# Gluster Introduction

Niels de Vos
GlusterFS co-maintainer
ndevos@redhat.com

Marcelo Barbosa
**GlusterFS contributor**
firemanxbr@fedoraproject.org

Agenda

- What is Gluster ?

- Architecture

- Quick start

- How to get involved ?

What is Gluster ?

Gluster is a distributed scale out filesystem that allows rapid provisioning of additional storage based on your storage consumption needs. It incorporates automatic failover as a primary feature. All of this is accomplished without a centralized metadata server.
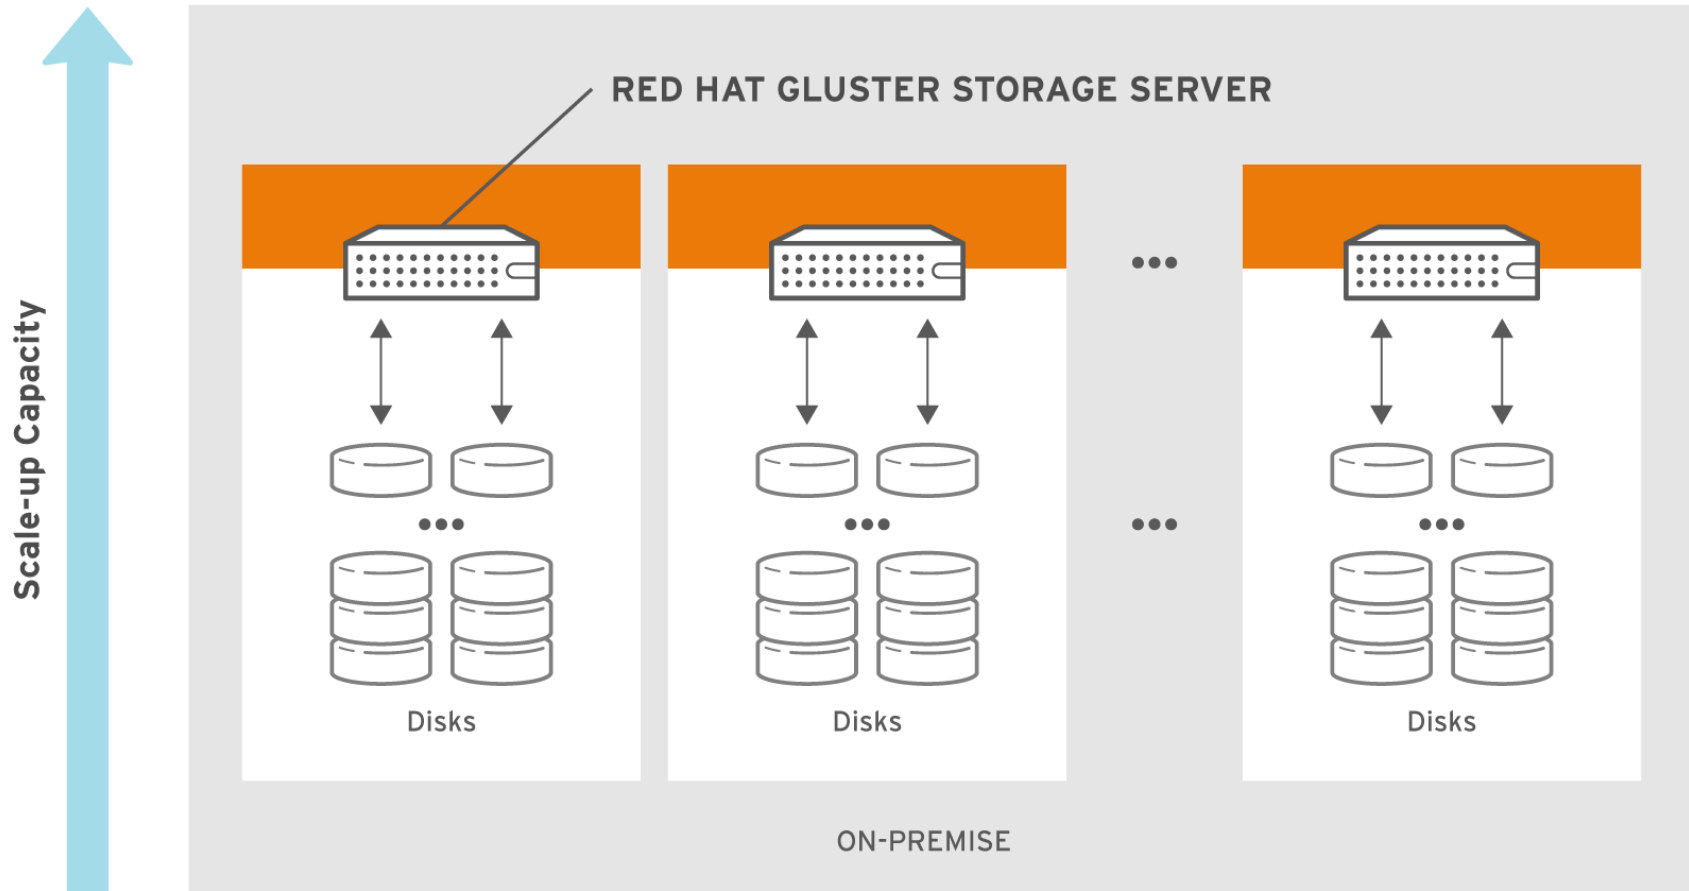
What is Gluster?

- 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

Scale-out and Scale-up
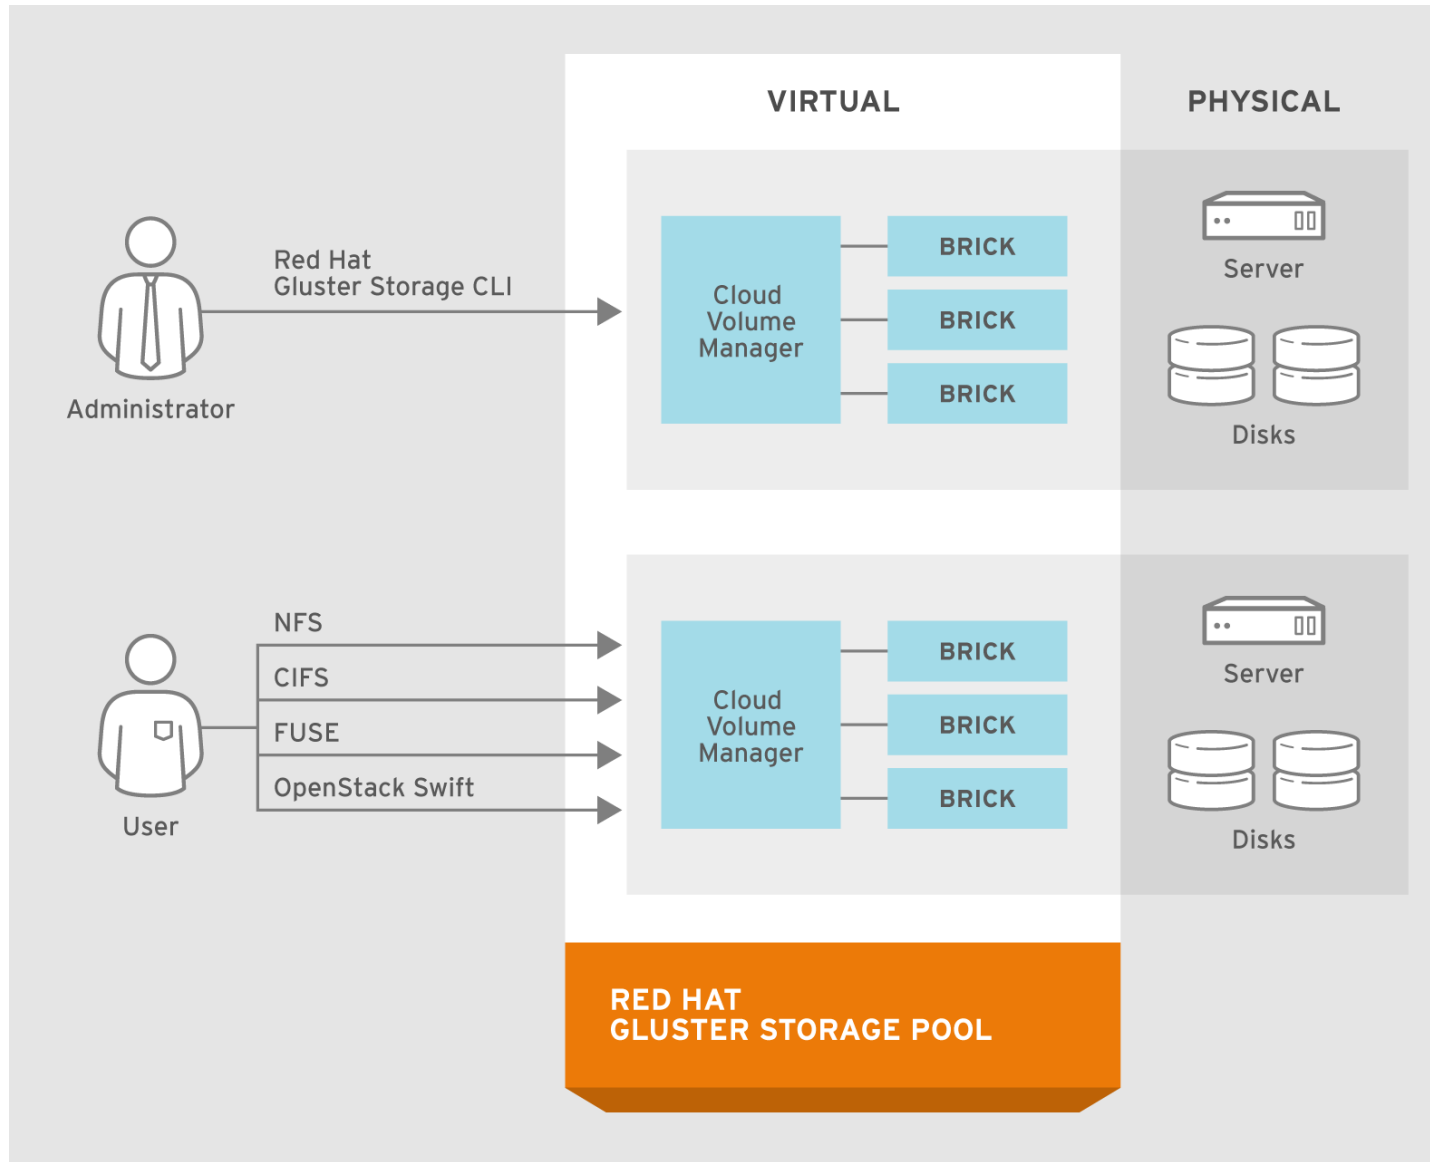
Data Access Overview

- 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)

- libgfapi flexible abstracted storage

  - Integrated with QEMU, Bareos and others
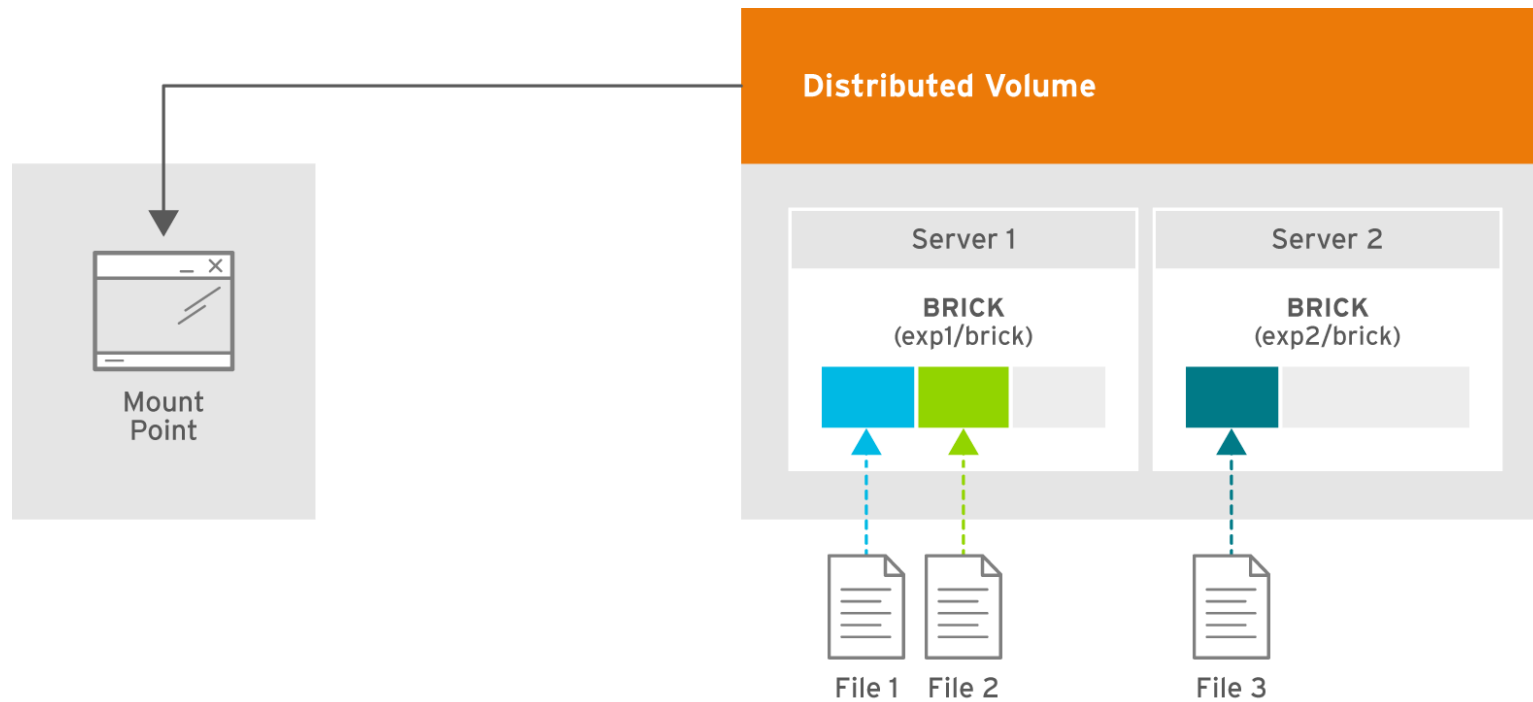
Architecture

Distributed Volume

- Files "evenly" spread across bricks
- *Similar* to file-level RAID 0
- Server/Disk failure could be catastrophic

Replicated Volume
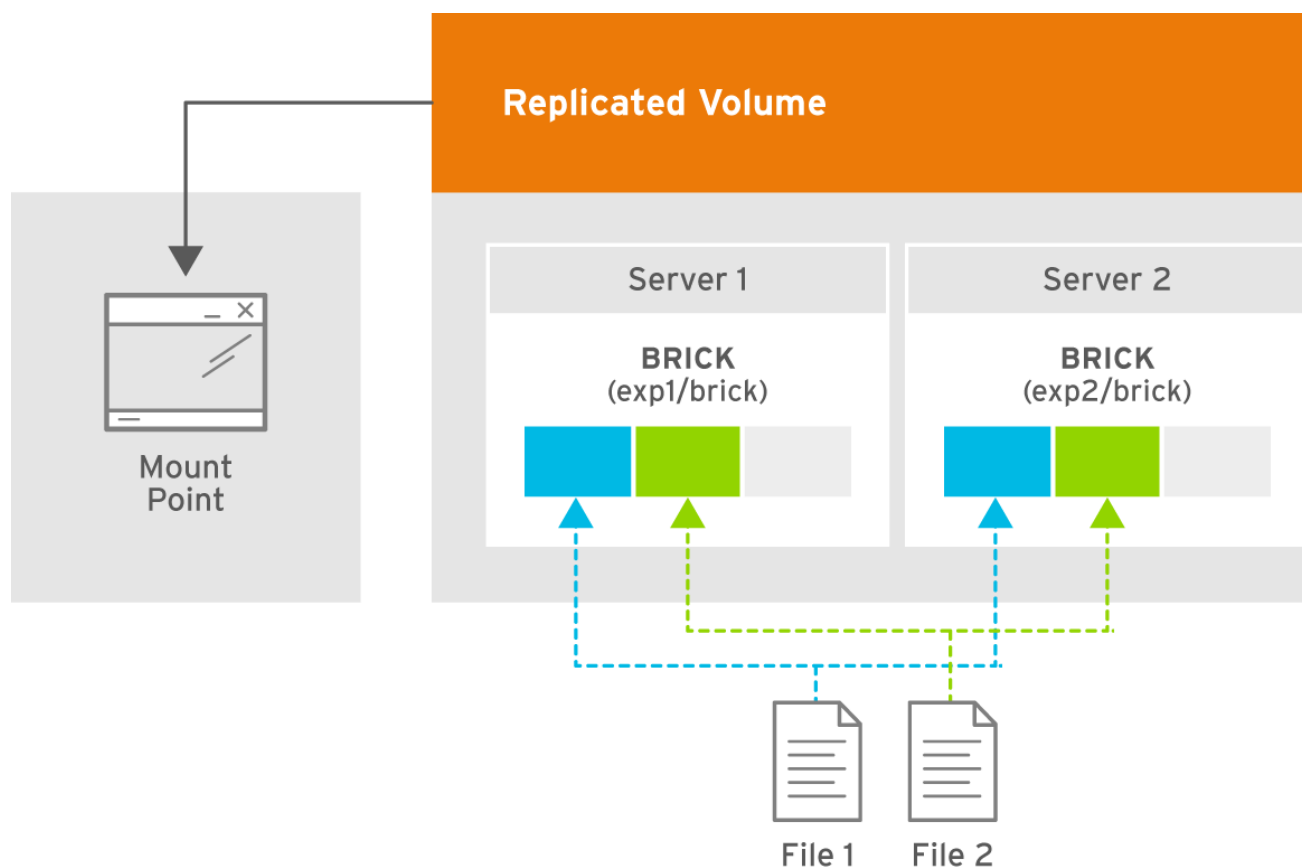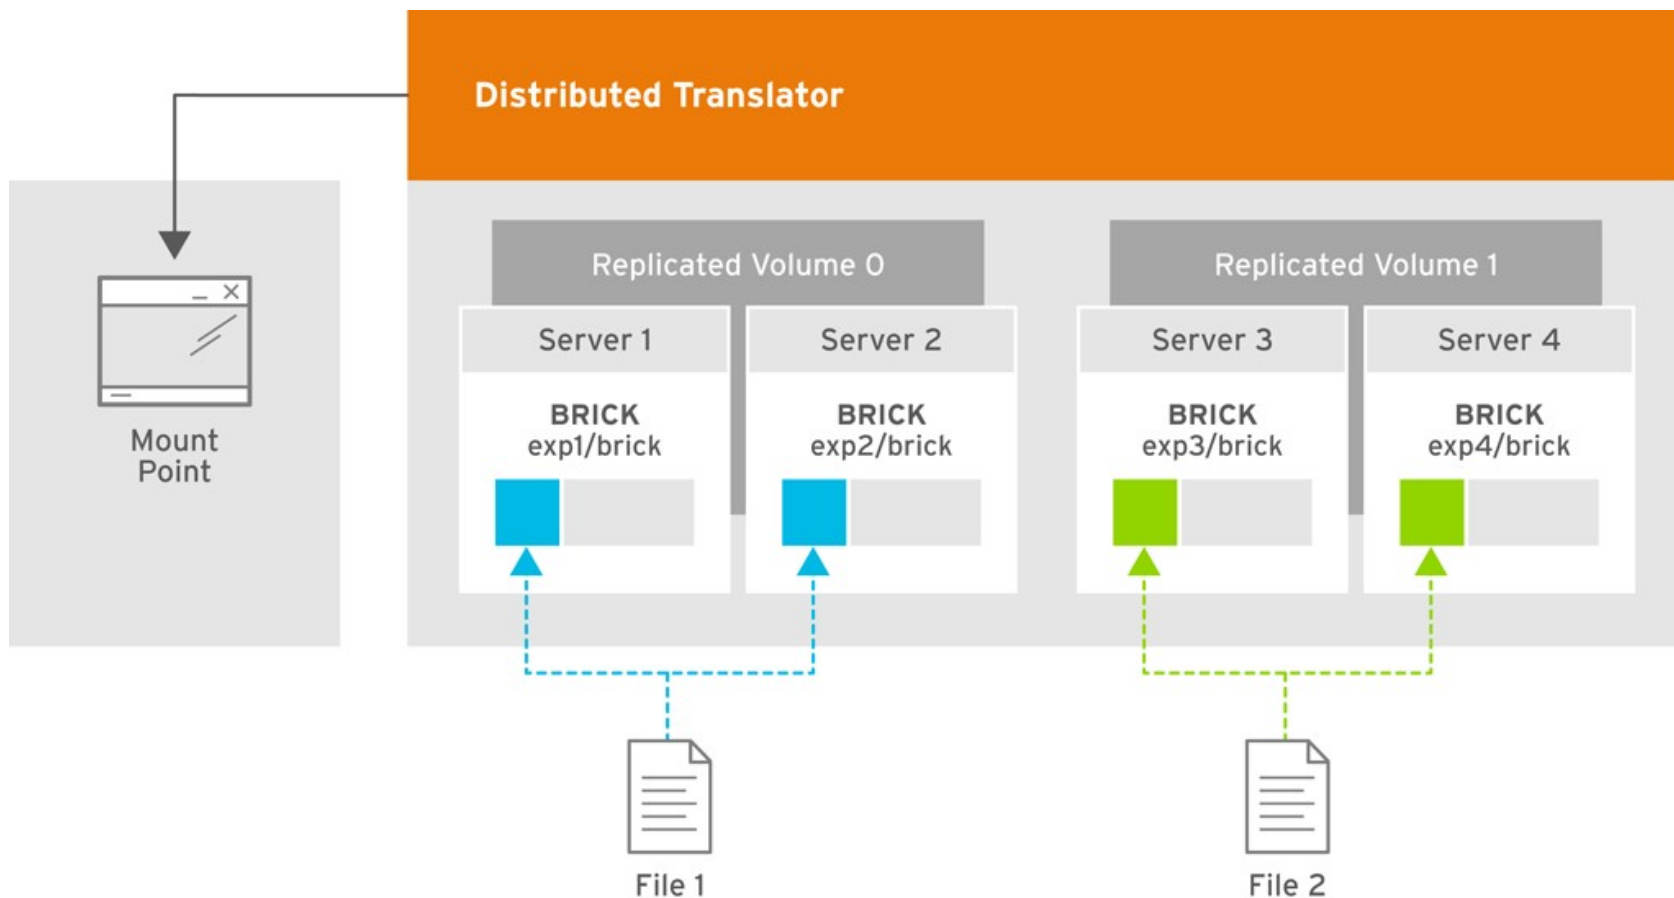
- Copies files to multiple bricks
- *Similar* to file-level RAID 1

Distributed Replicated Volume

- Distributes files across replicated bricks

Quick start

Assuming you have a disk at `/dev/sdb`:

```
# fdisk /dev/sdb
```

Format the partition:

```
# mkfs.xfs /dev/sdb1
```

Mount the partition as a Gluster "brick":

```
# mkdir -p /bricks/testvol
# mount /dev/sdb1 /bricks/testvol
```

Quick start

Add an entry to /etc/fstab:

```
# tail -n1 /proc/mounts >> /etc/fstab
```

Install Gluster packages on both nodes:

```
# yum install -y glusterfs-server (CentOS/RHEL 7)
```

```
# dnf install -y glusterfs-server (Fedora 22/23)
```

Quick start

Run the gluster peer probe command:

```
# gluster peer probe <ip or hostname of second host>
```

Configure your Gluster volume:

```
# gluster volume create testvol rep 2 \
    node01:/bricks/testvol/data \
    node02:/bricks/testvol/data
```

Test using the volume:

```
# mkdir /mnt/gluster
# mount -t glusterfs node01:/testvol
# cp -r /var/log /mnt/gluster
```

How to get involved ?

- Homepage
  http://gluster.org/

- Community IRC Chat (on Freenode)
  #gluster (for general topics)
  #gluster-dev (for developers)
  #gluster-meeting (meeting room)

- Mailing Lists
  http://www.gluster.org/mailman/listinfo/gluster-users
  http://www.gluster.org/mailman/listinfo/gluster-devel

- Documentation
  http://gluster.readthedocs.org/
  https://access.redhat.com/ - Red Hat Gluster Storage

Muito obrigado!