


Integration of Flexible Storage with the API of Gluster




Niels de Vos
GlusterFS co-maintainer

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


Agenda

- Introduction and Gluster basics
- Different approaches for automation
- The Gluster API
- Applications using the Gluster API




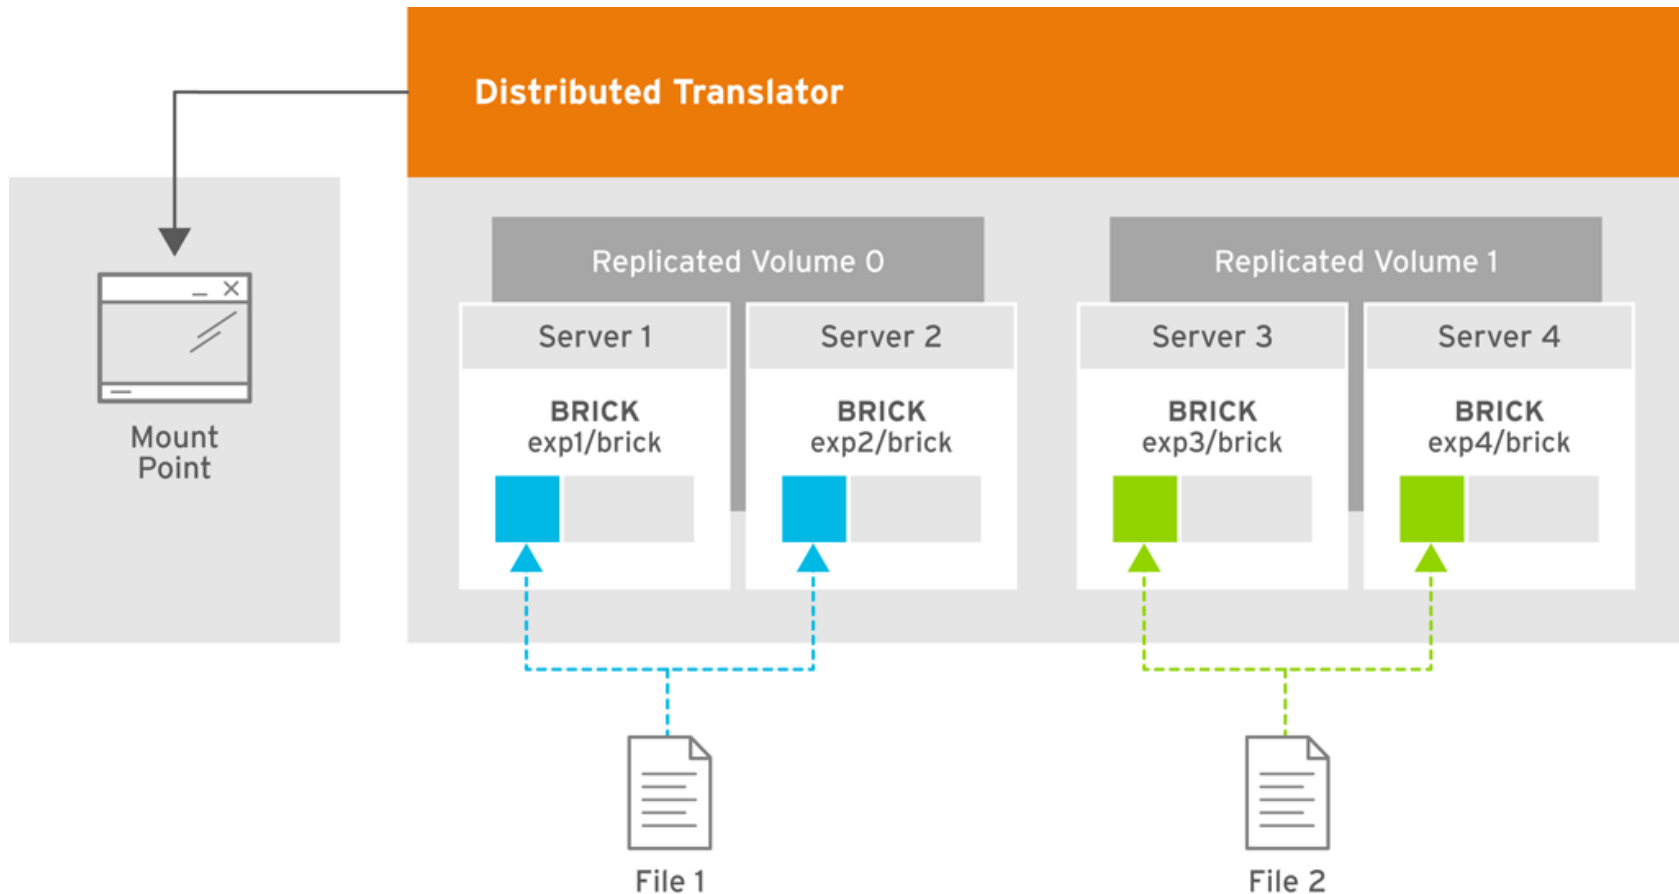
What is GlusterFS?

- Scalable, general-purpose storage platform
 - POSIX-y Distributed File System
 - Object Storage
 - Distributed block storage
 - Flexible storage accessible through an API
- No Metadata Server
- Heterogeneous Commodity Hardware
- Flexible and Agile Scaling
 - Capacity – Petabytes and beyond
 - Performance – Thousands of Clients



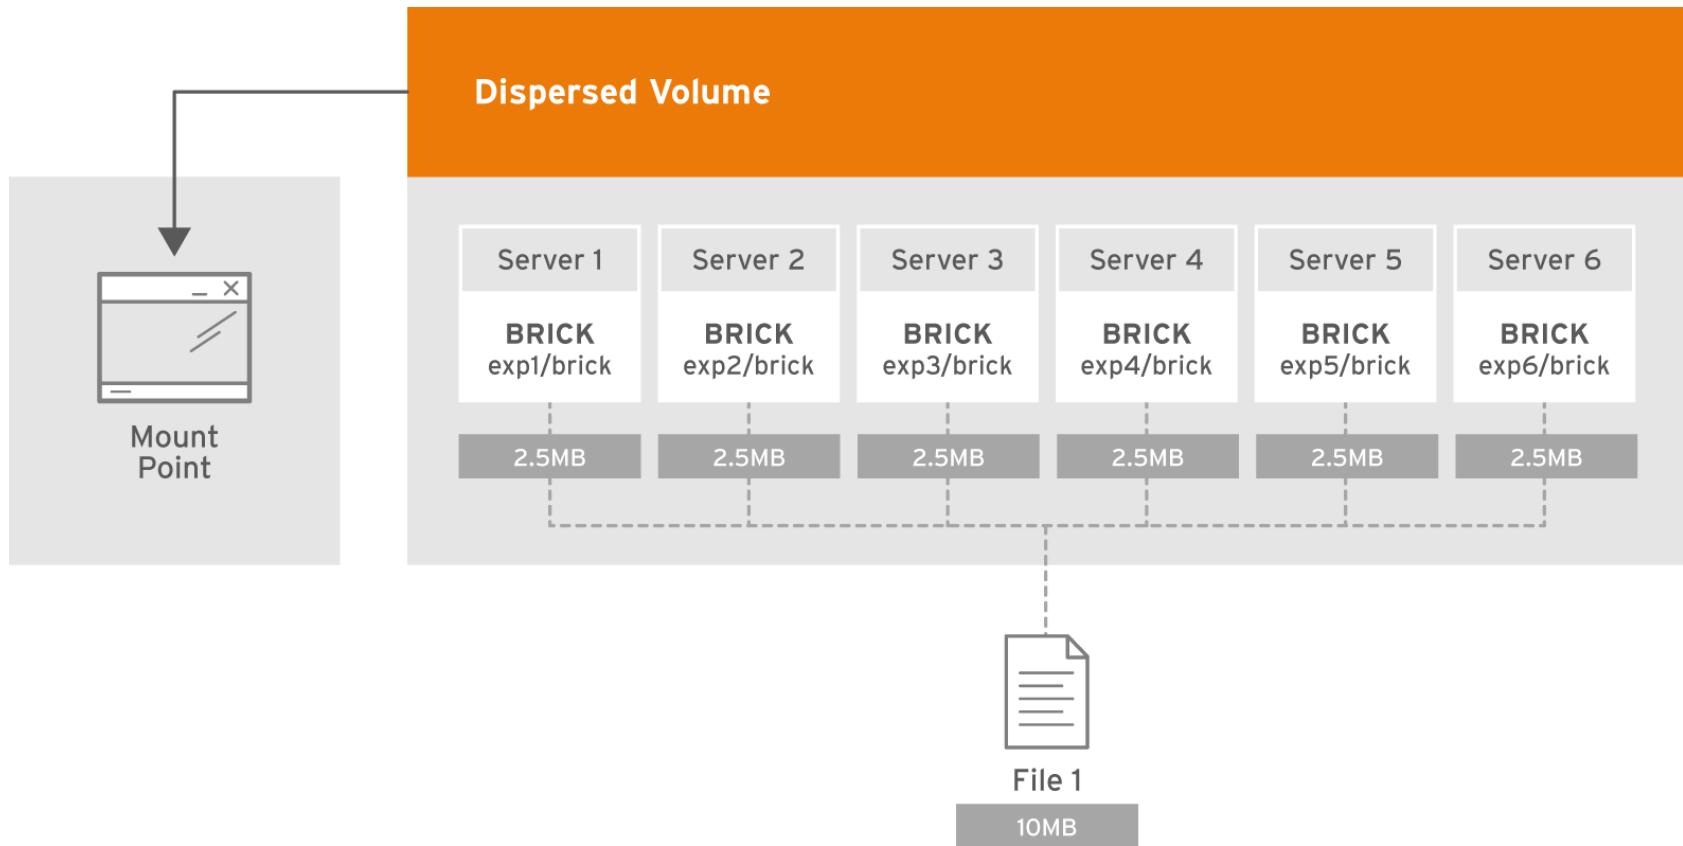
Distributed Replicated Volume

- 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



Automation Options




Automated Deployment and Management

- Ansible
- Gdeploy based on Ansible
- Kubernetes with Heketi
- oVirt for VMs on Gluster and Gluster management UI
- Storhaug for High-Availability NFS-Ganesha (+Samba)
- gluster-block for configuring iSCSI-targets




Example: Kubernetes and Heketi

- Dynamic persistent storage provisioning for containers
- Heketi handles storage requests from Kubernetes
 - Creation, snapshots, deletion, ...
- See presentation and demo from FOSDEM
<https://fosdem.org/2017/schedule/event/glustercontainer/>
- Try it with gluster-kubernetes (uses Vagrant)
<https://github.com/gluster/gluster-kubernetes>




Example: gluster-block

- Creates a disk image on a Gluster Volumes
- Configured (multipath) iSCSI
- Small file workloads are horrible for network filesystems
- Many network procedures cause for slow performance
- Pro: reduce the number of procedures
- Con: not a shared filesystem, single system/container




Gluster API introduction



Language Bindings


- Best performance for accessing data on Gluster
- No need to run applications as root
- libgfapi is a library written in C
- Python interface with libgfapi-python
- Java support through glusterfs-java-filessystem and libgfapi-jni
- New Golang bindings with gogfapi




Example: gfapi in Python

- Create a Volume to access the contents
- A volume acts like a directory
- Well documented API
<http://libgfapi-python.readthedocs.io/>

```
$ python
>>> from gluster import gfapi
>>> v = gfapi.Volume('vm122-180.example.com', 'objects')
>>> v.mount()
>>> v.mkdir('/idi2017')
>>> v.listdir('/')
['idi2017']
>>> fd = v.fopen('/idi2017/schedule.html', 'w')
>>> ret = fd.write('Welcome to Bologna!\n')
>>> fd.close()
```




Existing Use-Cases



Approaches to use Gluster

- Network filesystem
 - Standard applications, no modification needed
 - Much overhead due to POSIX filesystem API
- Object Storage
 - Simplified interface for GET/PUT data
- (virtual) block device
 - Large files for virtual machines
 - Loopback mounted filesystems for containers
- API for tight integration in other projects



Applications using the API


- Backup to and from Gluster
 - Bareos
- Virtual Machine
 - QEMU, libvirt, oVirt, OpenStack, CloudStack, Proxmox
- Network Filesystem
 - Samba, NFS-Ganesha
- Object Storage
 - gluster-swift with S3 option



Example: QEMU

- qemu-img with gluster:// URL


```
$ qemu-img create \  
  -f qcow2 \  
  gluster://vm122-180.example.com/vms/idi2017.qcow2 \  
  2G
```



Example: S3 Object Storage

- Object Storage
 - Proof of concept gluster-swift container with S3 option
<http://uyirpodiru.blogspot.com/2017/03/building-gluster-object-in-docker.html>
- s3curl example tool in Perl
<https://aws.amazon.com/code/128>

```
$ ./s3curl.pl --id demo \  
    -- http://172.17.0.2:8080/  
$ ./s3curl.pl --id demo --createBucket \  
    -- http://172.17.0.2:8080/idi2017  
$ ./s3curl.pl --id demo --put README \  
    -- http://172.17.0.2:8080/idi2017/README
```



Resources

Mailing lists:

integration@gluster.org

gluster-users@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