

GlusterFS for Sysadmins

David Nalley <david@gnsa.us> @ke4qqq

#whoami

- *Recovering sysadmin*
- *Apache CloudStack committer*
- *Member of the Gluster Advisory Board*
- *Distributed storage junkie*
- *Designer of ugly slides*
- *Poor imitation of Eco Willson*

My questions for you:

- *How many understand the relationship between LustreFS, GFS, GFS2, and GlusterFS?*
- *How many are using GlusterFS today?*
- *How many are using some distributed storage today, and if so, what is it?*

GlusterFS is...

- *a distributed/network/cluster filesystem...*
- *written in userspace*
- *proven in production at massive scale*
- *provides:*
 - *Scale-out NAS*
 - *Object storage*
 - *HDFS-compatible storage*
 - *and soon - decent block storage*

Volumes...

- *Made up of bricks on different nodes*
- *Each volume may have a different striped/replication/distribution strategy.*

*Distributed/replicated/striped/sm
othered/covered/cubed/*

- *LINK TO WAFFLE HOUSE HASH BROWN
MENU*

Menu options

- *Distributed*
- *Replicated*
- *Striped*
- *Distributed striped*
- *Distributed replicated*
- *Distributed striped replicated*

Distributed

- *Distributed volumes distributes files throughout the bricks in the volume. You can use distributed volumes where the requirement is to scale storage and the redundancy is either not important or is provided by other hardware/software layers.*

Replicated

- *Replicated volumes replicates files across bricks in the volume. You can use replicated volumes in environments where high-availability and high-reliability are critical.*

Striped

- *Striped volumes stripes data across bricks in the volume. For best results, you should use striped volumes only in high concurrency environments accessing very large files.*

Distributed Striped

- *Distributed striped volumes stripe data across two or more nodes in the cluster. You should use distributed striped volumes where the requirement is to scale storage and in high concurrency environments accessing very large files is critical.*

Distributed replicated

- *Distributed replicated volumes distributes files across replicated bricks in the volume. You can use distributed replicated volumes in environments where the requirement is to scale storage and high-reliability is critical. Distributed replicated volumes also offer improved read performance in most environments.*

Distributed striped replicated

- *Distributed striped replicated volumes distributes striped data across replicated bricks in the cluster. For best results, you should use distributed striped replicated volumes in highly concurrent environments where parallel access of very large files and performance is critical.*
- *Map Reduce only*

Striped replicated

- *Striped replicated volumes stripes data across replicated bricks in the cluster. For best results, you should use striped replicated volumes in highly concurrent environments where there is parallel access of very large files and performance is critical.*
- *Map Reduce only*

Ohhh yeah, don't forget georeplication...

- *Async*
- *Designed for remotely located storage*
- *Can replicate to more than two sites*

Userspace doesn't suck

- *Ignore your bias until you try it with GlusterFS*

The bad...

- *Log rotate not automated (wants you to call a command to do so)*
- *Rebalancing only partially automatic*

Demo setup

Presetup work done

- *GlusterFS installed*
(yum -y install glusterfs-server)
- *Filesystem created (XFS)*