redhat.

MANAGING UNSTRUCTURED DATA AT PETABYTE-SCALE

Joachim Schröder
Manager Solution Architects, DACH
Email: joachim.schroeder@redhat.com
November, 14th 2013

WHO'S RED HAT? - RED HAT PORTFOLIO

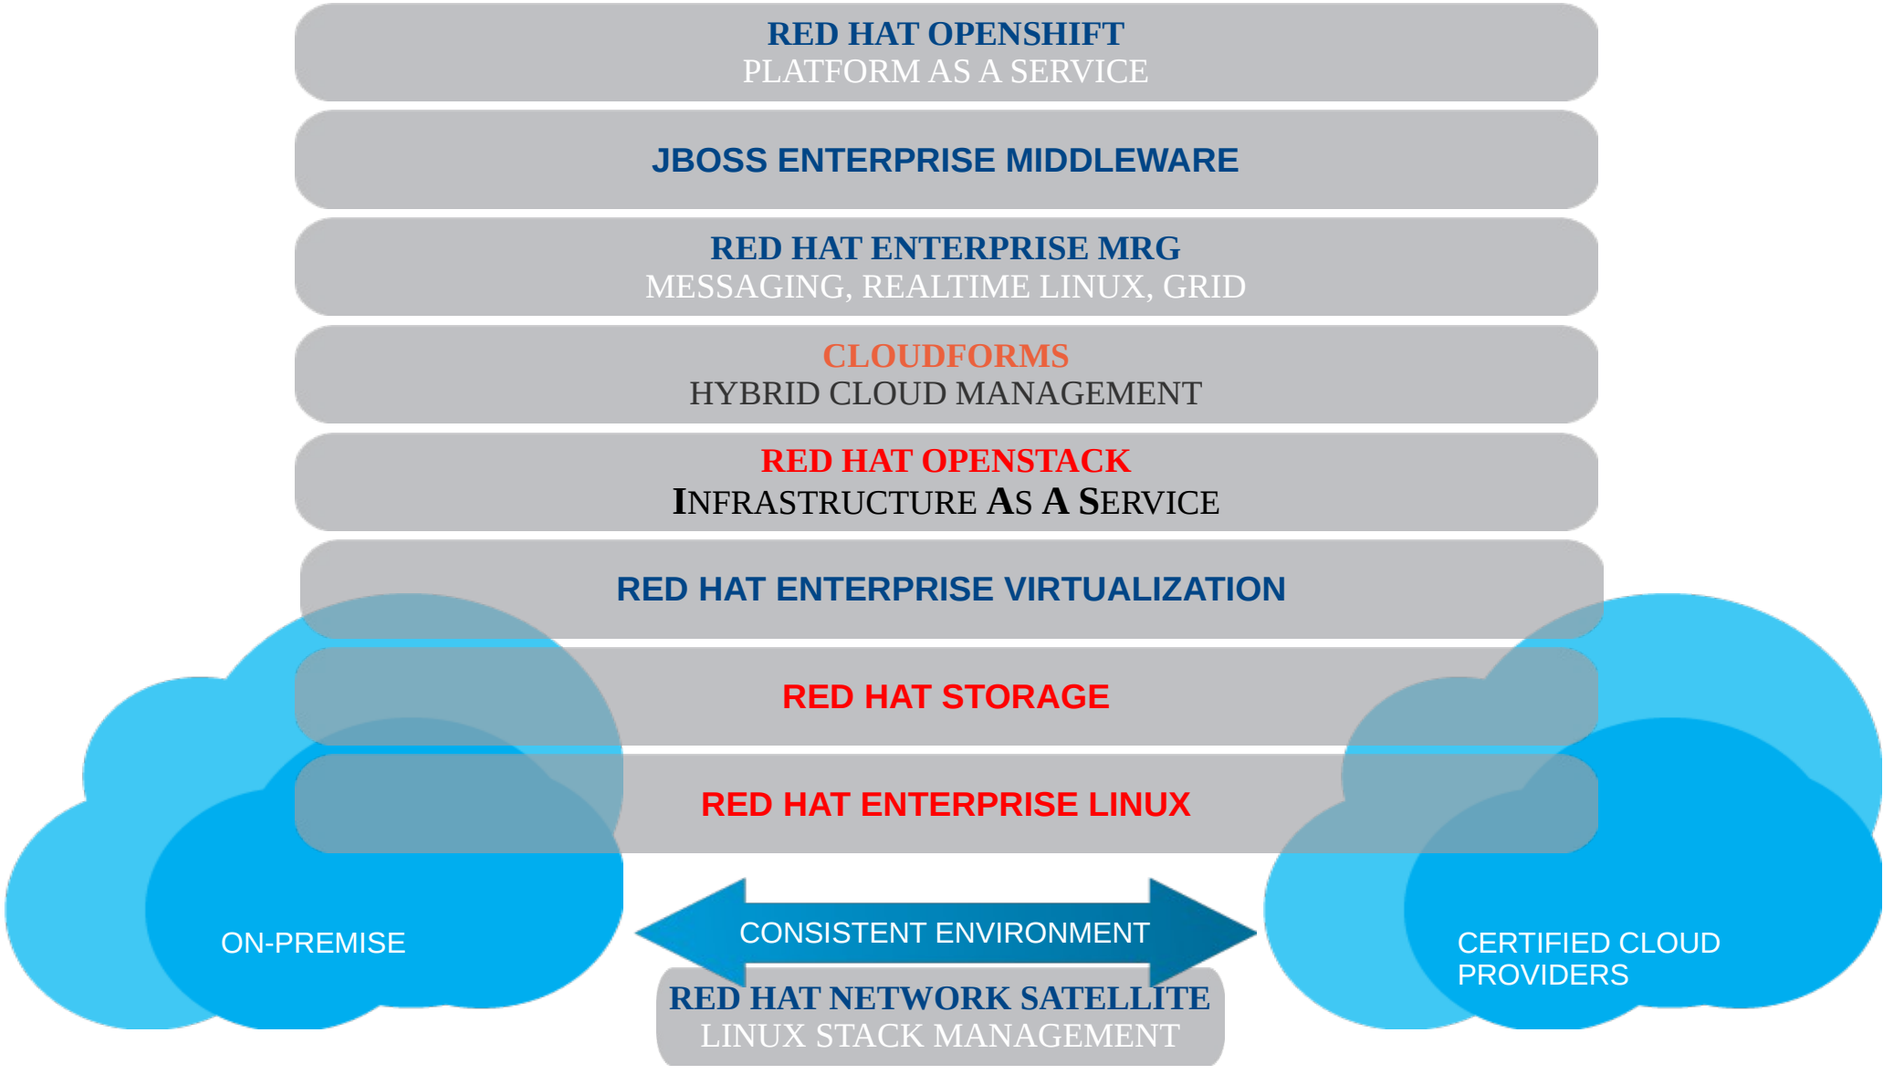

THE INFORMATION EXPLOSION

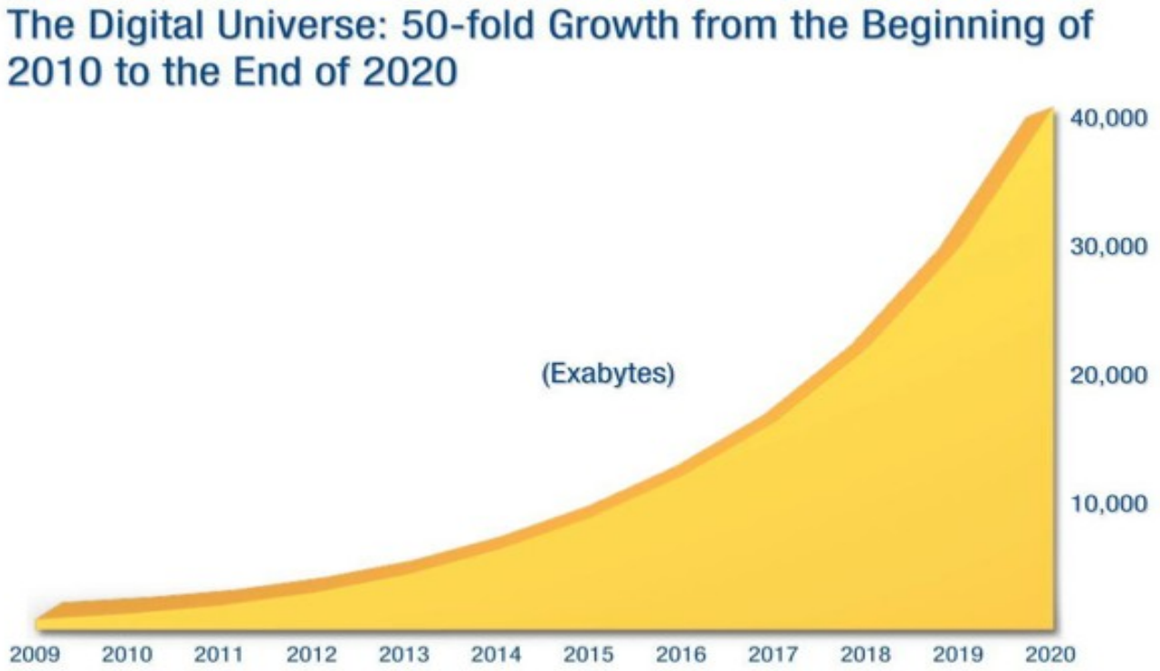

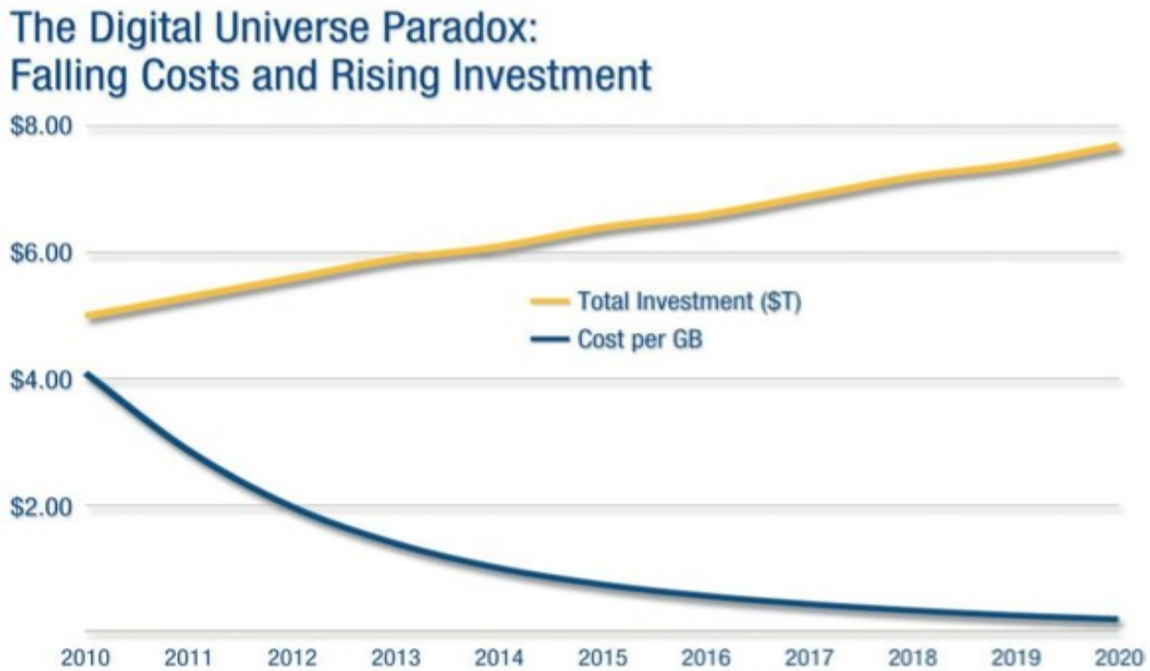

Source: IDC's Digital Universe Study, Dec 2012

Main growth drivers:
Virtualisation, Cloud, Mobile Computing and Big Data

RHS: Managing Unstructured Data

CORNERSTONE OF THE NEW SOFTWARE DEFINED DATACENTER

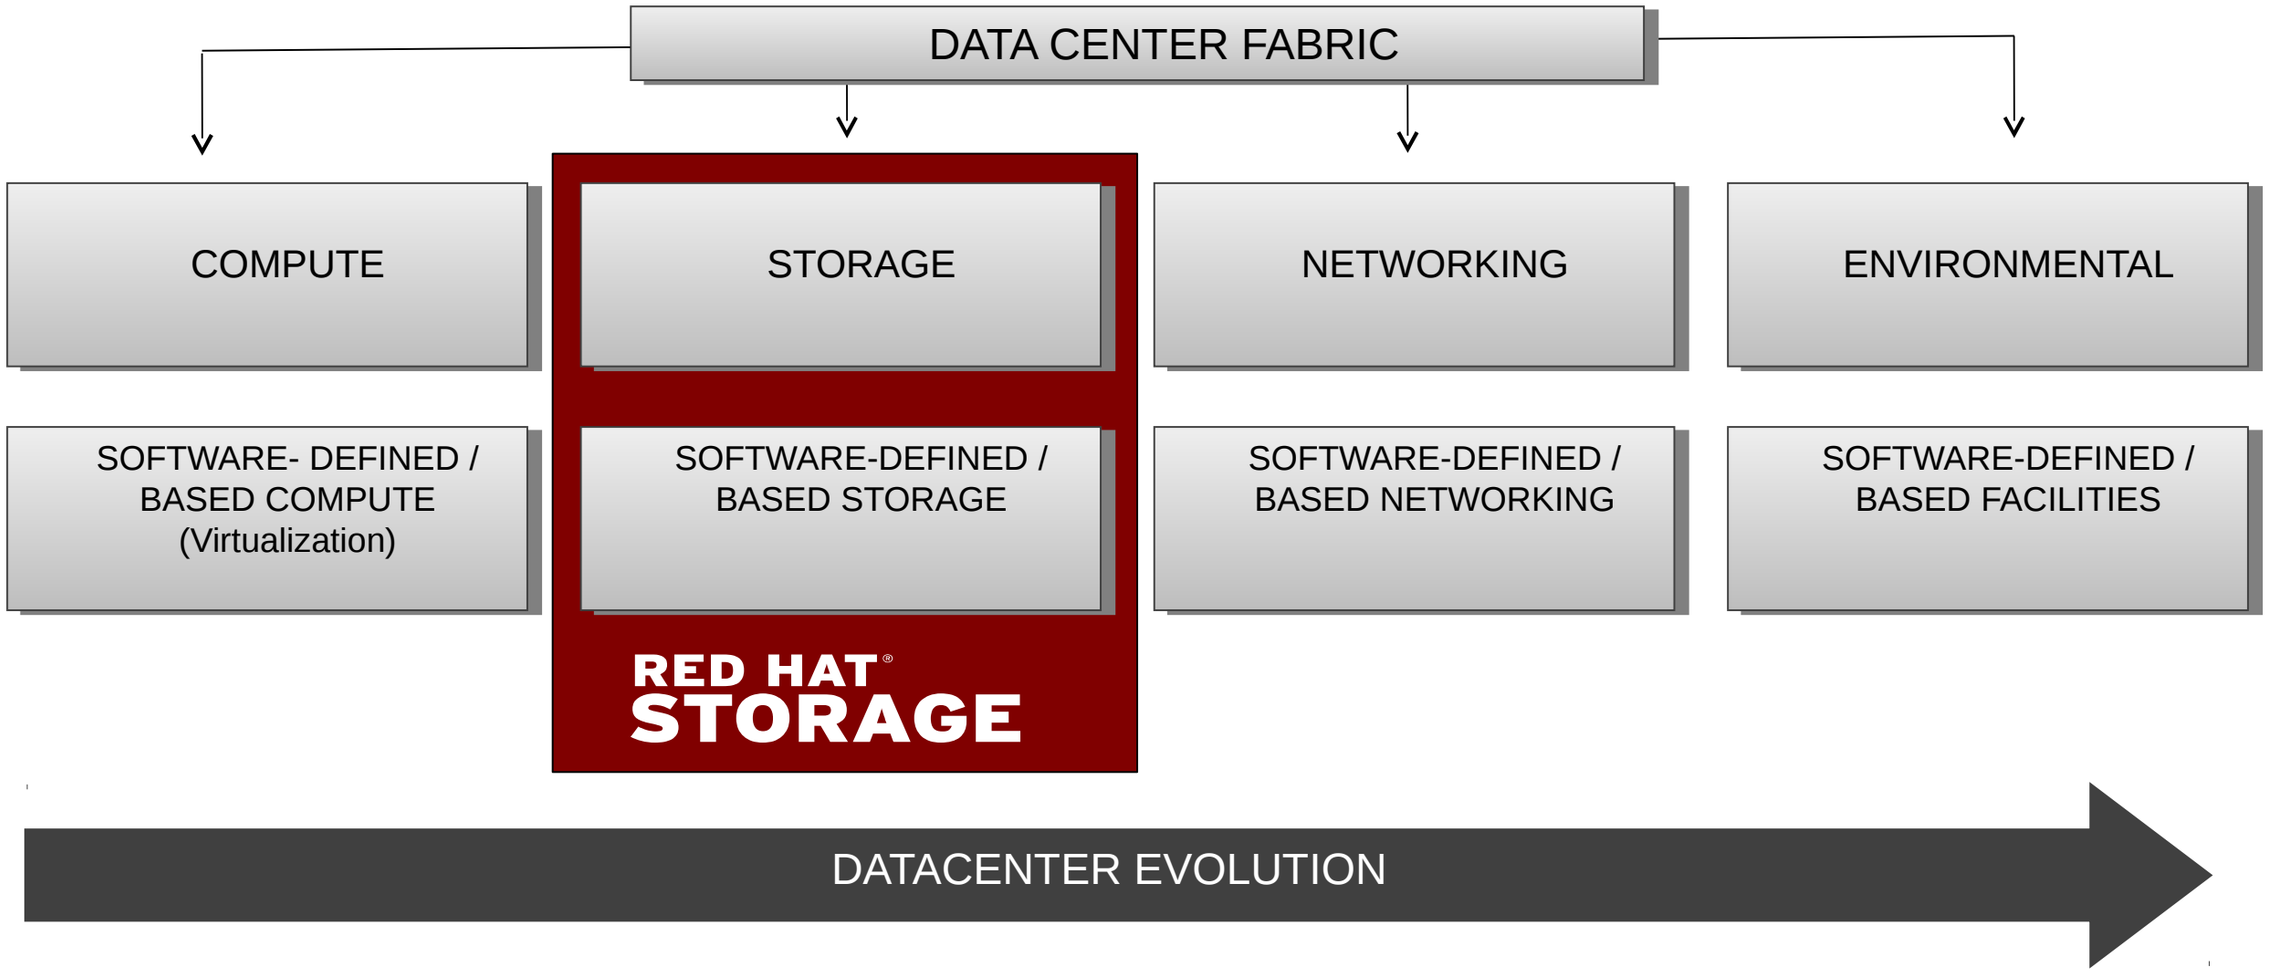

WHAT IS RED HAT STORAGE?

OpenSource

Scale-out NAS (Network Attached Storage)

deployable on

on-premise, virtualized and Cloud environments

based on GlusterFS

running on standard x86 Hardware

INCREASE DATA, APPLICATION AND INFRASTRUCTURE AGILITY

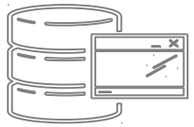

DATA SERVICES

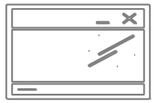

ENTERPRISE APPLICATIONS

BIG DATA WORKLOADS

CLOUD APPLICATIONS

ENTERPRISE MOBILITY

RED HAT® STORAGE

FILE SERVICES | **OPEN OBJECT APIs**

CONVERGED COMPUTE AND STORAGE

OPEN, SOFTWARE-DEFINED STORAGE PLATFORM

PHYSICAL

Standard x86 systems
Scale-out NAS solutions

VIRTUAL

Include idle or legacy resources

CLOUD

EBS EBS

SCALE-OUT STORAGE ARCHITECTURE

PERSISTENT DATA STORES

RED HAT STORAGE DEPLOYMENT ON-PREMISE

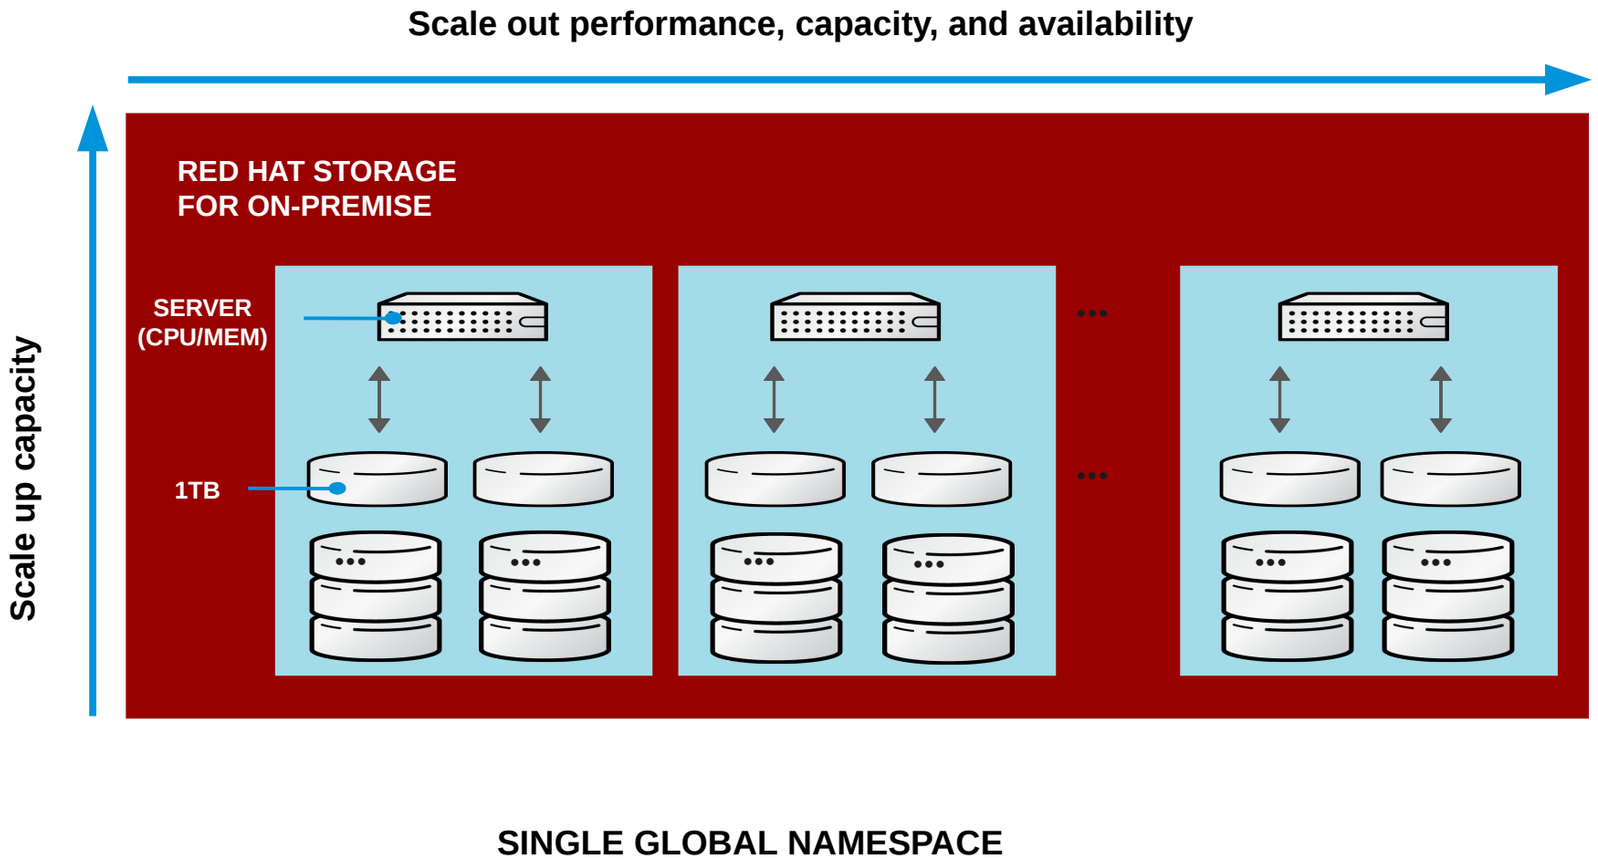

- Single namespace
- Aggregates CPU, memory, network capacity.
- Deploys on Red Hat-supported servers and underlying storage: DAS, JBOD.
- Scale out linearly.
- Scale out performance and capacity as needed.
- Replicate synchronously and asynchronously.

RED HAT STORAGE DEPLOYMENT ON AMAZON CLOUD

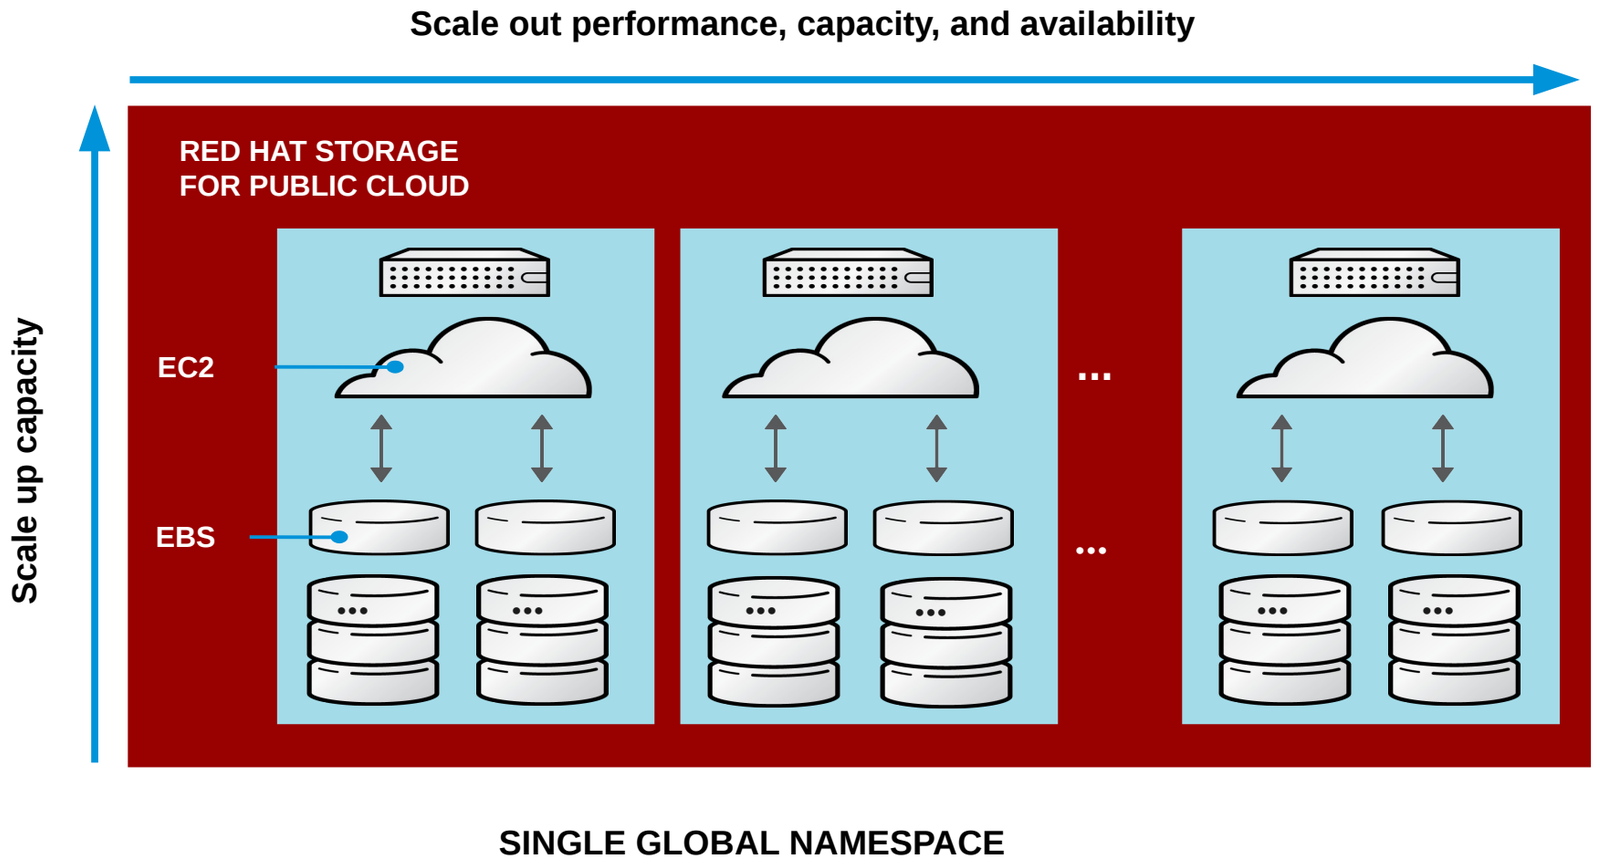

- GlusterFS Amazon Machine Images (AMIs)
- The only way to achieve high availability of Elastic Block Storage (EBS)
- Multiple EBS devices pooled
- POSIX compatible (no application to rewrite required to run on Amazon EC2)
- Scale out capacity and performance as needed

RED HAT STORAGE—50,000 FOOT OVERVIEW

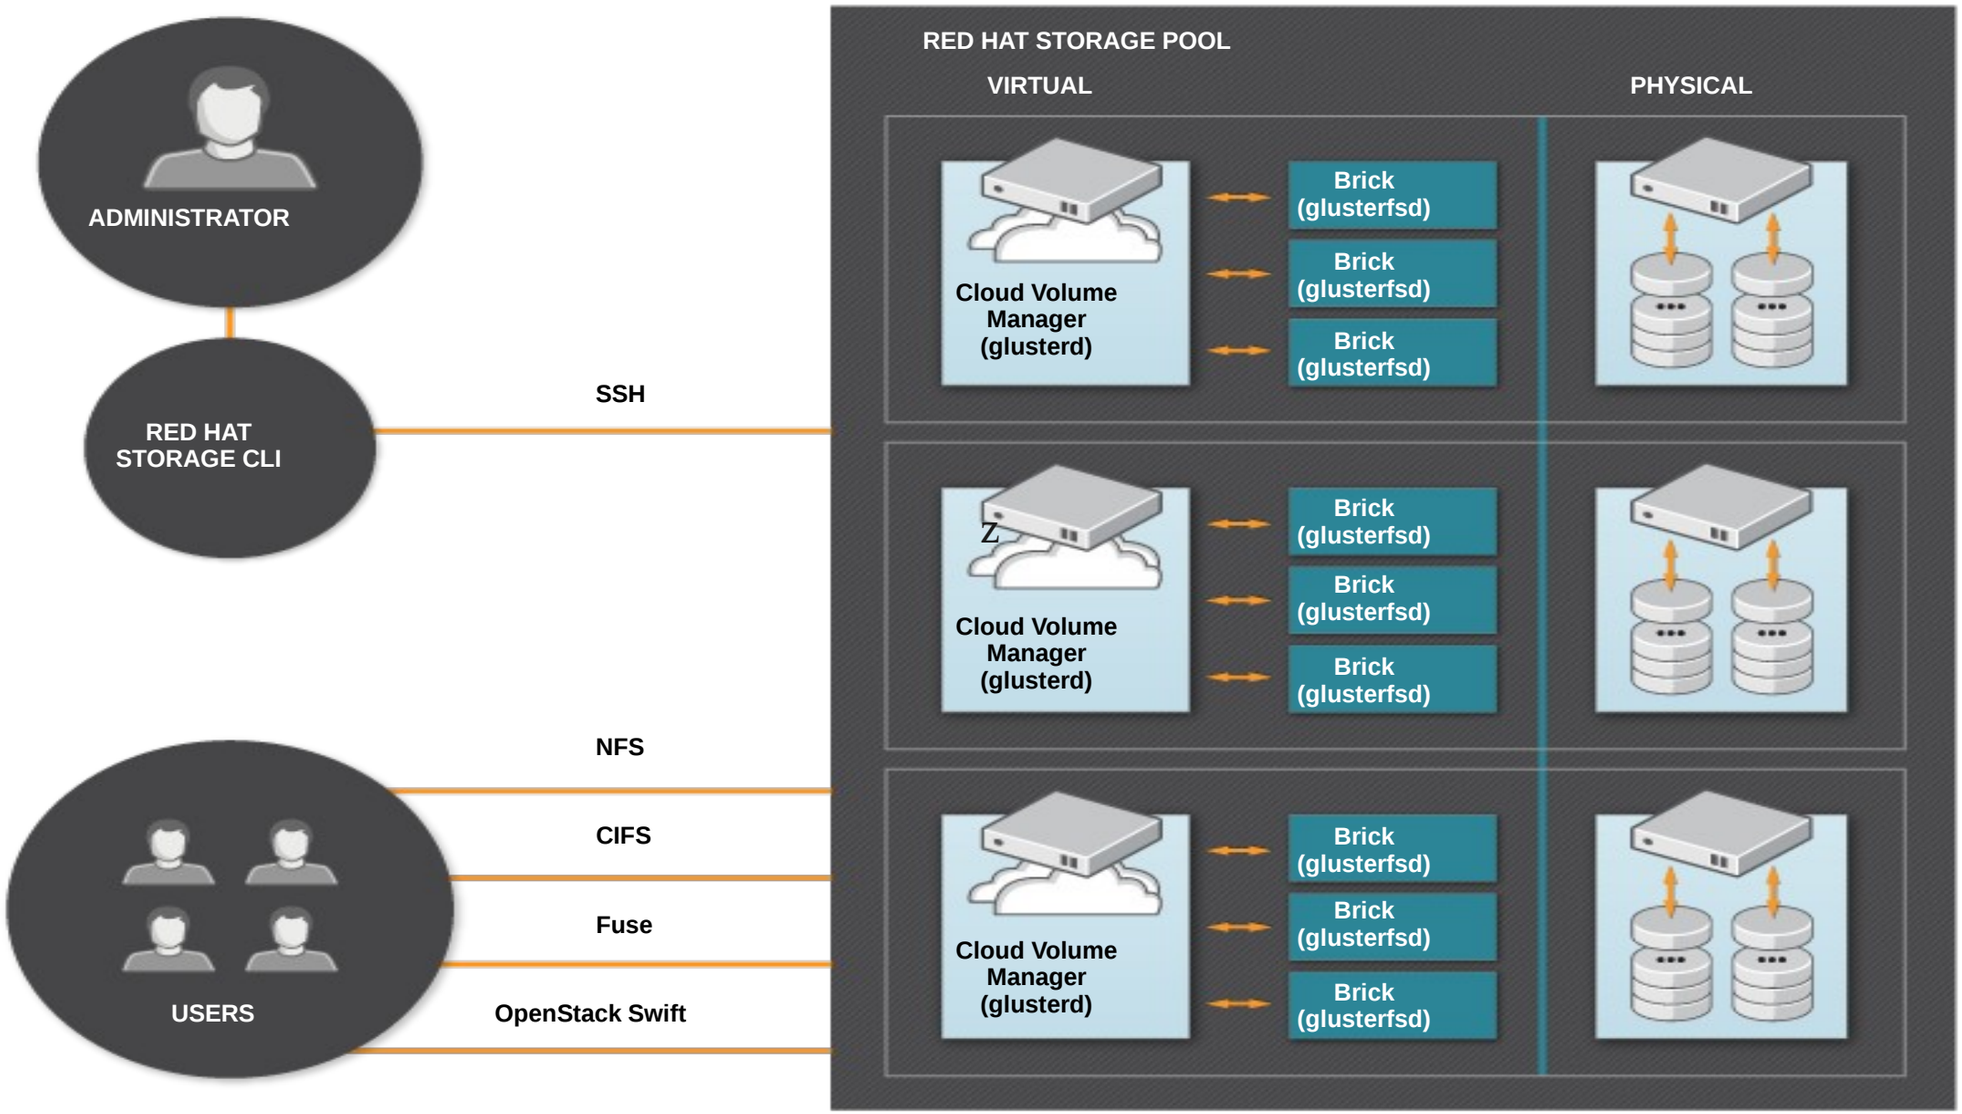

RED HAT STORAGE TECHNOLOGY STACK

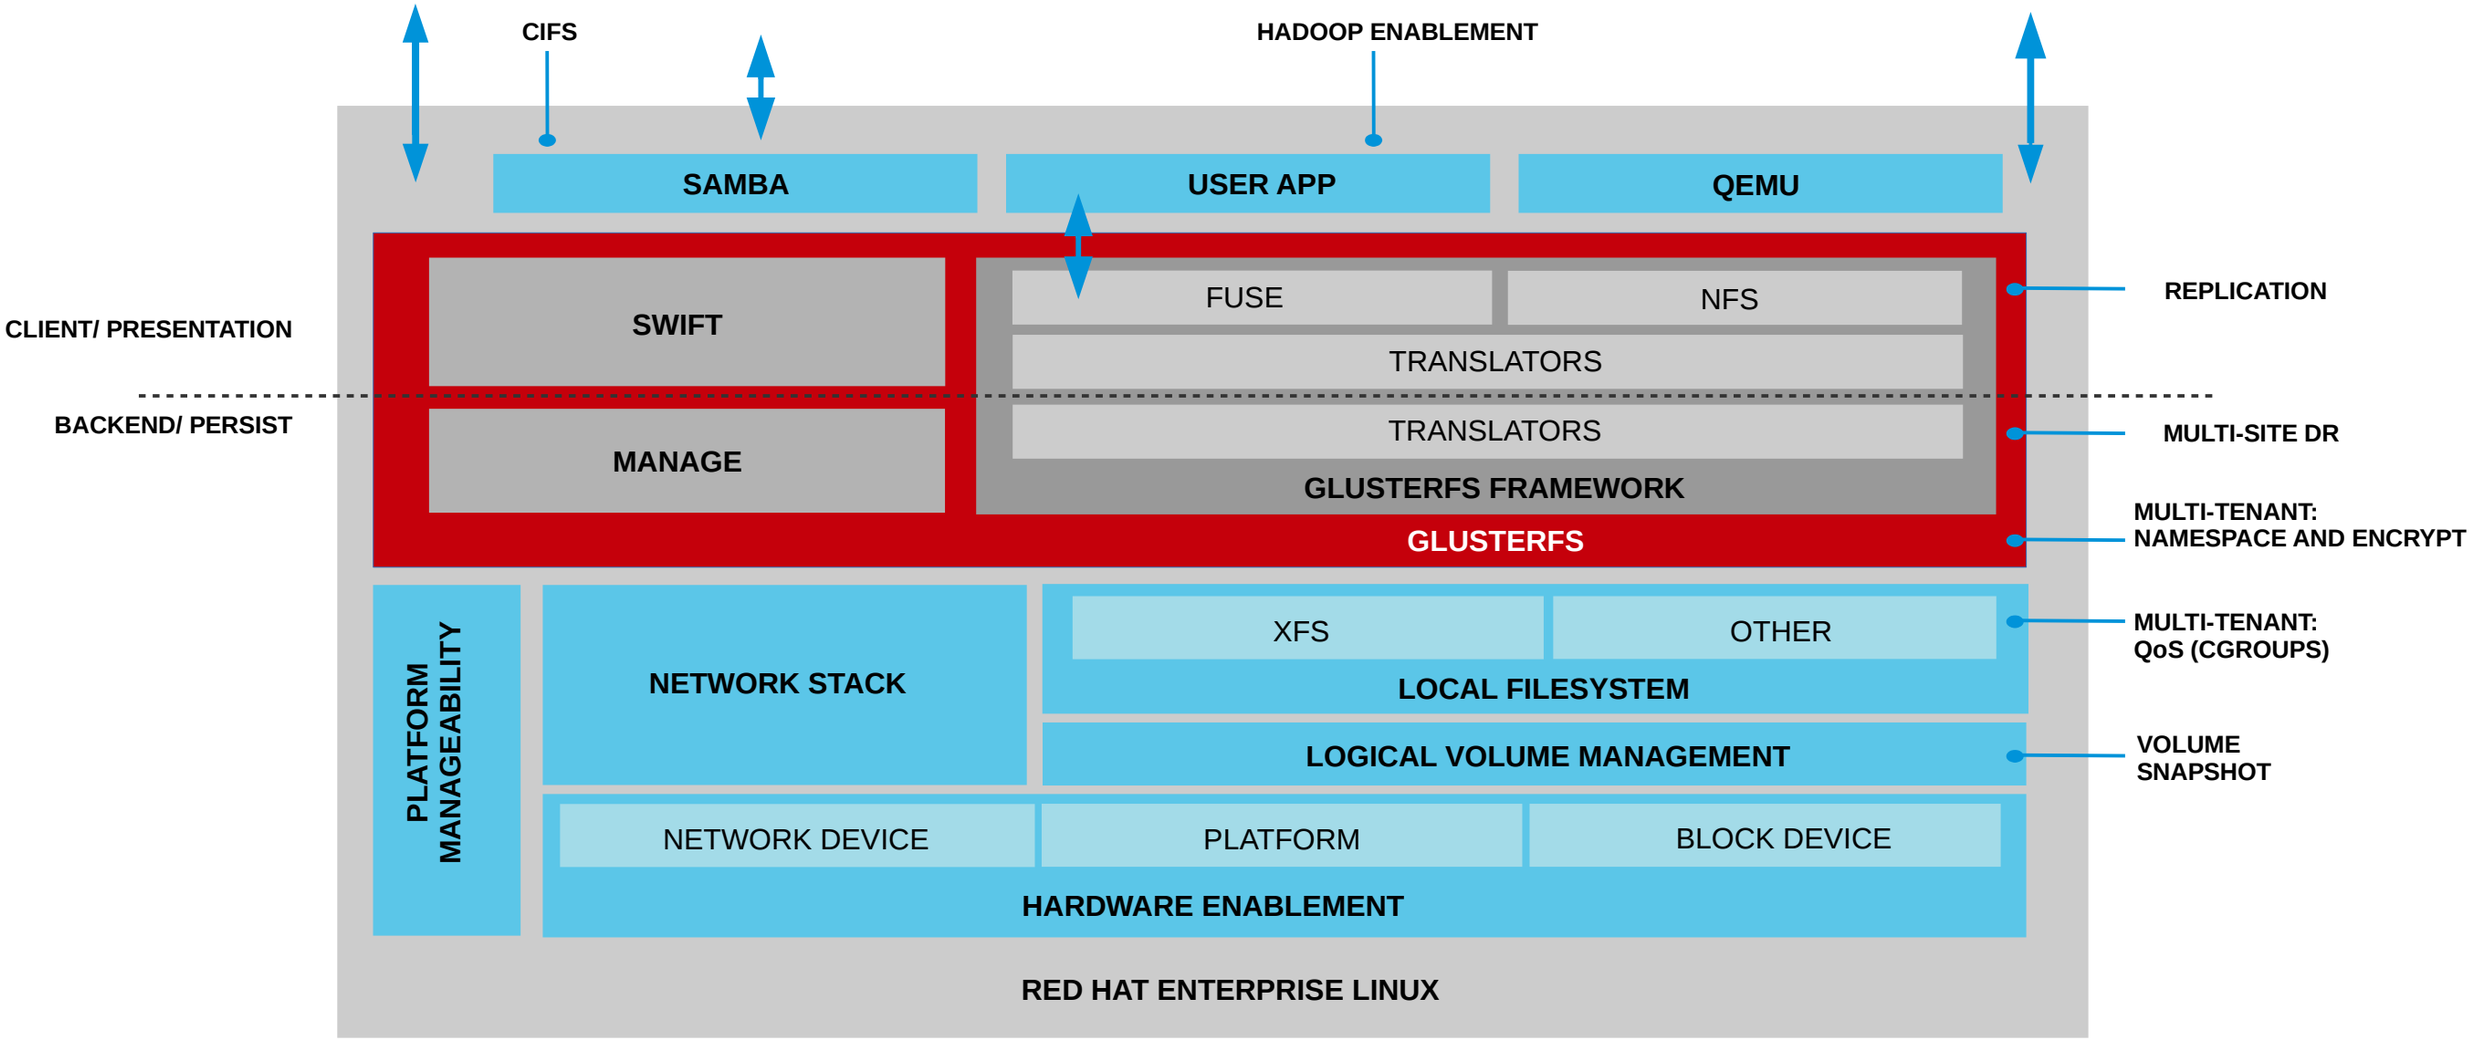

RED HAT STORAGE SCALABILITY

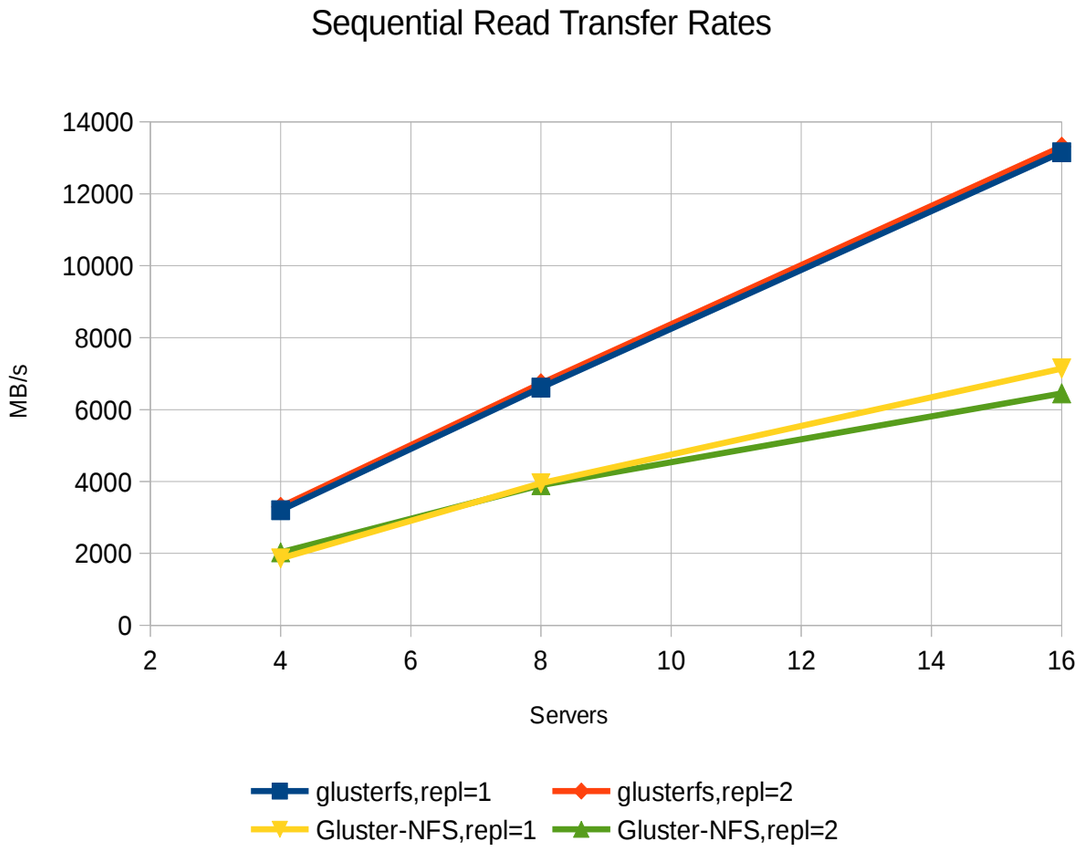

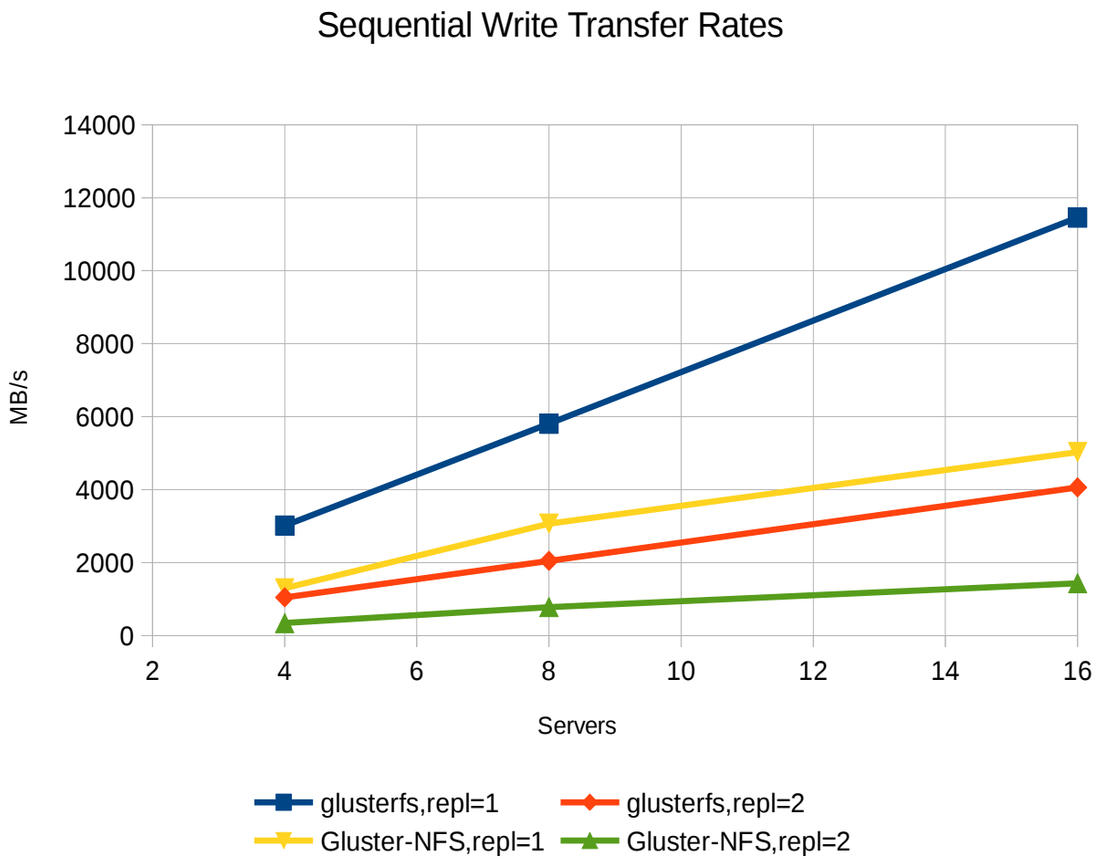

RHS-C Management Console

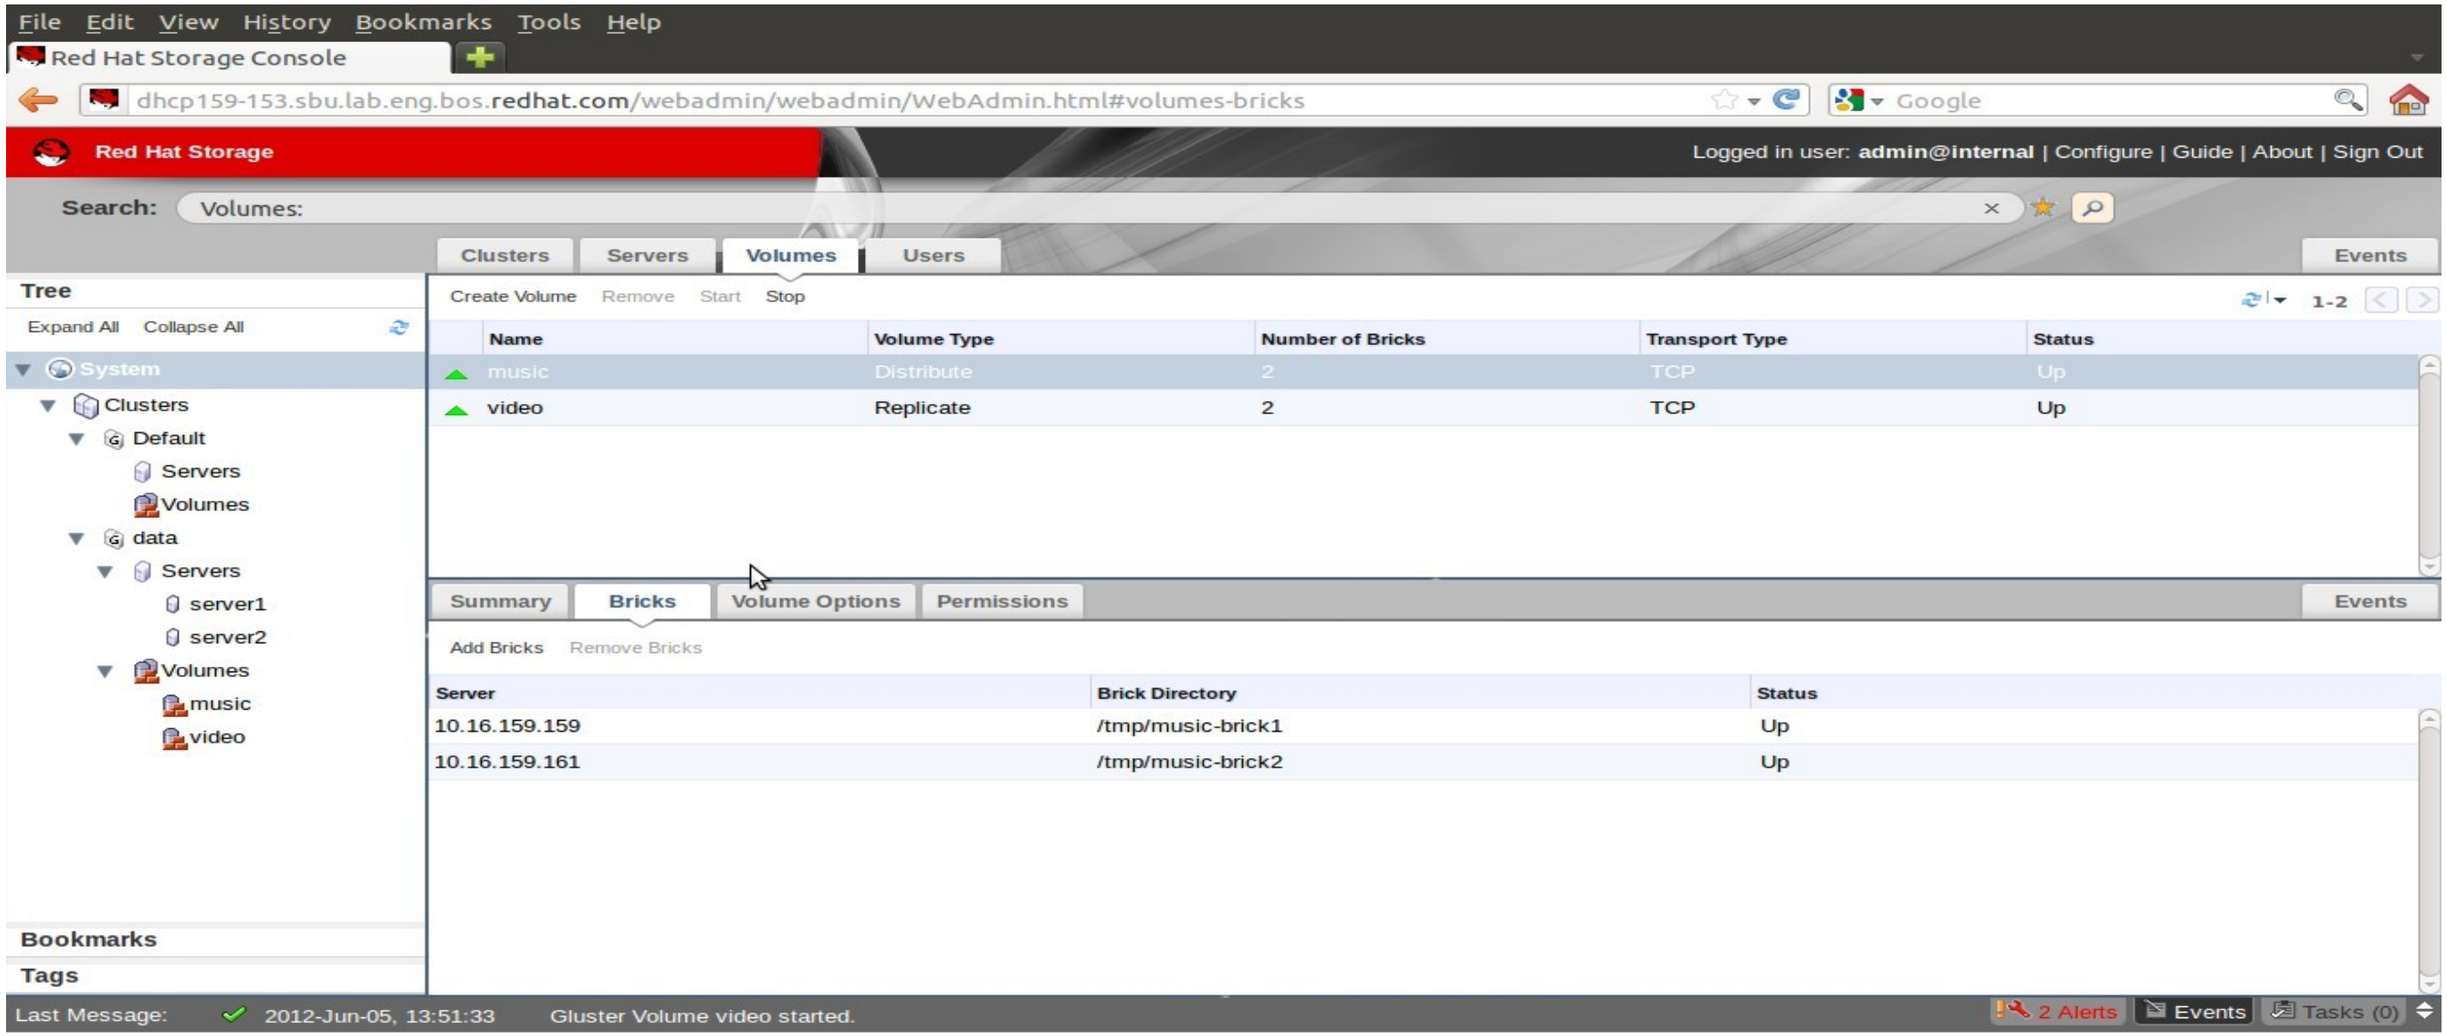

DESIGNED FOR MANAGING UNSTRUCTURED DATA

SUPPORTING A WIDE RANGE OF ENTERPRISE AND EMERGING WORKLOADS

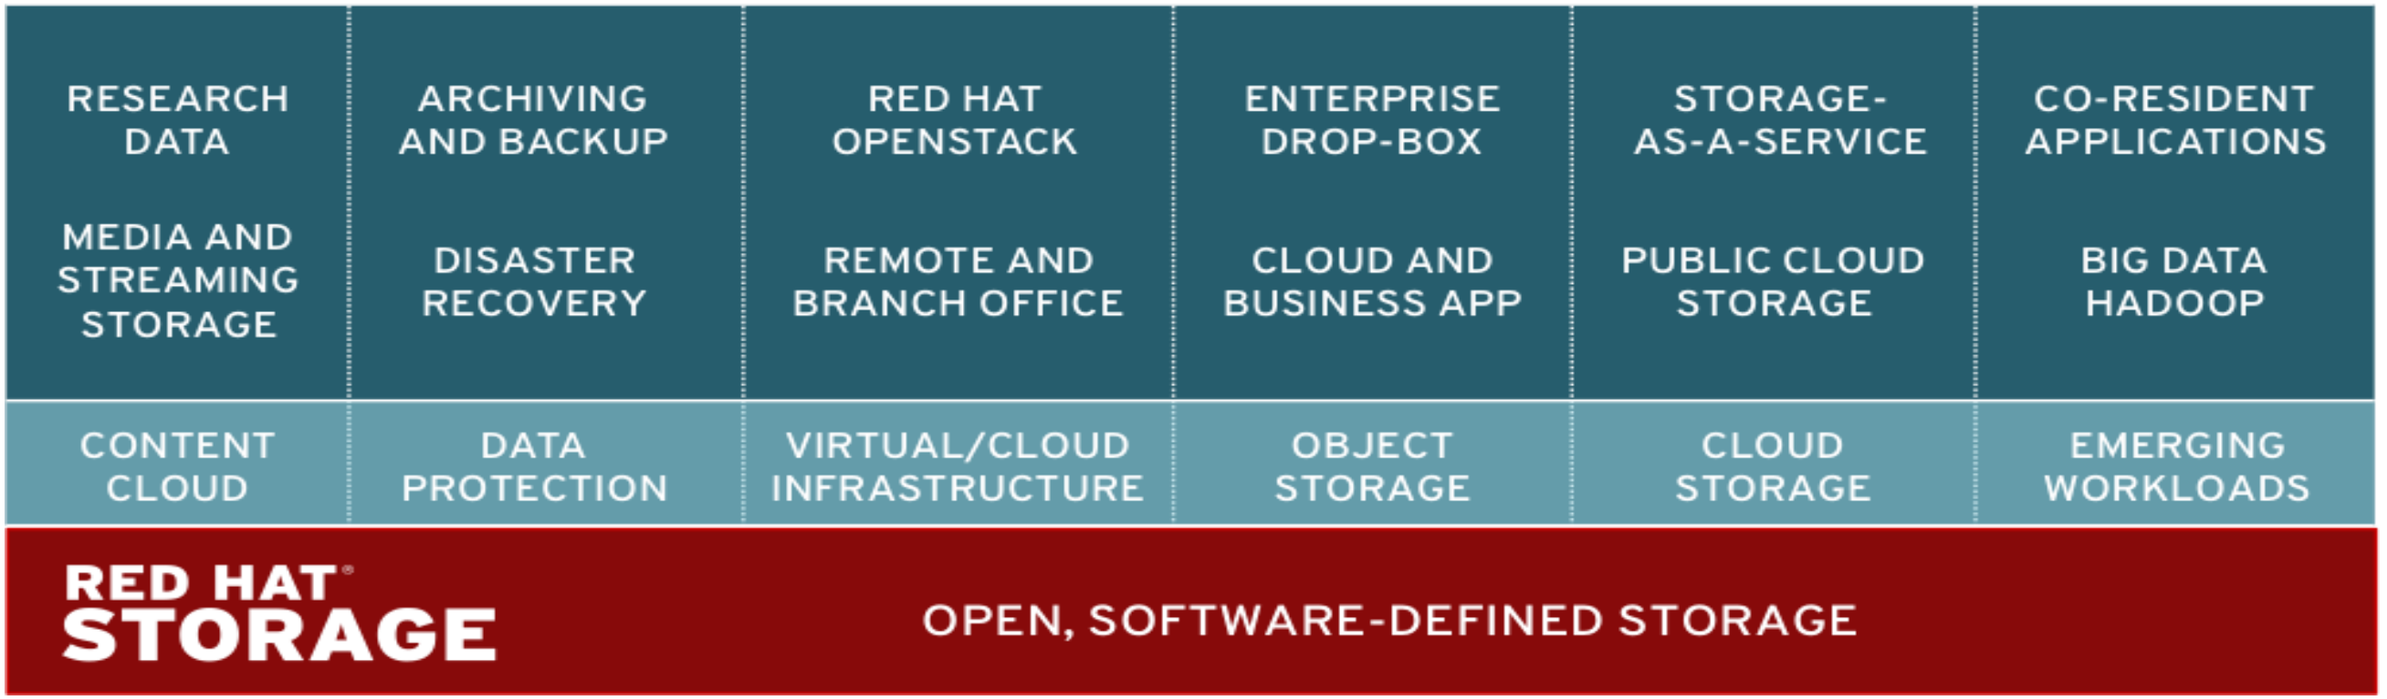

RED HAT STORAGE FOR OPENSTACK

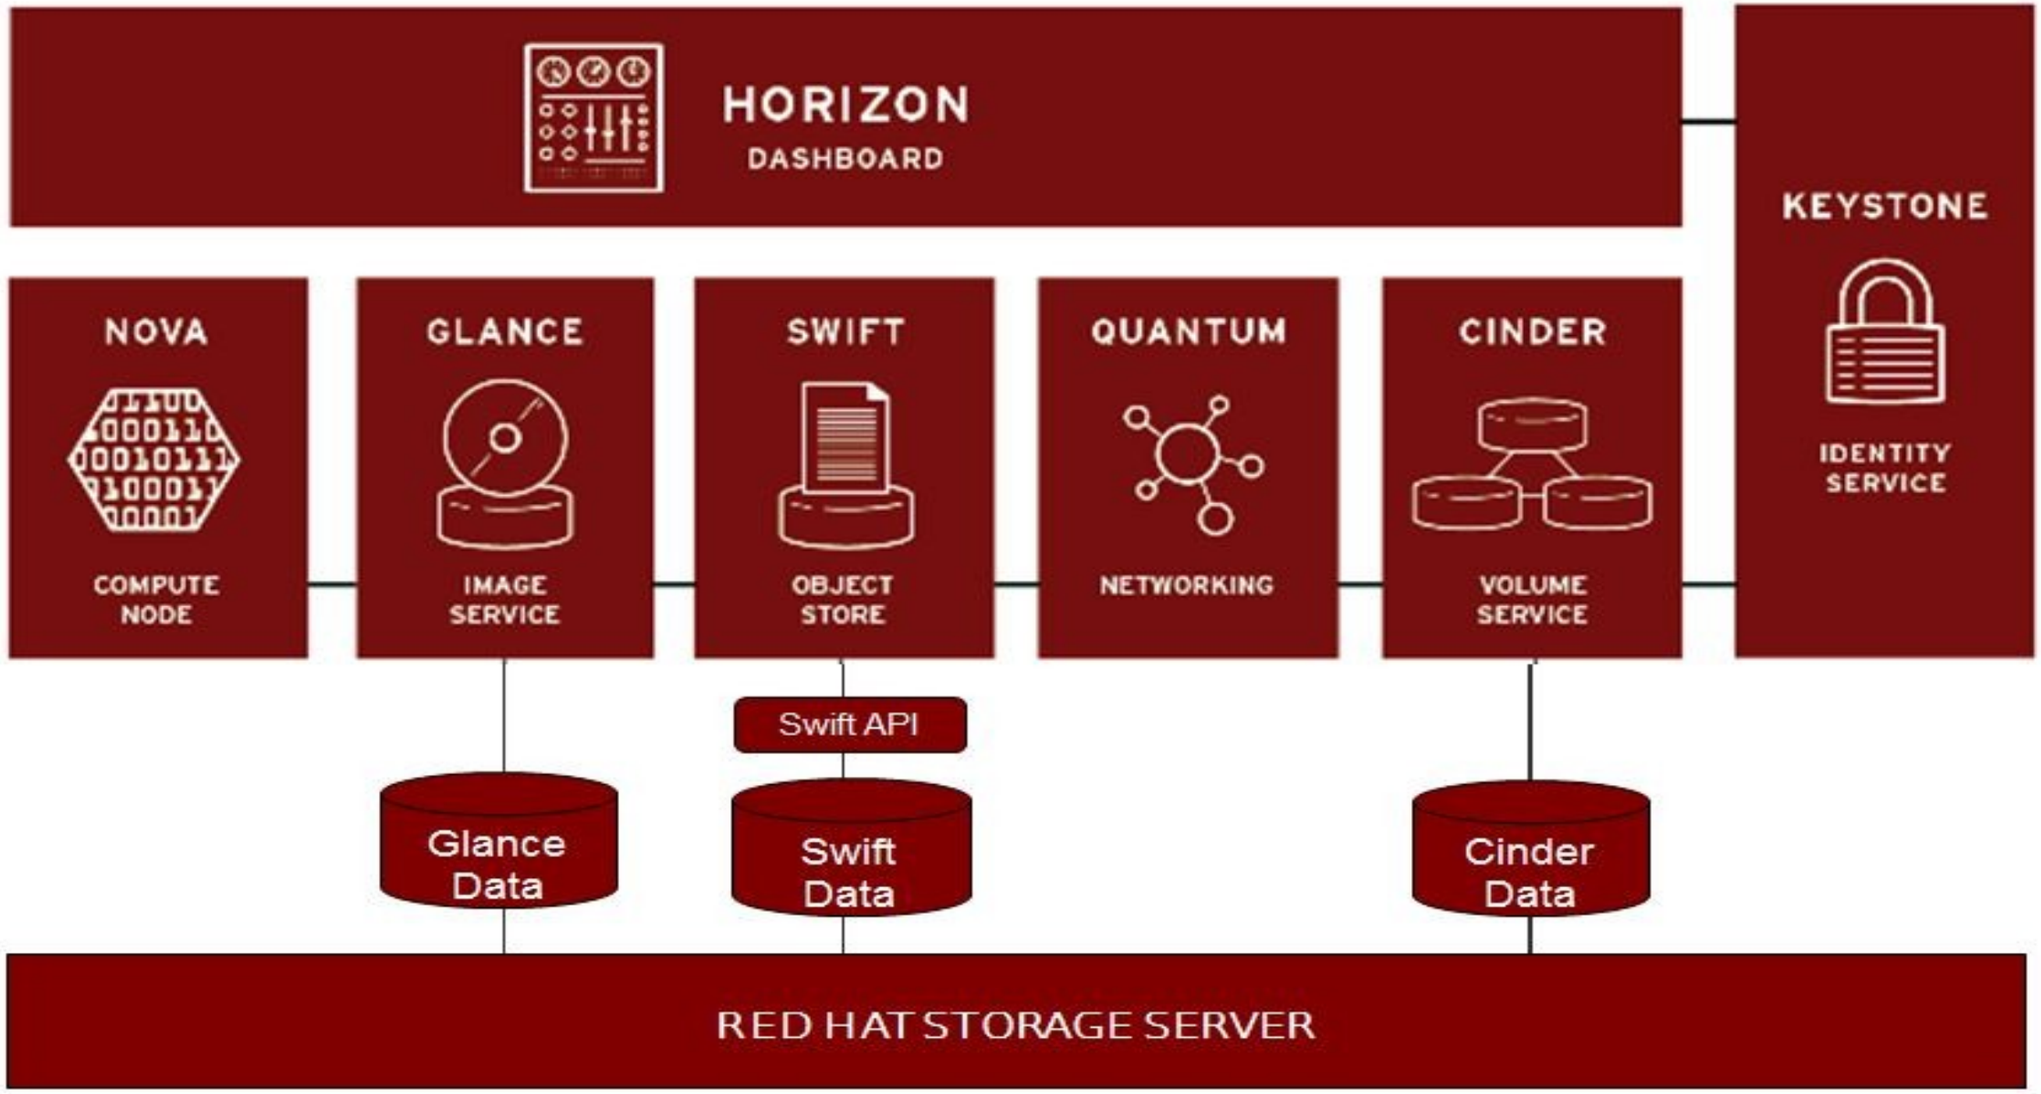

ENSURE GLOBAL DATA PROTECTION AND AVAILABILITY

TRANSPARENTLY DISTRIBUTE DATA GLOBALY

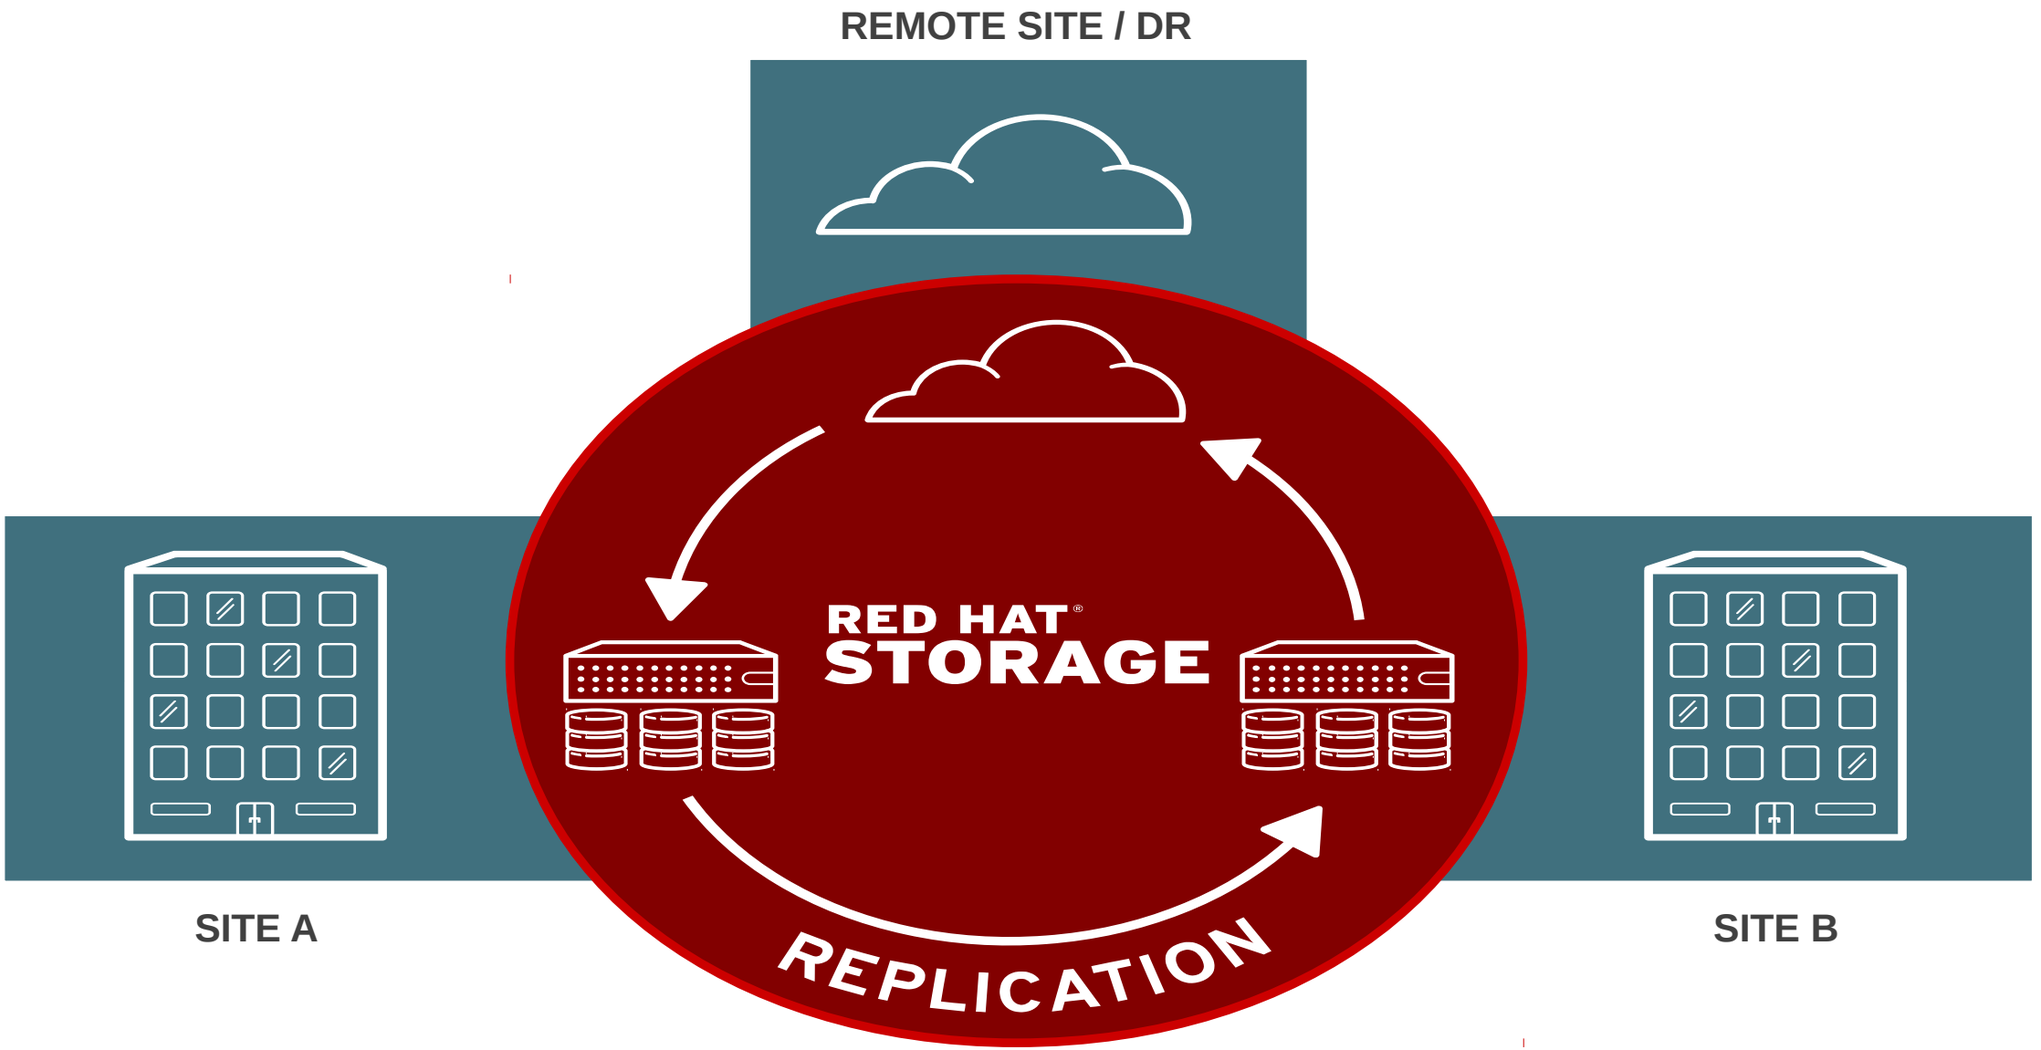

BRING APPLICATIONS CLOSER TO THE DATA

CONVERGING COMPUTE AND STORAGE

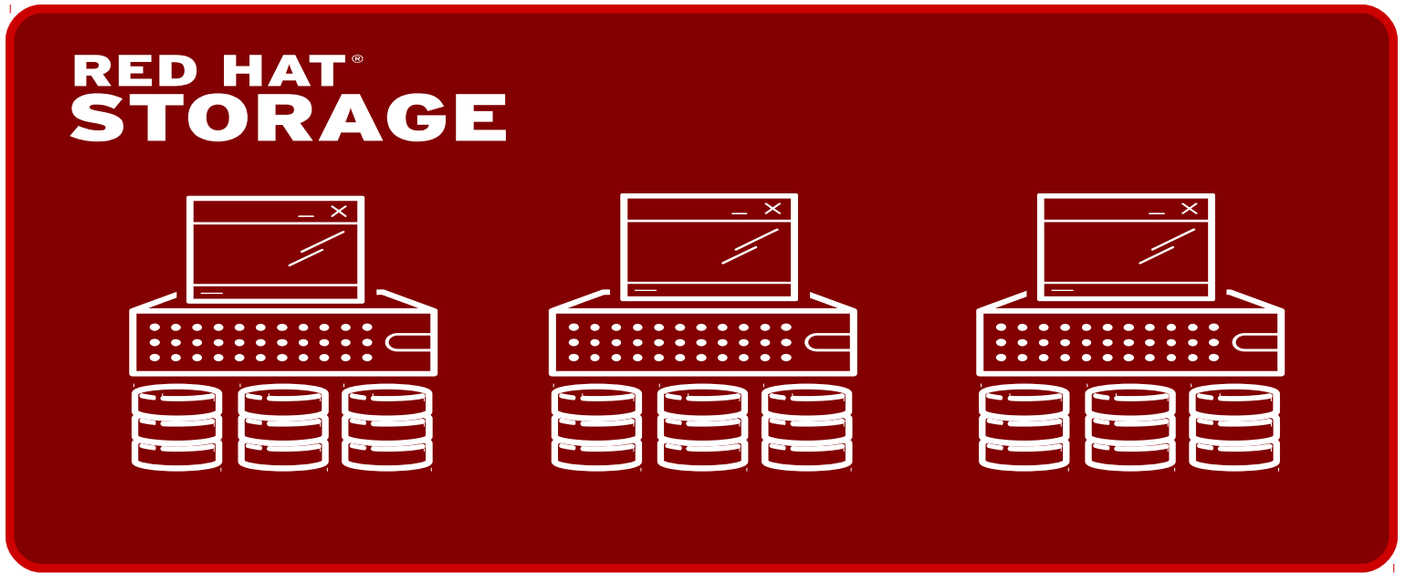

STORAGE RESIDENT APPLICATIONS

REDUCE LATENCY

PROCESS DATA LOCALLY

REDUCE COSTS

INCREASE AGILITY

HIGHLY AVAILABLE CLOUD STORAGE FOR AMAZON EC2

LEVERAGE THE ELASTICITY OF THE CLOUD WITHOUT RE-WRITING YOUR APPLICATIONS

CREATING HIGHLY AVAILABLE, SCALEABLE EBS STORAGE POOLS - ACROSS ZONES

Now available as AWS test-drive

DELIVER COST EFFECTIVE ELASTIC CAPACITY AND PERFORMANCE

53% - 78% REDUCTION IN COSTS

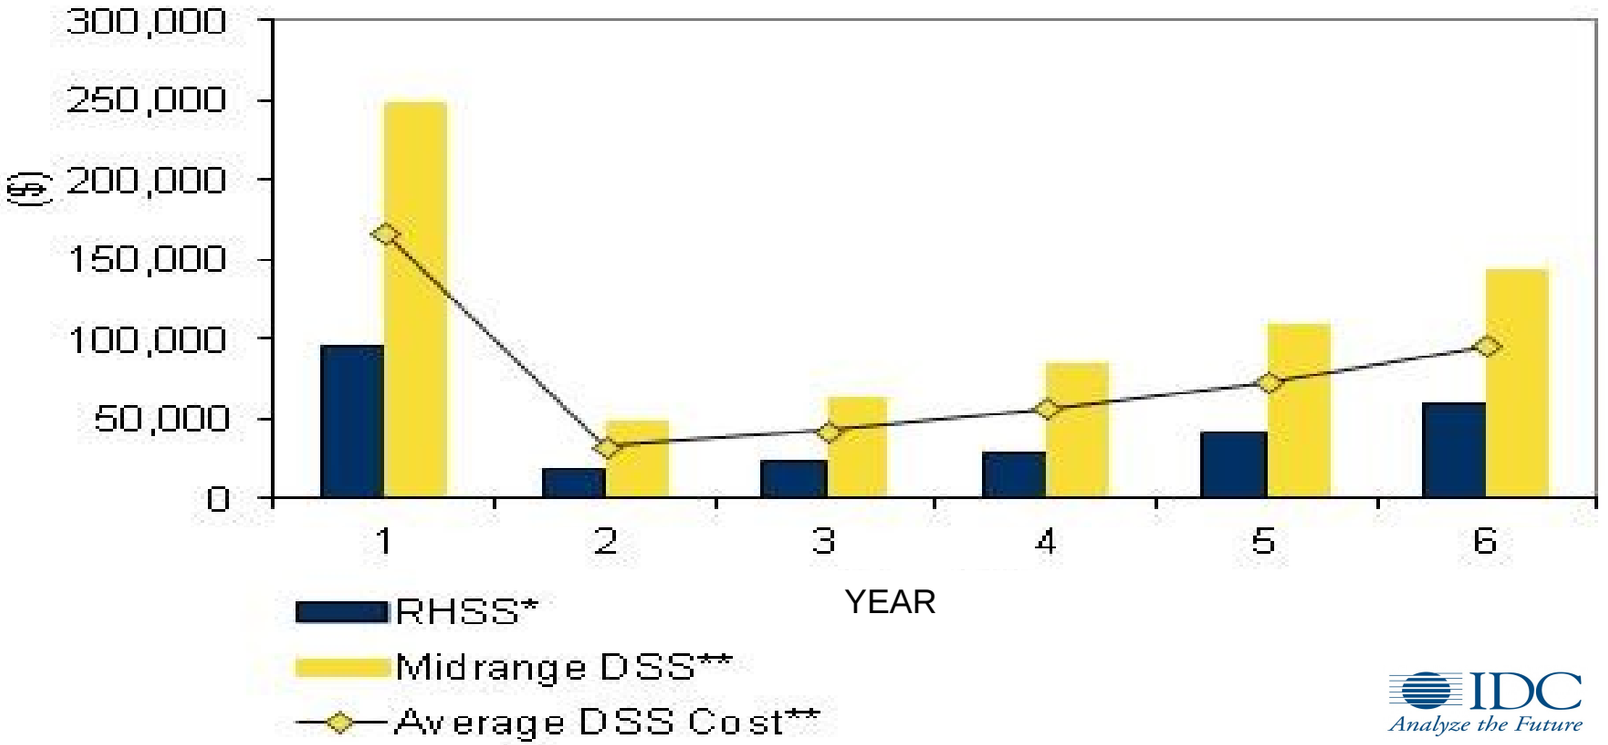

SOURCE: IDC REPORT – THE ECONOMICS OF SOFTWARE BASED STORAGE

MANAGING SPRAWLING UNSTRUCTURED FINANCIAL DATA

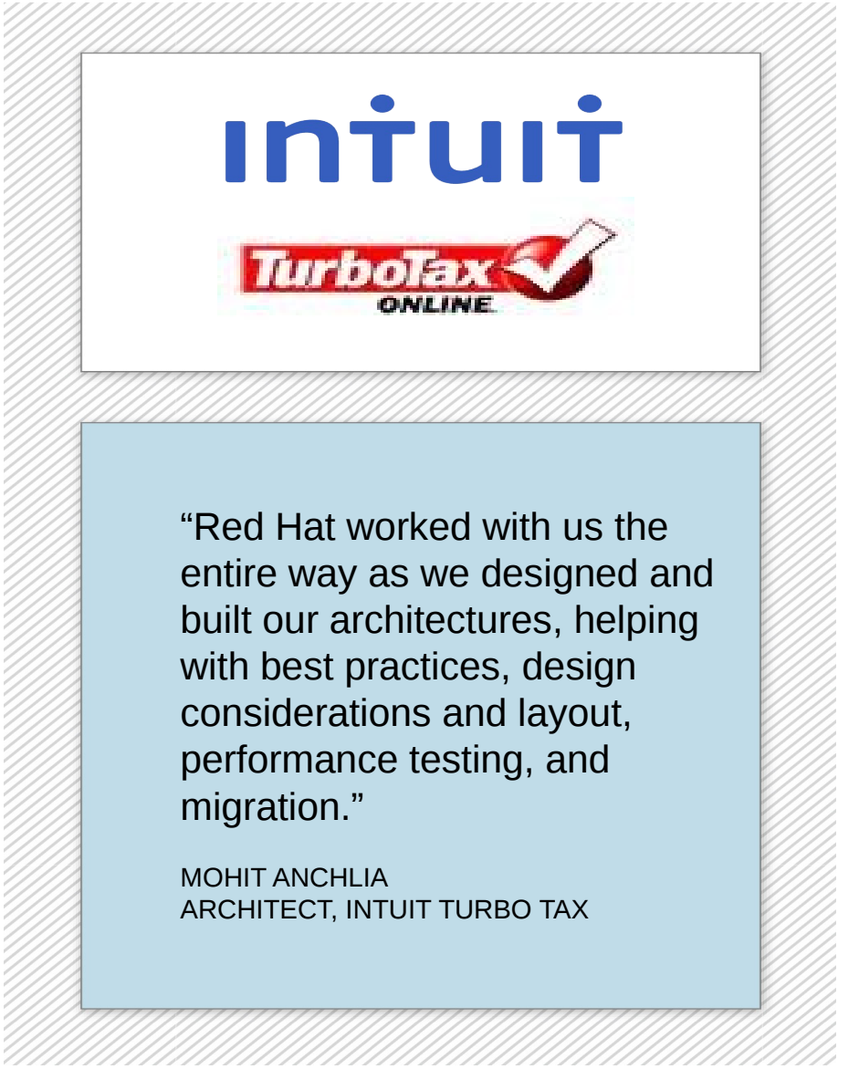

PROBLEM

- NEEDED A FAST, RELIABLE, AND COST-EFFECTIVE STORAGE SOLUTION TO MEET GROWING SAAS LINE OF BUSINESS
- TAX RETURNS AND OTHER DATA WERE BEING STORED AS BLOBS IN AN EXPENSIVE ORACLE DB

SOLUTION

- RED HAT STORAGE SERVER 2.O FOR ON-PREMISE OBJECT STORAGE
- HP DL2000s AND APACHE CASSANDRA

BENEFITS

- SCALEABLE ON-DEMAND STORAGE FOR UNSTRUCTURED DATA
- COST EFFECTIVE SOLUTION THAT LEVERAGES COMMODITY HARDWARE
- MEET GROWING CAPACITY AND PEAK PERFORMACE NEEDS
- ACHIEVE MULTI-SITE DISASTER RECOVERY

IS THE OPPORTUNITY REAL ?

SearchStorage

TechTarget

News | Premium Editorial | Multimedia | Storage Topics | Tutorials | Expert Advice | Vendor Content | Blogs | Storage Decisions Events

SEARCH GO

Home > Topics > Data Storage Management > Data management tools > Red Hat Gluster will transform storage market, IDC analyst predicts

Red Hat Gluster will transform storage market, IDC analyst predicts

Carol Sliwa
Published: 08 Aug 2013

Share

ESSENTIAL GUIDE RED HAT STORAGE SERVER SEEKS TO MIMIC SUCCESS OF ENTERPRISE LINUX

This article is part of an Essential Guide, our editor-selected collection of our best articles, videos and other content on this topic.Explore more in this guide:

1. - IDC'S NADKARNI PREDICTS MAINSTREAM ACCEPTANCE : READ MORE IN THIS SECTION

- **GlusterFS-based software will drive mainstream acceptance of software-based storage**

Explore other sections in this guide:

2. - Wikibon's Floyer: Manageability work needed

3. - Staimer foresees limited appeal

Red Hat Inc. will eventually turn the market "on its head" with its Red Hat GlusterFS-based storage software in the same way that it successfully challenged major server operating system vendors with its distribution of enterprise Linux, an IDC analyst predicted.

Ashish Nadkarni, a research director in the storage systems practice at Framingham, Mass.-based International Data Corp. (IDC), said Red Hat may need a few years to realize its vision of converting commodity hardware into a full-fledged storage platform.

But Nadkarni thinks Red Hat Gluster will challenge scale-out, file-based storage systems from major vendors as well as combined compute-storage platforms that cater to the high-performance computing market with data analytics software.

LISTEN TO THE ENTIRE PODCAST SERIES ABOUT RED HAT'S CHANCES WITH ITS STORAGE SERVER SOFTWARE

Red Hat's open source software-based storage holds appeal mainly for passive data

With push into software-based storage, Red Hat will help drive down cost of storage features

Latest News

- EMC Corp. prepares VNX2 launch
- Scality CEO: Object-oriented storage needs multiple access methods
- Coraid Inc. NAS gets more cache for heavy I/O loads
- Riverbed Granite storage box gets Fibre Channel support
- Emerging data storage trends embraced by BNY Mellon IT infrastructure

DELIVERING THE NEXT GENERATION OF OPEN SOFTWARE-DEFINED STORAGE TODAY

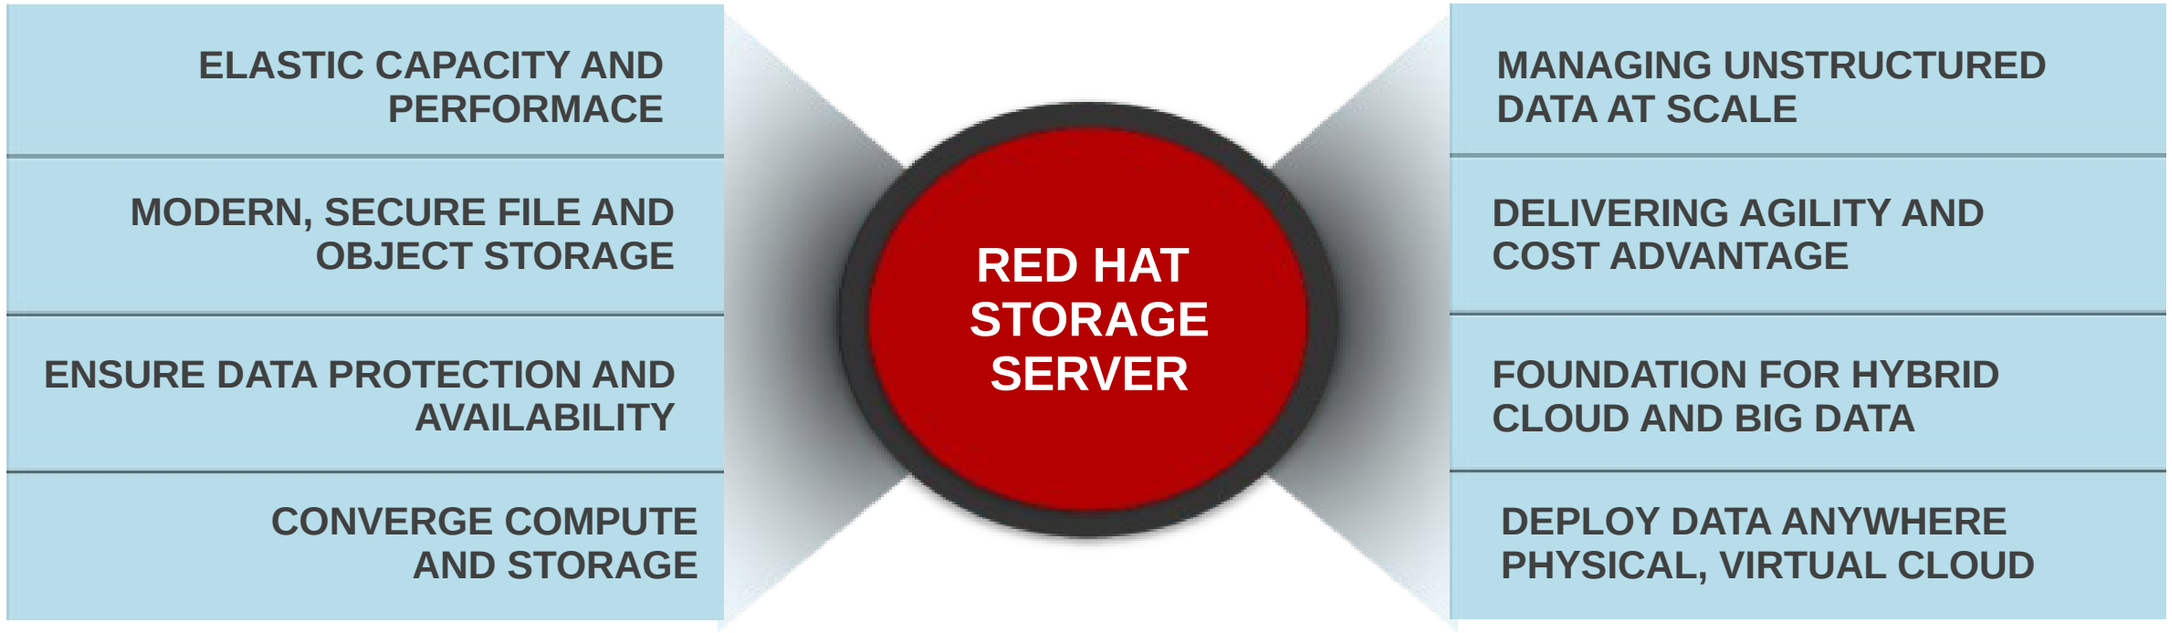

DESIGNED FOR TODAYS IT & DATA ECONOMICS

RED HAT STORAGE – DEVELOPING 3rd PARTY ECO-SYSTEM

RED HAT LEADS THROUGH OPEN INNOVATION

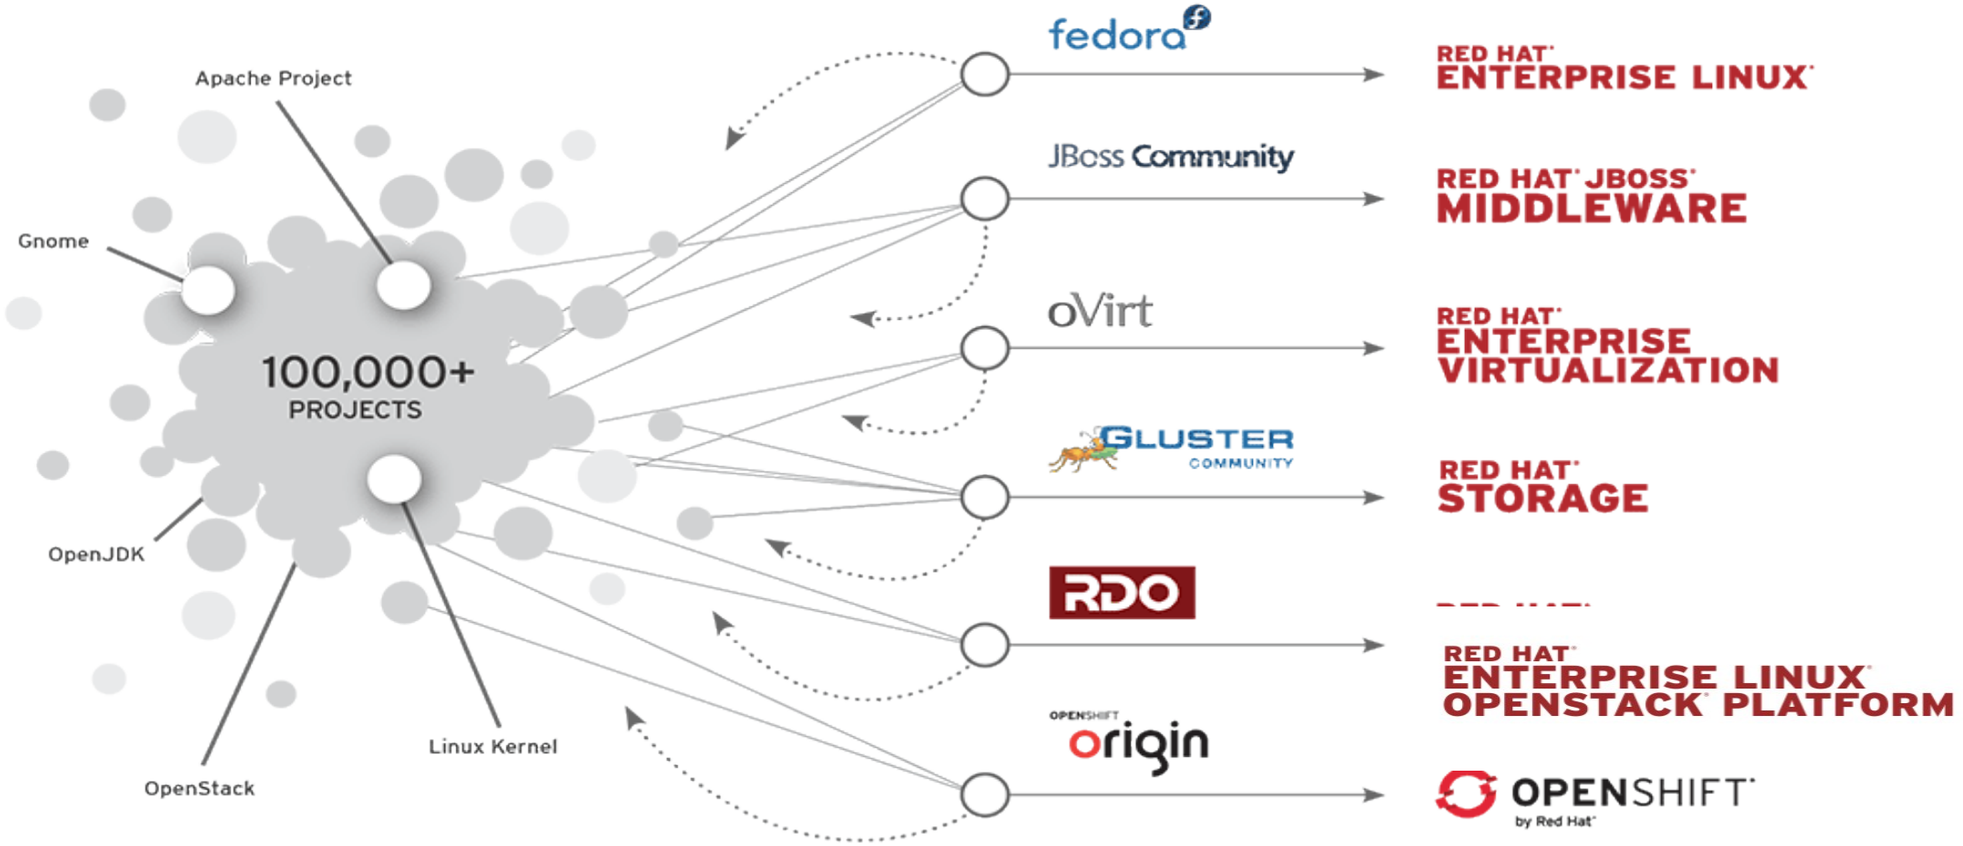

COMMUNITY INNOVATION

GLUSTER.ORG COMMUNITY FORGE ENHANCEMENTS AND PROJECTS

SNAPSHOTTING

CHANGE DETECTION

COMPRESSION

3-WAY REPLICATION

pNFS AND NFSv4 SUPPORT

FILE VERSIONING

ERASURE CODING

MULTI-MASTER GEO-REPLICATION

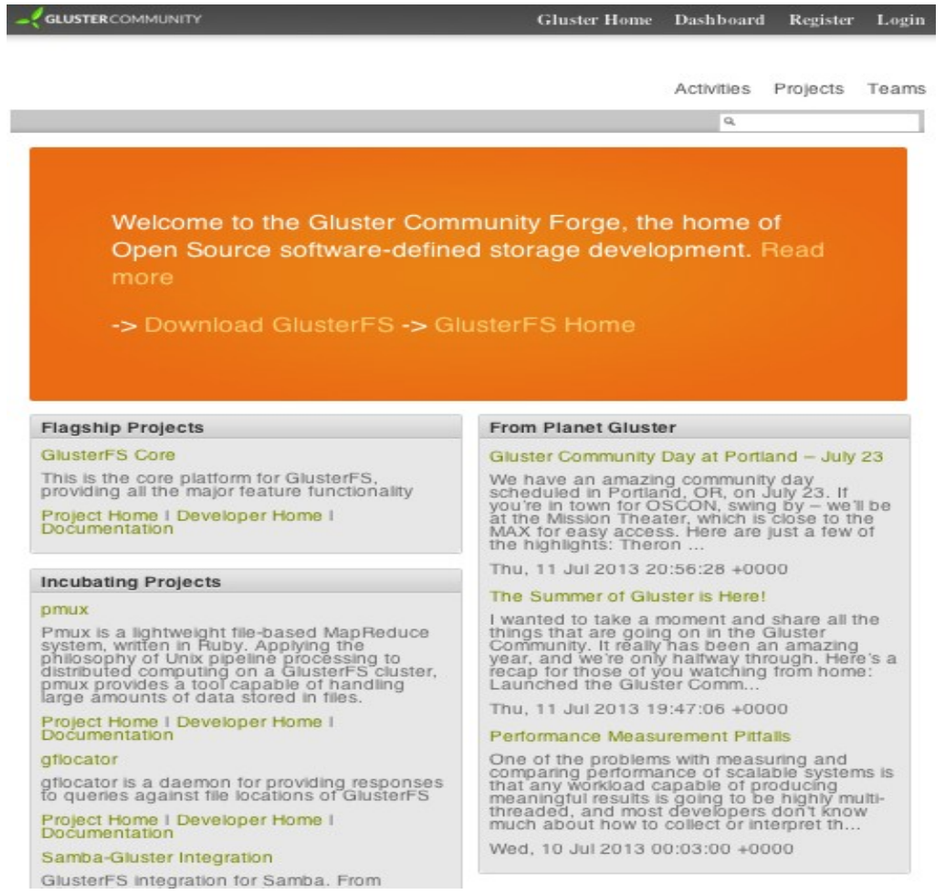

SMB 3.0 SUPPORT

NDMP SERVER

PUPPET MANAGEMENT MODULE

GTOP - MONITORING

GLUSTER PROFILING

SELINUX SUPPORT

PMUX – LIGHTWEIGHT MAP REDUCE

TRANSLATORS EXTENSION FOR PYTHON

RED HAT STORAGE INFORMATION RESOURCES

RED HAT STORAGE PRODUCT INFORMATION

HTTP://WWW.REDHAT.COM/PRODUCTS/STORAGE-SERVER/

RED HAT STORAGE SOLUTIONS

HTTP://WWW.REDHAT.COM/PROMO/LIBERATE/SOLUTIONS.HTML

RED HAT STORAGE CUSTOMER SUCCESS STORIES

HTTP://WWW.REDHAT.COM/PROMO/LIBERATE/RESOURCES.HTML

RED HAT STORAGE SERVICES

HTTP://WWW.REDHAT.COM/PROMO/LIBERATE/SERVICES.HTML

GLUSTER COMMUNITY

HTTP://WWW.GLUSTER.ORG

GLUSTER COMMUNITY FORGE

HTTP://FORGE.GLUSTER.ORG

BACKUP

How Does GlusterFS Work Without Metadata?

- Files are placed on a brick(s) in the cluster based on a calculation
- All native clients have an algorithm built-in
- All storage nodes have an algorithm built-in
- Files can then be retrieved based on the same calculation
- For non-native clients, the server handles retrieval and placement

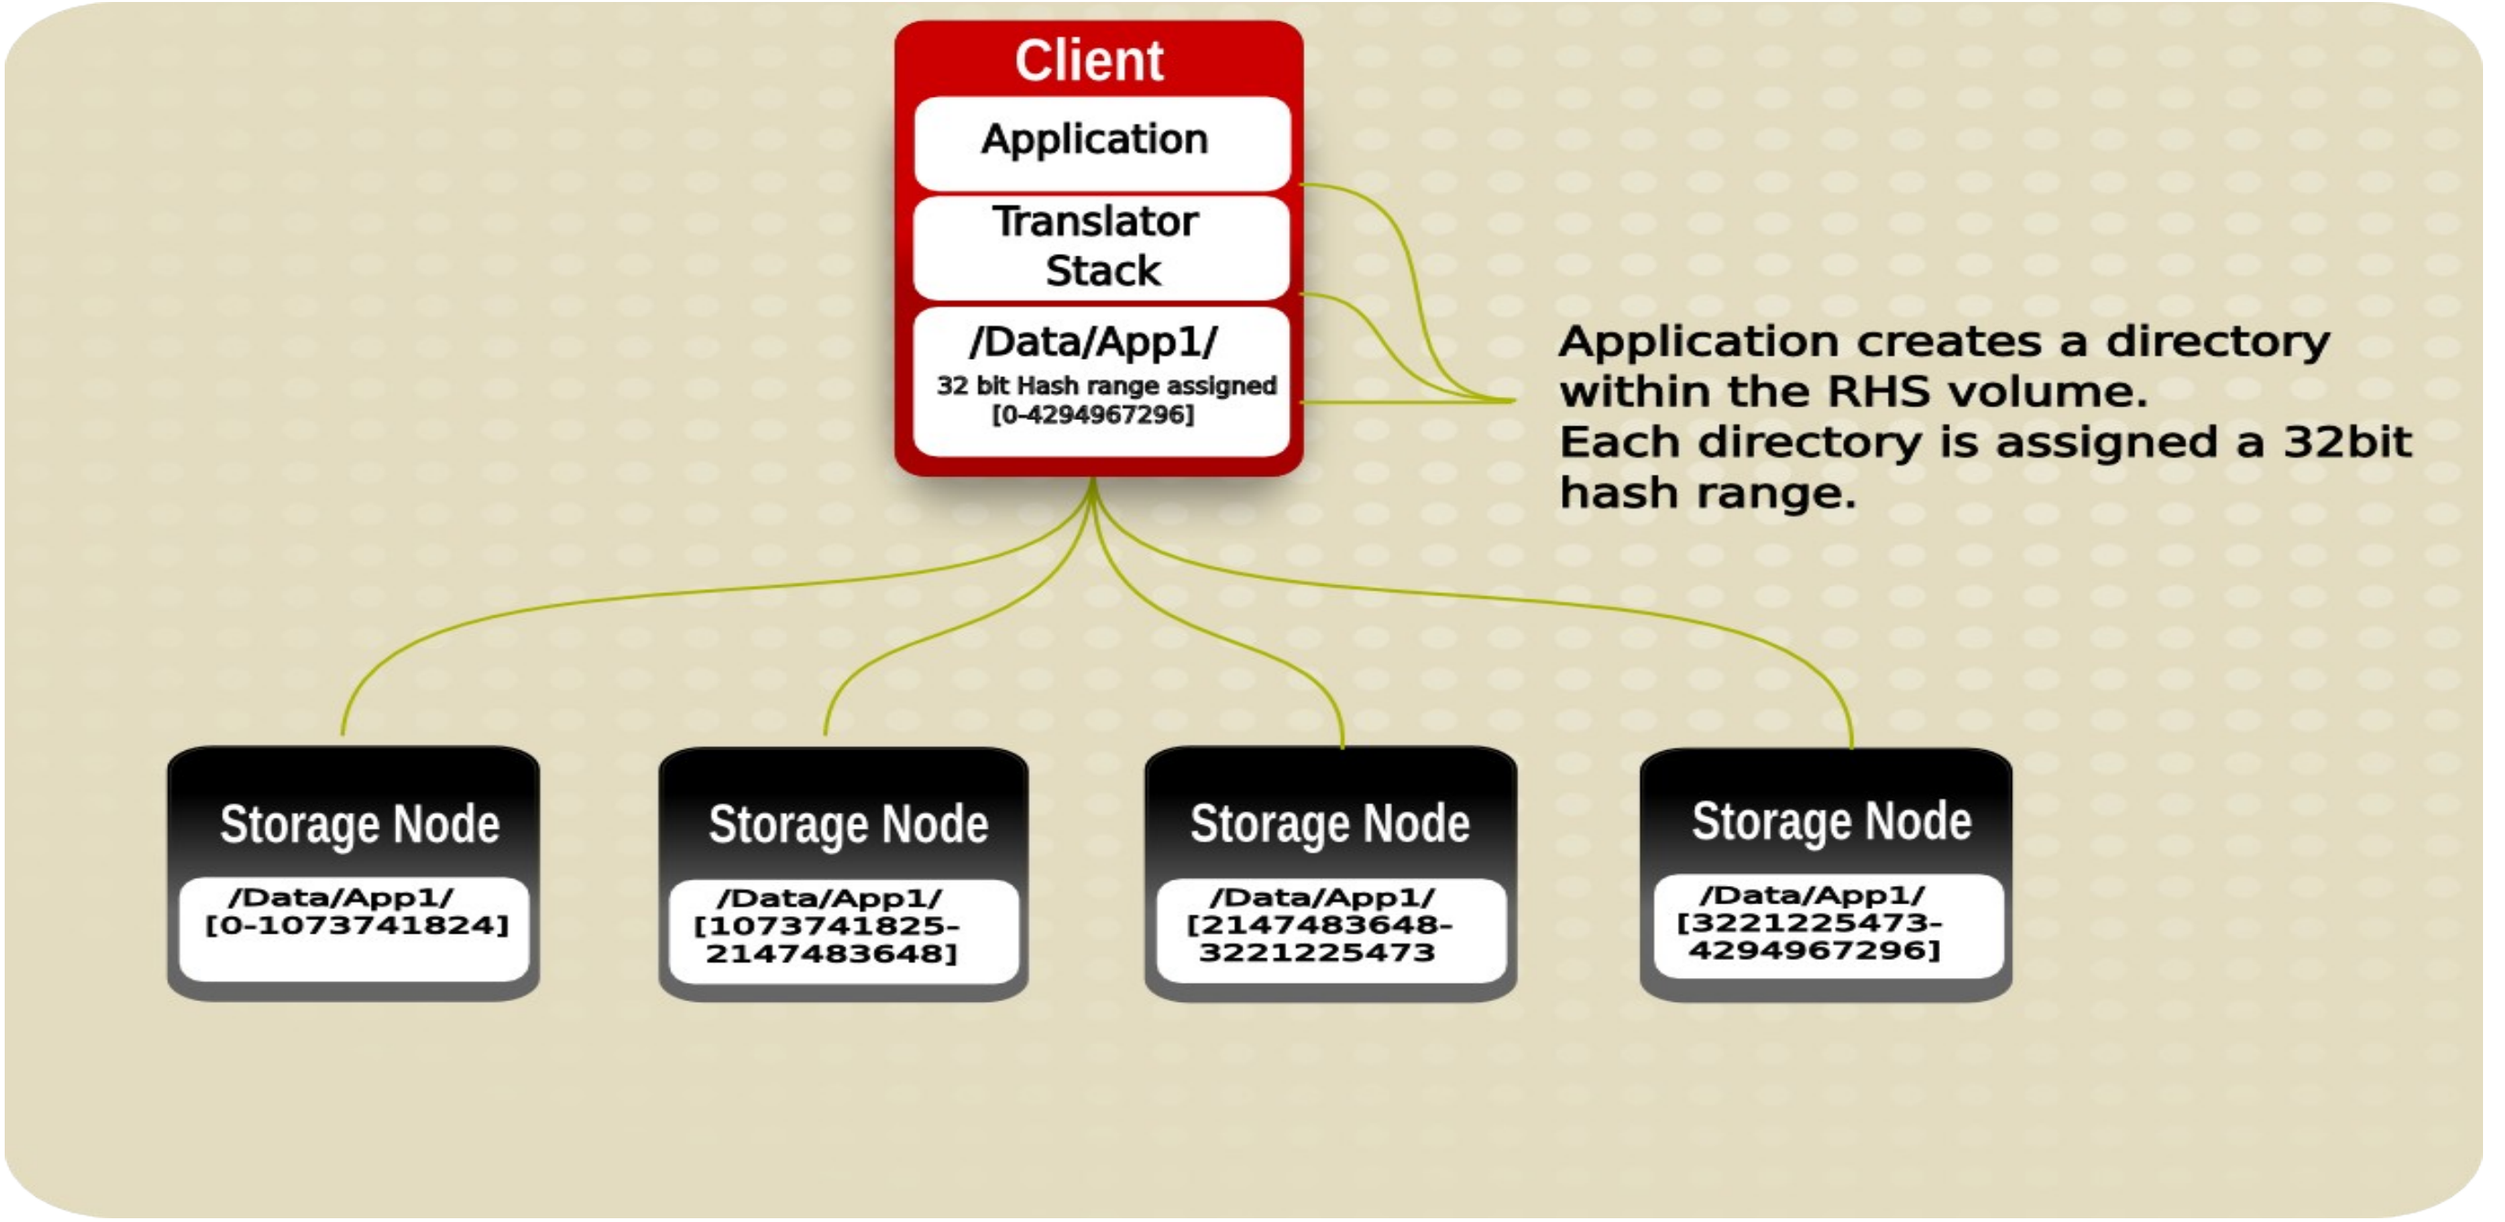
Client
Application
Translator Stack
/Data/App1/
32 bit Hash range assigned [0-4294967296]
Application creates a directory within the RHS volume.
Each directory is assigned a 32bit hash range.
Storage Node
/Data/App1/ [0-1073741824]
Storage Node
/Data/App1/ [1073741825-2147483648]
Storage Node
/Data/App1/ [2147483648-3221225473
Storage Node
/Data/App1/ [3221225473-4294967296]

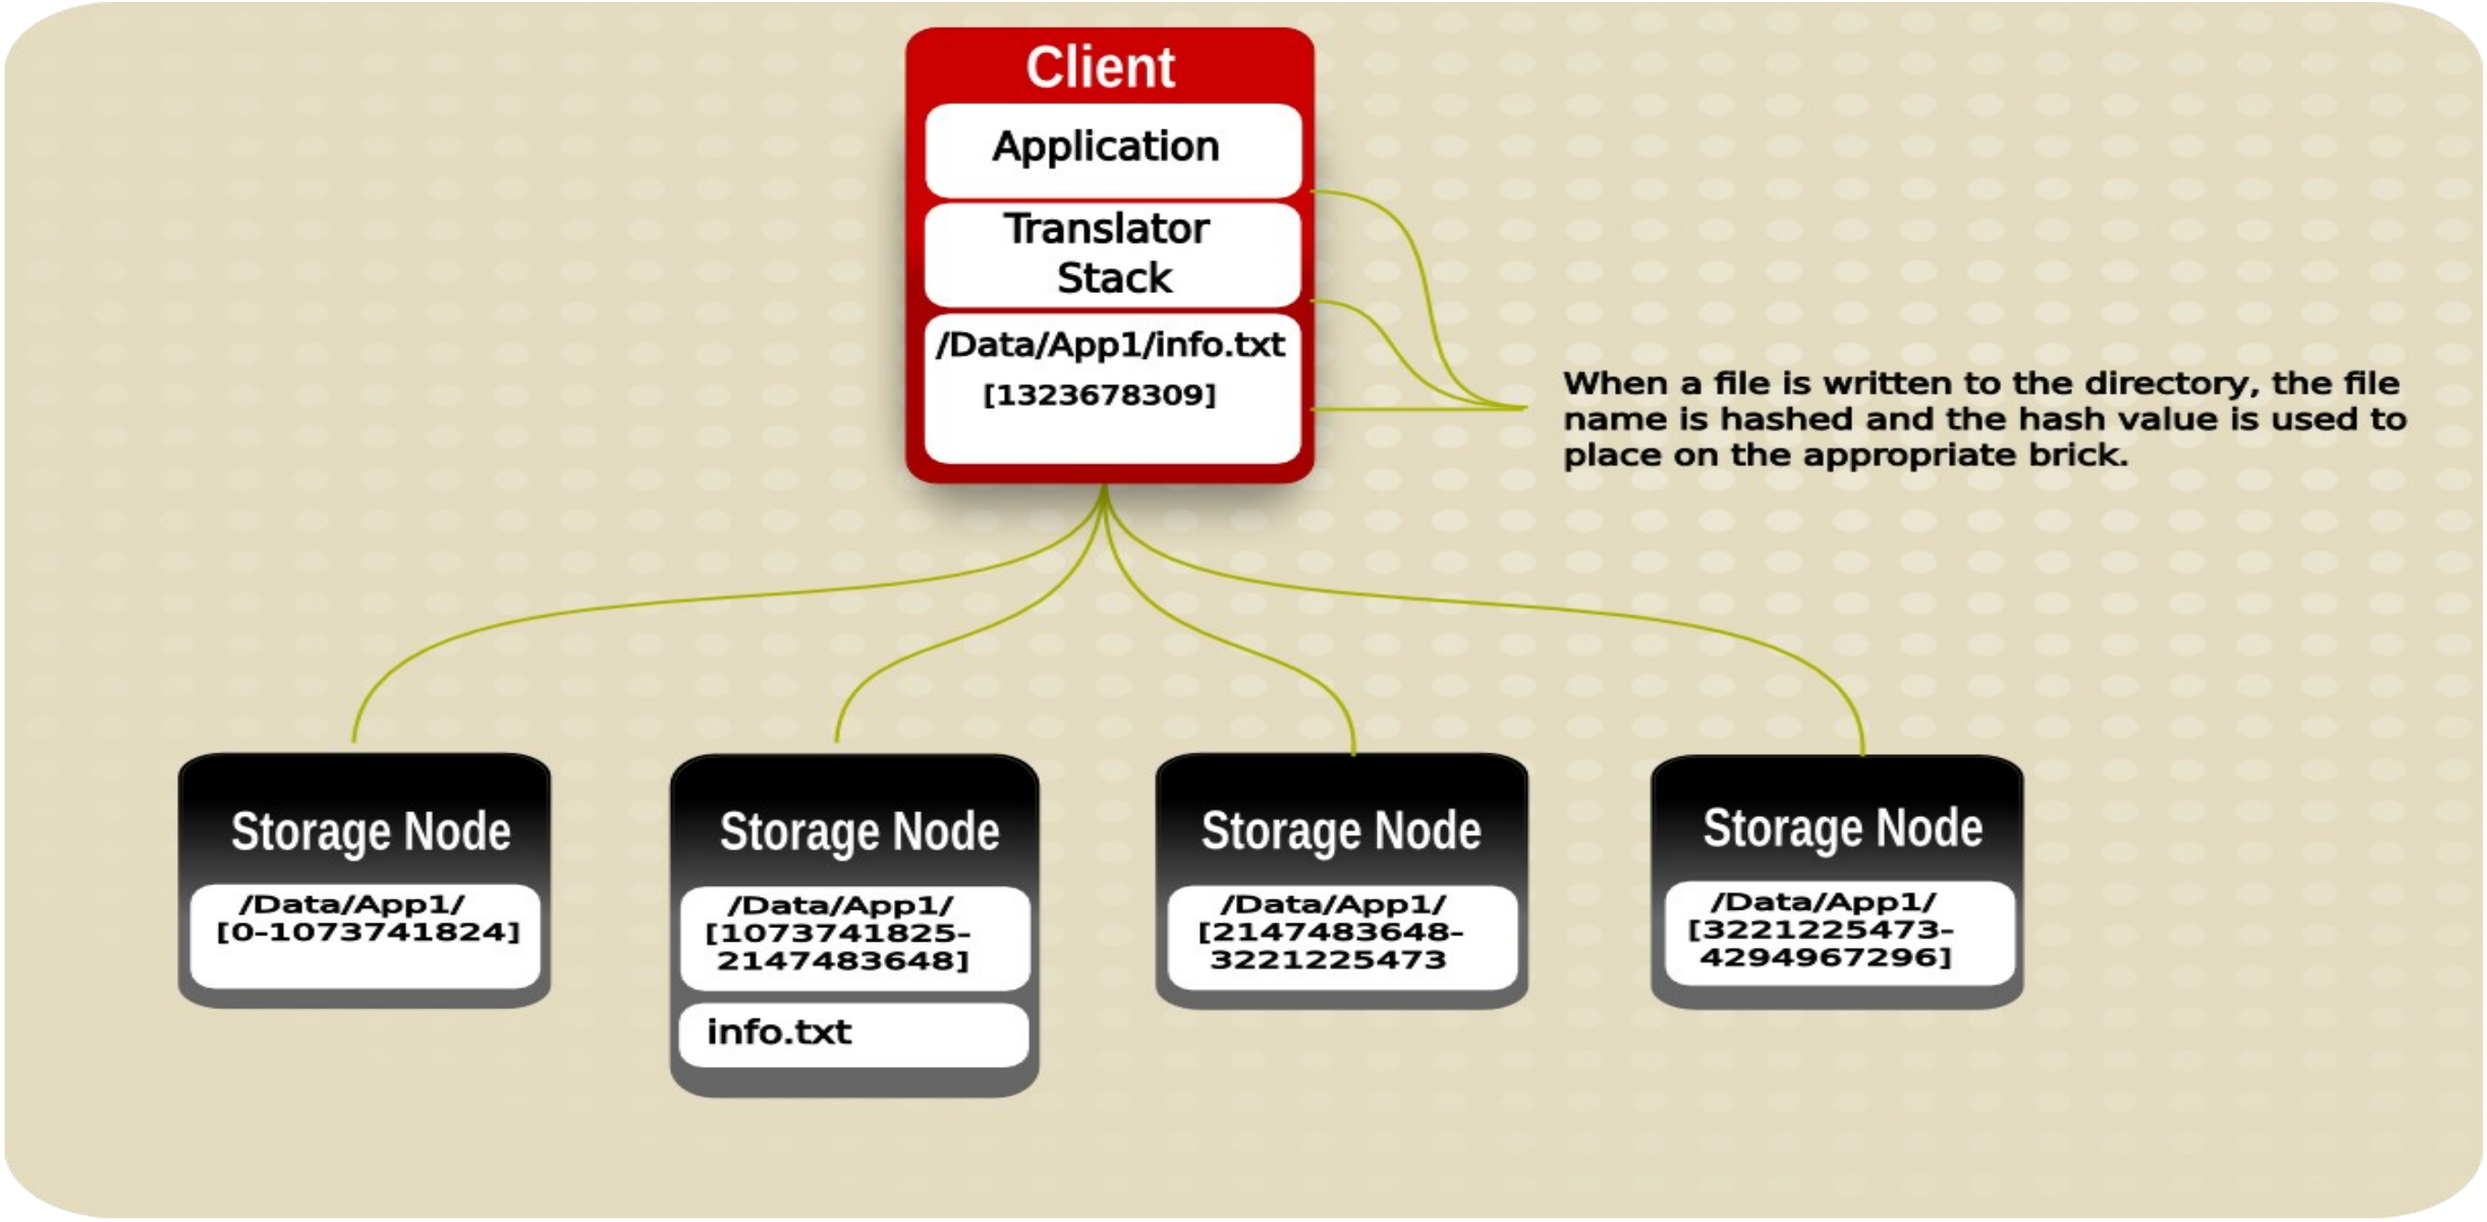
Client
Application
Translator Stack
/Data/App1/info.txt
[1323678309]
When a file is written to the directory, the file name is hashed and the hash value is used to place on the appropriate brick.
Storage Node
/Data/App1/
[0-1073741824]
Storage Node
/Data/App1/
[1073741825-
2147483648]
info.txt
Storage Node
/Data/App1/
[2147483648-
3221225473
Storage Node
/Data/App1/
[3221225473-
4294967296]

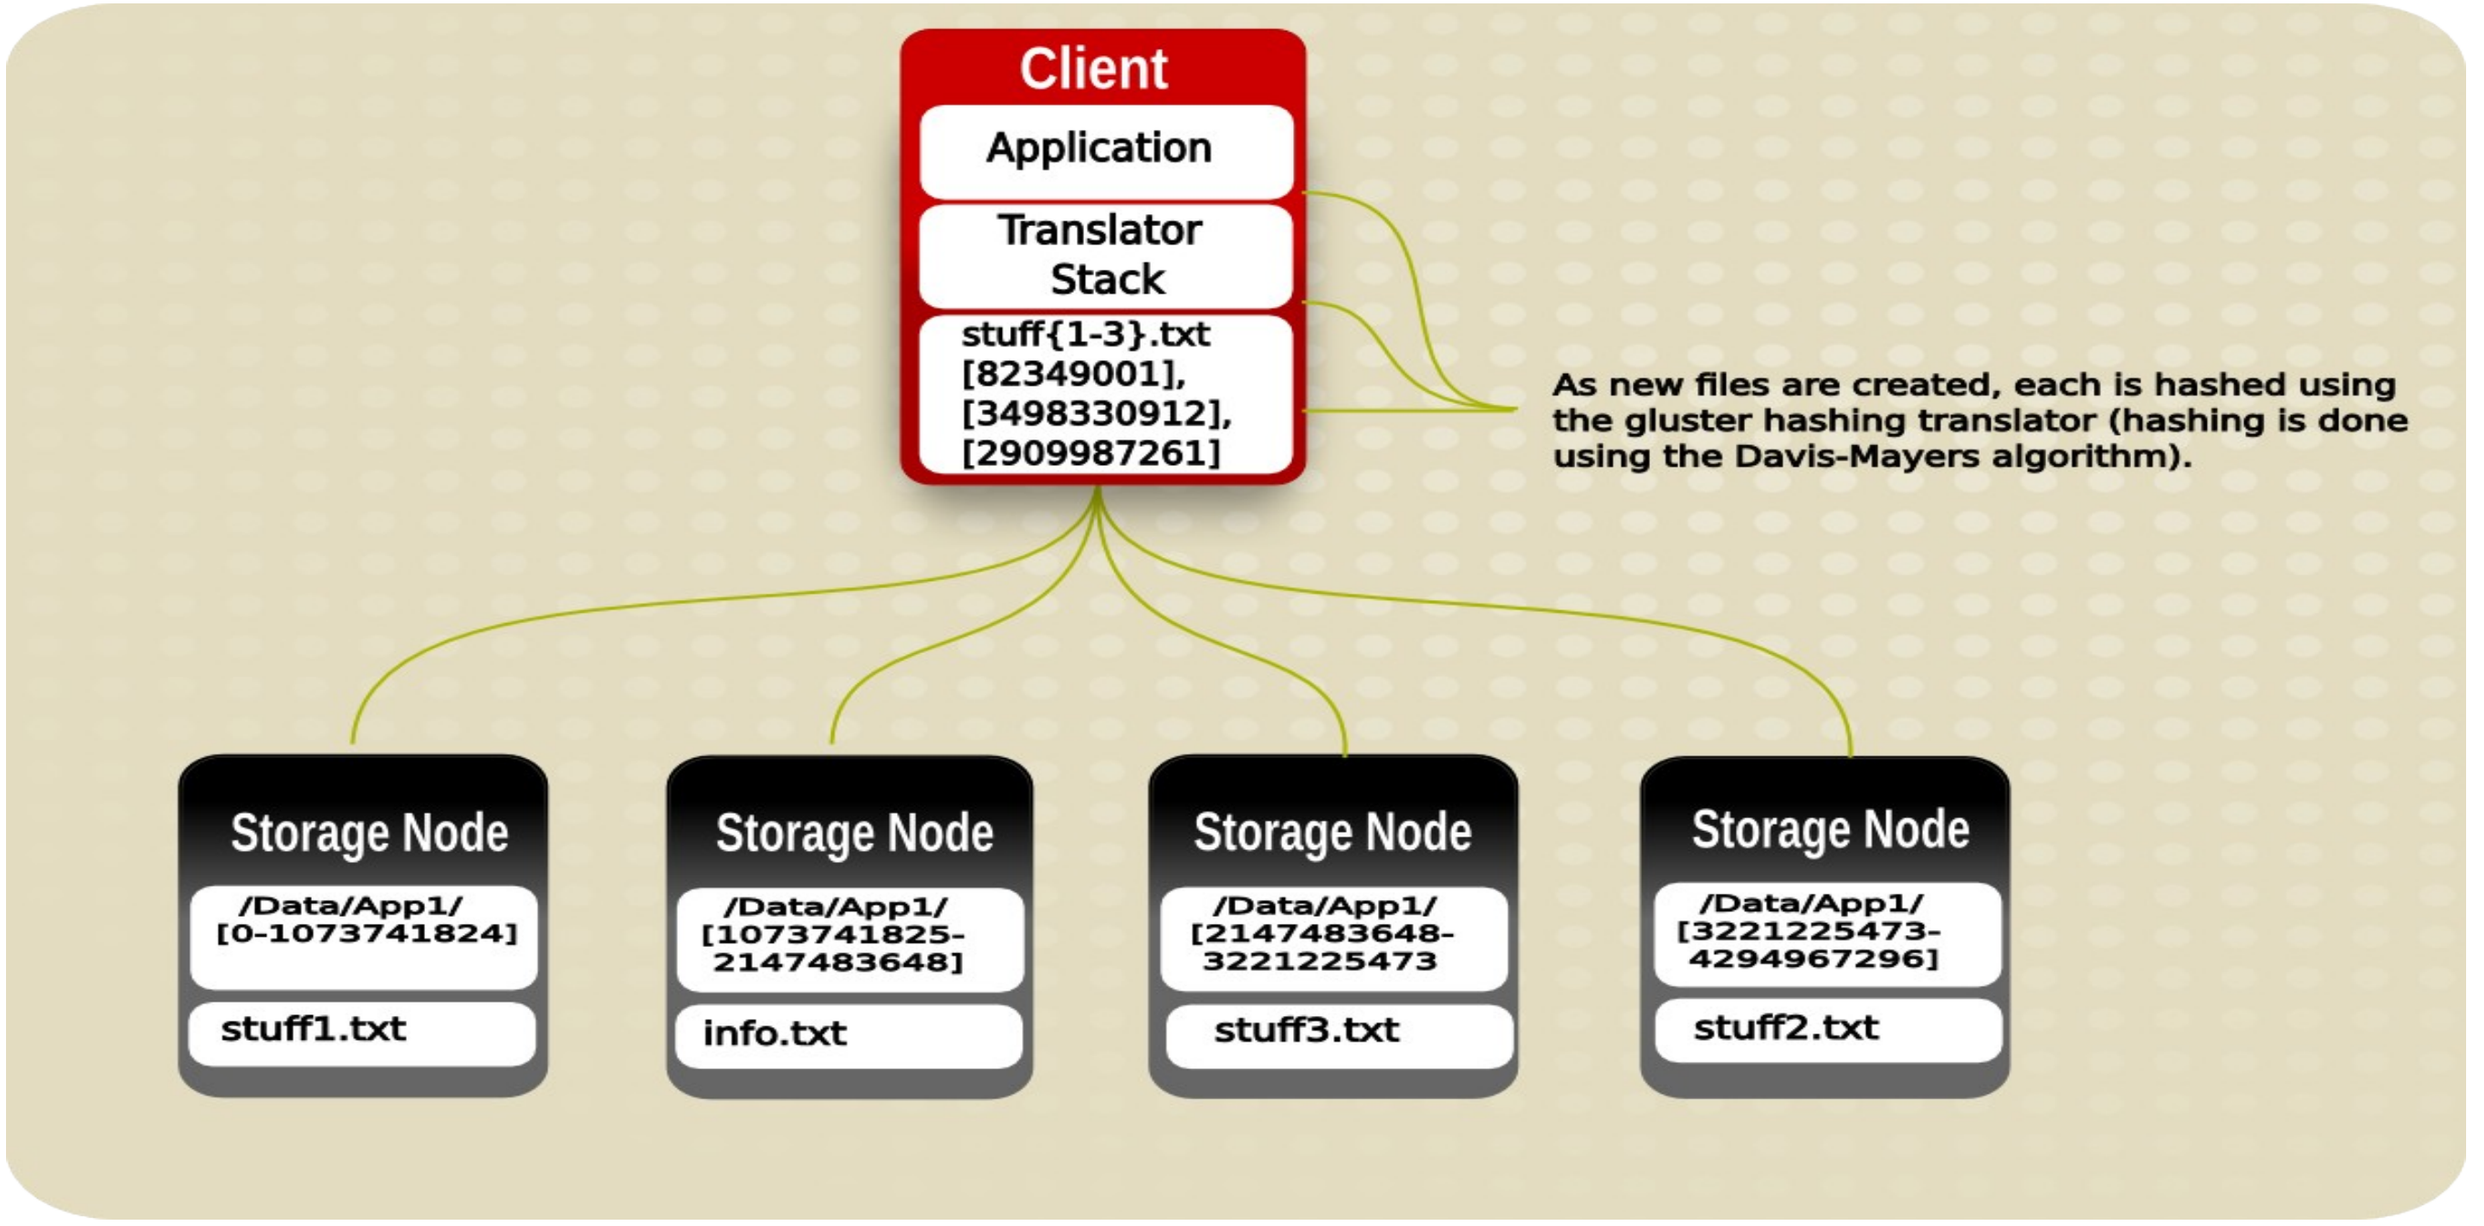
Client
Application
Translator Stack
stuff{1-3}.txt [82349001], [3498330912], [2909987261]
As new files are created, each is hashed using the gluster hashing translator (hashing is done using the Davis-Mayers algorithm).
Storage Node
/Data/App1/ [0-1073741824]
stuff1.txt
Storage Node
/Data/App1/ [1073741825-2147483648]
info.txt
Storage Node
/Data/App1/ [2147483648-3221225473
stuff3.txt
Storage Node
/Data/App1/ [3221225473-4294967296]
stuff2.txt

STACKABLE DESIGN–ELASTIC HASHING

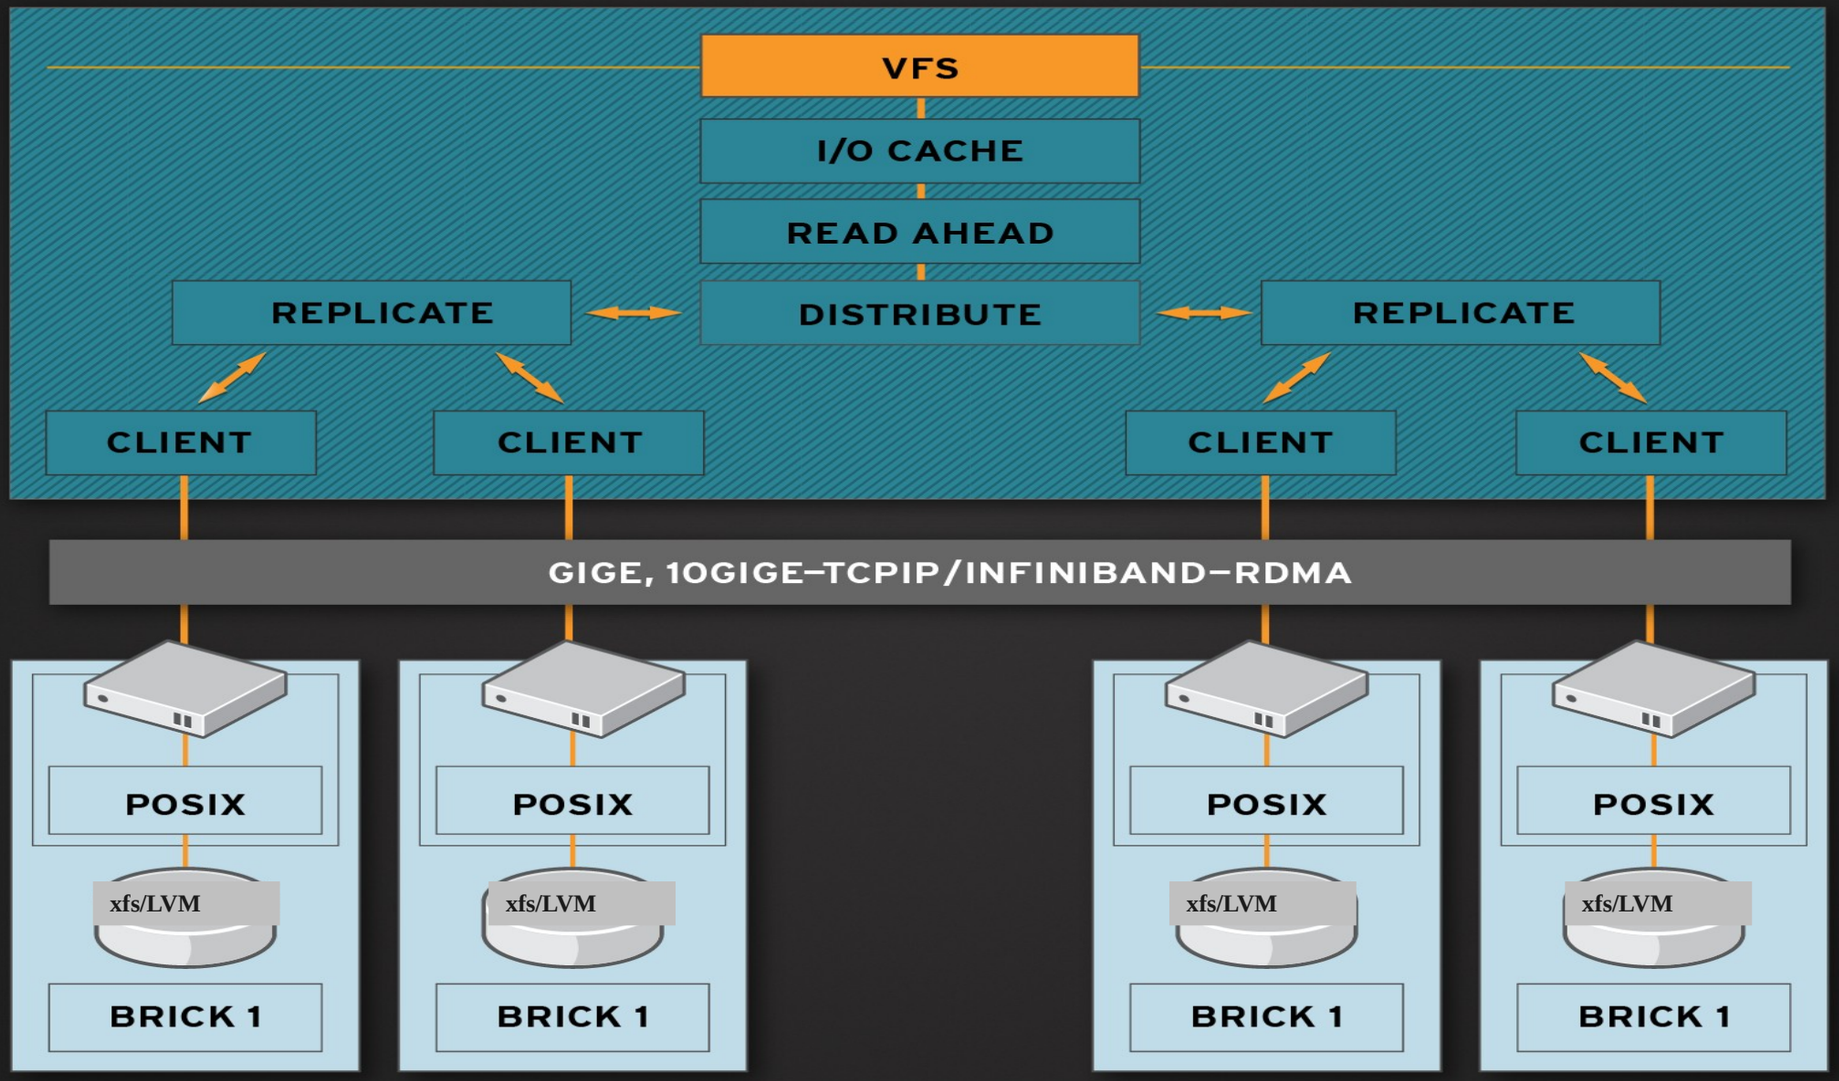

Bricks

- A brick is the combination of a node and a file system: hostname:/dir
- Each brick inherits limits of the underlying filesystem(xfs)
- RHS operates at the brick level, not at the node level
- Ideally, each brick in a cluster should be the same size

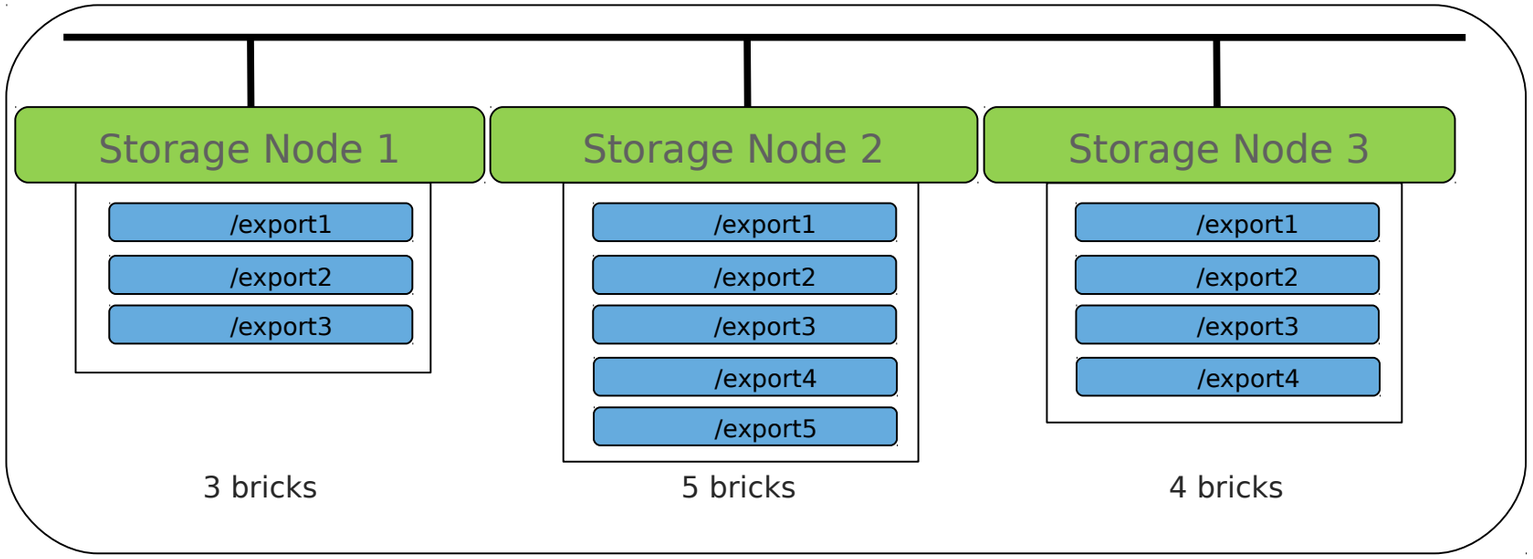

Volumes

- A volume consists of 1 or more bricks => exported with Gluster.
  - volumes have administrator assigned export names
  - a brick is a member of only one volume
- A namespace can have 1 or more volumes
  - A namespace can consist of replicated and distributed volumes
  - data in different volumes physically exists on different bricks
  - volumes can be mounted on clients using NFS, CIFS and/or GlusterFS clients (native FUSE client)

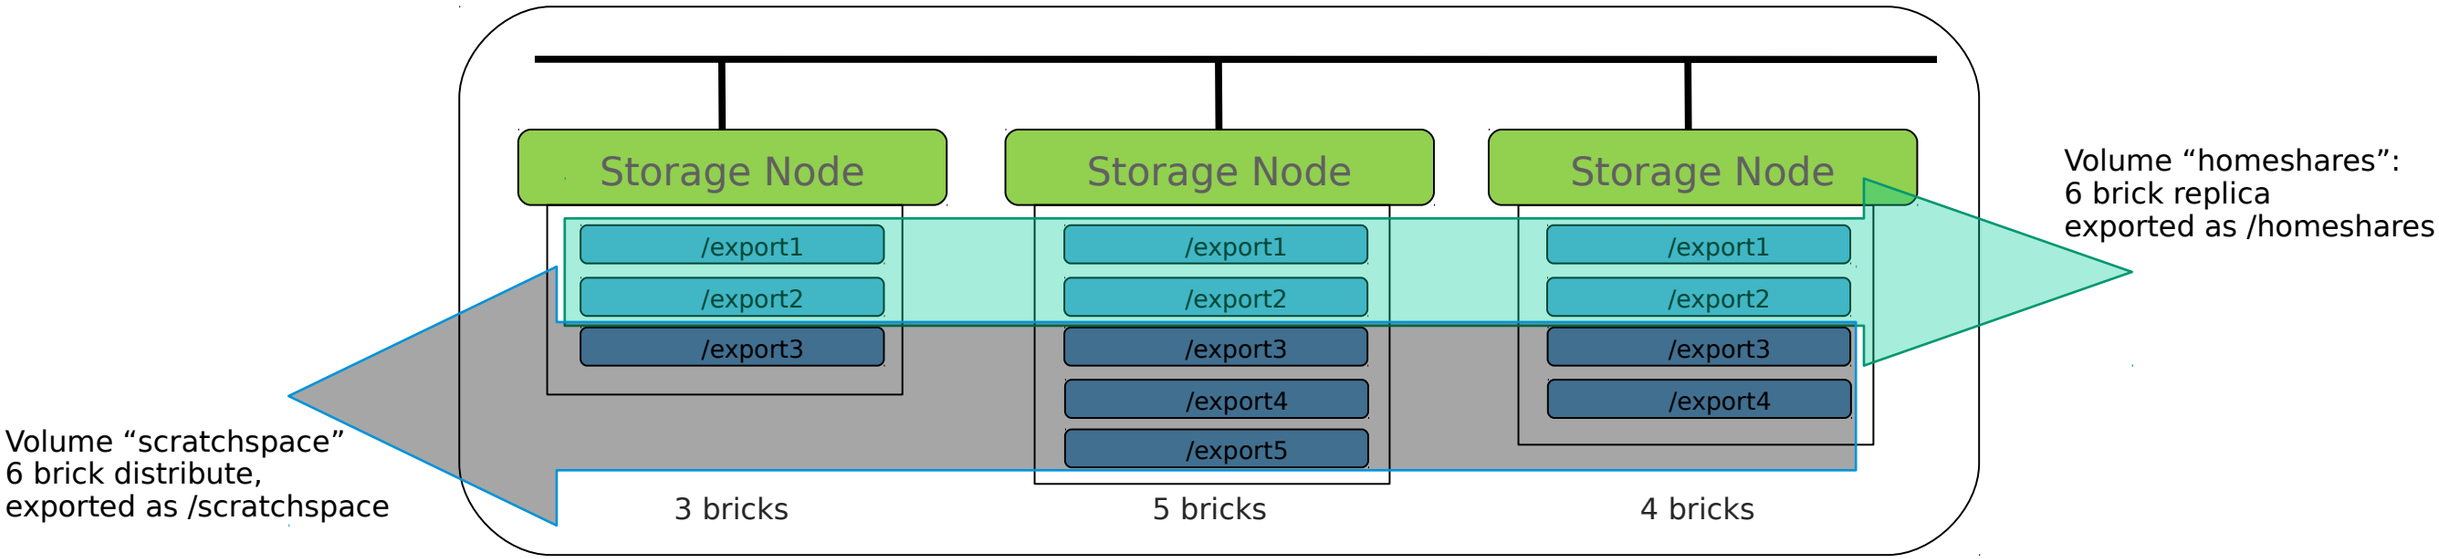

Data Flow with NFS/CIFS Client

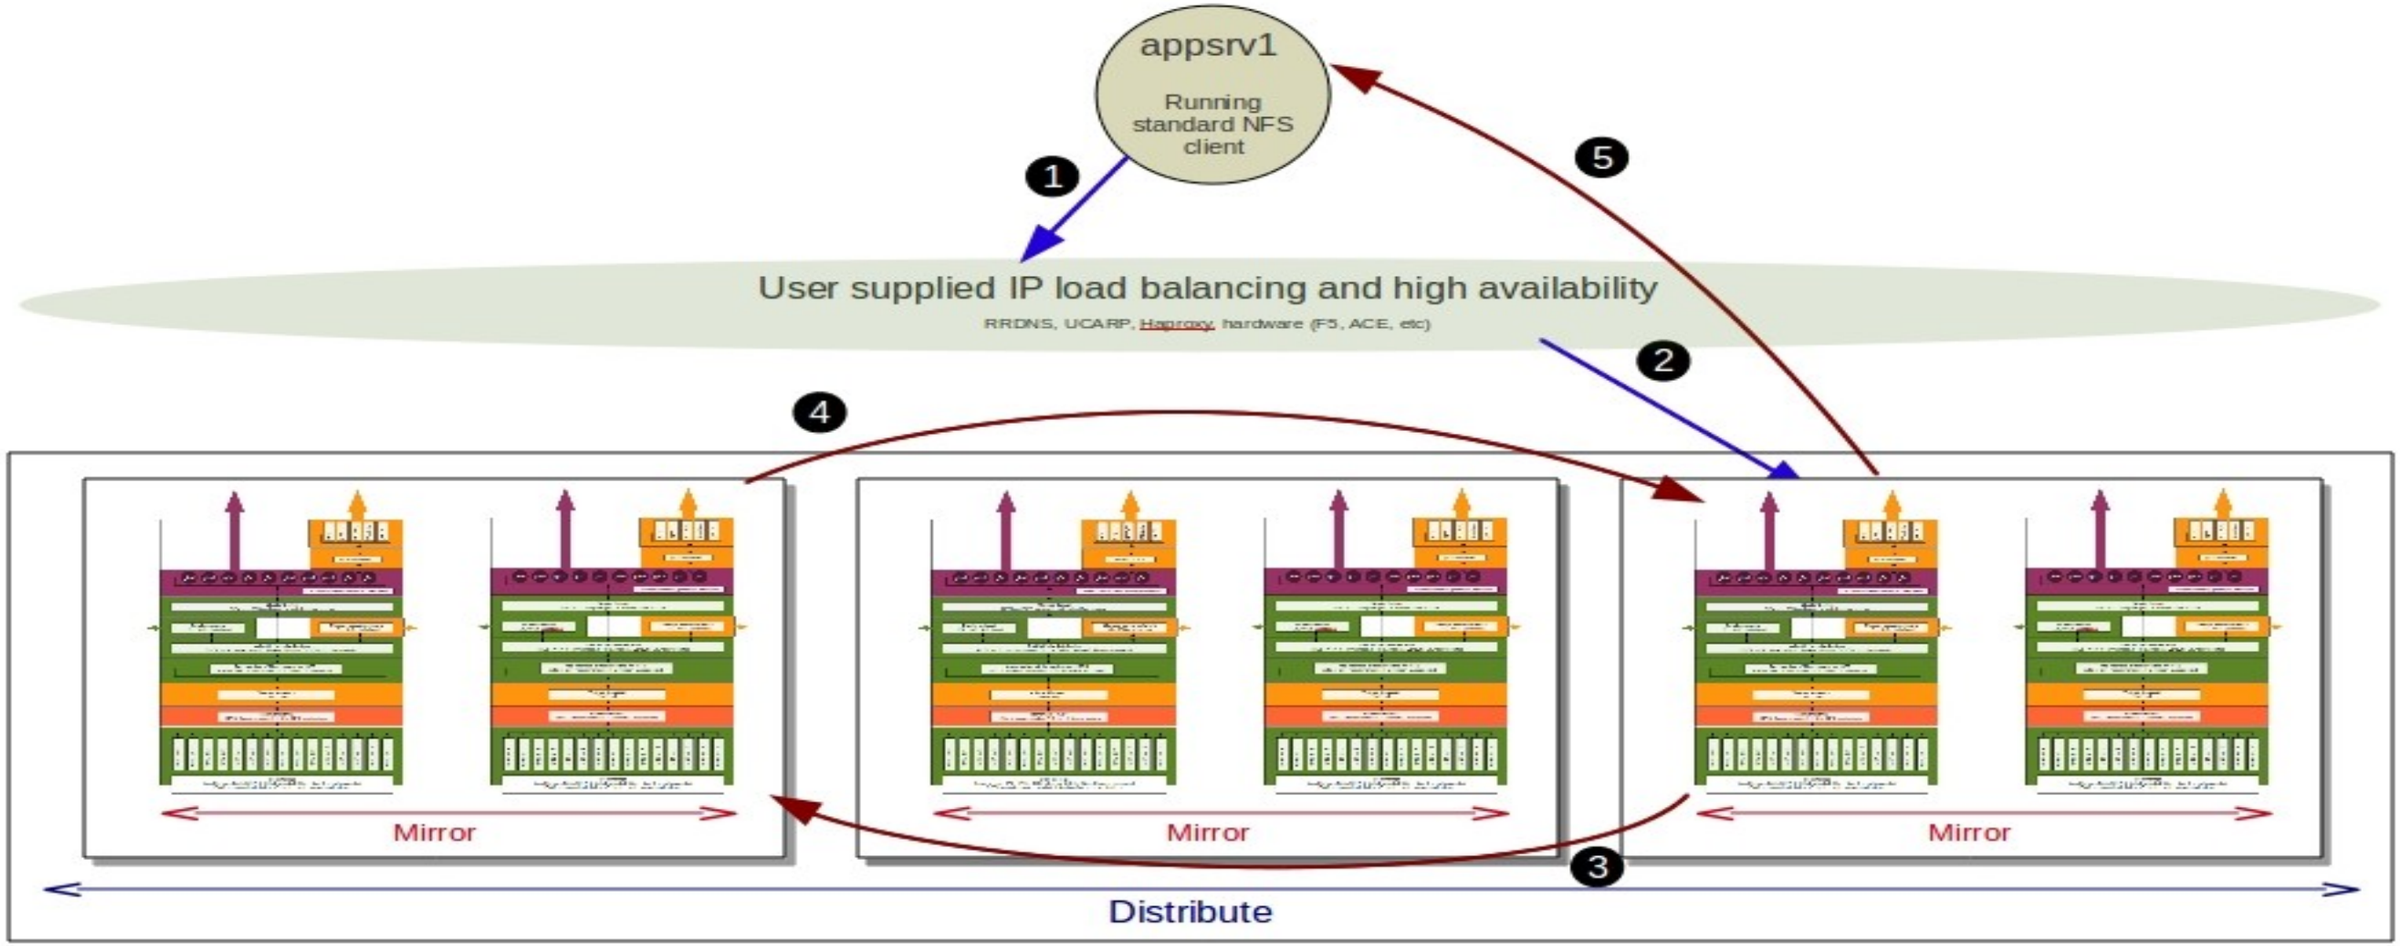

Data Flow with Native Client

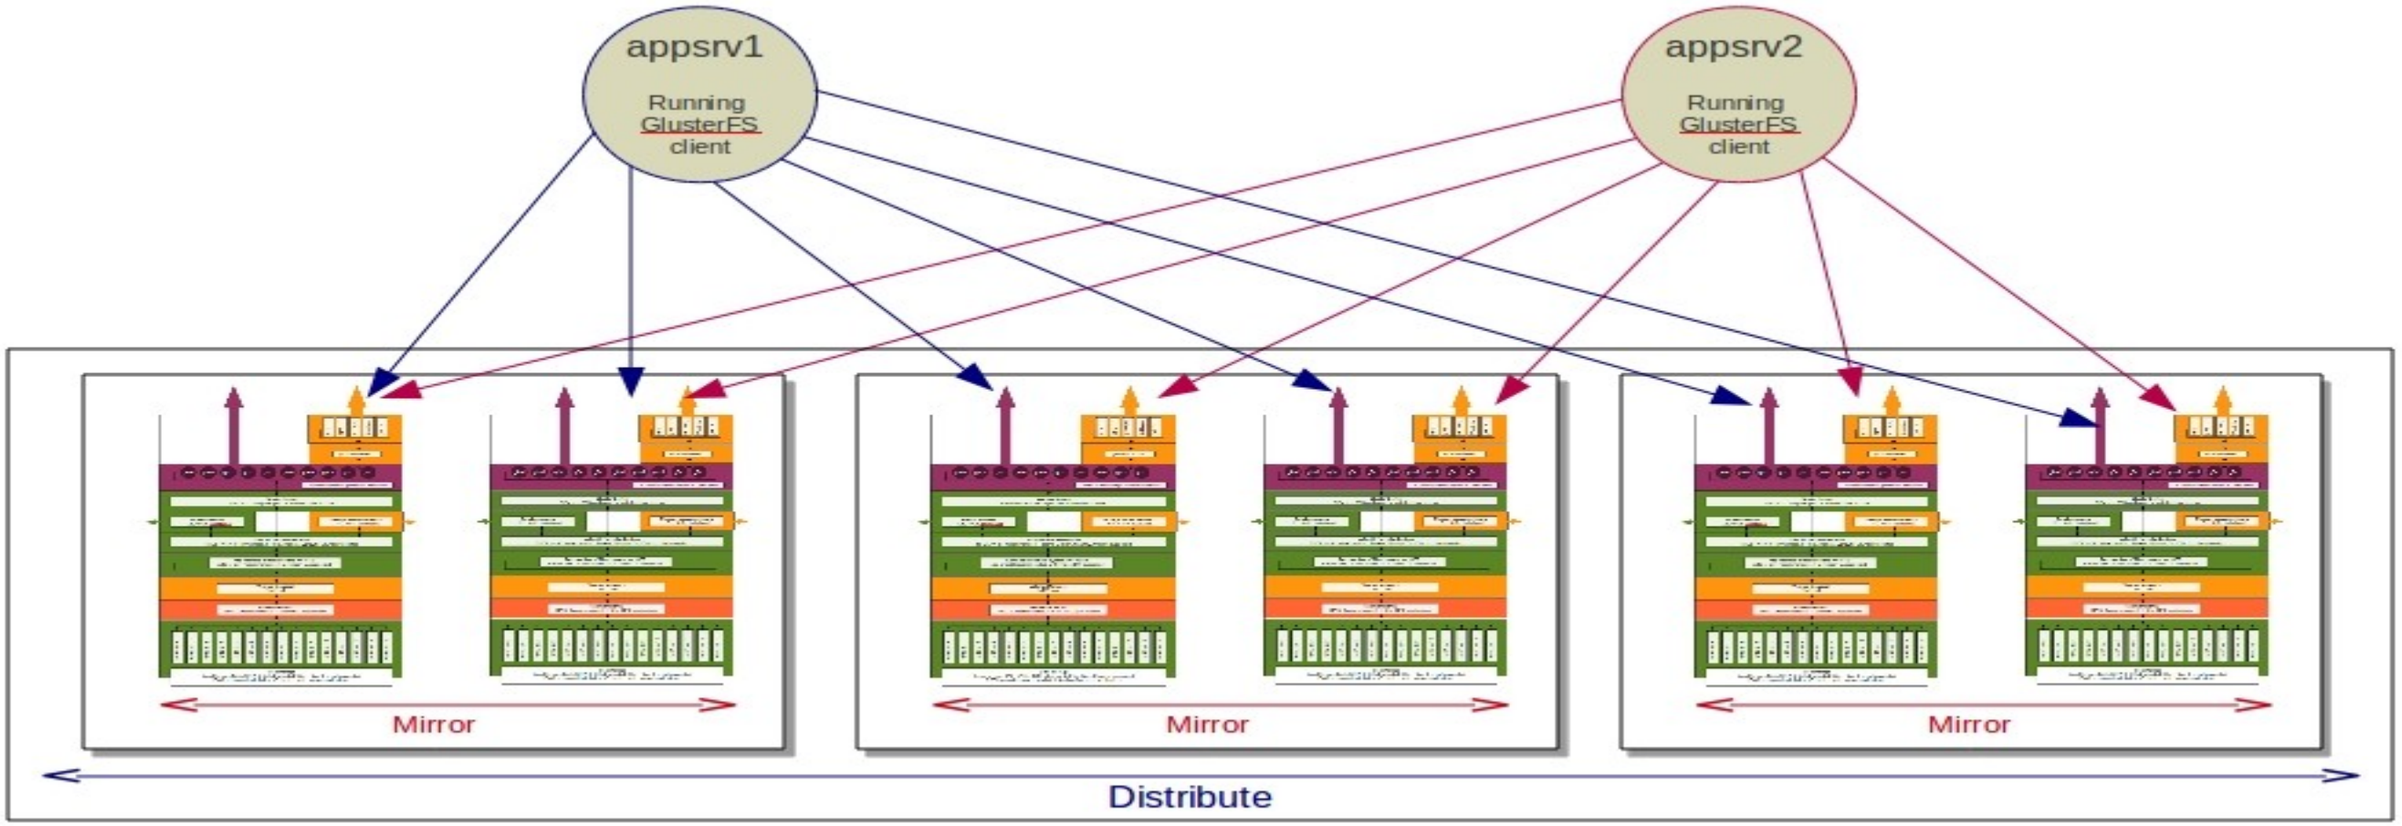

HOW DOES REPLICATION ACTUALLY WORK?

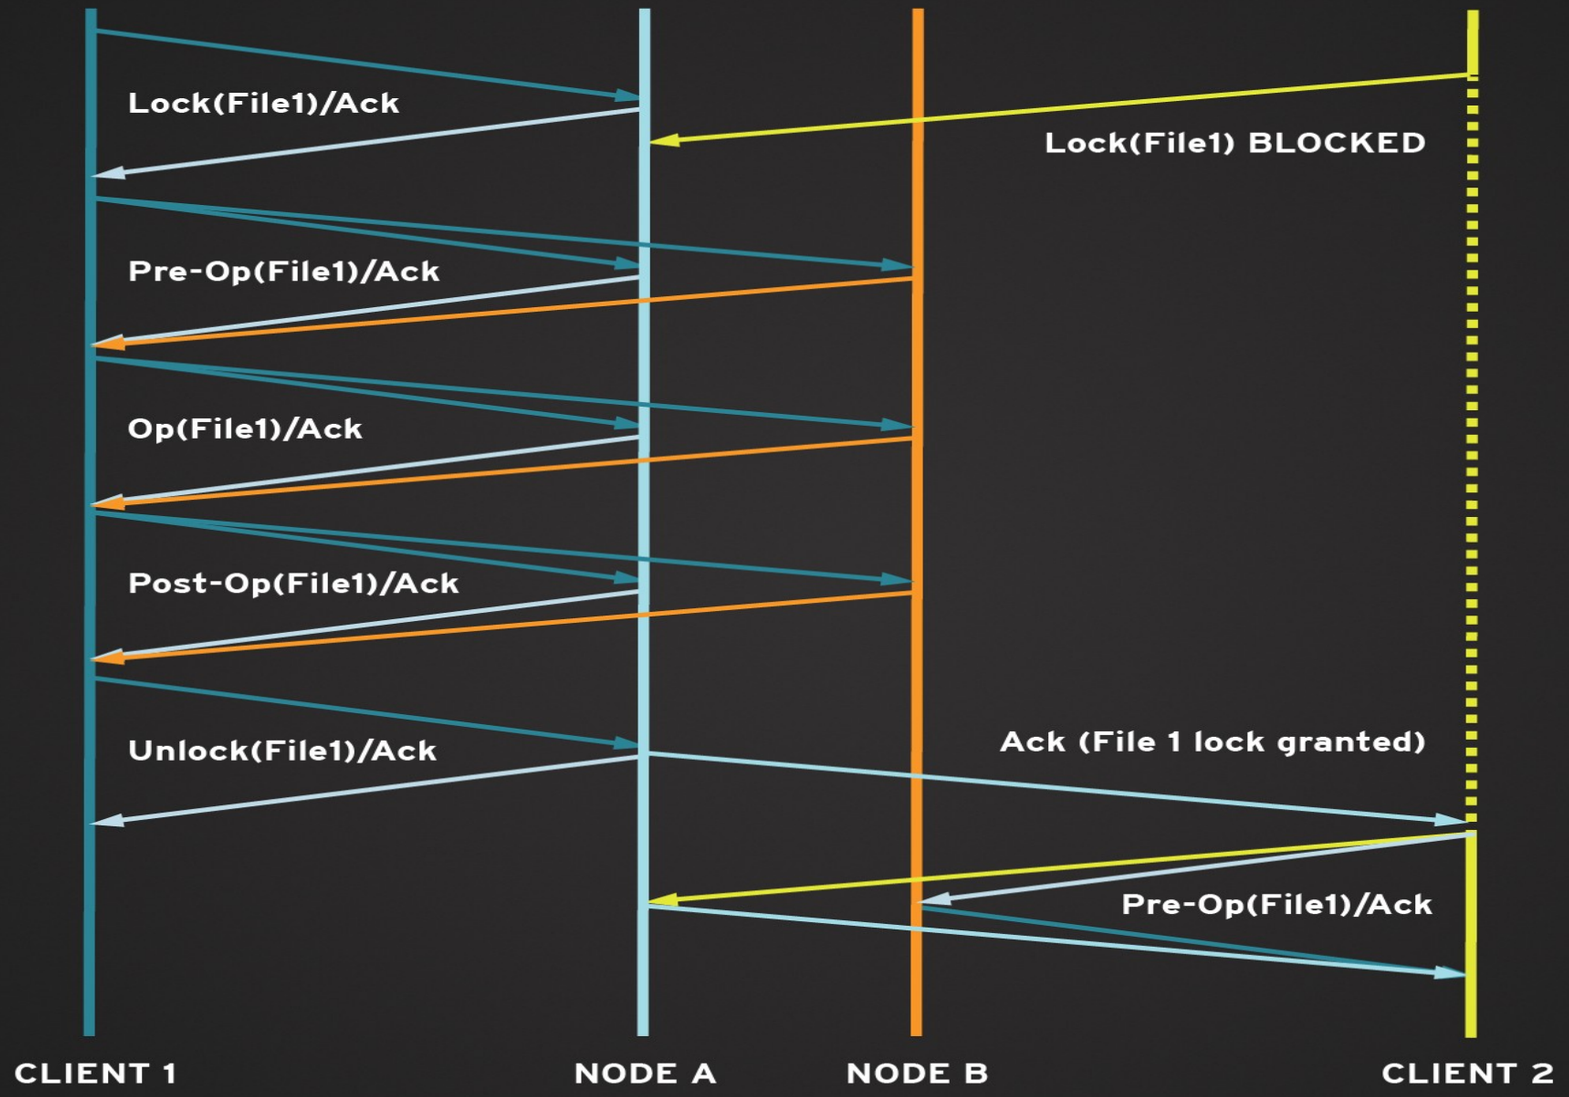

Seamless Integration for Hadoop Deployments

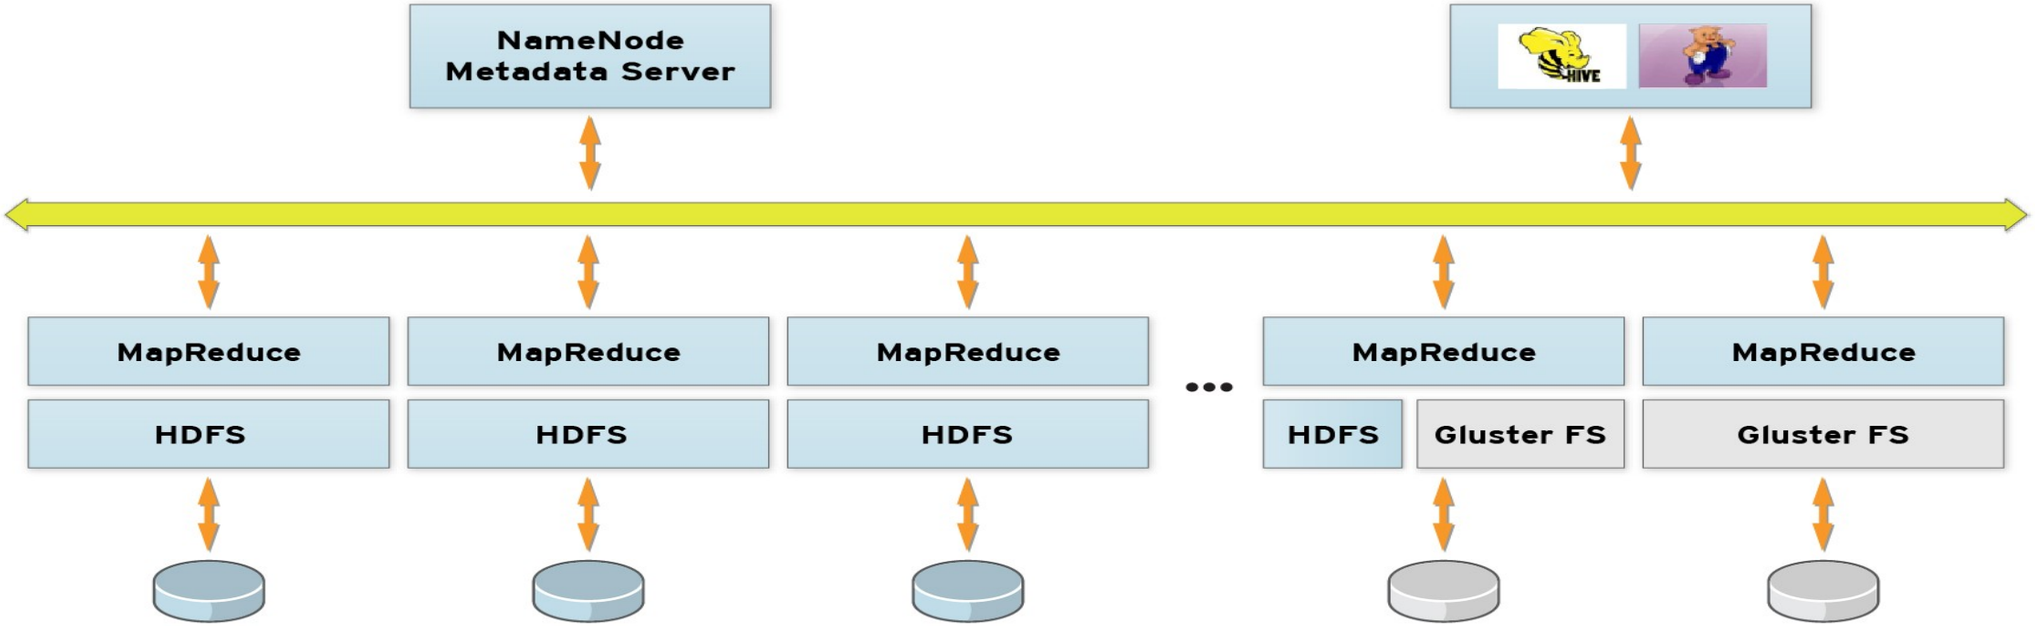

- GlusterFS can co-exist HDFS
- Does not use the NameNode metadata server

- Built using the Hadoop file system API
- Requires simple configuration file changes
- C Lib GlusterFS client enable GlusterFS direct access
- Provides Java binding for Hadoop compatibility

Hadoop architecture overview

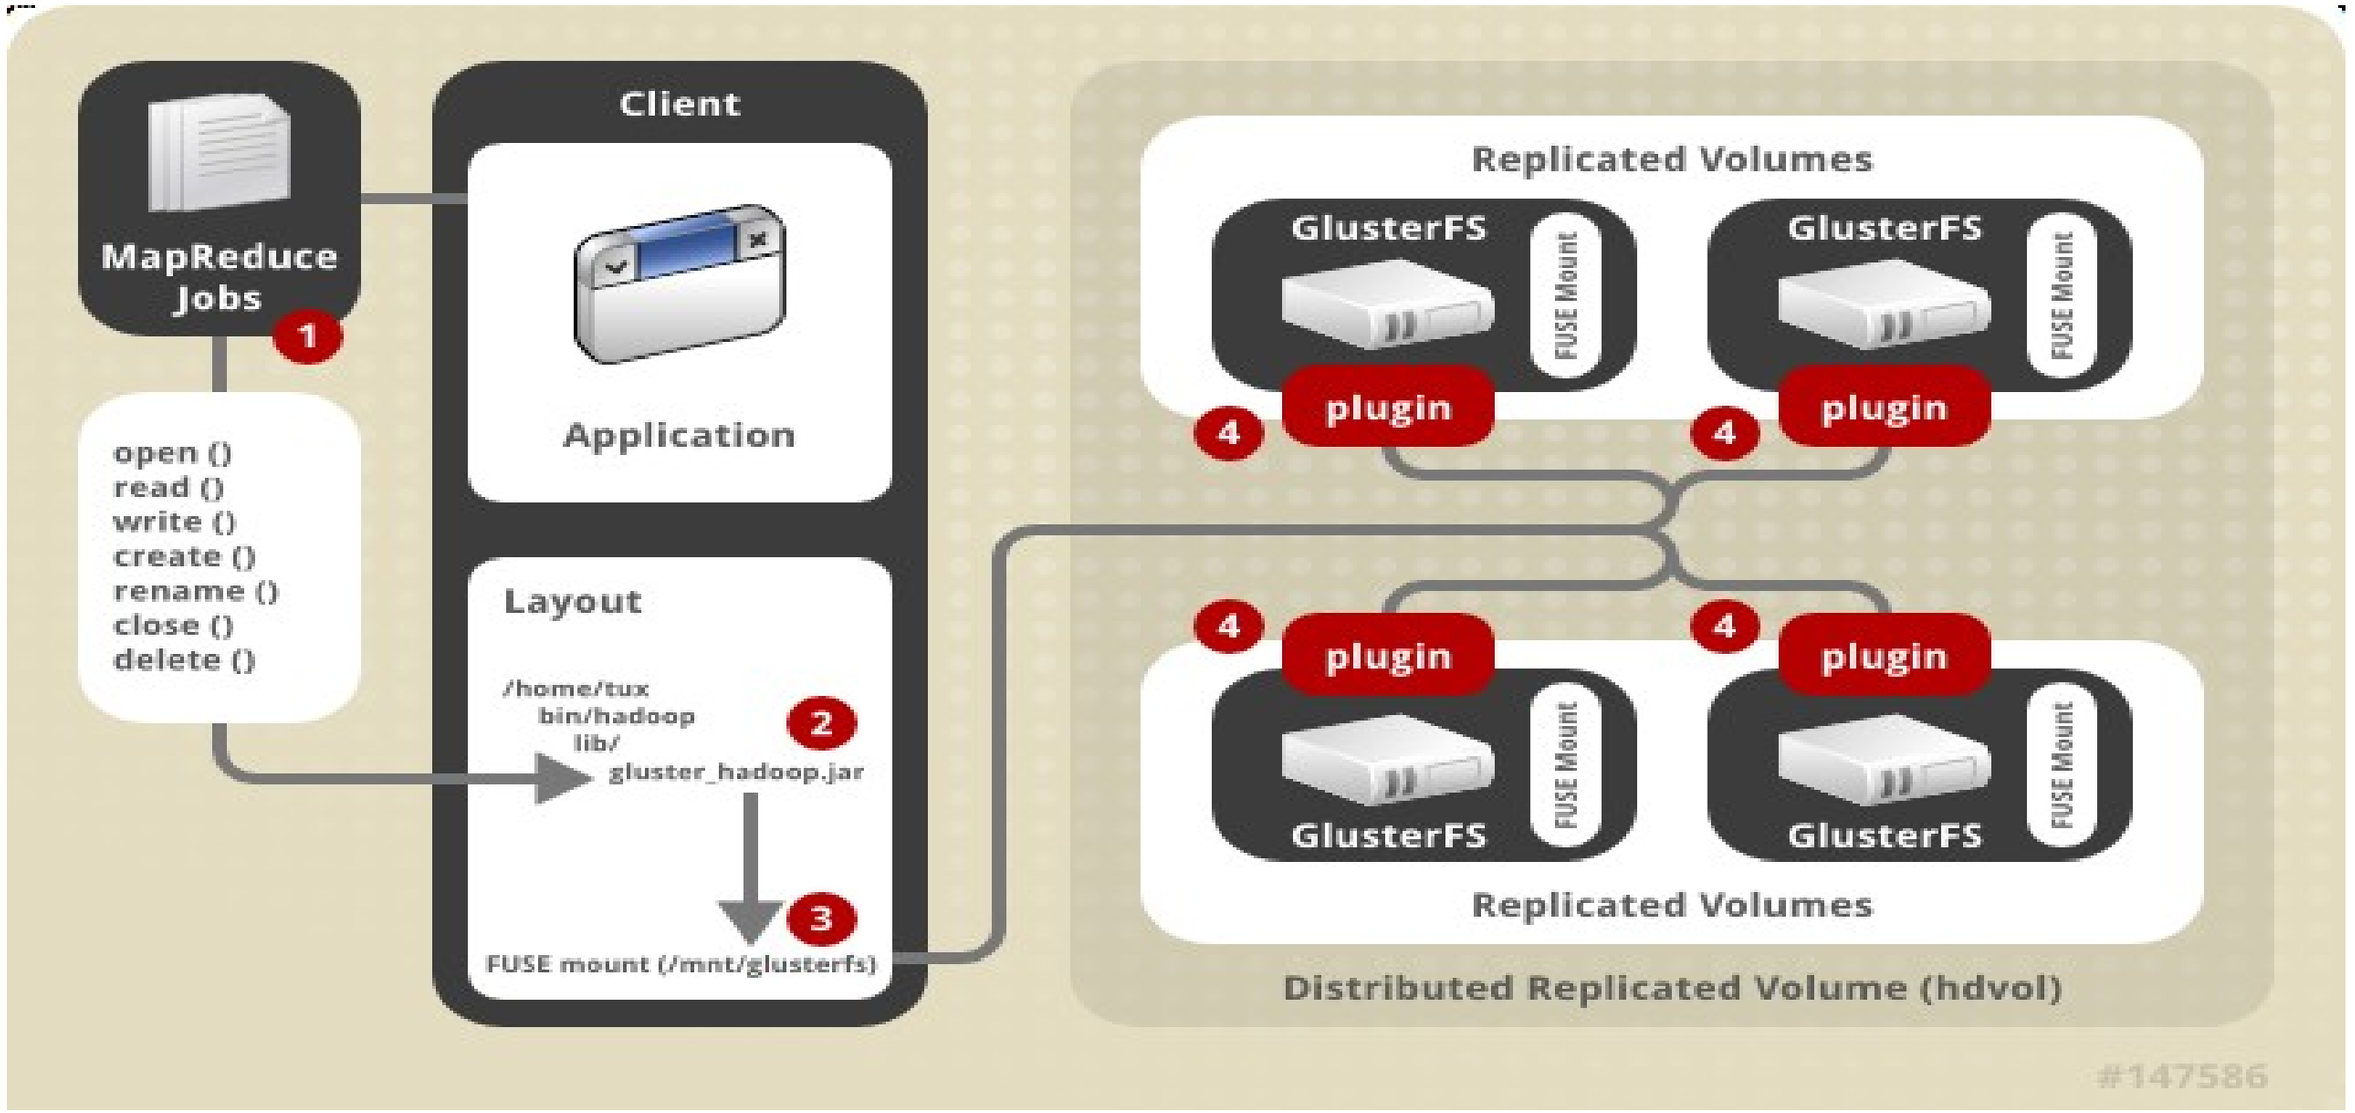

MANAGING UNSTRUCTURED DATA AT PETABYTE-SCALE

Joachim Schröder
Manager Solution Architects, DACH
Email: joachim.schroeder@redhat.com
November, 14th 2013

WHO'S RED HAT? - RED HAT PORTFOLIO

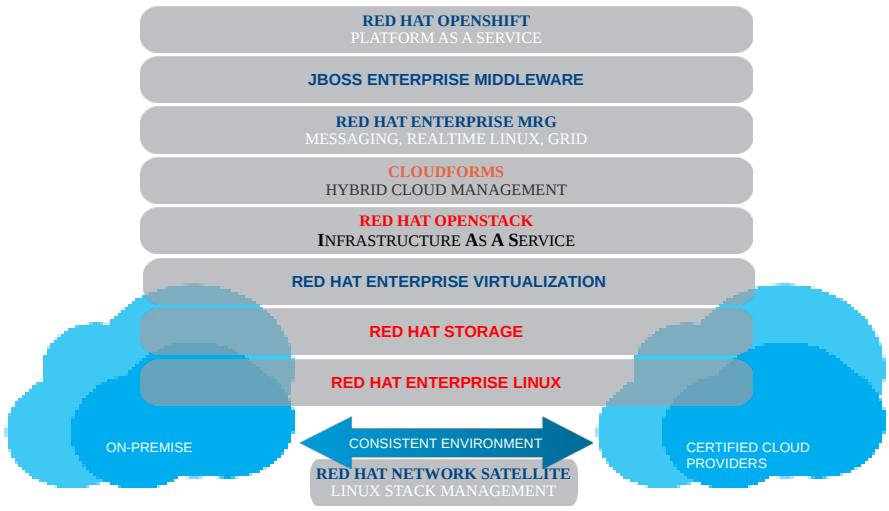

THE INFORMATION EXPLOSION

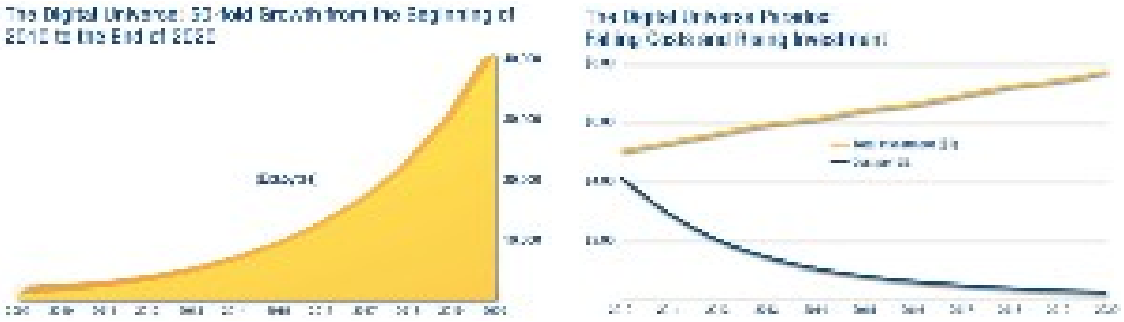

Source: IDC's Digital Universe Study, Dec 2012

Main growth drivers:
Virtualisation, Cloud, Mobile Computing and Big Data

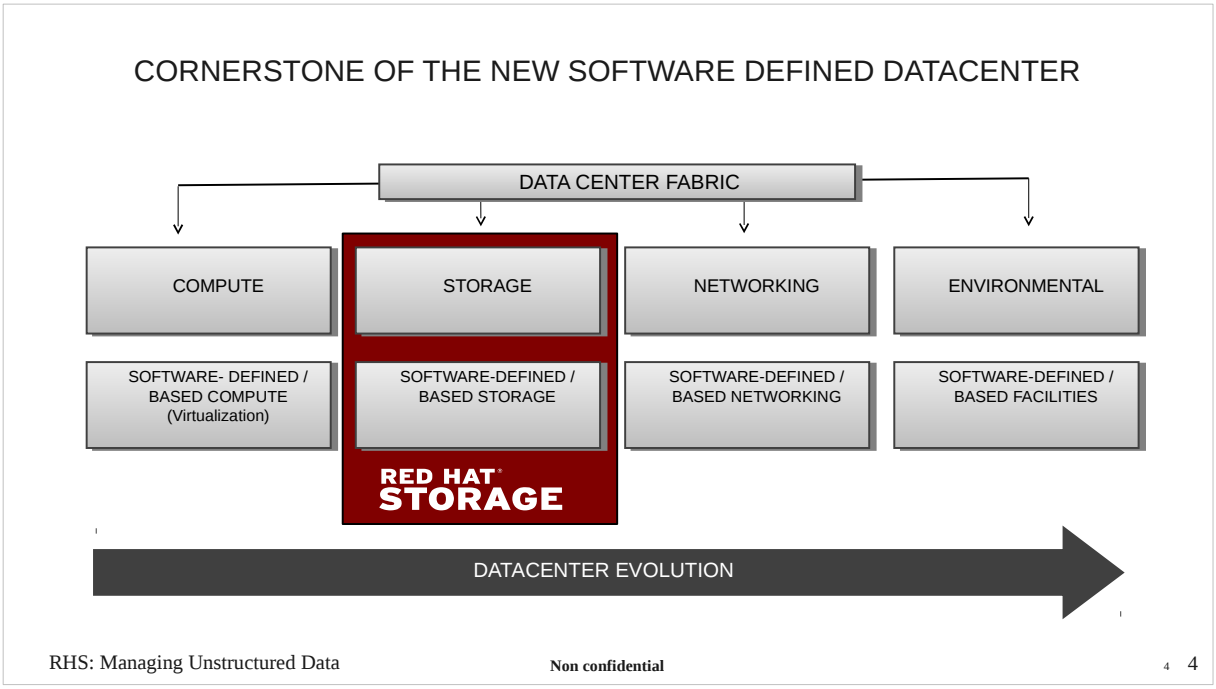

Todays Modern Data Center is Increasingly defined by and based on software.

Compute was the first – with virtualization – to begin to abstract data center resources aiding with

WHAT IS RED HAT STORAGE?

OpenSource

Scale-out NAS (Network Attached Storage)

deployable on

on-premise, virtualized and Cloud environments

based on GlusterFS

running on standard x86 Hardware

S
L
I
D
E

requiremnts fot he Big data challange:
available storage solution that can handle hardware failure.
ndards and possiblility to be replicated and access over geographical distance.
eal if a component would fail to maintan protection level.
gement or minimal manual management would be prefered.
ostic, so run your solution private, public or replicate in between, don’t be locked in
igrations when you lifecylcle hardware. Migrations might not be an option if you h

RED HAT STORAGE DEPLOYMENT ON-PREMISE

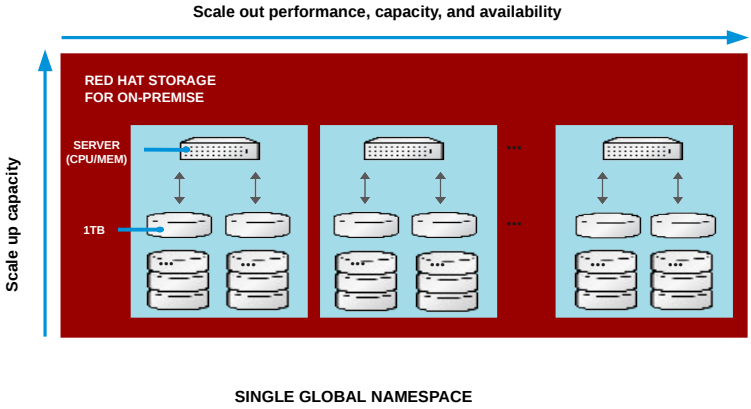

SINGLE GLOBAL NAMESPACE

- Single 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

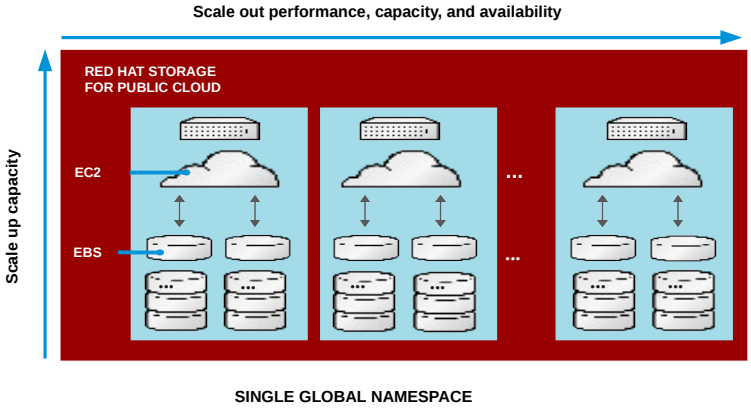

- GlusterFS Amazon Machine Images (AMIs)
- The only way to achieve high availability of Elastic Block Storage (EBS)
- Multiple EBS devices pooled
- POSIX compatible (no application to rewrite required to run on Amazon EC2)
- Scale out capacity and performance as needed

RED HAT STORAGE—50,000 FOOT OVERVIEW

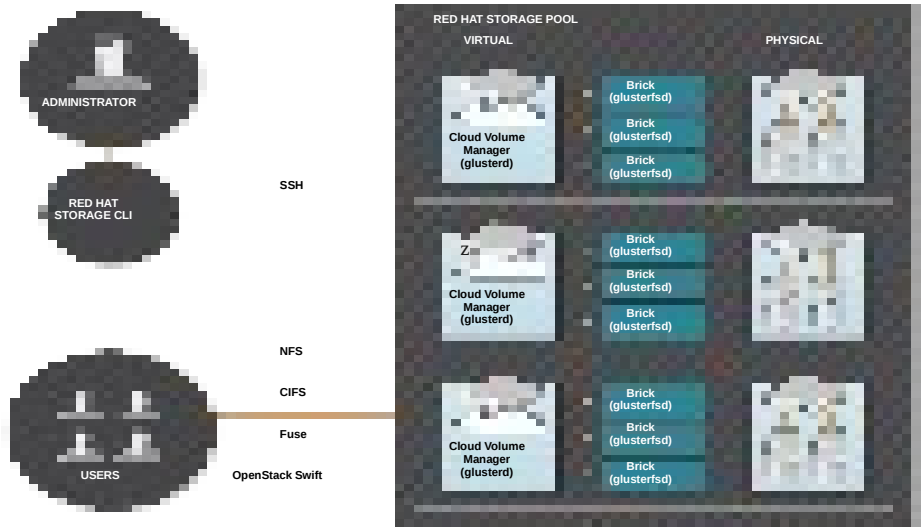

RED HAT STORAGE TECHNOLOGY STACK

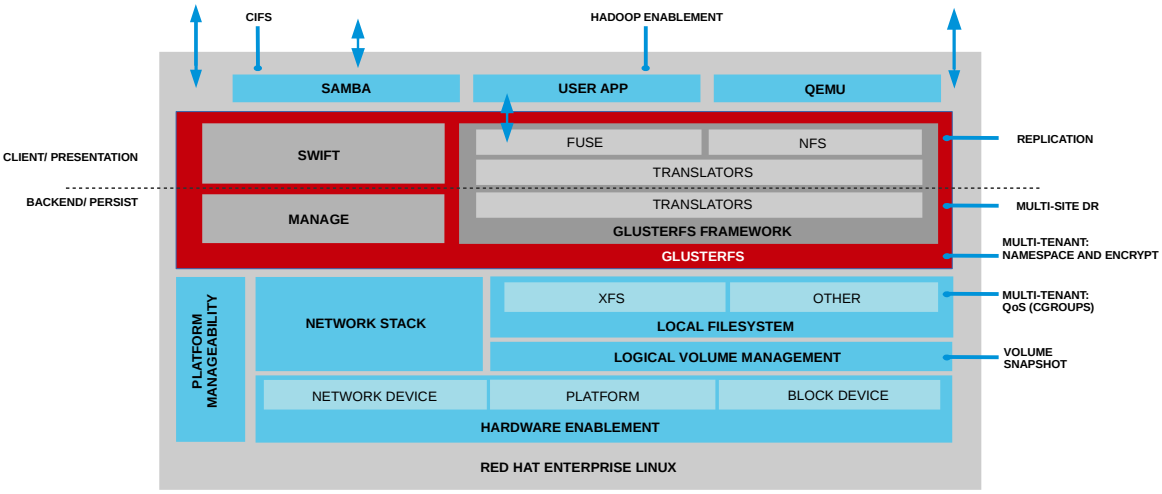

RED HAT STORAGE SCALABILITY

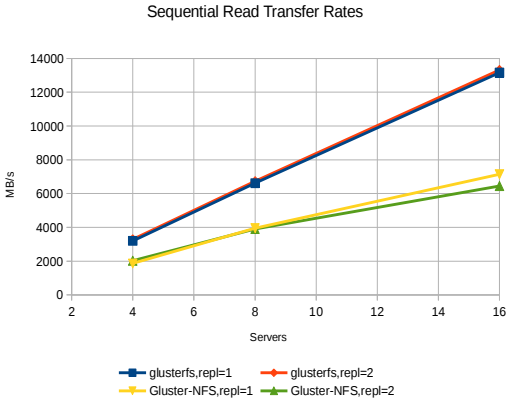

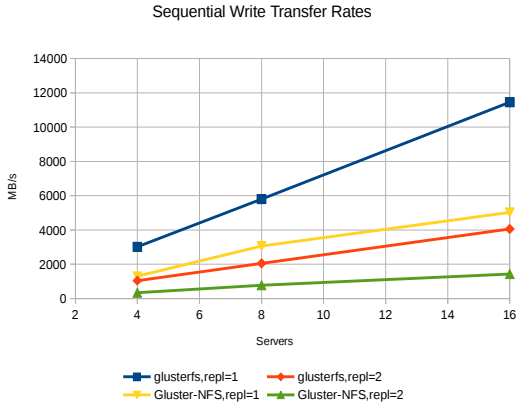

RHS-C Management Console

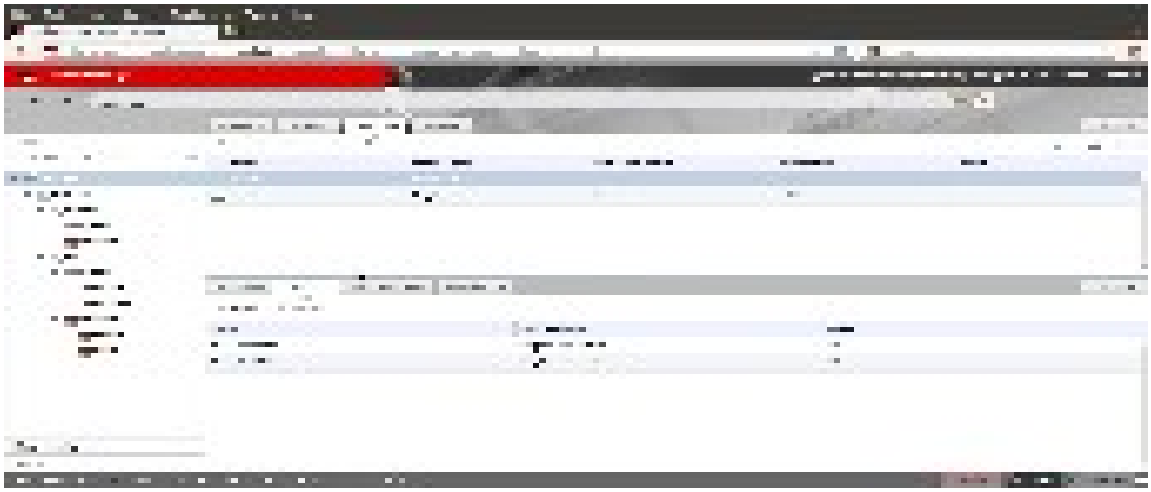

DESIGNED FOR MANAGING UNSTRUCTURED DATA

SUPPORTING A WIDE RANGE OF ENTERPRISE AND EMERGING WORKLOADS

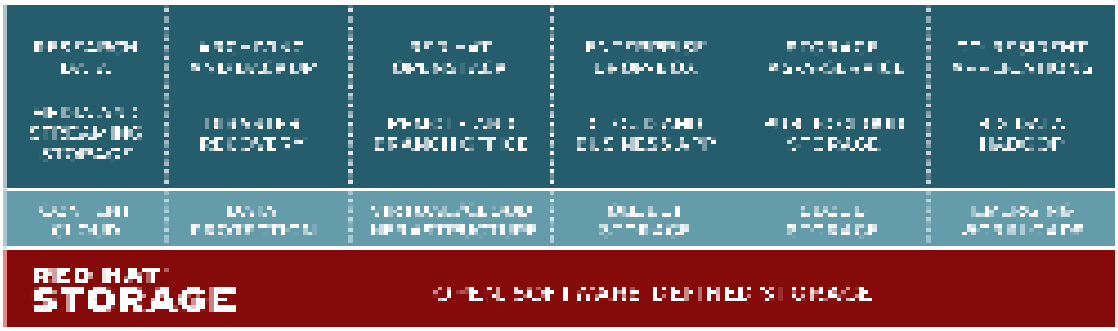

RED HAT STORAGE FOR OPENSTACK

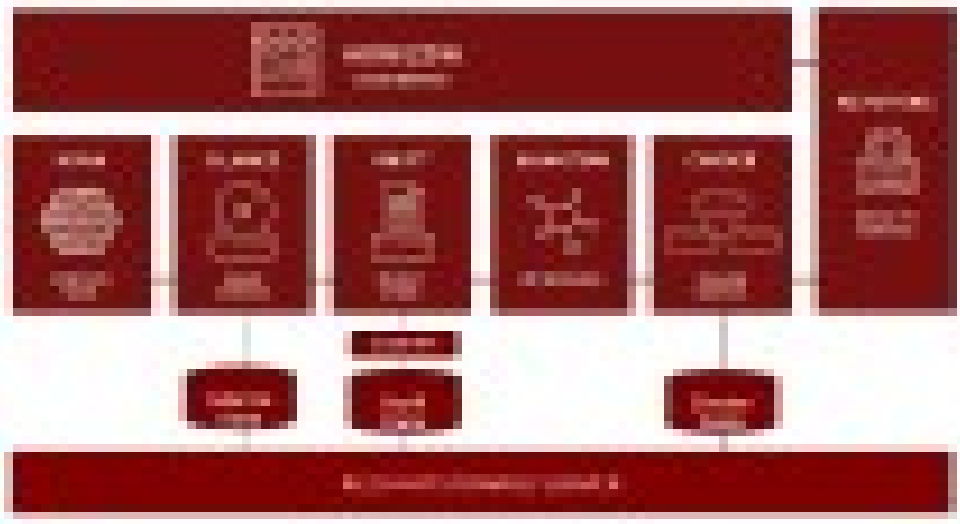

The different module components:

- Horizon : Management Dashboard
- Nova : Computing resources
- Glance : Image service
- Swift : Object Store
- Quantum : Networking module
- Cinder : Volume service
- Keystone : Authentication

Red Hat Storage fits in as an infrastructural component in below...

ENSURE GLOBAL DATA PROTECTION AND AVAILABILITY

TRANSPARENTLY DISTRIBUTE DATA GLOBALY

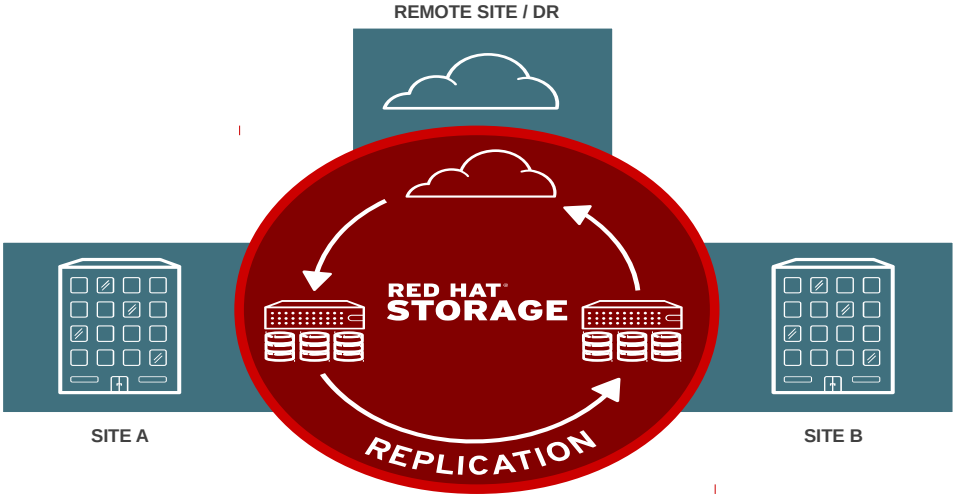

Red Hat Storage Allows You to Bring Application

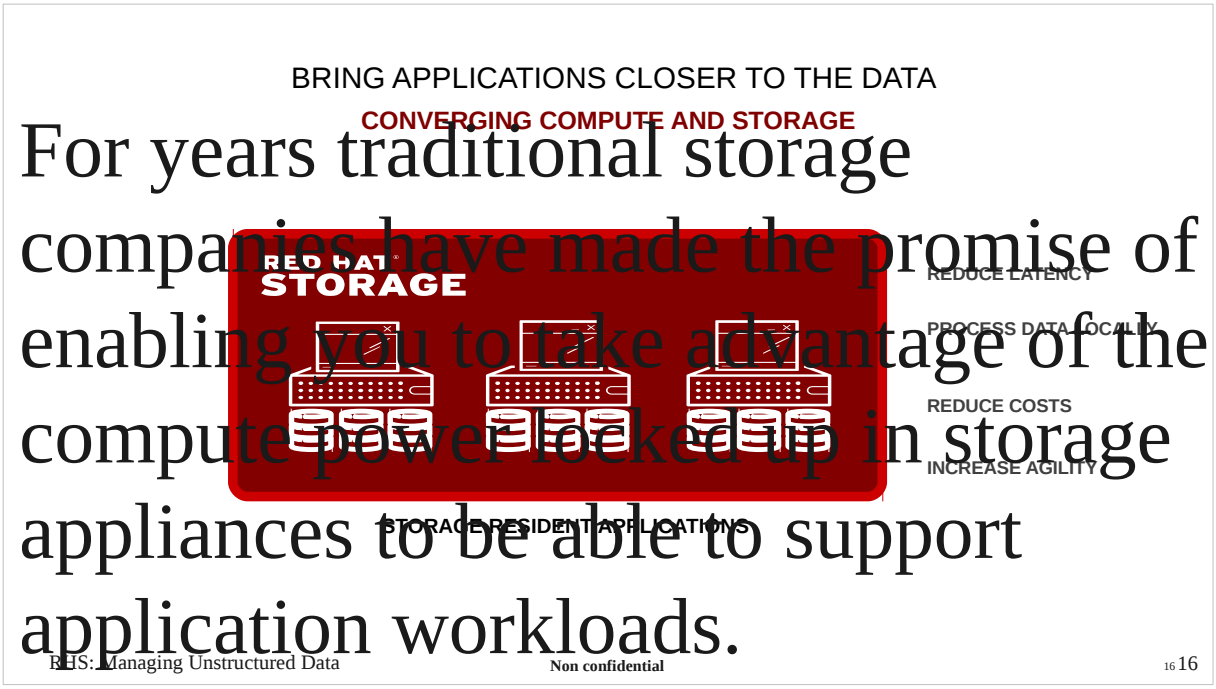

For years traditional storage companies have made the promise of enabling you to take advantage of the compute power locked up in storage appliances to be able to support application workloads.

With Todays new data landscape

REDUCE LATENCY

Gain increased performance for datasets by eliminating the network hop introduced by traditional architectures

PROCESS DATA LOCALLY

REDUCE COSTS

HIGHLY AVAILABLE CLOUD STORAGE FOR AMAZON EC2

LEVERAGE THE ELASTICITY OF THE CLOUD WITHOUT RE-WRITING YOUR APPLICATIONS

CREATING HIGHLY AVAILABLE, SCALEABLE EBS STORAGE POOLS - ACROSS ZONES

Now available as AWS test-drive

CUSTOMER USE CASE IntelliTEK

DELIVER COST EFFECTIVE ELASTIC CAPACITY AND PERFORMANCE

53% - 78% REDUCTION IN COSTS

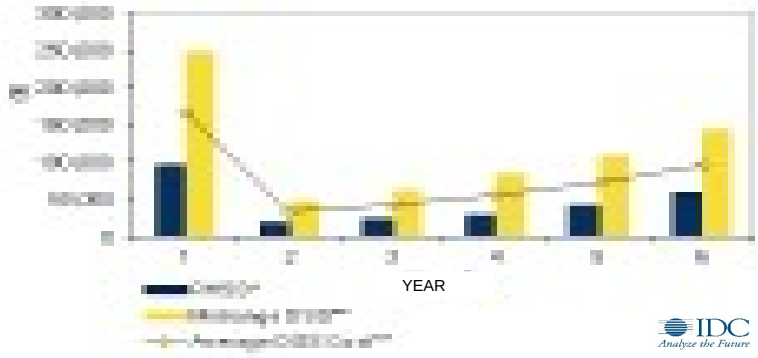

SOURCE: IDC REPORT – THE ECONOMICS OF SOFTWARE BASED STORAGE

MANAGING SPRAWLING UNSTRUCTURED FINANCIAL DATA

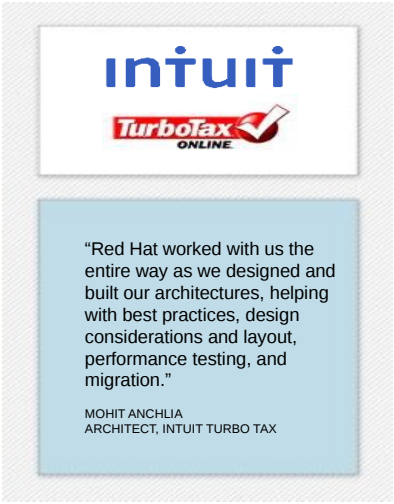

PROBLEM

- NEEDED A FAST, RELIABLE, AND COST-EFFECTIVE STORAGE SOLUTION TO MEET GROWING SAAS LINE OF BUSINESS
- TAX RETURNS AND OTHER DATA WERE BEING STORED AS BLOBS IN AN EXPENSIVE ORACLE DB

SOLUTION

- RED HAT STORAGE SERVER 2.O FOR ON-PREMISE OBJECT STORAGE
- HP DL2000s AND APACHE CASSANDRA

BENEFITS

- SCALEABLE ON-DEMAND STORAGE FOR UNSTRUCTURED DATA
- COST EFFECTIVE SOLUTION THAT LEVERAGES COMMODITY HARDWARE
- MEET GROWING CAPACITY AND PEAK PERFORMACE NEEDS
- ACHIEVE MULTI-SITE DISASTER RECOVERY

RHS: Managing Unstructured Data

Presentation Path:

Pandora serves up all of its music files through Red Hat Storage.

Imagine the scalability challenges Pandora faces. Each store song needs to be transcoded into 12 different file formats, depending on the device (phone, tablet, computer, etc.) accessing it.

Pandora needs to scale up immediately to accommodate a peak in traffic and, at the same time, accommodate long tail content access as well.

There is a publicly referenceable case study related to this customer. There is no formal write-up available.

IS THE OPPORTUNITY REAL ?

SearchStorage — TechTarget

News | Premium Editorial | Multimedia | Storage Topics | Tutorials | Expert Advice | Vendor Content | Blogs | Storage Decisions Events | SEARCH | GO

Home > Topics > Data Storage Management > Data management tools > Red Hat Gluster will transform storage market, IDC analyst predicts

Red Hat Gluster will transform storage market, IDC analyst predicts

Carol Sliwa
Published: 08 Aug 2013

Share

ESSENTIAL GUIDE — **RED HAT STORAGE SERVER SEEKS TO MIMIC SUCCESS OF ENTERPRISE LINUX**

This article is part of an Essential Guide, our editor-selected collection of our best articles, videos and other content on this topic.Explore more in this guide:

1. - IDC'S NADKARNI PREDICTS MAINSTREAM ACCEPTANCE : READ MORE IN THIS SECTION
   - GlusterFS-based software will drive mainstream acceptance of software-based storage

Explore other sections in this guide:

2. - Wikibon's Floyer: Manageability work needed
3. - Stalmer foresees limited appeal

Red Hat Inc. will eventually turn the market "on its head" with its Red Hat GlusterFS-based storage software in the same way that it successfully challenged major server operating system vendors with its distribution of enterprise Linux, an IDC analyst predicted.

Ashish Nadkarni, a research director in the storage systems practice at Framingham, Mass.-based International Data Corp. (IDC), said Red Hat may need a few years to realize its vision of converting commodity hardware into a full-fledged storage platform.

But Nadkarni thinks Red Hat Gluster will challenge scale-out, file-based storage systems from major vendors as well as combined compute-storage platforms that cater to the high-performance computing market with data analytics software.

LISTEN TO THE ENTIRE PODCAST SERIES ABOUT RED HAT'S CHANCES WITH ITS STORAGE SERVER SOFTWARE

Red Hat's open source software-based storage holds appeal mainly for passive data

With push into software-based storage, Red Hat will help drive down cost of storage features

Latest News

- EMC Corp. prepares VNX2 launch
- Scality CEO: Object-oriented storage needs multiple access methods
- Coraid Inc. NAS gets more cache for heavy I/O loads
- Riverbed Granite storage box gets Fibre Channel support
- Emerging data storage trends embraced by BNY Mellon IT infrastructure

DELIVERING THE NEXT GENERATION OF OPEN SOFTWARE-DEFINED STORAGE TODAY

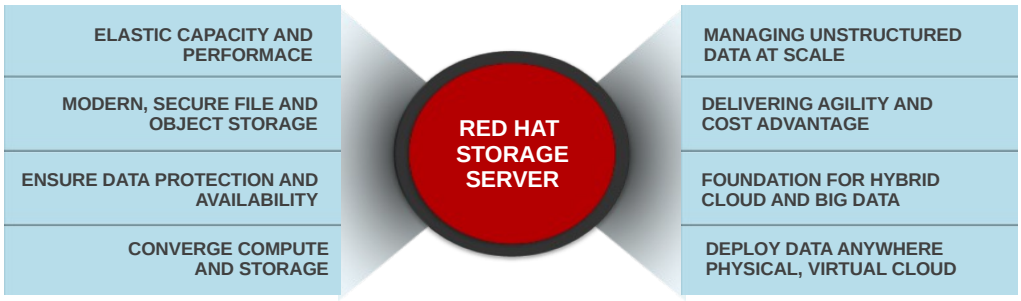

DESIGNED FOR TODAYS IT & DATA ECONOMICS

RHS: Managing Unstructured Data

RED HAT STORAGE – DEVELOPING 3rd PARTY ECO-SYSTEM

RED HAT LEADS THROUGH OPEN INNOVATION

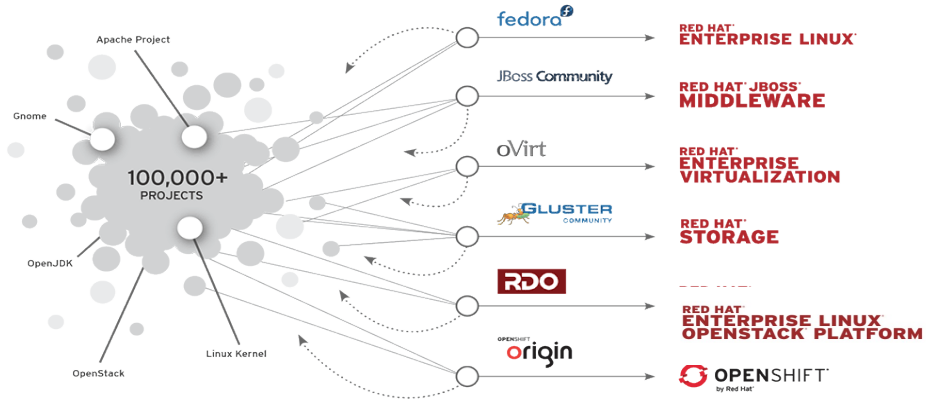

Gerry

COMMUNITY INNOVATION

GLUSTER.ORG COMMUNITY FORGE ENHANCEMENTS AND PROJECTS

SNAPSHOTTING

CHANGE DETECTION

COMPRESSION

3-WAY REPLICATION

pNFS AND NFSv4 SUPPORT

FILE VERSIONING

ERASURE CODING

MULTI-MASTER GEO-REPLICATION

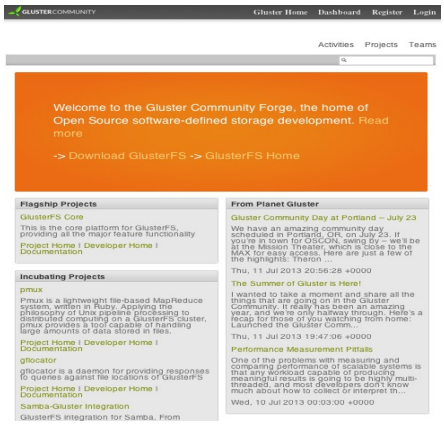

SMB 3.0 SUPPORT

NDMP SERVER

PUPPET MANAGEMENT MODULE

GTOP - MONITORING

GLUSTER PROFILING

SELINUX SUPPORT

PMUX – LIGHTWEIGHT MAP REDUCE

TRANSLATORS EXTENSION FOR PYTHON

BACKUP

How Does GlusterFS Work Without Metadata?

- Files are placed on a brick(s) in the cluster based on a calculation
- All native clients have an algorithm built-in
- All storage nodes have an algorithm built-in
- Files can then be retrieved based on the same calculation
- For non-native clients, the server handles retrieval and placement

redhat

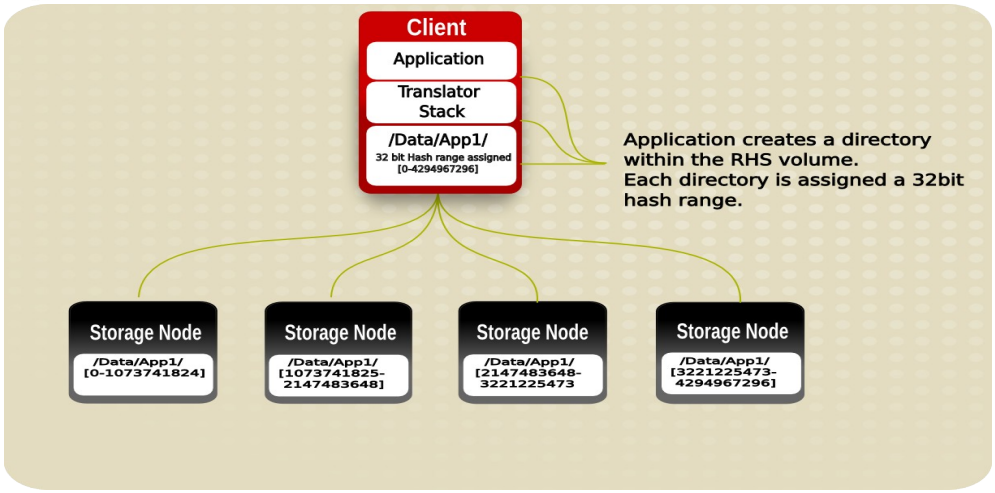

redhat

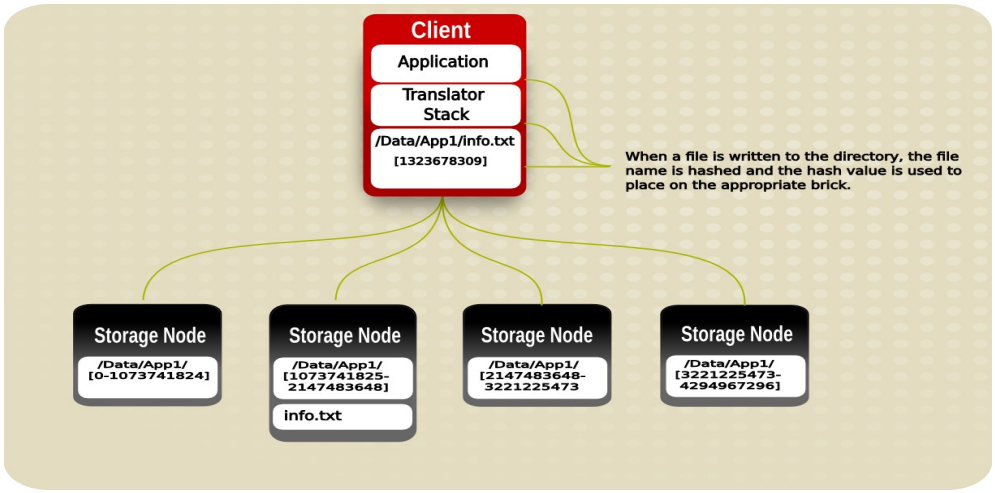

redhat

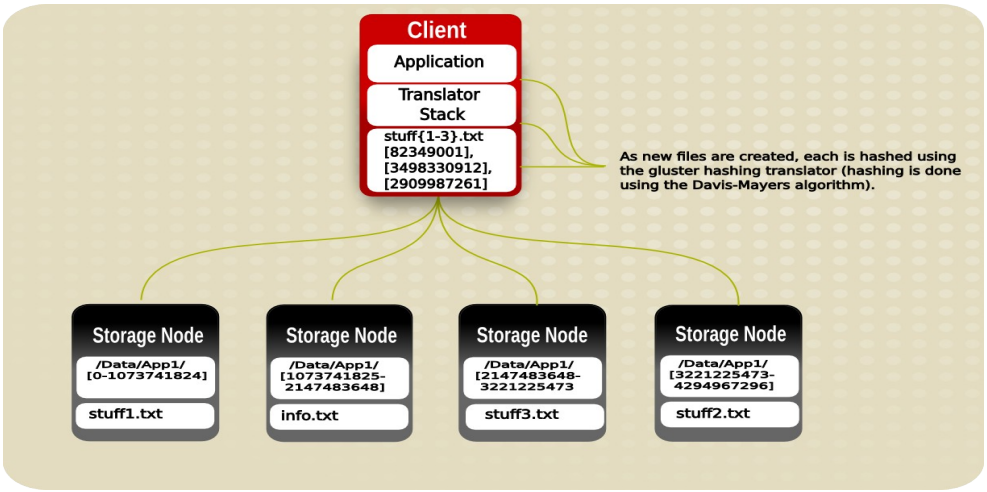

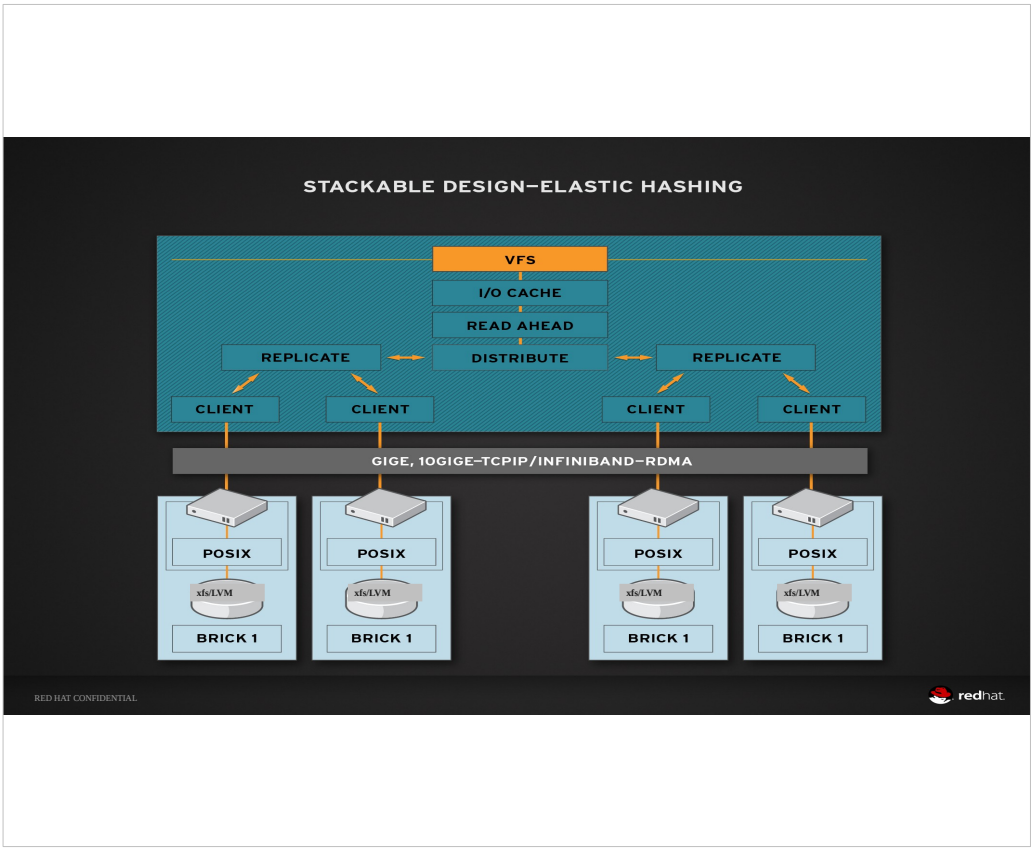

Past decade: Linux + volume x86 servers transformed the *server* market

Displaced costly proprietary RISC/UNIX systems
Enabled new classes of workloads
Showed superior economics

Current decade: Open-source-based storage + volume x86 servers tranform *storage* market

Displacing costly proprietary SAN and NAS systems
Cost 1/3 to 1/2 the price of alternative proprietary solutions
Enabling new classes of workloads
Helping realize the true potential of hybrid clouds

Bricks

- A brick is the combination of a node and a file system: hostname:/dir
- Each brick inherits limits of the underlying filesystem(xfs)
- RHS operates at the brick level, not at the node level
- Ideally, each brick in a cluster should be the same size

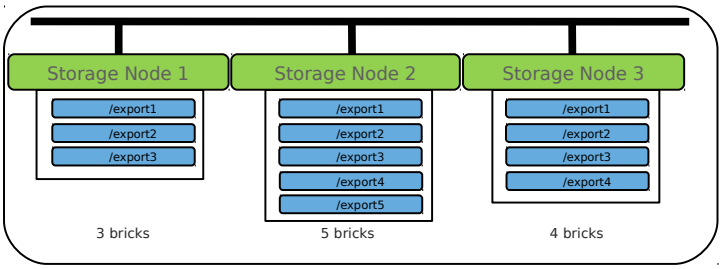

Volumes

- A volume consists of 1 or more bricks => exported with Gluster.
  - volumes have administrator assigned export names
  - a brick is a member of only one volume
- A namespace can have 1 or more volumes
  - A namespace can consist of replicated and distributed volumes
  - data in different volumes physically exists on different bricks
  - volumes can be mounted on clients using NFS, CIFS and/or GlusterFS clients (native FUSE client)

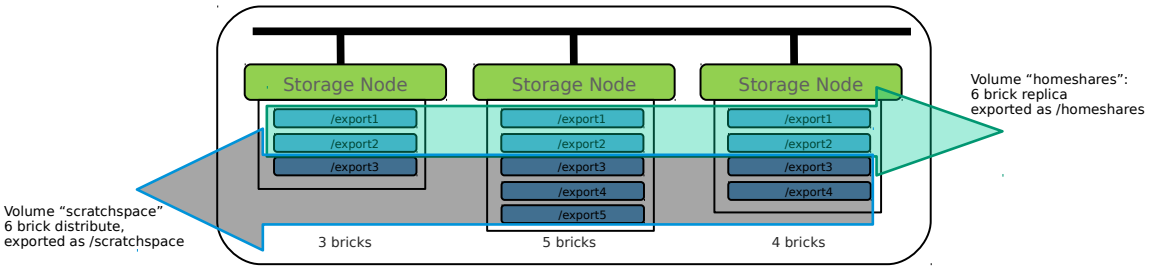

Data Flow with NFS/CIFS Client

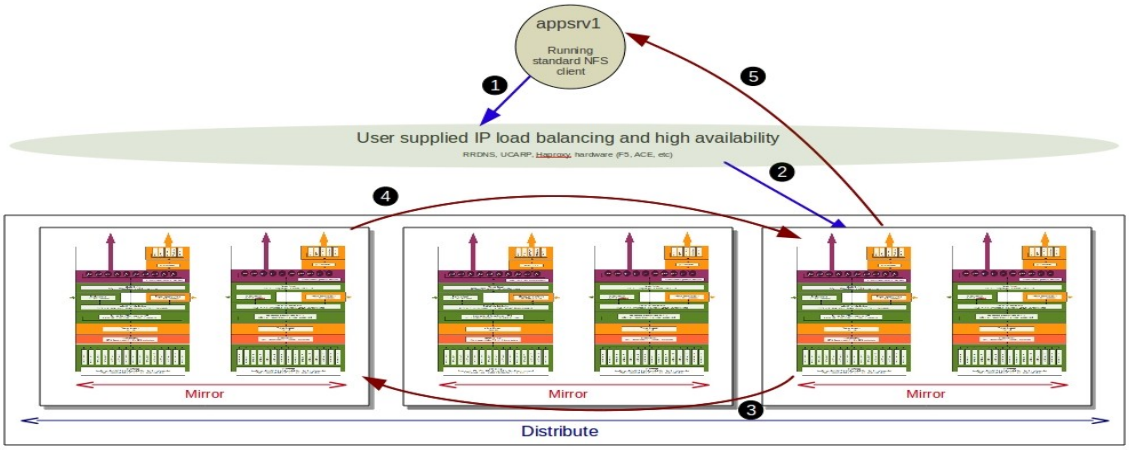

Data Flow with Native Client

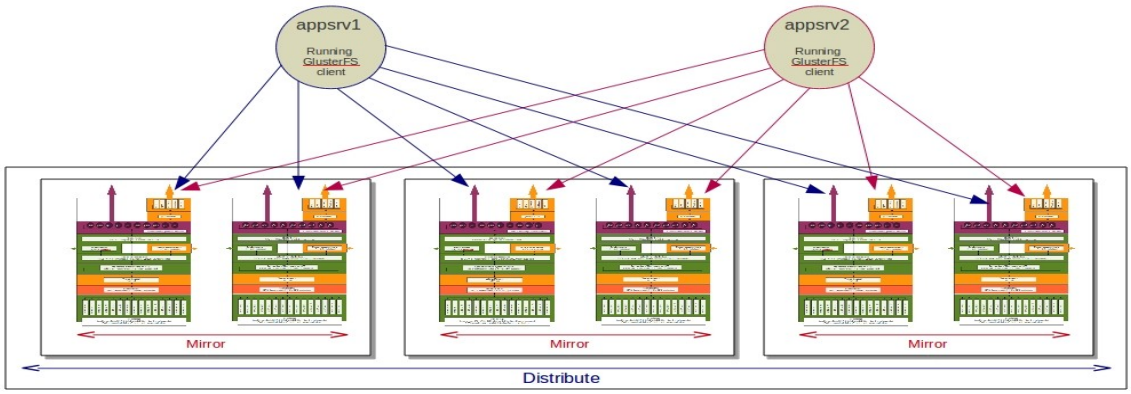

HOW DOES REPLICATION ACTUALLY WORK?

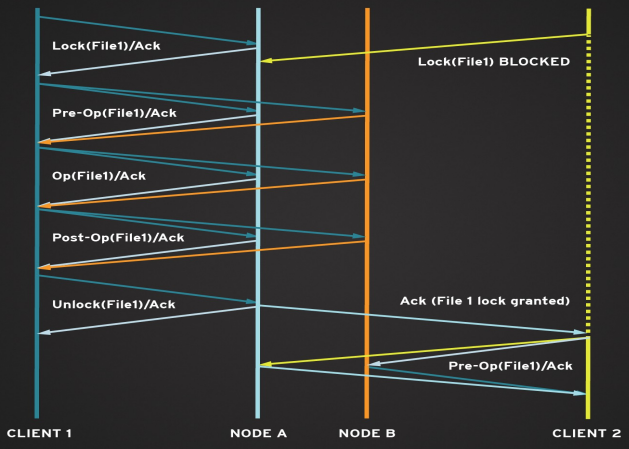

redhat.

Seamless Integration for Hadoop Deployments

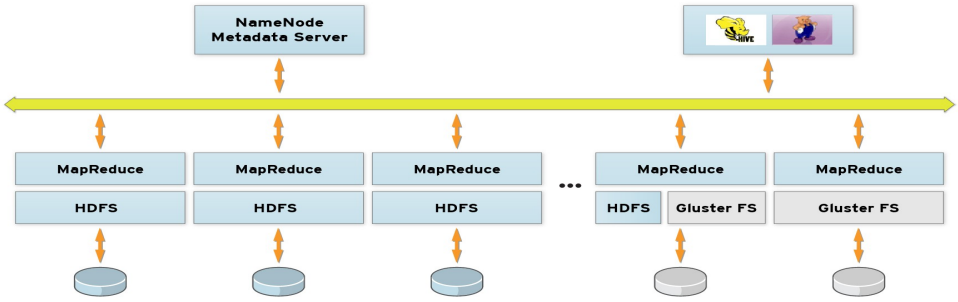

- GlusterFS can co-exist HDFS
- Does not use the NameNode metadata server

- Built using the Hadoop file system API
- Requires simple configuration file changes
- C Lib GlusterFS client enable GlusterFS direct access
- Provides Java binding for Hadoop compatibility

Hadoop architecture overview