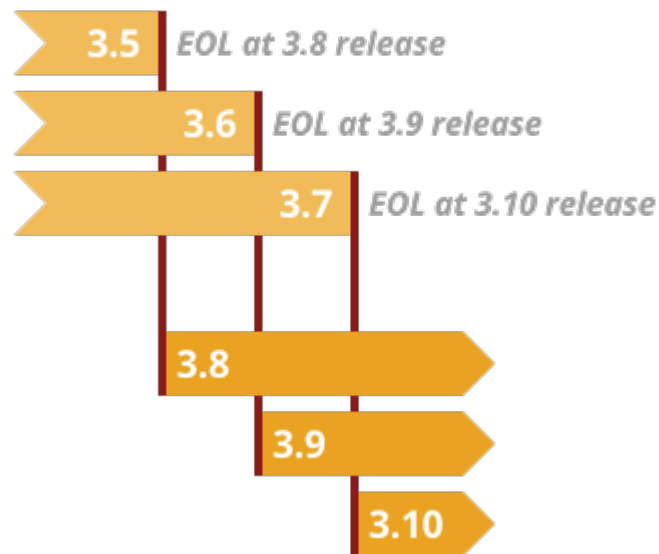


Release Schedule




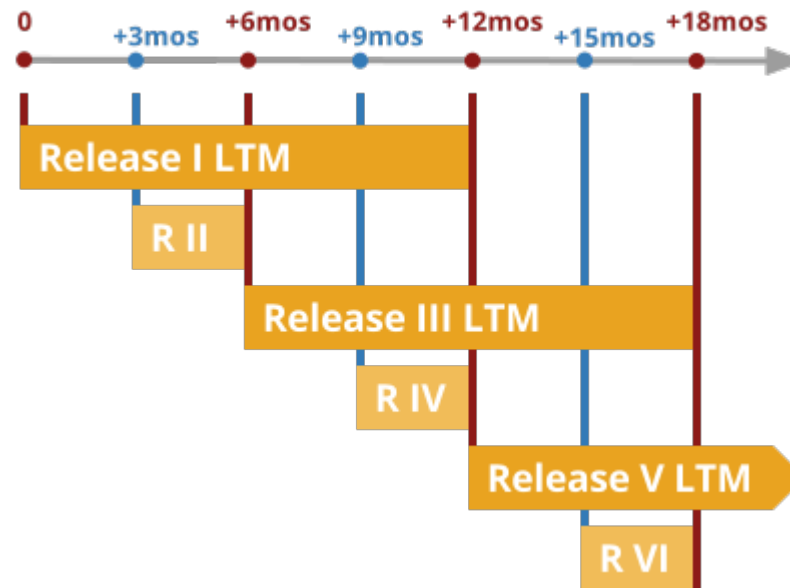
Release Schedule upto version 3.7

- A release every 6 months
- Provide monthly updates for three version




New Release Schedule from version 3.8

- A release every 3 months
- Monthly updates for all maintained releases
- Alternating between Short-Term-Maintenance (STM) and Long-Term-Maintenance (LTM) versions




Recent Feature Additions




Features included in the 3.9/STM release

- reset-brick command to ease replacing bricks
- Multi-threaded self-heal for Disperse volumes
- Hardware extension acceleration for Disperse volumes
- Lock revocation through CLI
- On-demand scrubbing for Bitrot Detection
- Event API for realtime cluster notifications
- Simplified configuration for geo-replication




Features included in the 3.8/LTM release

- Base for REST Management APIs for Gluster
- Manageable by Heketi
 - Easier integration in OpenStack, Kubernetes, ...
- SEEK_DATA/SEEK_HOLE for sparse files
- Geo replication improvements
 - Tiering aware
 - Sharding support
- More options for policy based split-brain resolution
- Multi-threaded self heal




Features included in version 3.7

- Small-file performance enhancements
- Tiering for hot and cold contents
- Trash translator making undelete of files possible
- Netgroups and advanced exports configuration (NFS)
- BitRot detection
- Upcall infrastructure to notify client processes
- Support for NFS Ganesha clusters
- ... and more




Upcoming Feature Additions




Features planned for the 3.10 release

- Brick-multiplexing
- md-cache enhancements
- Converged High-Availability with Storhaug
 - Pacemaker managed NFS-Ganesha and Samba
- Progress and time estimation of rebalance
- Tier daemon as a service managed by GlusterD
- Improved debugability for gfapi applications




Expected for the 3.11 release

- Scale out/in support with Tiering
- Subdirectory mounting for the FUSE client
- SELinux contexts on Gluster Volumes
- Server-side DHT to improve REaddir operations



Plans for the next 4.0 release

- Scalability and manageability improvements
 - Journal Based Replication
 - Improved Distributed Hashing Translator
 - GlusterD 2.0
- inotify like functionality
- Kerberos for the Gluster protocols
- Improved WORM, Retention and Compliance
- ... and much more



Resources

Mailing lists:


gluster-users@gluster.org
gluster-devel@gluster.org

IRC:


#gluster and #gluster-dev on Freenode

Links:

<http://gluster.org/>
<http://gluster.readthedocs.io/>
<https://github.com/gluster/>



Thank you!



Niels de Vos
ndevos@redhat.com
ndevos on IRC
@nixpanic on Twitter