



# Gluster Overview

Scott McCarty  
5/2012

# Red Hat Storage Technology Highlights

## High Performance

- Fast file access - no metadata server
- Scale-out with linear performance
- Petabyte scale storage
- 1GbE, 10GbE

## Architecture

- Software only—you choose the infrastructure
- Fully distributed architecture—no metadata server
- Simultaneous file and object storage
- Stackable userspace design
- No kernel dependencies
- Simple installation
- Early maturity and rich functionality

## High-Availability

- No single point of failure
- Replication to survive hardware failure
- Self-healing
- Data stored in NFS-like native format
- Business Continuity and Disaster Recover
- Replicate from datacenters to public clouds

## No Hot-Spots/Bottlenecks

- File placement and retrieval is algorithmically controlled
- Files are spread evenly throughout the file system

## Single Global Namespace

- Single mount point – easier storage management
- Spans data center/private cloud or public cloud
- Eases and accelerates movement to the cloud



# Why It's Different?

## No metadata server

- No single point of failure, automated self heal and failover
- No performance bottleneck on data lookups for fast file access

## Built in replication

- Synchronous for inter-node replication
- Asynchronous for geo-replication

## No block size restrictions

- Ideal for small and large files

## POSIX compatible file system

- Out of the box NFS, CIFS and GlusterFS native access

## Expanded data access options

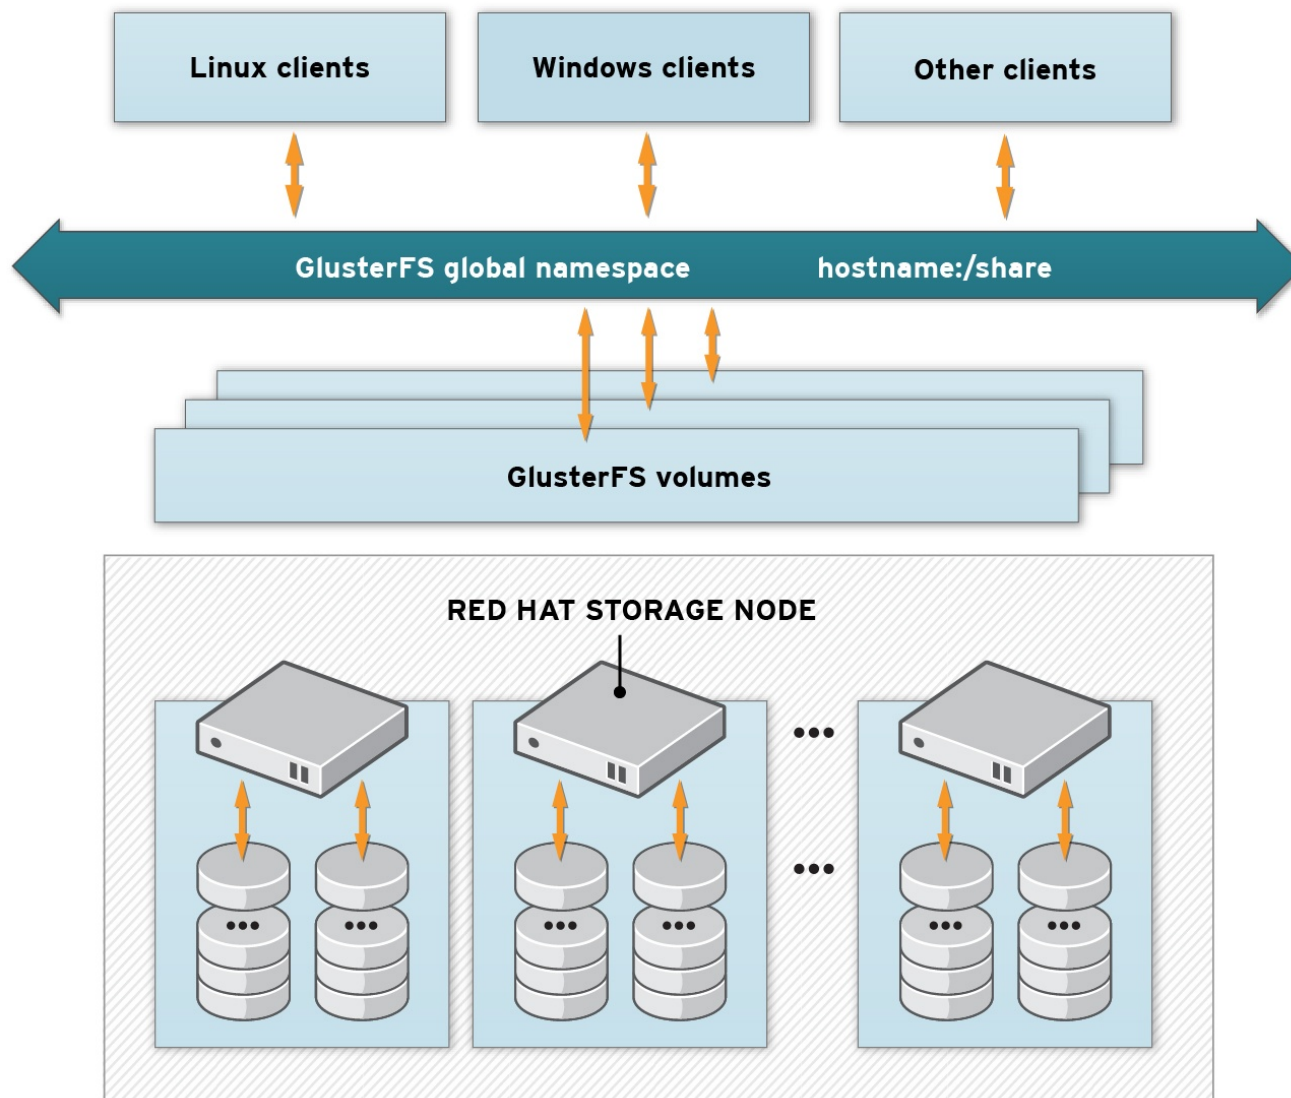
- File and object access to data
- Access files from your object interface and access data within objects as files
- File based applications can access data without modification

## Reduces requirement for replicated files from 3 to 2

- 33% capacity savings



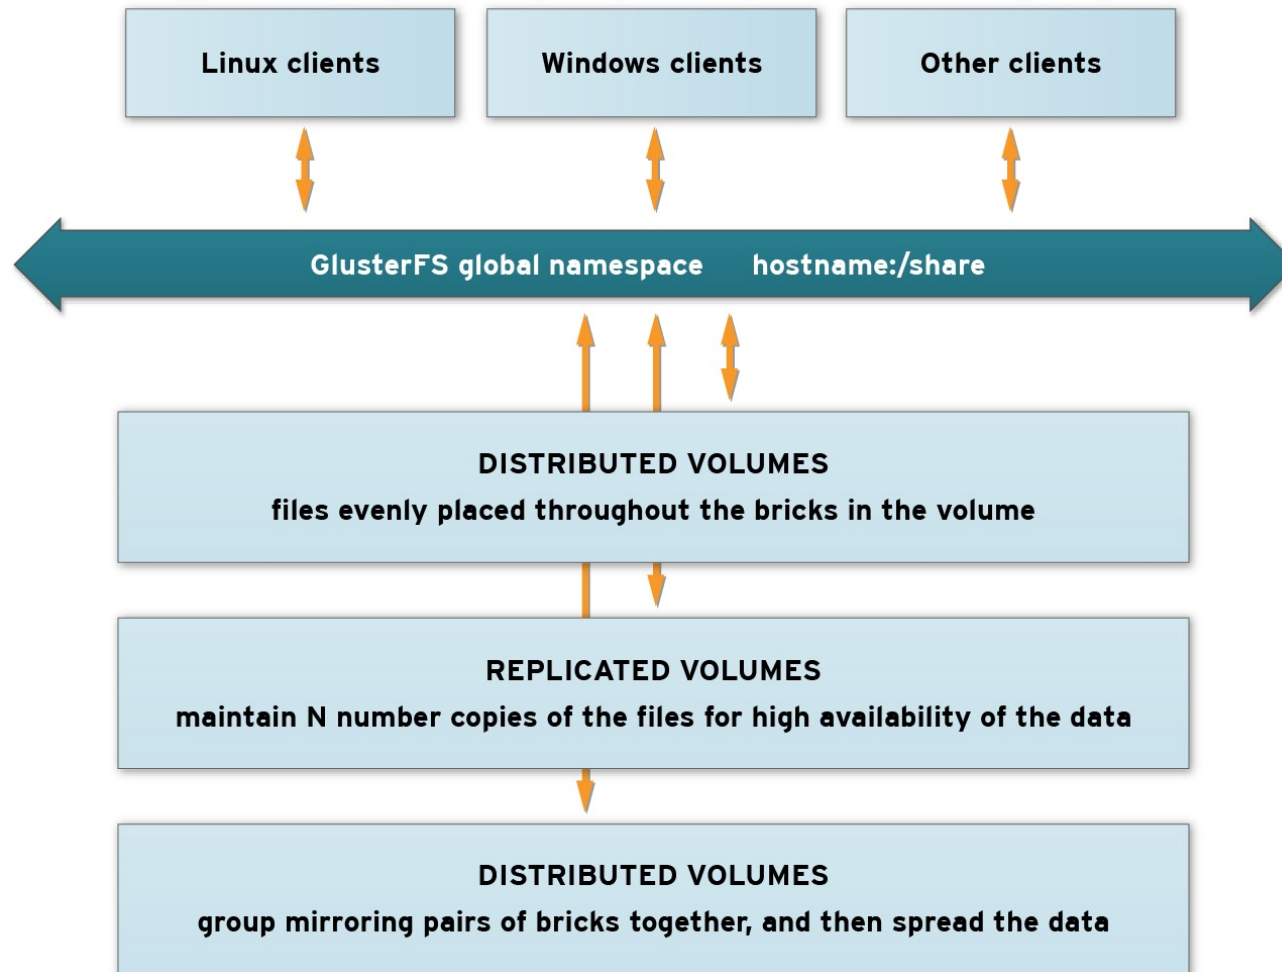
# Providing Access to Your Data



- GlusterFS enables you to create a Global Namespace
- On that namespace you can create volumes where data resides
- Clients access data from the volumes
- GlusterFS handles all volume-level policies
  - Distribute
  - Replicate
  - Geo-Rep
  - And more...



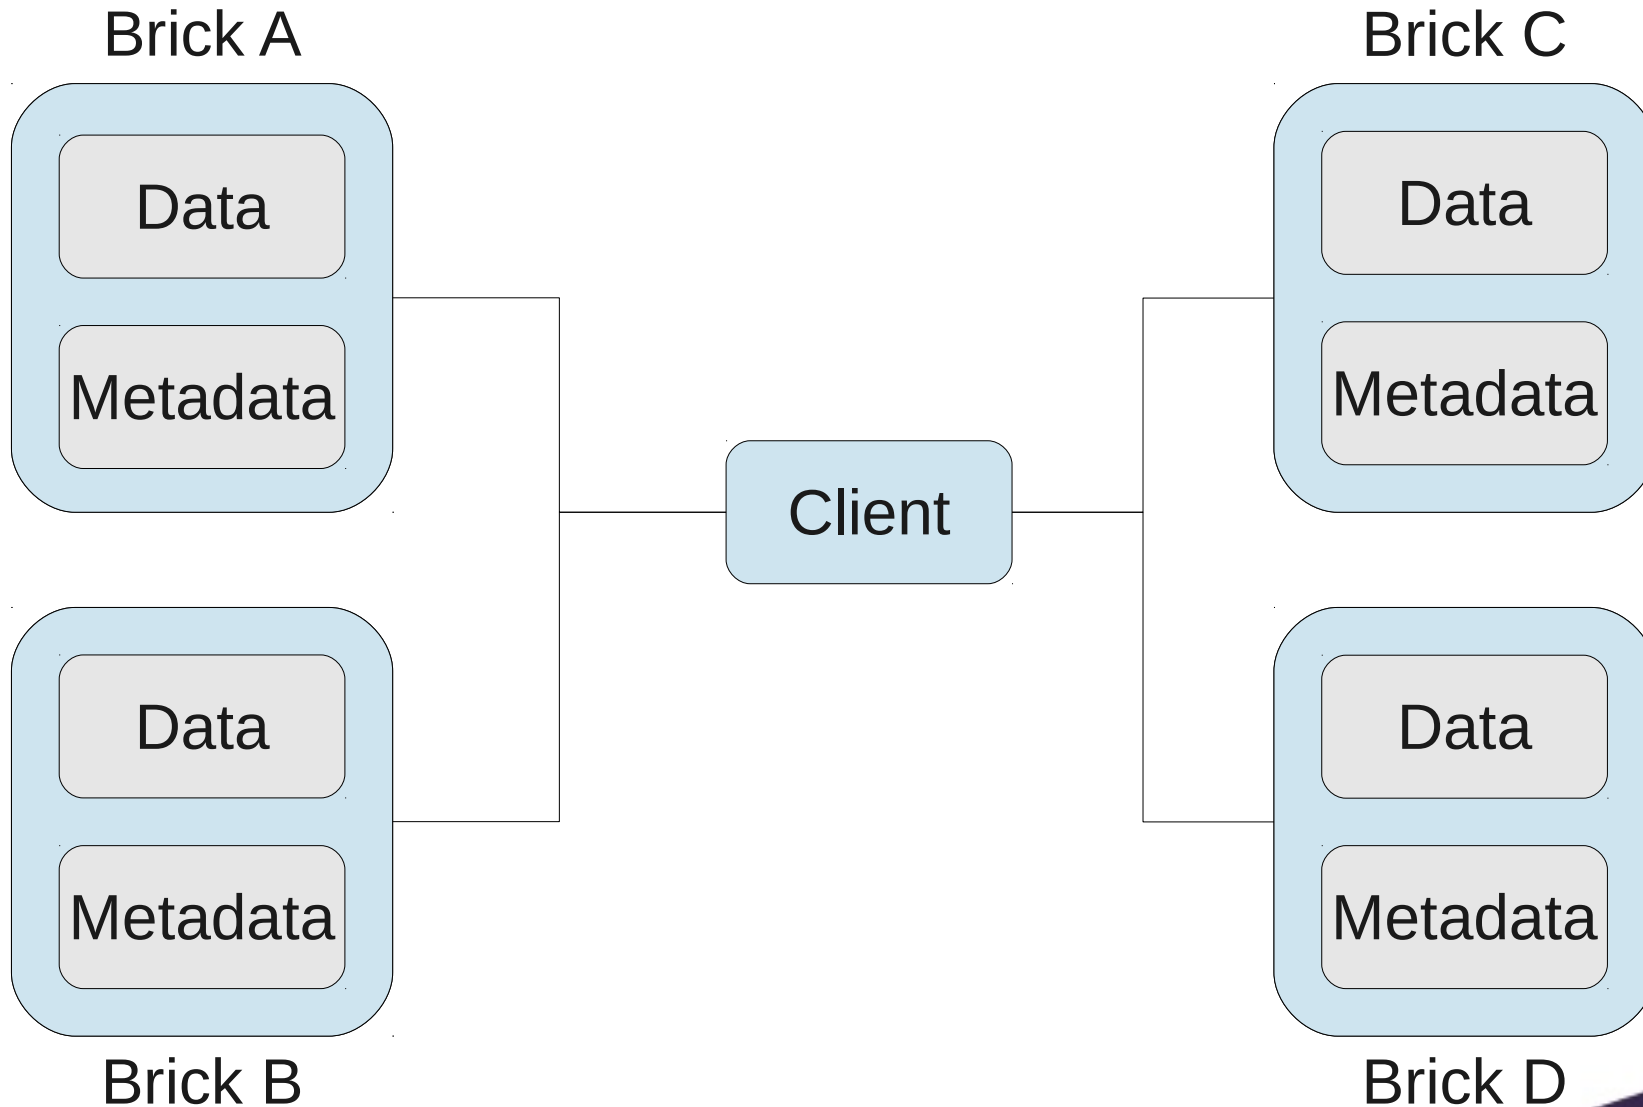
# Gluster FS Handles Everything Else From There



- Any GlusterFS node can handle client requests
- GlusterFS handles distributing, replicating, and remotely replicating the data
- Clients perceive volumes as being one share that they can read and write the data
- Everything that GlusterFS does behind that is transparent to the client



# GlusterFS Diagram





# Questions