

# Building a fileserver-cluster with GlusterFS

Stefan Kania

12. Mai 2016

# Introduction

Why use Clusters?

# Introduction

## Why use Clusters?

- Eliminate *single point of failure*

# Introduction

## Why use Clusters?

- Eliminate *single point of failure*
- *failover* in case of a server outage

# Introduction

## Why use Clusters?

- Eliminate *single point of failure*
- *failover* in case of a server outage
- *Load balancing* under heavy load

# Introduction

## Why use Clusters?

- Eliminate *single point of failure*
- *failover* in case of a server outage
- *Load balancing* under heavy load
- Maintenance during normal working hours

# GlusterFS

## Characteristics of GlusterFS

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume



# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network
- Expandable

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network
- Expandable
- PosixACL Support

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network
- Expandable
- PosixACL Support
- Many different configurations possible

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network
- Expandable
- PosixACL Support
- Many different configurations possible
- self healing in case on knode is down

# GlusterFS

## Characteristics of GlusterFS

- merge space from many knodes to one volume
- mount via *fusemount* over the network
- Expandable
- PosixACL Support
- Many different configurations possible
- self healing in case on knode is down
- support of filesystem-snapshots (LVM2), starting with version 3.6

# GlusterFS

## Different configurations

# GlusterFS

## Different configurations

- Replicated Volume



# GlusterFS

## Different configurations

- Replicated Volume
- Distributed Volume

# GlusterFS

## Different configurations

- Replicated Volume
- Distributed Volume
- Striped Volume

# GlusterFS

## Different configurations

- Replicated Volume
- Distributed Volume
- Striped Volume
- Replicated-Distributed Volume

# GlusterFS

## Configuring GlusterFS

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks
- mount the filesystem on all knodes



# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks
- mount the filesystem on all knodes
- create the volume

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks
- mount the filesystem on all knodes
- create the volume
- start the volume

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks
- mount the filesystem on all knodes
- create the volume
- start the volume
- mount the volume

# GlusterFS

## Configuring GlusterFS

- Configure your system to use the packages from [gluster.org](http://gluster.org)
- Install the packages
- format the filesystem on all bricks
- mount the filesystem on all knodes
- create the volume
- start the volume
- mount the volume
- use the volume and be happy

# GlusterFS

## Gluster and Snapshots

# GlusterFS

## Gluster and Snapshots

- all bricks have to be configured with LVM2

# GlusterFS

## Gluster and Snapshots

- all bricks have to be configured with LVM2
- Gluster-version must be at least 3.6

# GlusterFS

## Gluster and Snapshots

- all bricks have to be configured with LVM2
- Gluster-version must be at least 3.6
- All bricks have to be on it's own partition



# GlusterFS

## Gluster and Snapshots

- all bricks have to be configured with LVM2
- Gluster-version must be at least 3.6
- All bricks have to be on it's own partition
- All LVM2-volumes must be *thinly provisioned Volumes*

## how it works



Let's do some clustering