



# Open Storage in the Enterprise

## With GlusterFS and Red Hat Storage

Dustin L. Black, RHCA  
Sr. Technical Account Manager & Team Lead  
Red Hat Global Support Services

LinuxCon Europe -- 2013-10-23





Dustin L. Black, RHCA  
Sr. Technical Account Manager  
Red Hat, Inc.

dustin@redhat.com  
@dustinblack





# #what<sub>i</sub>s TAM


- Premium named-resource support
- Proactive and early access
- Regular calls and on-site engagements
- Customer advocate within Red Hat and upstream
- Multi-vendor support coordinator
- High-touch access to engineering
- Influence for software enhancements
- **NOT** Hands-on or consulting

# What is GlusterFS?

- 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
- Standards-Based – Clients, Applications, Networks
- Flexible and Agile Scaling
  - Capacity – Petabytes and beyond
  - Performance – Thousands of Clients


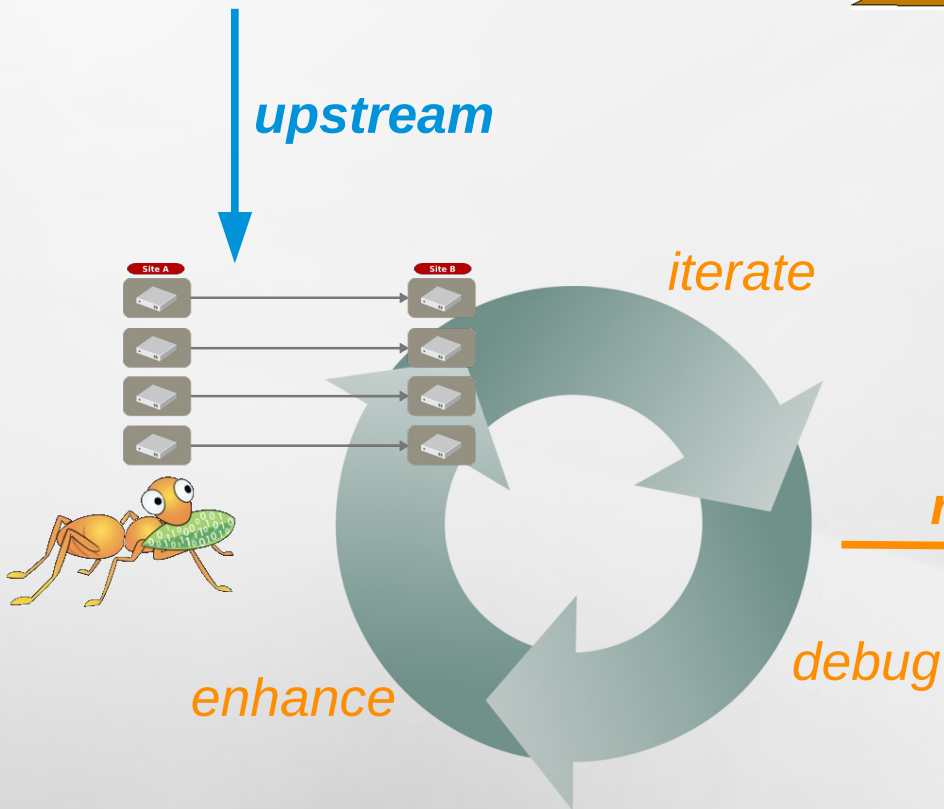



**Network Interconnect**



# What is Red Hat Storage?

- Enterprise Implementation of GlusterFS
- Software Appliance
- Bare Metal Installation
- Built on RHEL + XFS
- Subscription Model
- Storage Software Appliance
  - Datacenter and Private Cloud Deployments
- Virtual Storage Appliance
  - Amazon Web Services Public Cloud Deployments



# Use Case: Media Storage via Object Interface

**Open Storage in the Enterprise**

With GlusterFS and Red Hat Storage

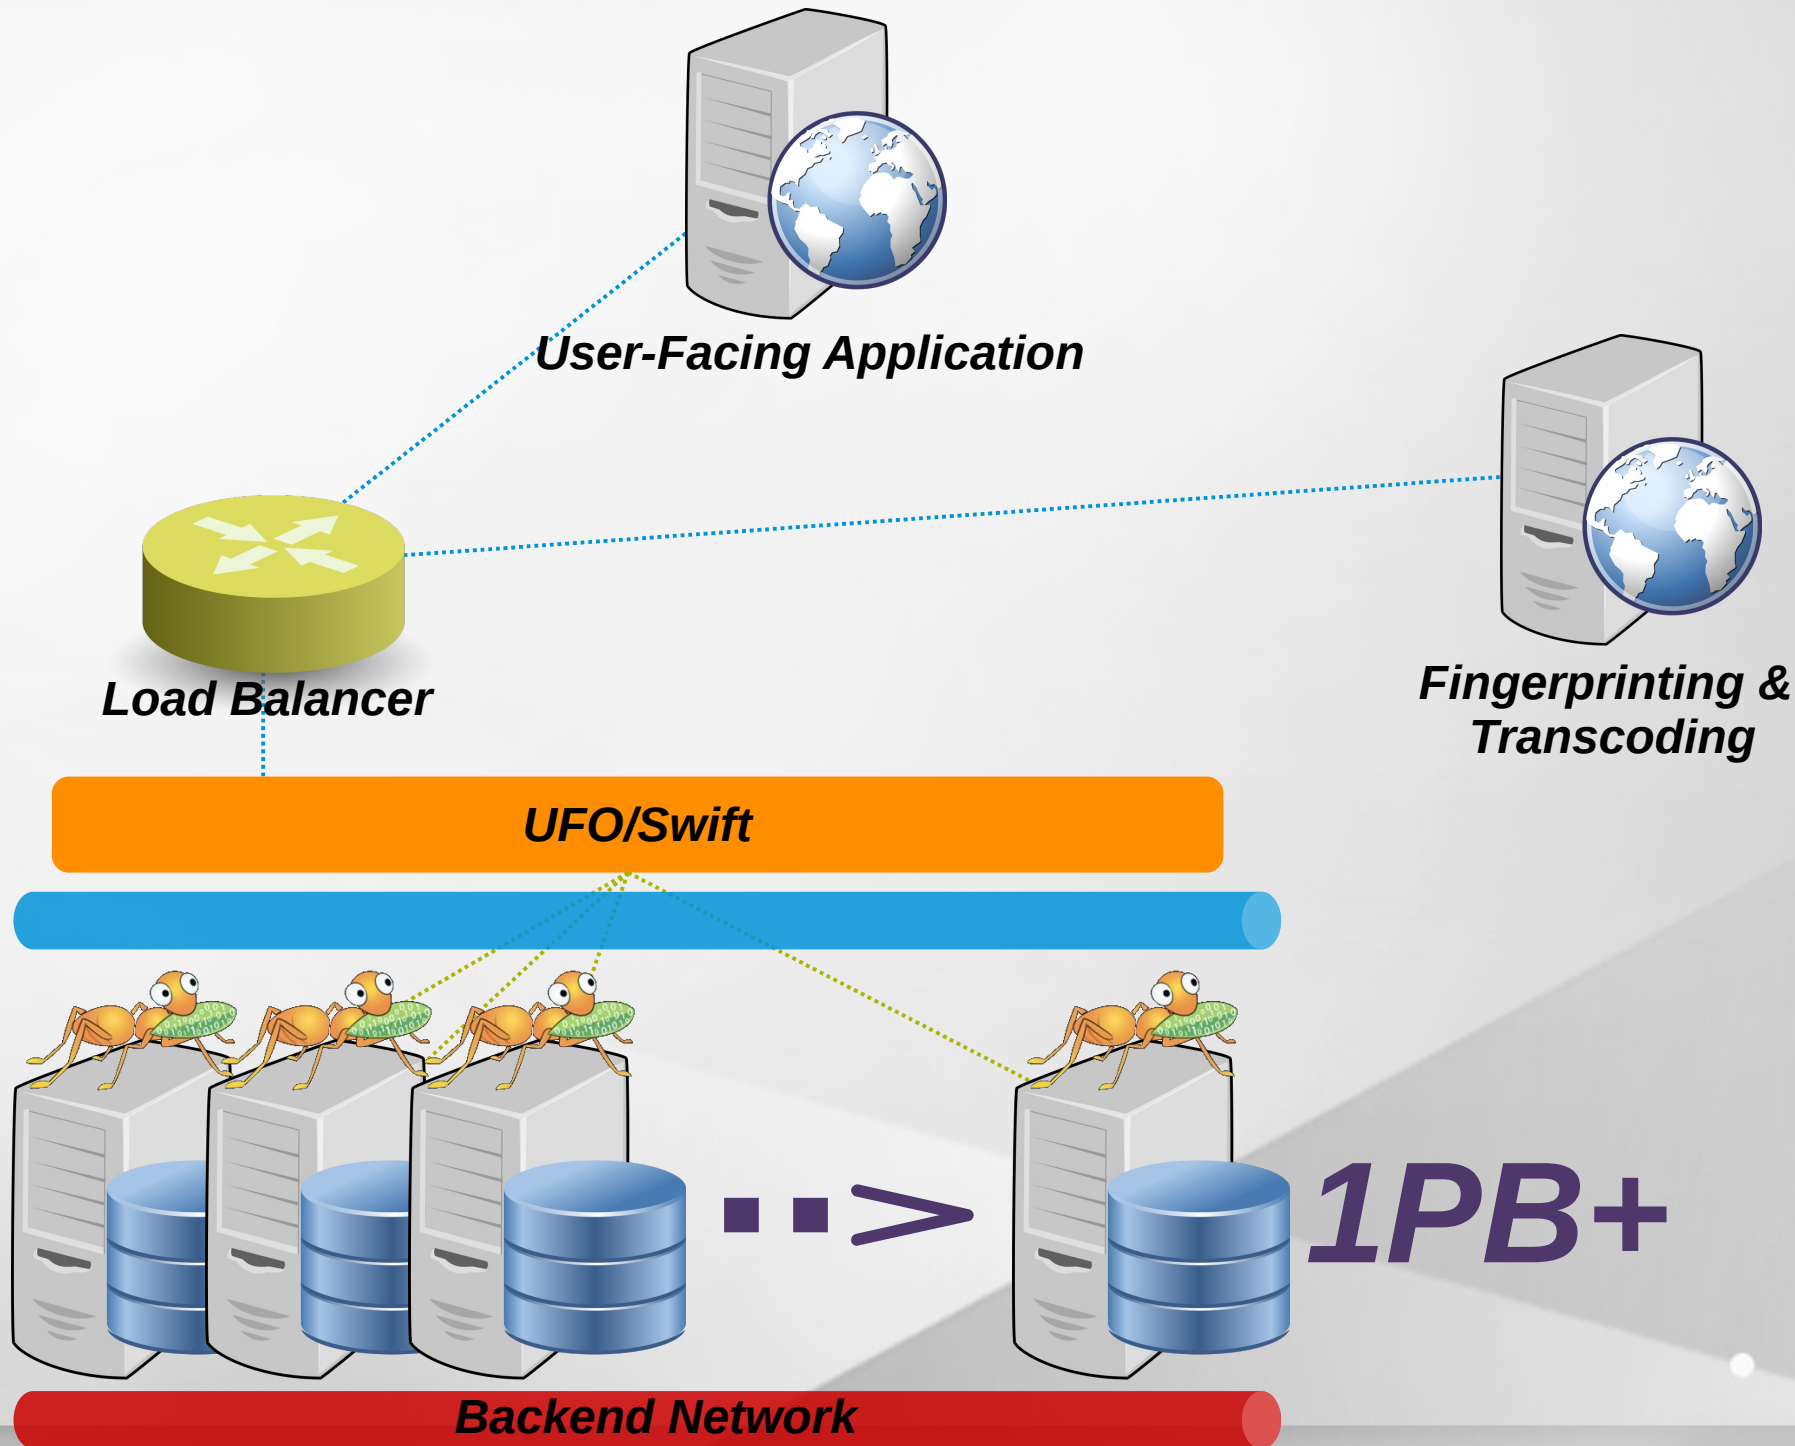


# Goals

- Media file storage for customer-facing app
- Drop-in replacement for legacy object backend
- 1PB plus 1TB/day growth rate
- Minimal resistance to increasing scale
- Multi-protocol capable for future services
- Fast transactions for fingerprinting and transcoding

# Implementation

- 12 Dell R710 nodes + MD1000/1200 DAS
  - Growth of 6 -> 10 -> 12 nodes
- ~1PB in total after RAID 6
- GlusterFS Swift interface from OpenStack
- Built-in file+object simultaneous access
- Multi-GBit network with segregated backend



# Use Case: Self-Service Provisioning with Accounting and Chargeback

**Open Storage in the Enterprise**

With GlusterFS and Red Hat Storage

# Goals

- Add file storage provisioning to existing self-service virtualization environment
  - Automate the administrative tasks
- Multi-tenancy
  - Subdivide and limit usage by corporate divisions and departments
  - Allow for over-provisioning
  - Create a charge-back model
- Simple and transparent scaling

# Implementation

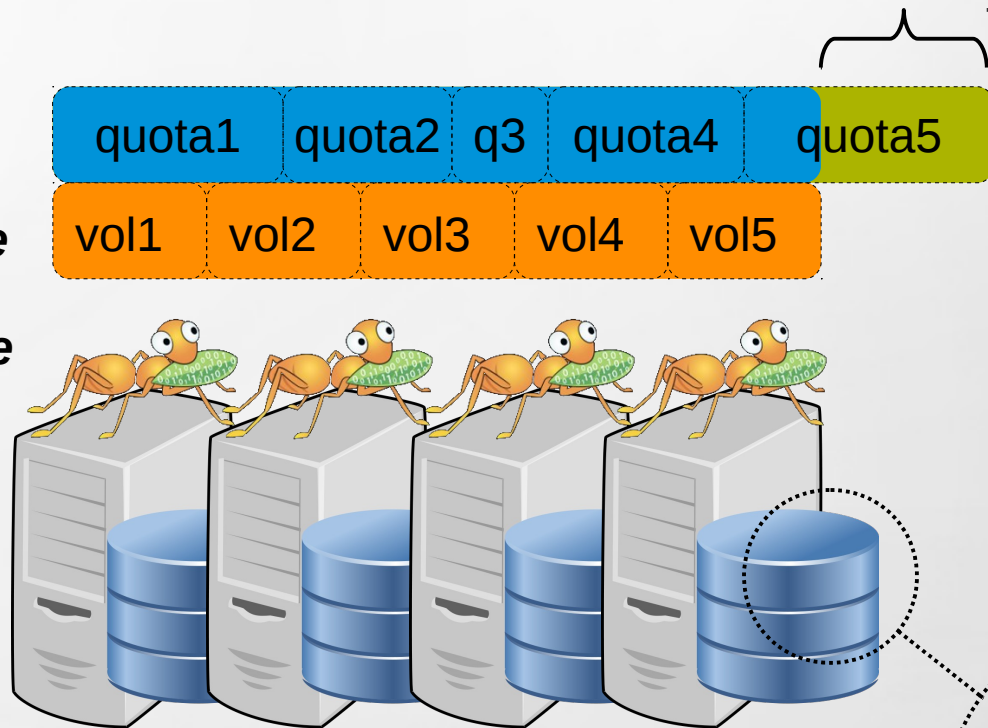
- Dell R510 nodes with local disk
- ~30TB per node as one XFS filesystem
- Bricks are subdirectories of the parent filesystem
  - Volumes are therefore naturally over-provisioned
- Quotas\* placed on volumes to limit usage and provide for accounting and charge-back
- Only 4 gluster commands needed to allocate and limit a new volume; Easily automated

*\*New feature in RHS 2.1; Coming in GlusterFS 3.5*

**Quotas are set on a volume level and are independent of each other and of the underlying available space**

**Over-Provisioned**

**Volumes created from the bricks share a total amount of available space**



**One backend filesystem is sub-divided into directories that are used as bricks**

**XFS Filesystem**

```
/xfs/brick1
/xfs/brick2
/xfs/brick3
/xfs/brick4
/xfs/brick5
```

```
# gluster volume create
# gluster volume start
# gluster volume quota enable
# gluster volume quota limit-usage
```

# Use Case:

NoSQL Backend with SLA-Bound  
Geo-Replication

**Open Storage in the Enterprise**

With GlusterFS and Red Hat Storage



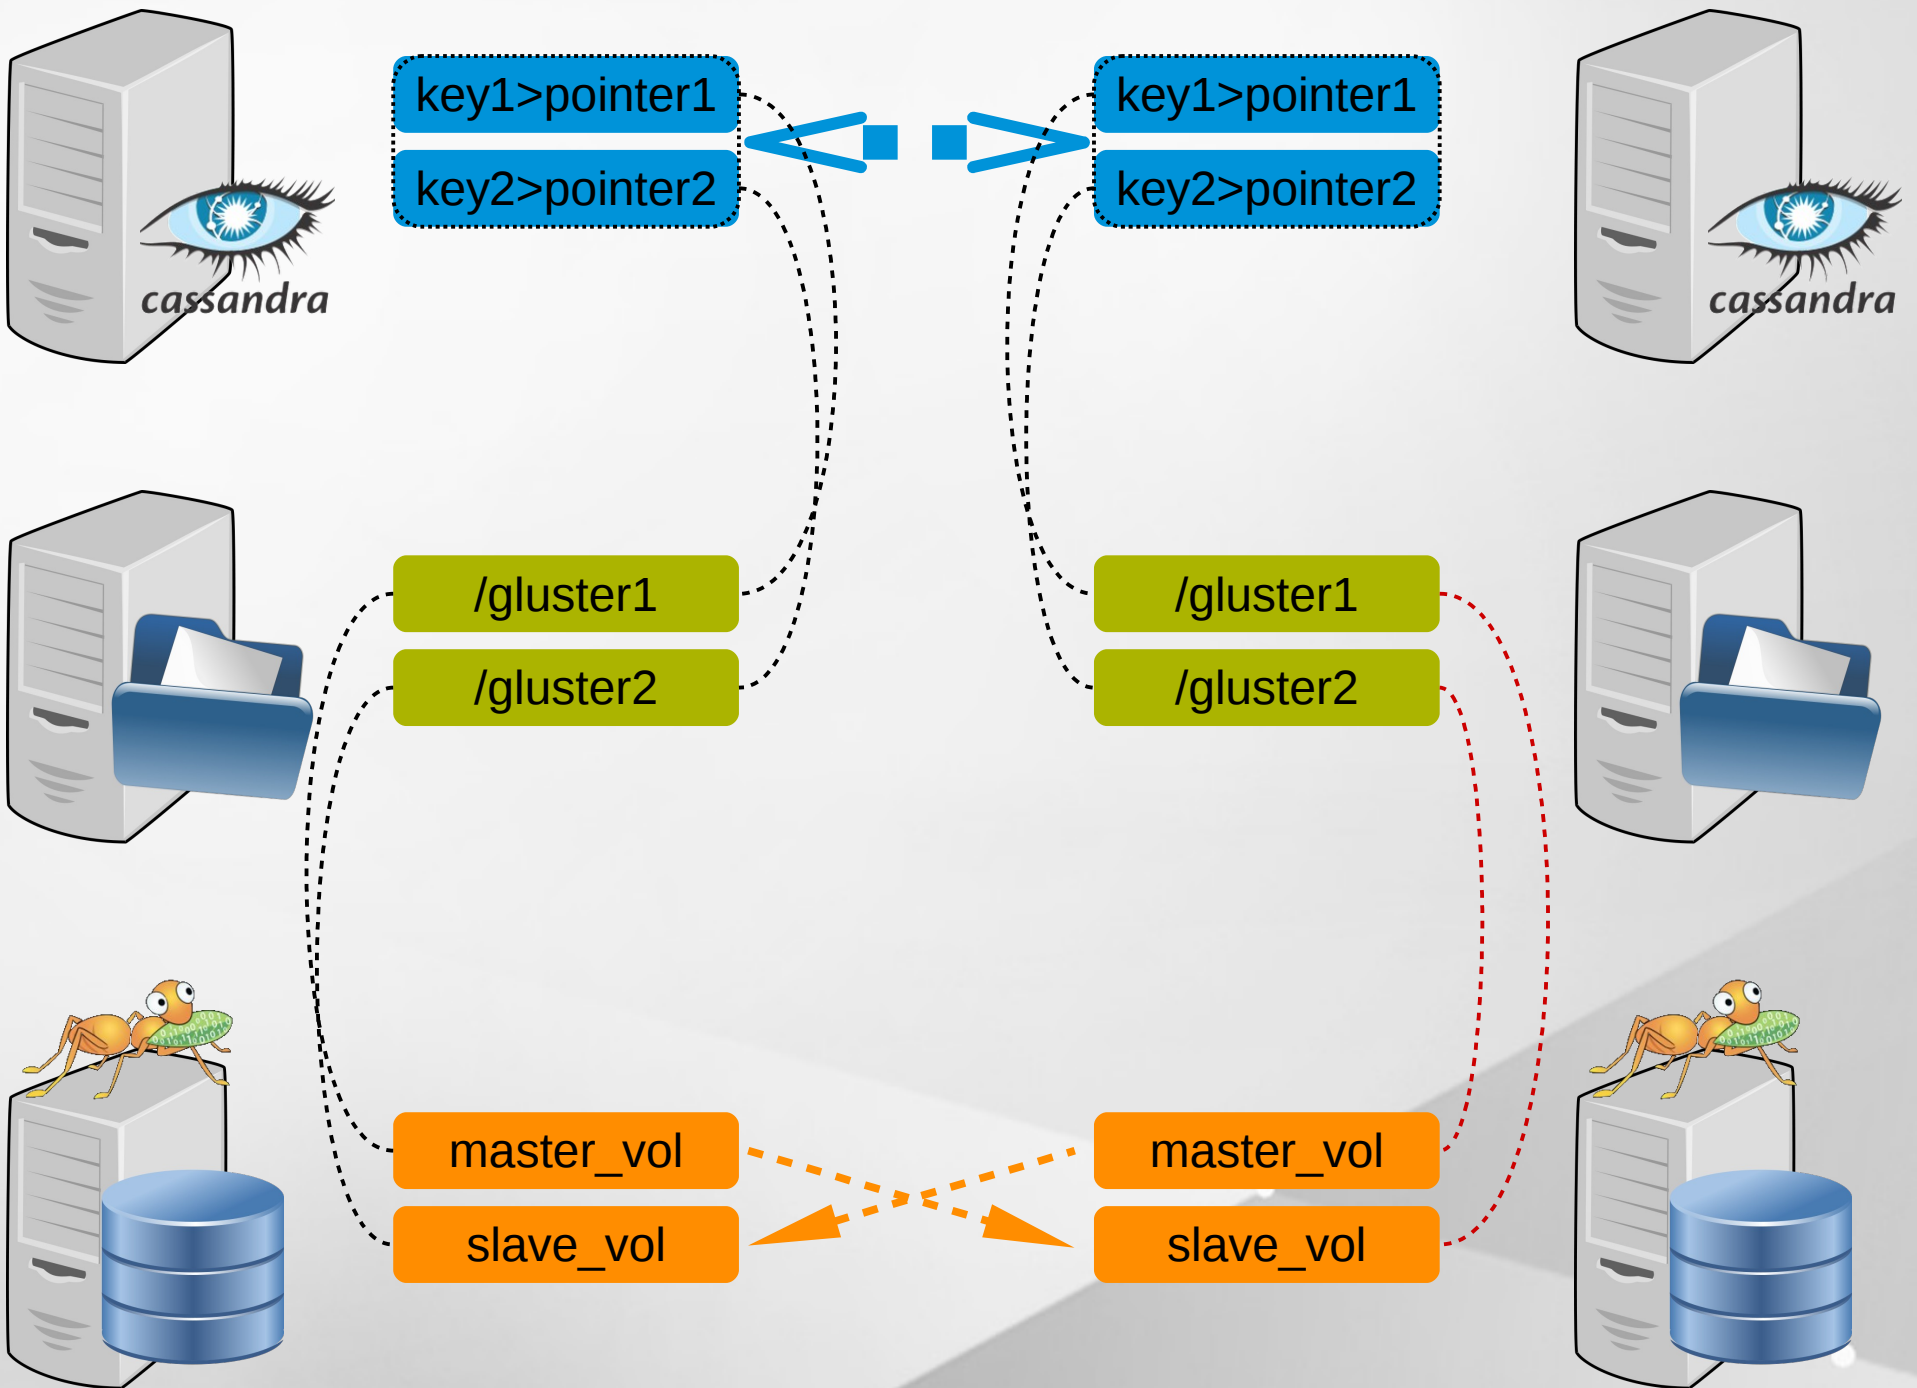
# Goals

- Replace legacy database key/blob architecture
- Divide and conquer
  - NoSQL layer for key/pointer
  - Scalable storage layer for blob payload
- Active/Active sites with 30-minute replication SLA
- Performance tuned for small-file WORM patterns

# Implementation

- HP DL170e nodes with local disk
- ~4TB per node
- Cassandra replicated NoSQL layer for key/pointer
- GlusterFS parallel geo-replication\* for data payload site copy exceeding SLA standards
- Worked with Red Hat Engineering to modify application data patterns for better small-file performance

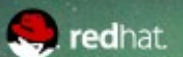
*\*New feature in RHS 2.1; Coming in GlusterFS 3.5*



# Questions?

## Open Storage in the Enterprise

With GlusterFS and Red Hat Storage



# Do it!

- Build a test environment in VMs in just minutes!
- Get the bits:
  - Fedora 19 has GlusterFS packages natively
  - RHS 2.1 ISO available on the Red Hat Portal
  - Go upstream: [www.gluster.org](http://www.gluster.org)



# Thank You!

Slides Available at: <http://people.redhat.com/dblack>

- [dustin@redhat.com](mailto:dustin@redhat.com) / [@dustinblack](#)

- [storage-sales@redhat.com](mailto:storage-sales@redhat.com)

- **RHS:**

[www.redhat.com/storage/](http://www.redhat.com/storage/)

- **GlusterFS:**

[www.gluster.org](http://www.gluster.org)

- **TAM:**

[access.redhat.com/support/offerings/tam/](http://access.redhat.com/support/offerings/tam/)



 [@Glusterorg](#)

 [@RedHatStorage](#)



 [Gluster](#)

 [Red Hat Storage](#)

## Open Storage in the Enterprise

With GlusterFS and Red Hat Storage