



# RED HAT STORAGE SERVER

## TECHNICAL OVERVIEW

Ingo Börnig

Solution Architect,

Red Hat

24.10.2013

# NEW STORAGE REQUIREMENTS FOR THE MODERN HYBRID DATACENTER

**DESIGNED FOR THE NEW DATA LANDSCAPE** – PETABYTE SCALE

**LINEAR SCALABILITY** – PERFORMANCE AND CAPACITY

**HIGHLY AVAILABLE** – WITHOUT BREAKING THE BANK

**ACCESSIBLE** – FROM ANY APPLICATION, THROUGH ANY DEVICE

**OPEN AND INTEROPERABLE** – STANDARDS AND OPEN TECHNOLOGIES

**SELF-HEALING, SELF-MANAGING** – REDUCING OPERATIONAL OVERHEAD

**HYBRID DATACENTER FOUNDATION** - PRIVATE, PUBLIC, AND HYBRID CLOUDS

**EXTENSIBLE** – INNOVATE TO MEET YOUR UNIQUE BUSINESS REQUIREMENTS

**DESIGNED FOR TODAY'S IT ECONOMICS** – DO MORE WITH WHAT YOU HAVE

# RED HAT STORAGE DESIGN GOALS

## Scale out

- Elimination of metadata servers
- Effective distribution of data to achieve scalability and flexibility

## Linear Scaling

- Capacity – scale up vertically
- Performance – scale out horizontally

## Elasticity

- Flexibly adapt to the growth or reduction of data in the enterprise.
- Add or remove resources to/from storage pool with zero application disruption.

## Deployment agnostic

- Deploy on-premise, in the public cloud, or a hybrid setup.

## Must run on commodity hardware

- Industry standard servers
- No-purpose built hardware


# RED HAT STORAGE ATTRIBUTES

- Based on GlusterFS technology
  - Distributed file system on top of local file system
  - Red Hat Storage uses XFS
  - No metadata server with elastic hashing algorithm
  - Spreads files across XFS directories called bricks
  - Stores metadata in extended attributes of XFS
  - Modular stackable architecture based on user-space FUSE\*
- \*Filesystem in user space


# INCREASE DATA, APPLICATION AND INFRASTRUCTURE AGILITY



DATA SERVICES



ENTERPRISE APPLICATIONS



BIG DATA WORKLOADS



CLOUD APPLICATIONS



ENTERPRISE MOBILITY

FILE SERVICES

OPEN OBJECT APIs

CONVERGED COMPUTE AND STORAGE

OPEN, SOFTWARE-DEFINED STORAGE PLATFORM

**RED HAT<sup>®</sup>  
STORAGE**

PHYSICAL



Standard x86 systems  
Scale-out NAS solutions

VIRTUAL



Include idle or  
legacy resources

CLOUD



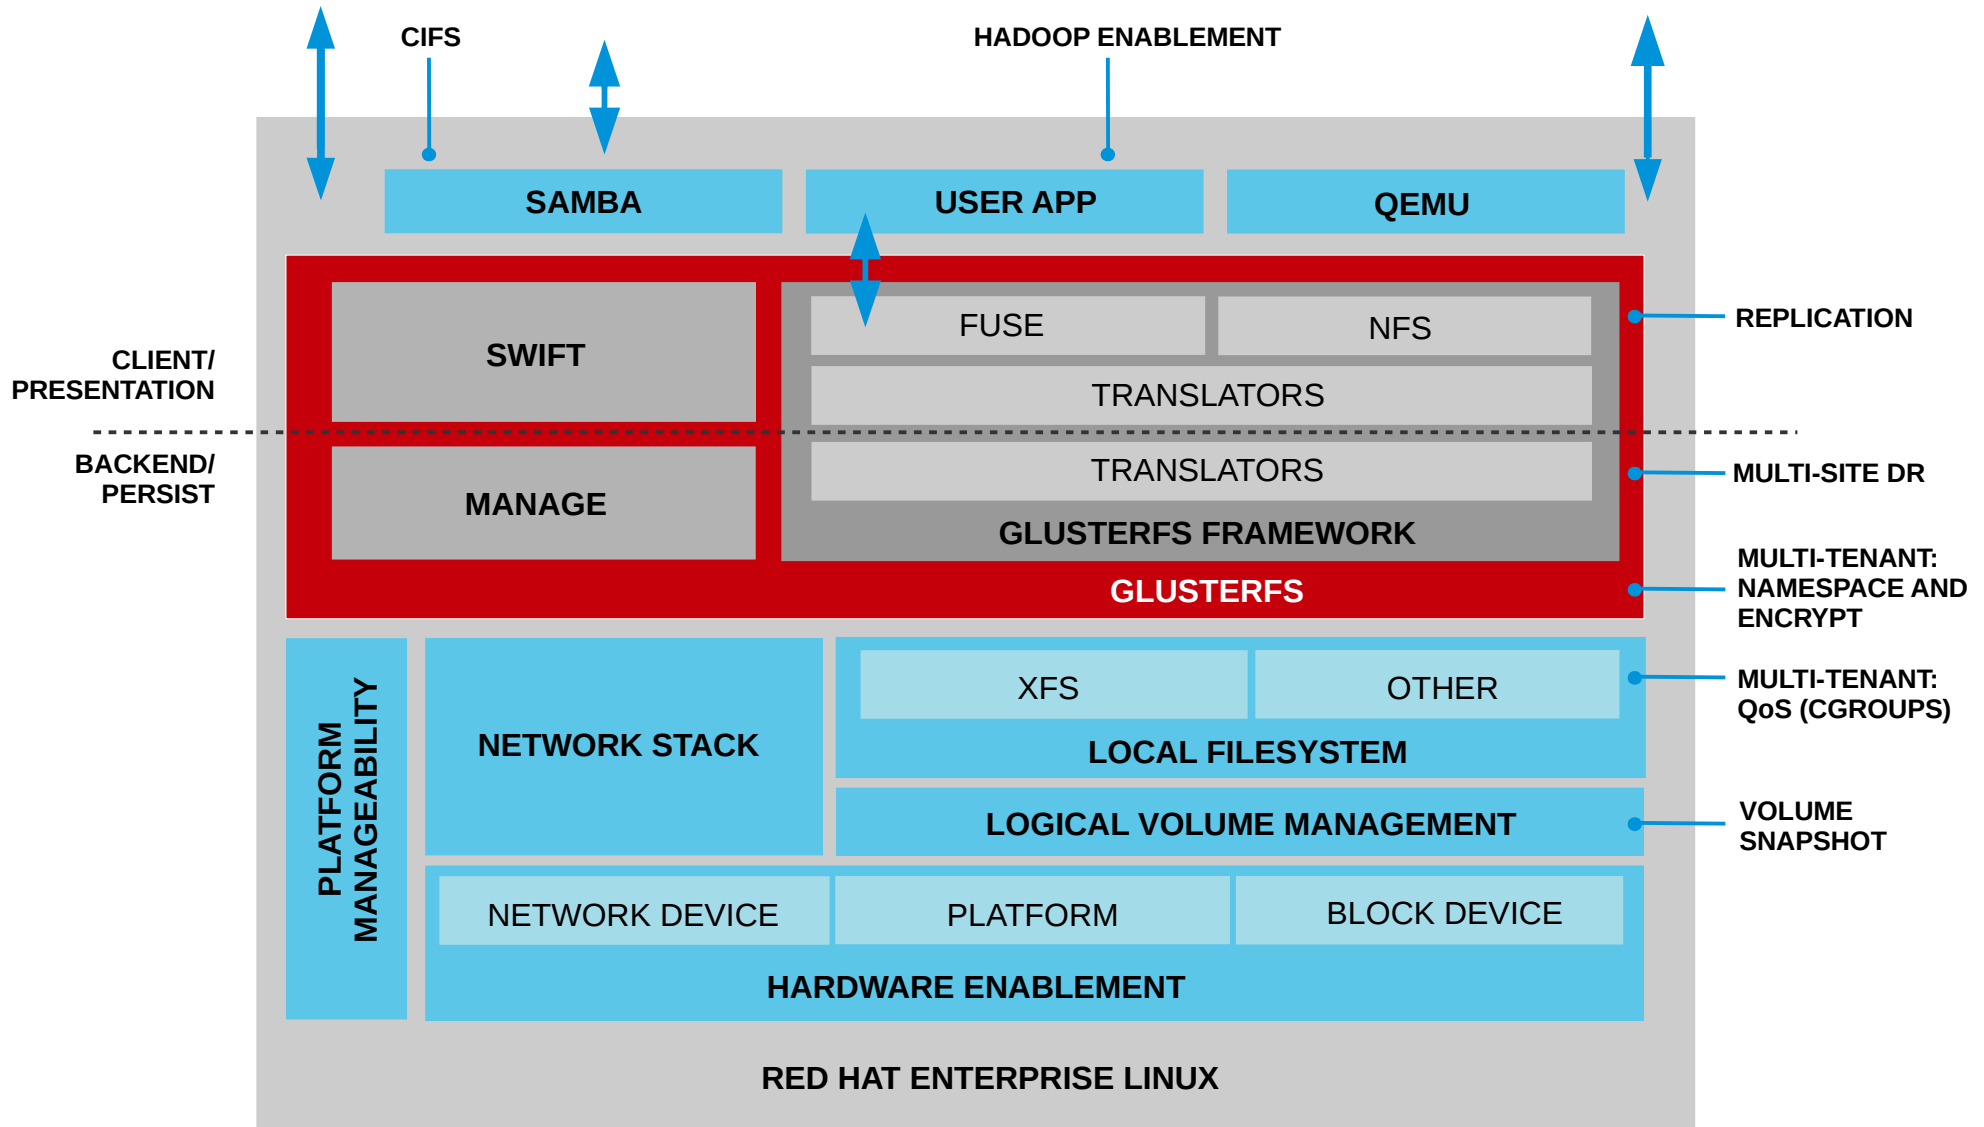
EBS EBS



SCALE-OUT STORAGE  
ARCHITECTURE

PERSISTANT DATA STORES

# Red Hat Storage technology stack



# Red Hat Storage concepts



## VOLUME

A namespace presented as a POSIX mount point and is comprised of bricks.



## BRICK

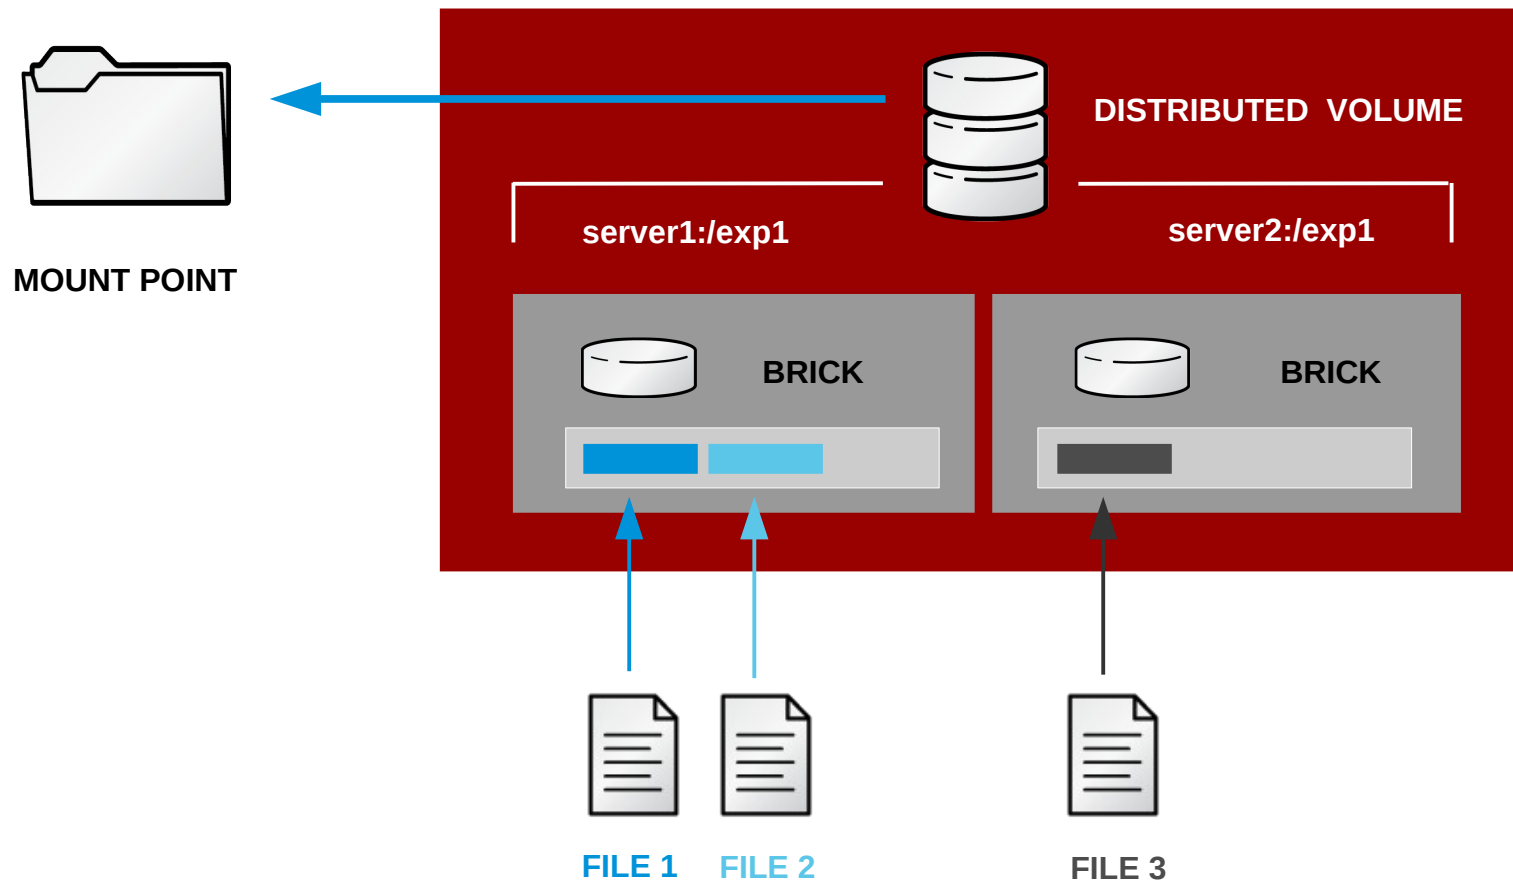
The basic unit of storage, represented by an export directory on a server



## SERVER/NODES

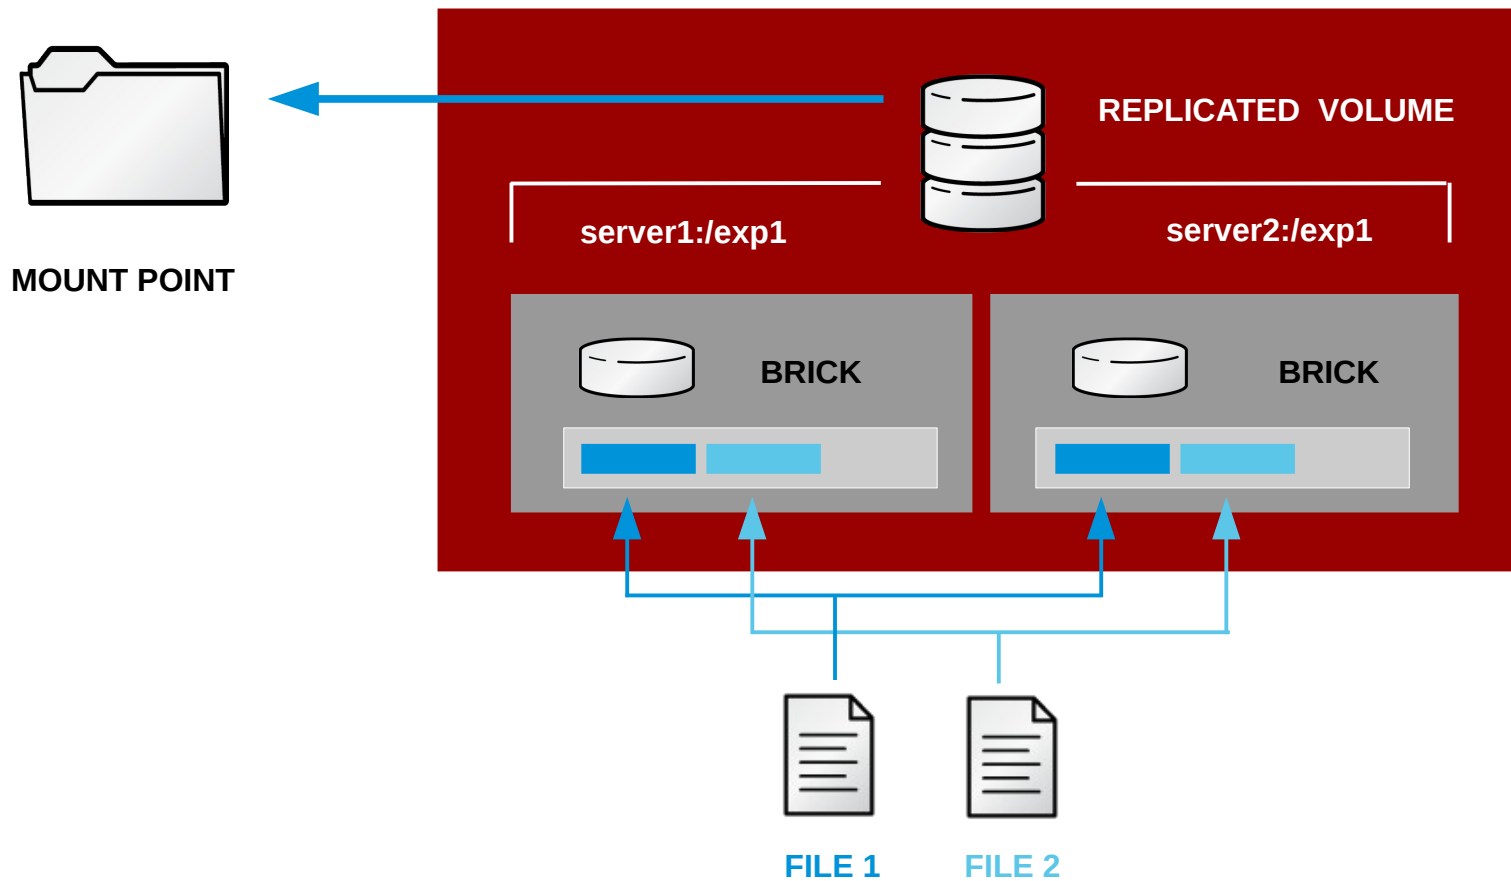
Contain the bricks

# Red Hat Storage user perspective (distributed volumes)






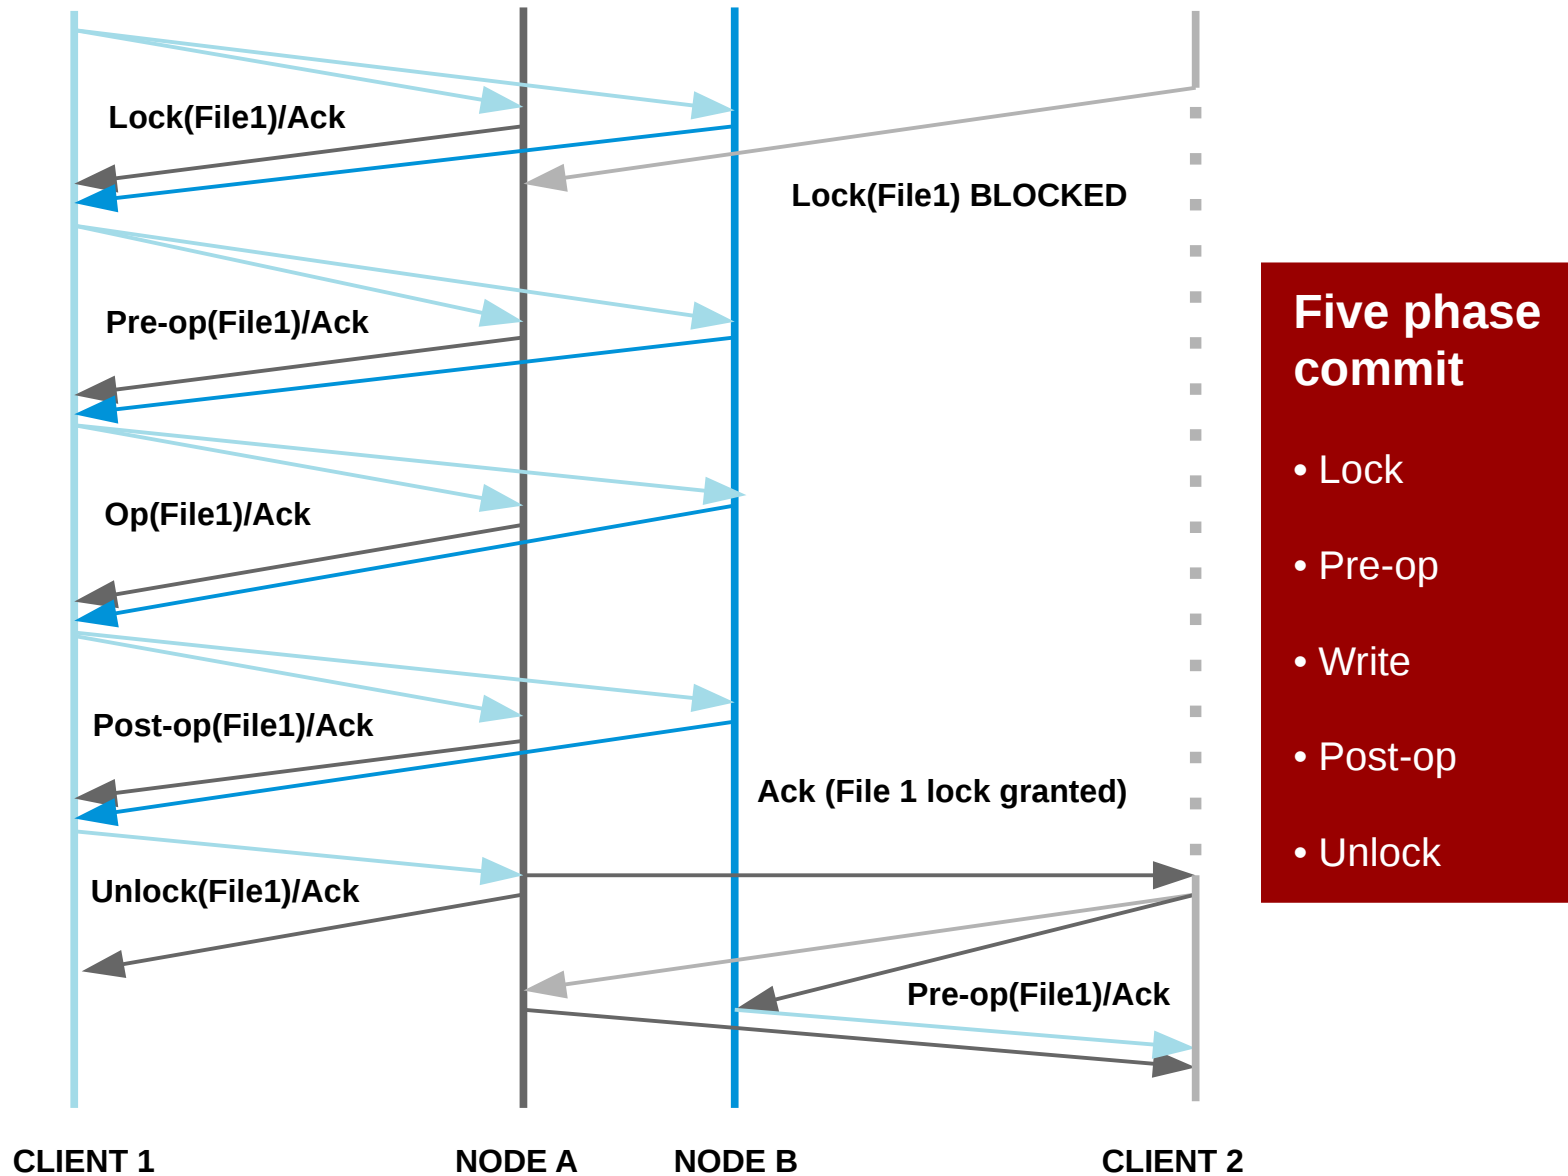
# Red Hat Storage user perspective (replicated volumes)



# Red Hat Storage user perspective (distributed replicated volumes)



# How does replication actually work?



# SIMPLIFIED AND UNIFIED STORAGE MANAGEMENT

## SINGLE PANE OF GLASS FOR CONVERGED STORAGE AND COMPUTE

### STORAGE OPERATIONS

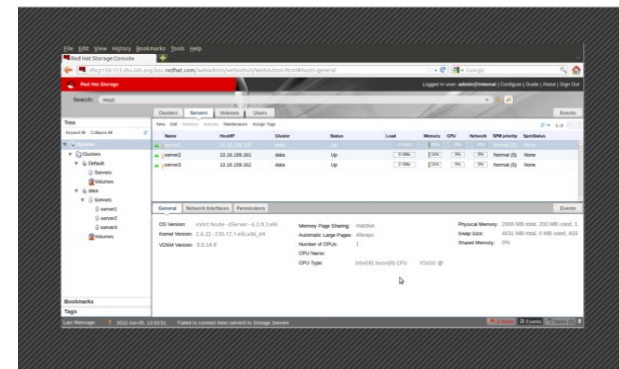
INTUITIVE USER INTERFACE

VOLUME MANAGEMENT

ON-PREMISE AND PUBLIC CLOUD

### VIRTUALIZATION AND STORAGE

SHARED MANAGEMENT WITH RHEV-M



### PROVISIONING

INSTALLATION AND CONFIGURATION

UPDATE MANAGEMENT

LIFECYCLE MANAGEMENT

FAMILIAR RHEL TOOLS

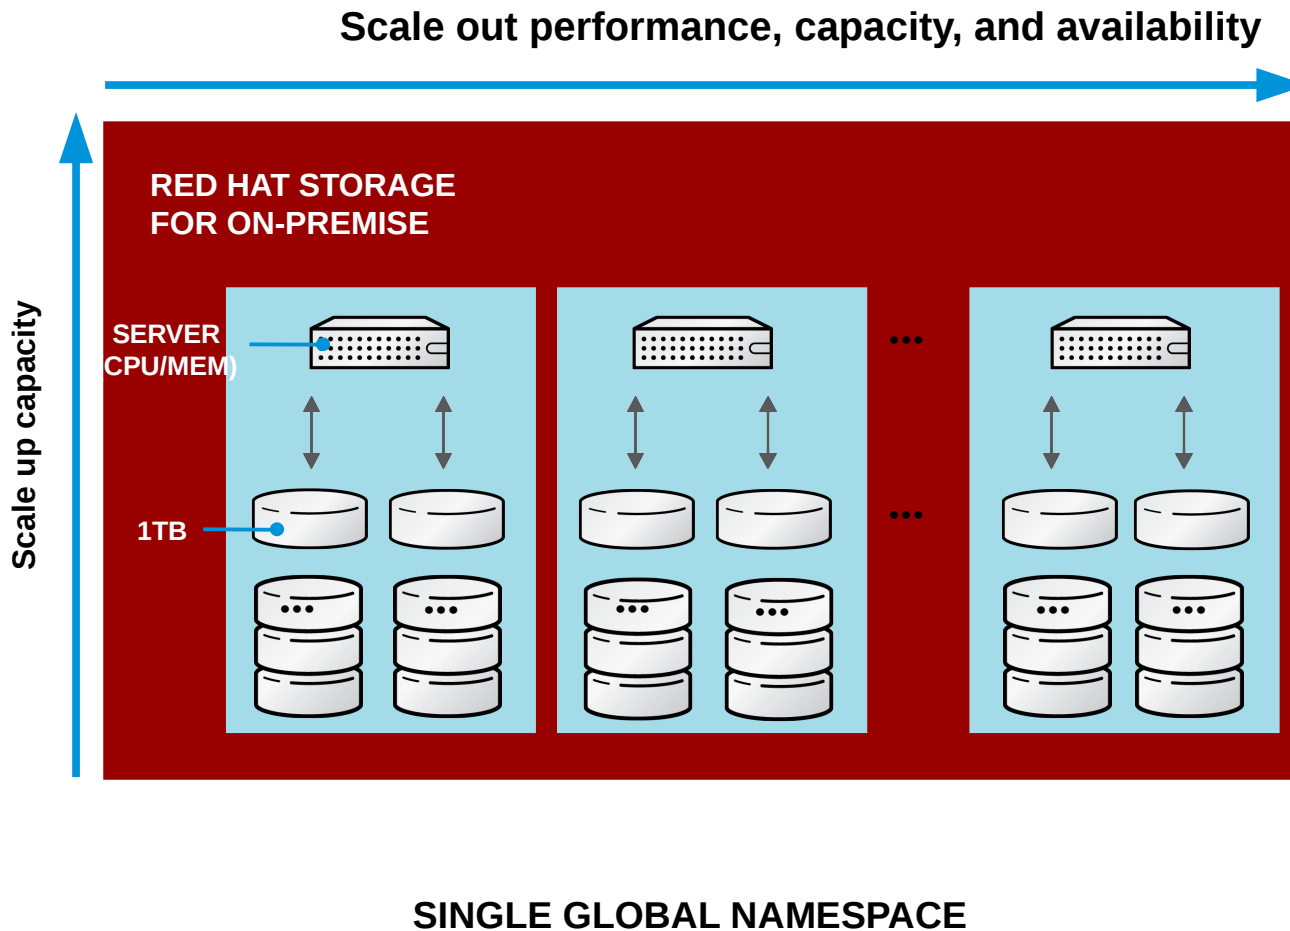
### MANAGEMENT TOOLS & FRAMEWORK

**RHS  
CONSOLE  
STORAGE**

**RHEV  
MANAGER  
VIRTUALIZATION**

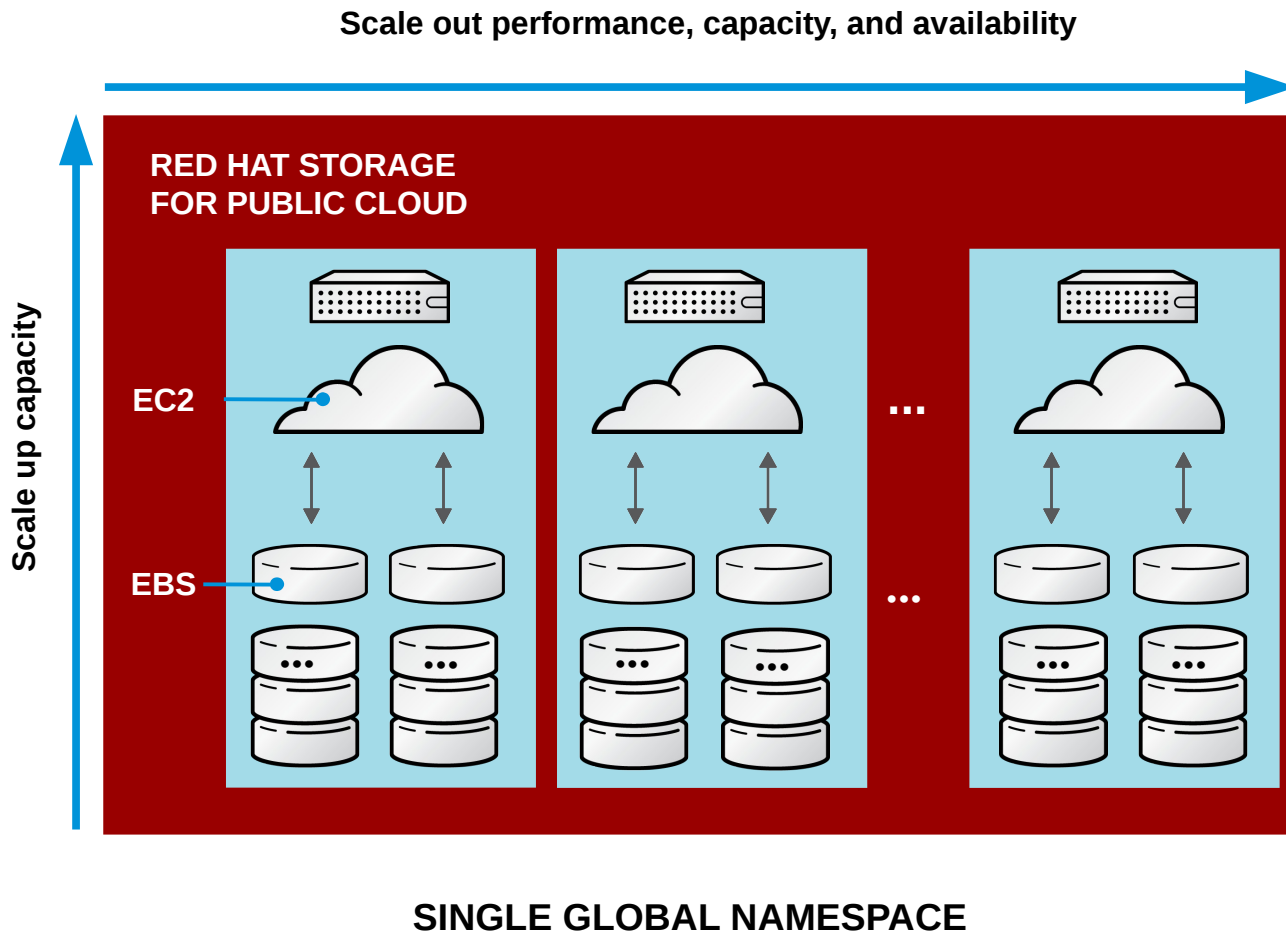
**RED HAT NETWORK  
SATELLITE  
PROVISIONING & LIFECYCLE MGMT**

# Red Hat Storage Deployment on Physical Servers



- Global namespace
- Aggregates CPU, memory, network capacity.
- Deploys on Red Hat-supported servers and underlying storage: DAS, JBOD.
- Scale out linearly.
- Scale out performance and capacity as needed.
- Replicate synchronously and asynchronously.

# Red Hat Storage Deployment on Amazon Cloud



- GlusterFS Amazon Machine Images (AMIs)
- The only way to achieve high availability of Elastic Block Storage (EBS)
- Multiple EBS devices pooled
- POSIX compatible (no application to rewrite required to run on Amazon EC2)
- Scale out capacity and performance as needed

# What is productized in Red Hat Storage?

## RED HAT® STORAGE

- Red Hat Enterprise Linux
- XFS
- GlusterFS
- Red Hat Storage console management station

Physical Server: 2 socket x86 with 12-36 disks

or

Virtual Server: Amazon, AWS, Red Hat Virtualization, or VMware

**A pre-integrated, pre-verified and ready to run software platform**

**Sourced by customer**

## Generic requirements for a Red Hat Storage server

- Must be 2-socket (4-core or 6-core) servers from HP, Dell, Fujitsu, IBM, Cisco, NEC and Hitachi ONLY (no 1-socket, 4-socket servers, or 8-socket servers, for example)
- Intel Xeon Nehalem-EX and beyond (CPU)
- Equivalent AMD processors also approved
- Minimum RAM requirements are use case specific, ranging from 16 - 48 GB
- Reliable backplane RAID controller shipped by server vendors or from OEM manufacturers
- RAID 6 support in hardware RAID controller
- RAID controller card must be flash-backed or battery-backed
- 2 X 10GigE (copper or optical) preferred. 2 X 1GigE also supported
- Drive LED call-out (optional)
- Interrogate and control write-back state of drives (optional)


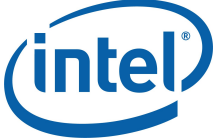






# Red Hat Storage compatible hardware vendors




# RED HAT STORAGE PARTNER SOLUTIONS

## JOINTLY CREATED BUSINESS SOLUTIONS

BACK-UP AND ARCHIVE	 commvault® <i>solving forward</i>	SIGNIFICANTLY REDUCE BACK-UP WINDOWS VS. LEGACY STORAGE SOLUTIONS.
CONTENT CLOUD		SPEED ABILITY TO RAPIDLY IMPLEMENT CONTENT CLOUDS.
BIG DATA STORAGE		PURPOSE-BUILT STORAGE SOLUTION FOR LARGE DATA WORKLOADS
INDUSTRY BLUEPRINTS		DELIVER INDUSTRY-SPECIFIC BLUEPRINTS TO MINIMIZE THE TIME VALUE FOR DATA SOLUTIONS
ENTERPRISE DROP BOX		STORAGE RESIDENT ENTERPRISE DROP BOX THAT SCALES WITH YOUR DATA
ENHANCED SECURITY CONTROLS		ACTIVE DIRECTORY AUTHENTICATION AND SINGLE SIGN-ON

# Content Cloud

- Used for applications that aggregate large quantities of data that build up into large files
- Massive simultaneous consumption of multimedia content by thousands of users



## Segment:

- Verticals that require data analysis such as oil and gas or patient records management
- Content providers, CDNs

**Workload:** Files written once and read extensively

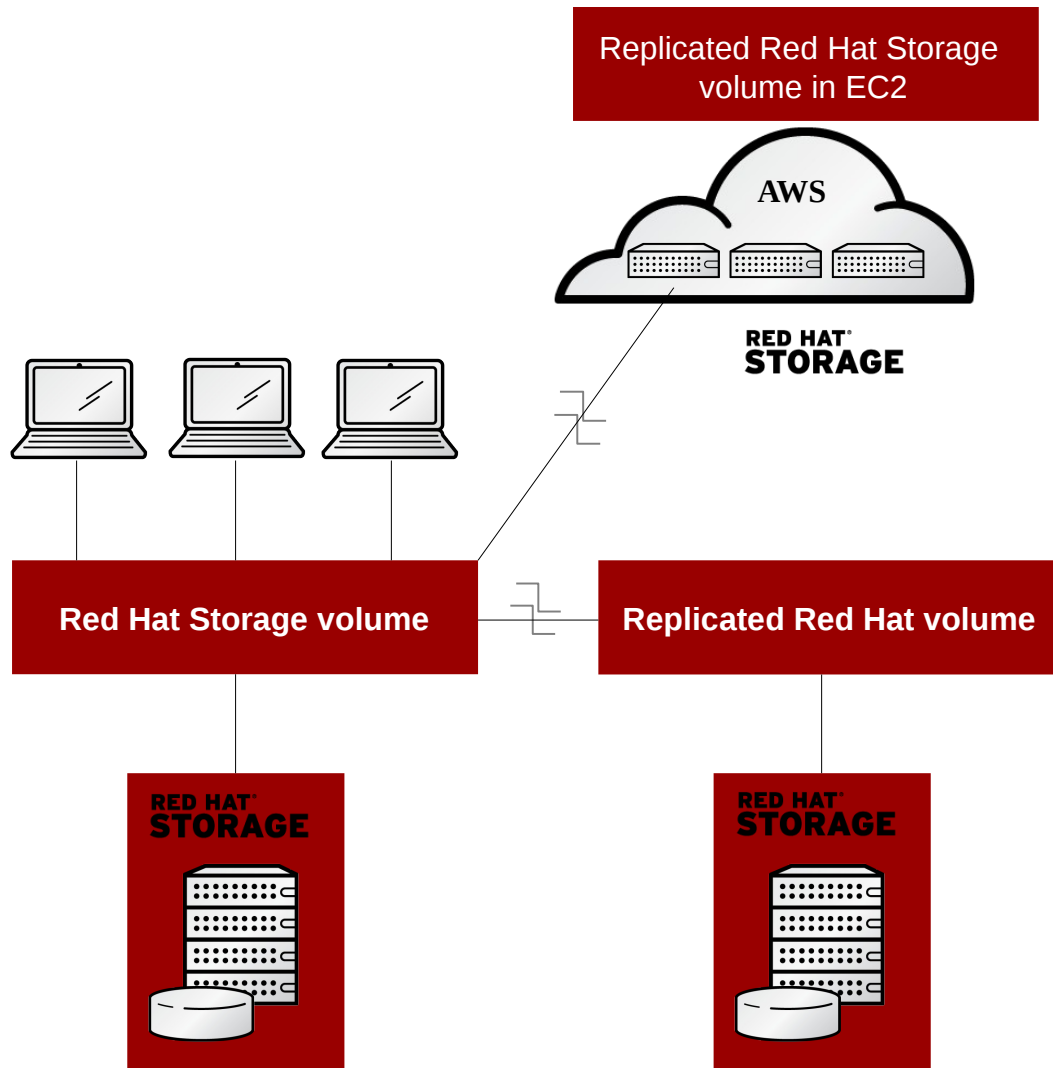
- Data needs to be shared and load can be distributed across servers
- Often scale at petabyte size

## Why Red Hat Storage:

- TCO effective vs. potential NAS growth
- Massive scalability at cost

# Data Protection

- Medium to large enterprises with DR requirements and multiple datacenters to remain synchronized
- Nearline archival, often used as tape replacement for faster and cost effective access



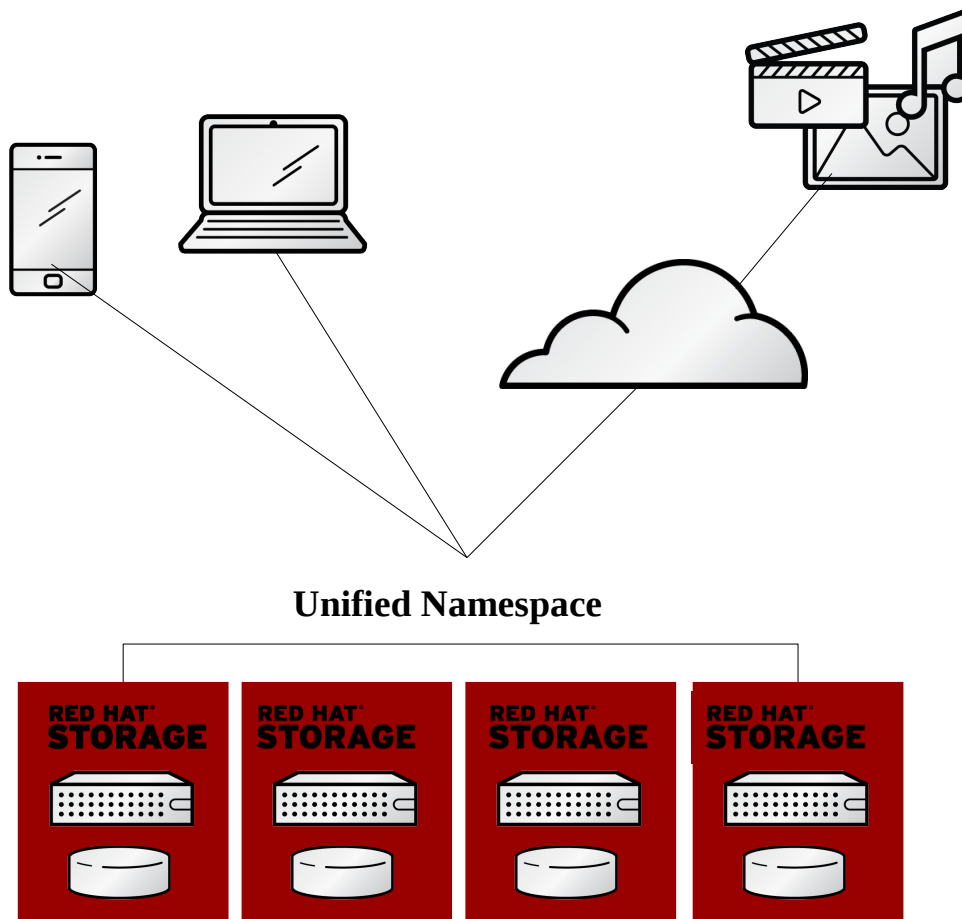
**Segment:** Medium / large enterprises – compliance / business-model driven

**Workload:** Files written once and replicated to cloud storage or on-premise

**Why Red Hat Storage:**  
Cost effective solution when compared to competition

# Object Storage

- Service Providers offer storage to end users to store content accessible from a variety of devices
- Enterprise drop-box and cloud storage for service providers



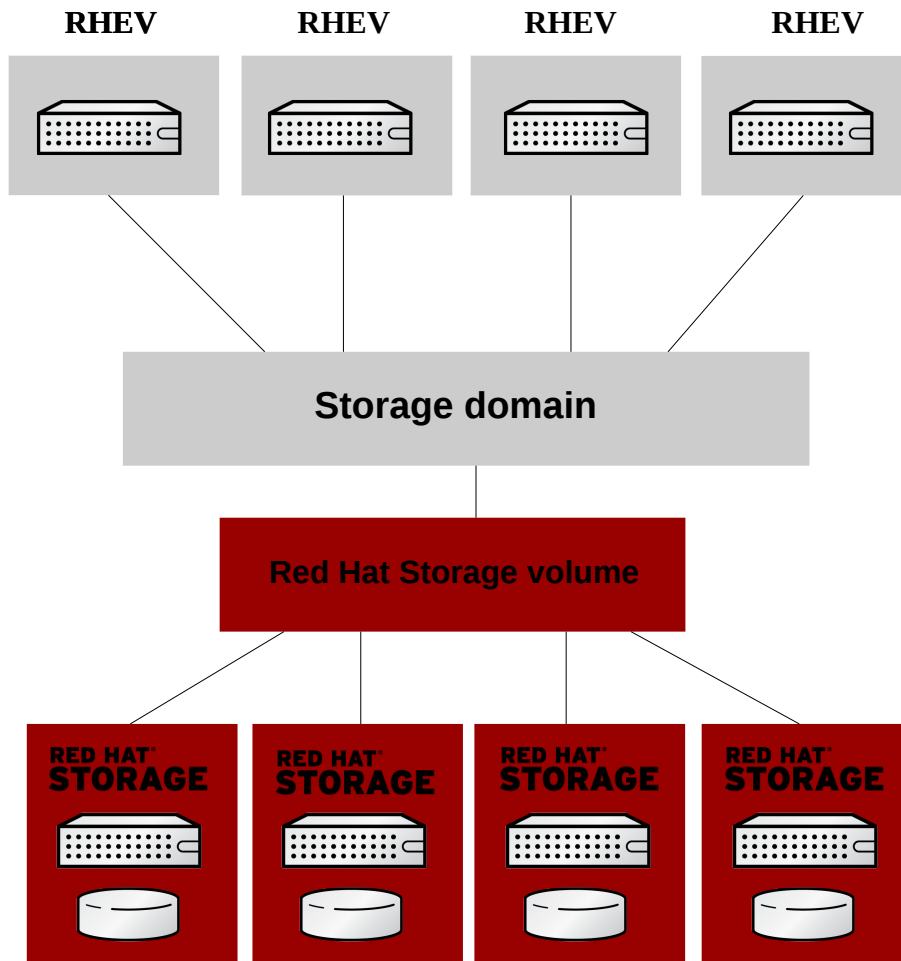
**Segment:** SPs and verticals looking for new business models

**Workload:** Balance reads and writes wide range of file sizes, high throughput

**Why Red Hat Storage:** Simultaneous storage and retrieval of files and objects interchangeably

# Live VM Image Store

Red Hat Storage provides a cost effective and reliable alternative for VM storage



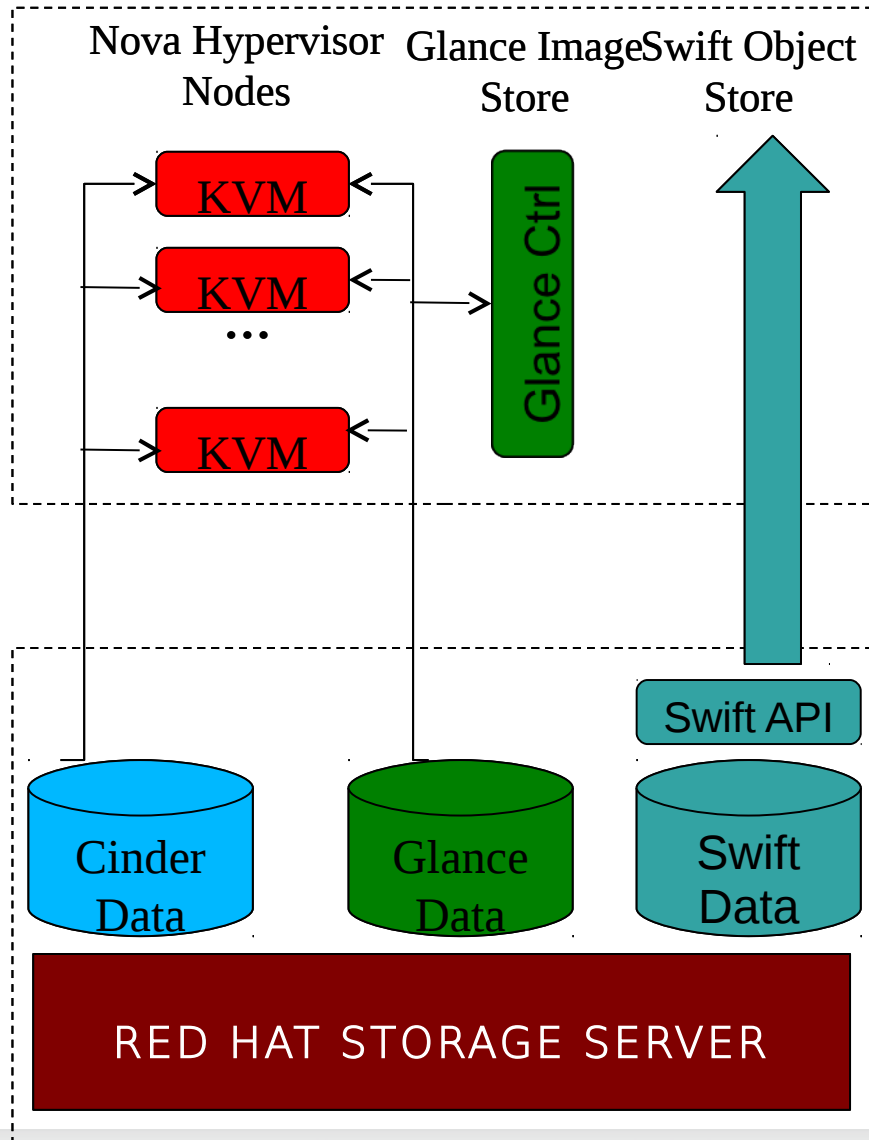
**Segment:** Wherever virtualization is used

## Why Red Hat Storage:

- Cost effective solution when compared to competition.
- Integrated solution from one vendor without vendor lock-in.
- Virtualization without a SAN on Linux.

# CONVERGED PRIVATE CLOUD WITH RHEL OPENSTACK PLATFORM AND RED HAT STORAGE

## RHELOSP

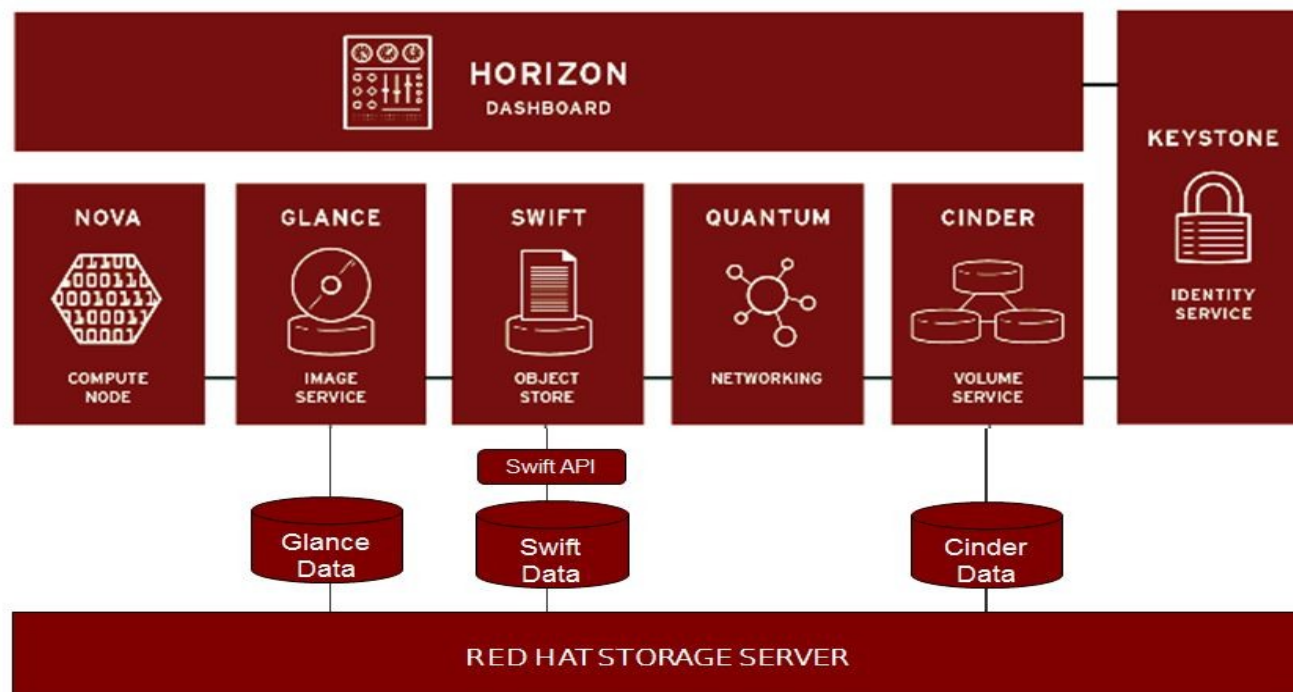


### RHS provides:

- Cinder Block Service
- Glance Image Service
- Swift Object Service
- Integration with LDAP and Active Directory (AD)

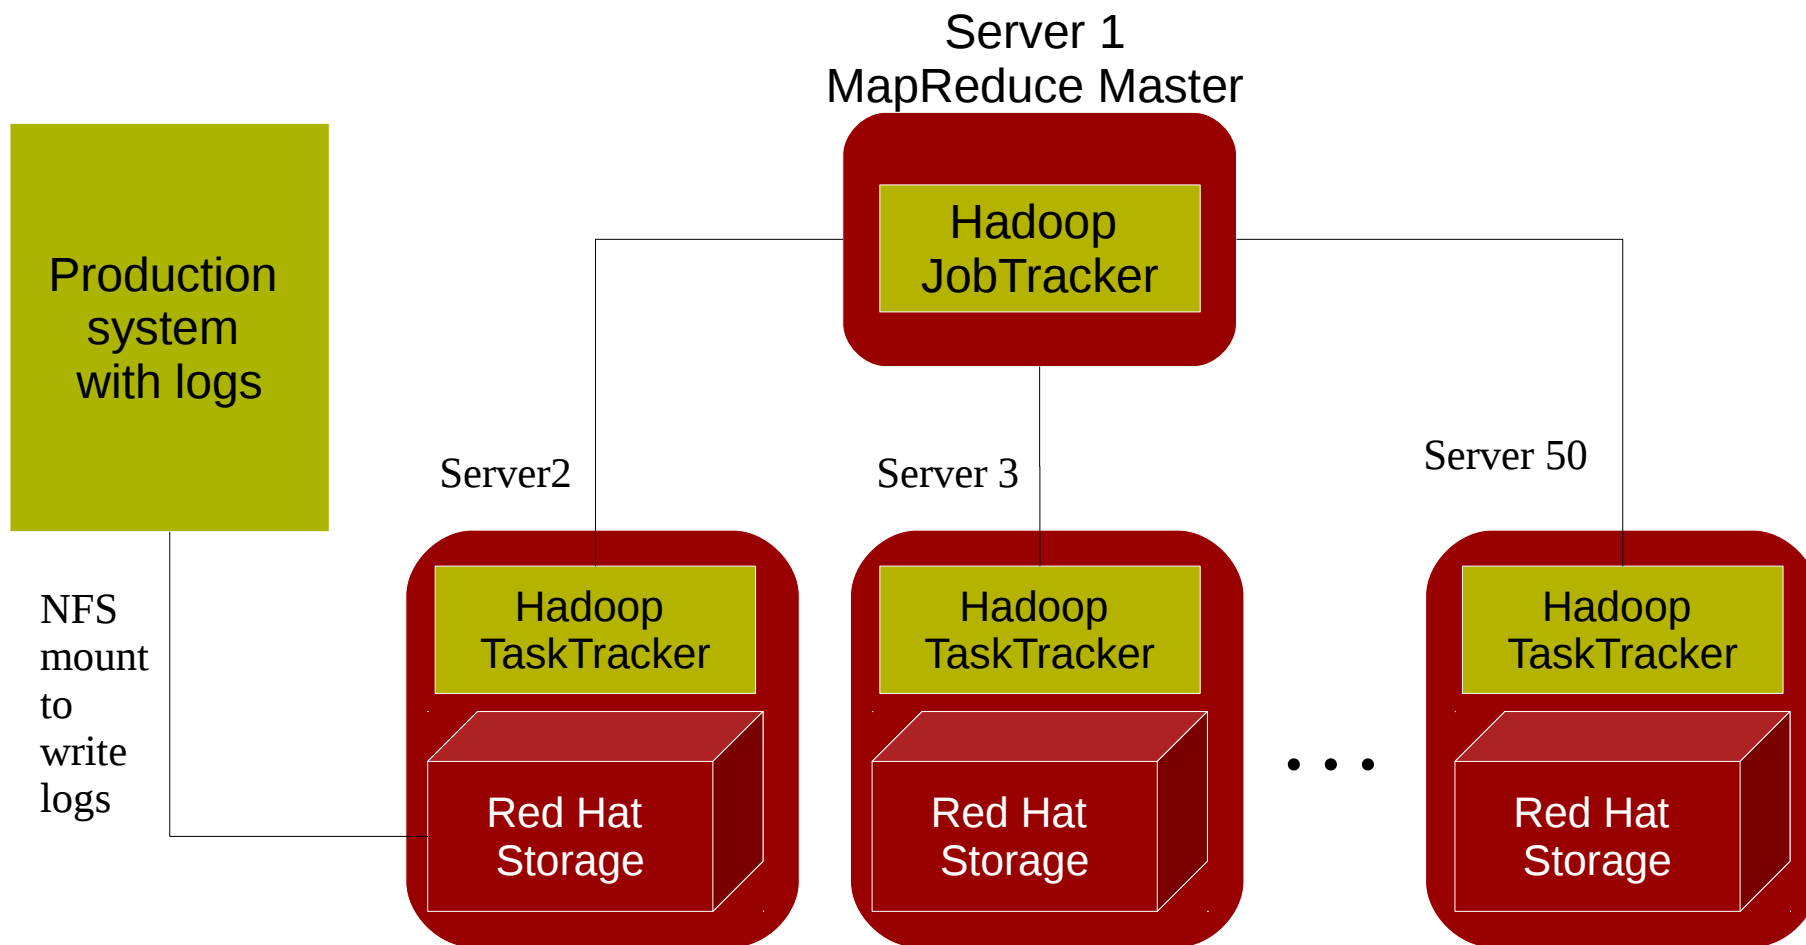
## Storage provider for OpenStack

- Unified OpenStack storage today
  - Support Cinder, Glance, and Swift with a common view
- Designed for future OpenStack innovation
  - File storage, Elastic Hadoop (Savanna), Local storage in virtual machines
- Red Hat Global Enterprise support for OpenStack and the open hybrid cloud
  - RHEL OpenStack Platform + Red Hat Storage + Red Hat Satellite





# Apache Hadoop with Red Hat Storage



- Eliminates NameNode bottlenecks and SPOF
- Disaster recovery built-in with geo-replication
- POSIX compliance for data ingest/export
- Maintains data locality as cluster scales
- Out-of-the-box HDFS API compatibility
- Management tools reduce scale-out costs

# Red Hat Storage Server roadmap summary

CY 2013

CY 2014

		CY 2013				CY 2014	
Releases	Anshi		Big Bend	Corbett	Denali		
	Q1	Q2	Q3	Q4	Q1	Q2	
Key Features	<p><b>Theme: Red Hat Enterprise Virtualization images store, robustness (GA)</b></p> <ul style="list-style-type: none"> <li>• Red Hat Enterprise Linux 6.2 based</li> <li>• Red Hat Enterprise Virtualization image store</li> <li>• Live migration of VMs</li> <li>• Root squash support</li> <li>• ~100 bug fixes</li> </ul>		<p><b>Theme: High Performance Geo-Replication (RC)</b></p> <ul style="list-style-type: none"> <li>• Red Hat Enterprise Linux 6.4 based</li> <li>• High performance Geo-replication</li> <li>• OpenStack Grizzly Support for object API</li> <li>• Complete integration and support for RHEL Openstack platform</li> <li>• SMB 2.0 support, performance and AD integration</li> <li>• Quota Support</li> <li>• Red Hat Satellite support</li> </ul>	<p><b>Theme: Console, Windows Performance (RC)</b></p> <ul style="list-style-type: none"> <li>• Red Hat Enterprise Linux 6.4 based</li> <li>• Storage Console full support</li> </ul> <p><b>Features</b></p> <ul style="list-style-type: none"> <li>• NFSv3 ACL support</li> </ul>	<p><b>Theme: Snapshots (RC)</b></p> <ul style="list-style-type: none"> <li>• Red Hat Enterprise Linux 6.5 based</li> <li>• Snapshots</li> </ul> <p><b>Features</b></p> <ul style="list-style-type: none"> <li>• Better support for small files performance</li> </ul>		

# Red Hat Storage Server Road Map

- Anshi (v2.0 U 4 CY 13 Q1)

Theme: Virtual Machine Image Store

- RHEL 6.2 & Gluster FS 3.3 based
- RHEV image store
- Enables Cinder plug-in for OpenStack
- Root Squashing

- Big Bend (v2.1 CY 13 Q3)

Theme: High Performance Geo replication

- Based on RHEL 6.4 & GlusterFS 3.4
- Parallel, distributed asynchronous geo-replication
- OpenStack Grizzly Support for object API
- Complete integration and support for RHEL Openstack platform
- SMB 2.0 support & performance enhancements for small file and write intensive workloads. AD integration.
- Quota Support
- Red Hat Satellite support

# Red Hat Storage Server Road Map

## Corbett (CY 13 Q4)

### Theme: Manageability

- RHEL 6.4 based
- Red Hat Storage Console support
- NFS ACL support and other enhancements
- RDMA over Infiniband (under consideration)

## Denali (CY 14 Q1/2)

### Theme: Volume Snapshots

- RHEL 6.5 based
- Leverage dm-thinp based snapshots
- Integration with commercial backup toolsets
- Better performance for small file workload



**THANK YOU**