



GlusterFS

Petascale Cloud Filesystem

GlusterFS Simple Commands

gluster peer probe HOSTNAME

gluster volume info

gluster volume create VOLNAME [stripe COUNT] [replica COUNT] [transport tcp | rdma] BRICK ...

gluster volume delete VOLNAME

gluster volume add-brick VOLNAME NEW-BRICK ...

gluster volume rebalance VOLNAME start



1st meeting room




1st Office US

Story of Gluster

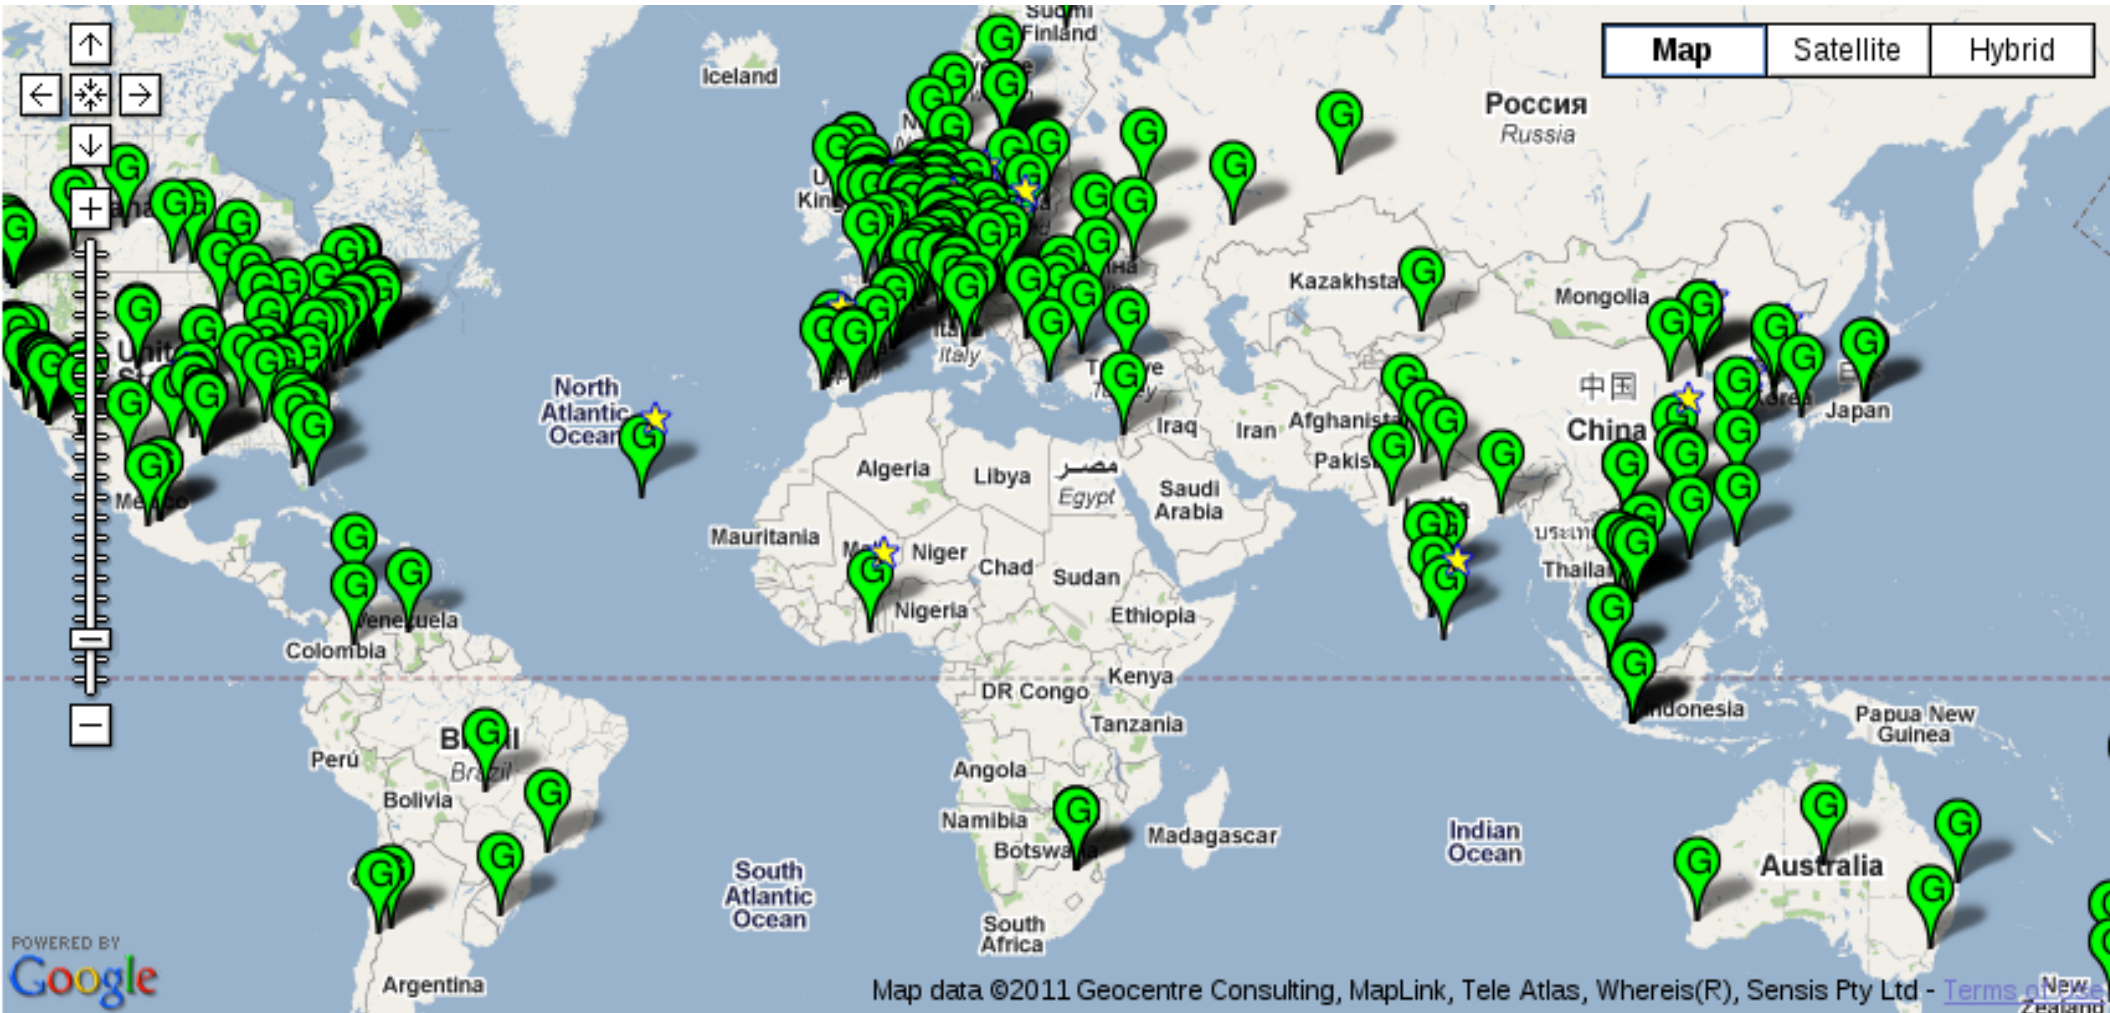


1st Office Bengalooru



1st Office Bengalooru

Registered Gluster Deployments





Deutsche Bank



GlusterFS

Commercial Deployments

Gluster Architecture Advantages

'Google Storage' for Everyone

- Intelligence in the SW
- Leverage commodity HW
- Scale-out elastically
- Replication for reliability
- Software enables virtualization

❖ Software only

❖ No metadata server

- Fully distributed architecture, no bottleneck
- Gluster Elastic Hash

❖ High performance global namespace

- Scale out with linear performance
- Hundreds of petabytes
- 1 GbE, 10GbE and InfiniBand (native RDMA)

❖ High availability

- Replication to survive hardware failure
- Self-healing
- Data stored in NFS-like native format

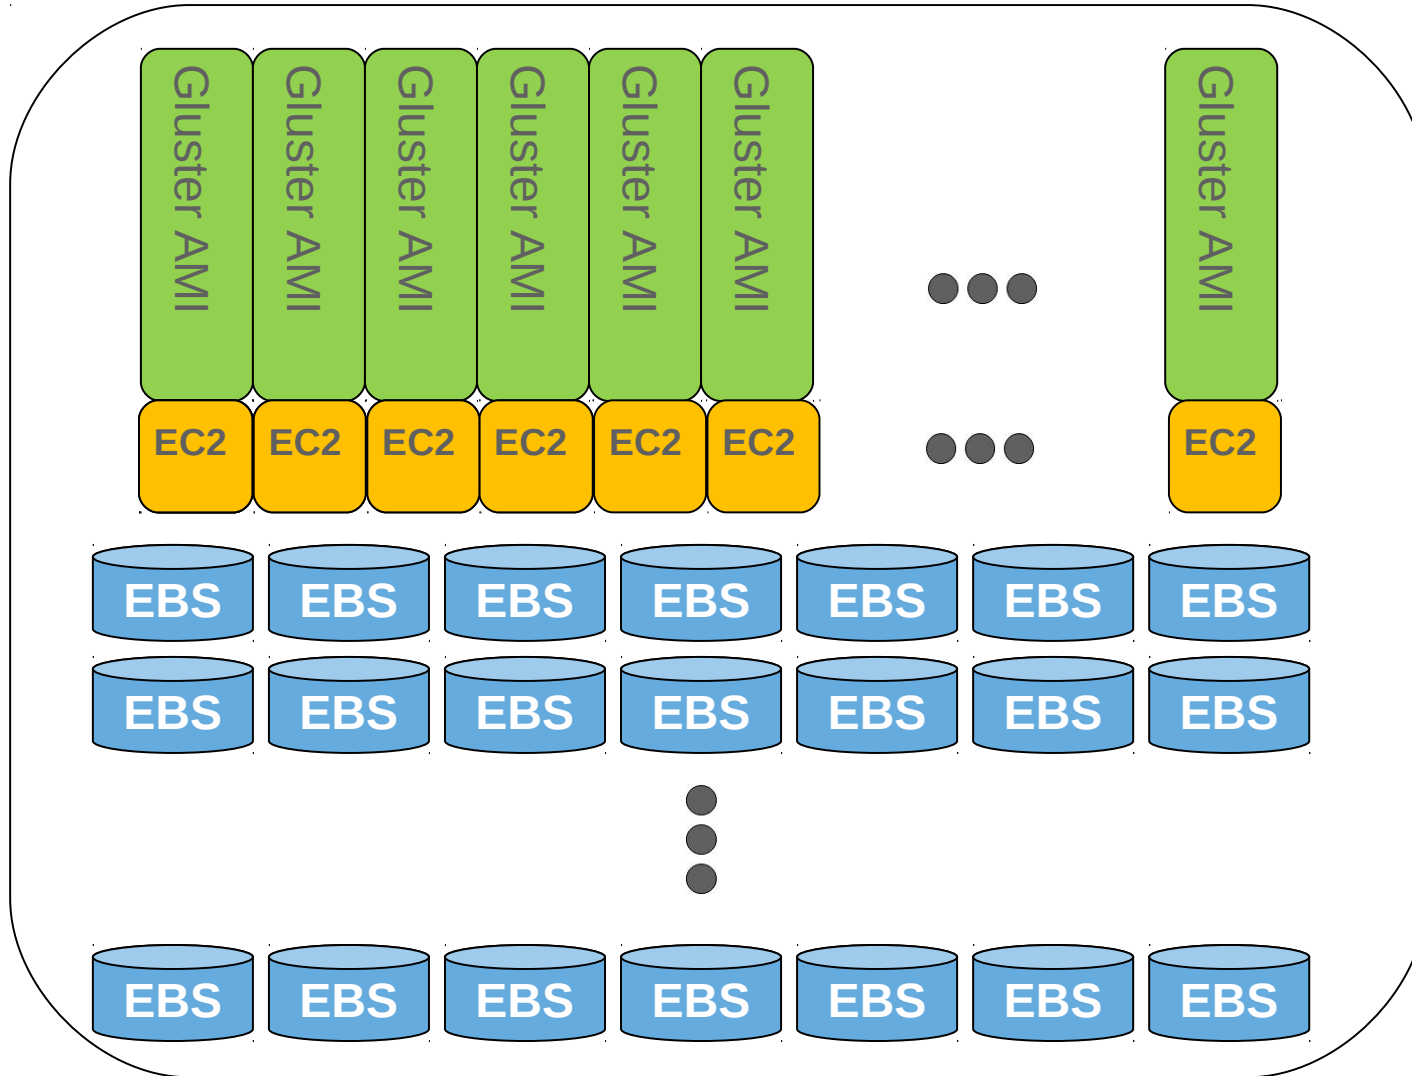
❖ Stackable userspace design

- No kernel dependencies, simple install
- Match specific workload profiles
- Early maturity and rich functionality

Gluster Elastic Capabilities

- ❖ **Volumes abstracted from the hardware**
- ❖ **Grow, shrink, migrate across physical systems**
 - Like vMotion for storage volumes
 - Automatic rebalancing
- ❖ **No application interruption**
- ❖ **File system configuration changes accepted at run time**
 - React to changing workloads, live performance tuning
- ❖ **Native NFS**

Anatomy of a storage pool: Public Cloud (e.g. AWS)



Pandora Internet Radio



- 1.2 PB of audio served per week
- 13 million files
- Over 50 GB/sec peak traffic

❖ Problem

- Explosive user & title growth
- As many as 12 file formats for each song
- 'Hot' content and long tail

❖ Solution

- Three data centers, each with a six-node GlusterFS cluster
- Replication for high availability
- 250+ TB total capacity

❖ Benefits

- Easily scale capacity
- Centralized management; one administrator to manage day-to-day operations
- No changes to application
- Higher reliability

Partners Healthcare

Private Cloud: Centralized Storage as a Service



- Over 500 TB
- 9 Sun “Thumper” systems in cluster

❖ Problem

- Capacity growth from 144TB to 1+PB
- Multiple distributed users/departments
- Multi OS access - Windows, Linux and Unix

❖ Solution

- GlusterFS Cluster
- Solaris/ZFS/x4500 w/ InfiniBand
- Native CIFS/ NFS access

❖ Benefits

- Capacity on demand / pay as you grow
- Centralized management
- Higher reliability
- OPEX decreased by 10X



Join our Gluster Community

[Community.gluster.org](https://community.gluster.org)