redhat.

STORAGE FOR OPENSHIFT CONTAINERS feat. RED HAT GLUSTER STORAGE

Shawn Houston
Cloud Storage Solutions Architect

LINUX CONTAINERS:

Software packaging concept that typically includes an application and all of its runtime dependencies

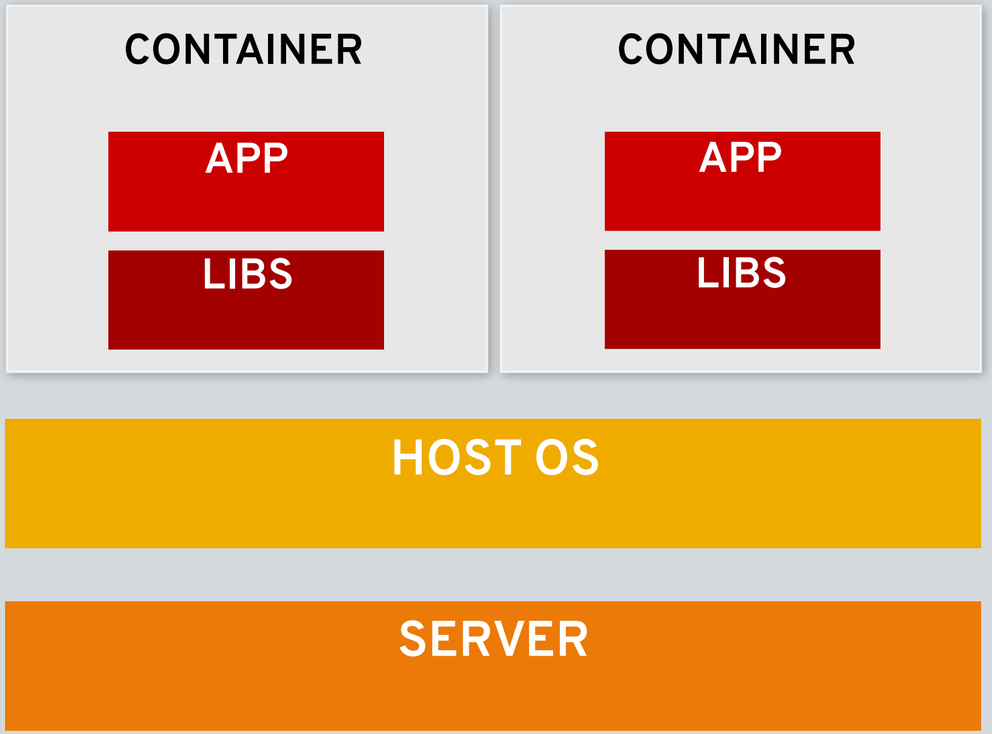

BENEFITS

- **HIGHER** quality software releases
- **SHORTER** test cycles
- **EASIER** application management

CONTAINERS V.S. VIRTUALIZATION

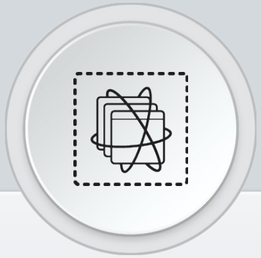

Containers

- Abstracts OS Kernel
- Limited to Linux
- One CPU and memory mgr
- Up in seconds
- 100s or 1000s
- Multiple copies of single app

VS

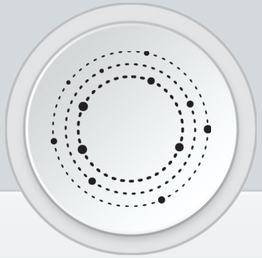

Virtualization

- Abstracts entire device
- Any Operating System
- Two CPU and memory mgrs
- Up in hours or days
- 10s or 100s
- Multiple apps

WHY PERSISTENT STORAGE FOR CONTAINERS?

"For which workloads or application use cases have you used/do you anticipate to use containers?"

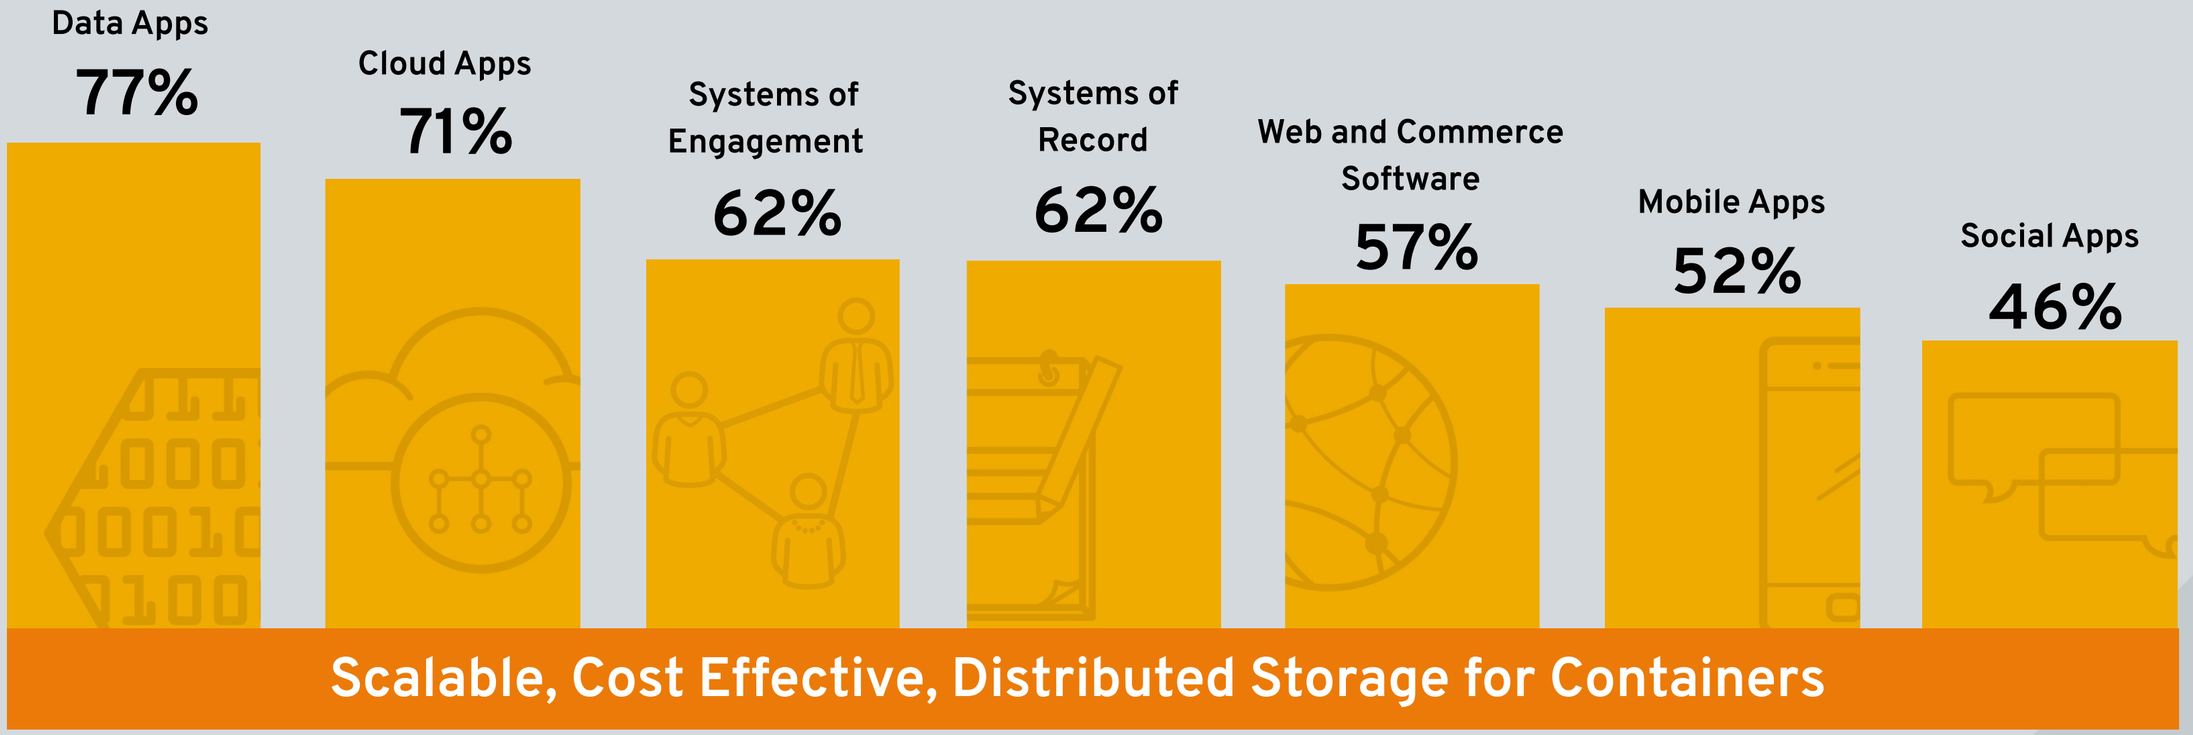

Base: 194 IT operations and development decision-makers at enterprise in APAC, EMEA, and North America Source: A commissioned study conducted by Forrester Consulting on behalf of Red Hat, January 2015

redhat.

THE ROAD TO STORAGE AS A SERVICE

Development Model	Application Architecture	Deployment & Packaging	Application Infrastructure	Storage
Waterfall	Monolithic	Bare Metal	Data Center	Scale Up
Agile	N-tier	Virtual Servers	Hosted	Scale Out
DevOps	**MicroServices**	**Containers**	**Hybrid Cloud**	**Storage as a Service**

STORAGE INNOVATION FOR CONTAINERIZED APPLICATIONS

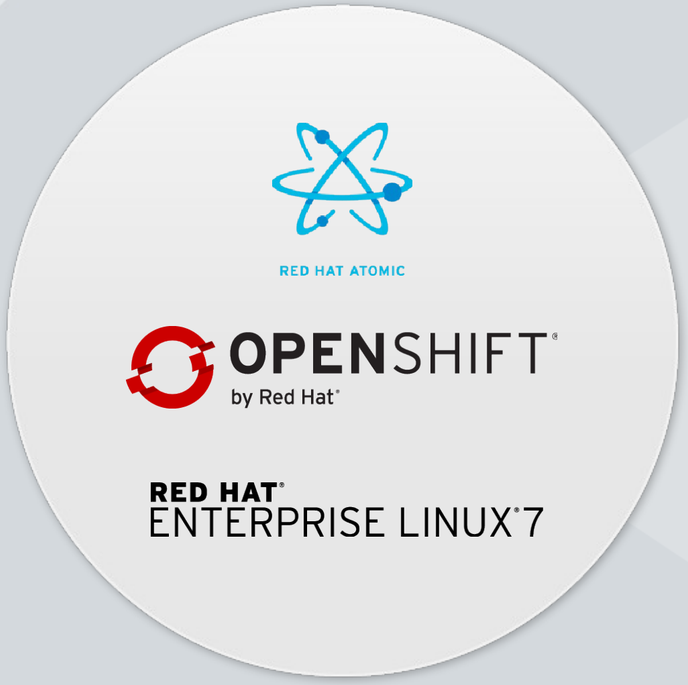

- Ceph RBD
- Amazon EBS
- Fiber Channel
- GCE
- iSCSI
- NFS
- GlusterFS

AUTOMATED CONFIGURATION

SINGLE CONTROL PANEL

CHOICE OF PERSISTENT STORAGE

CONTAINER READY STORAGE

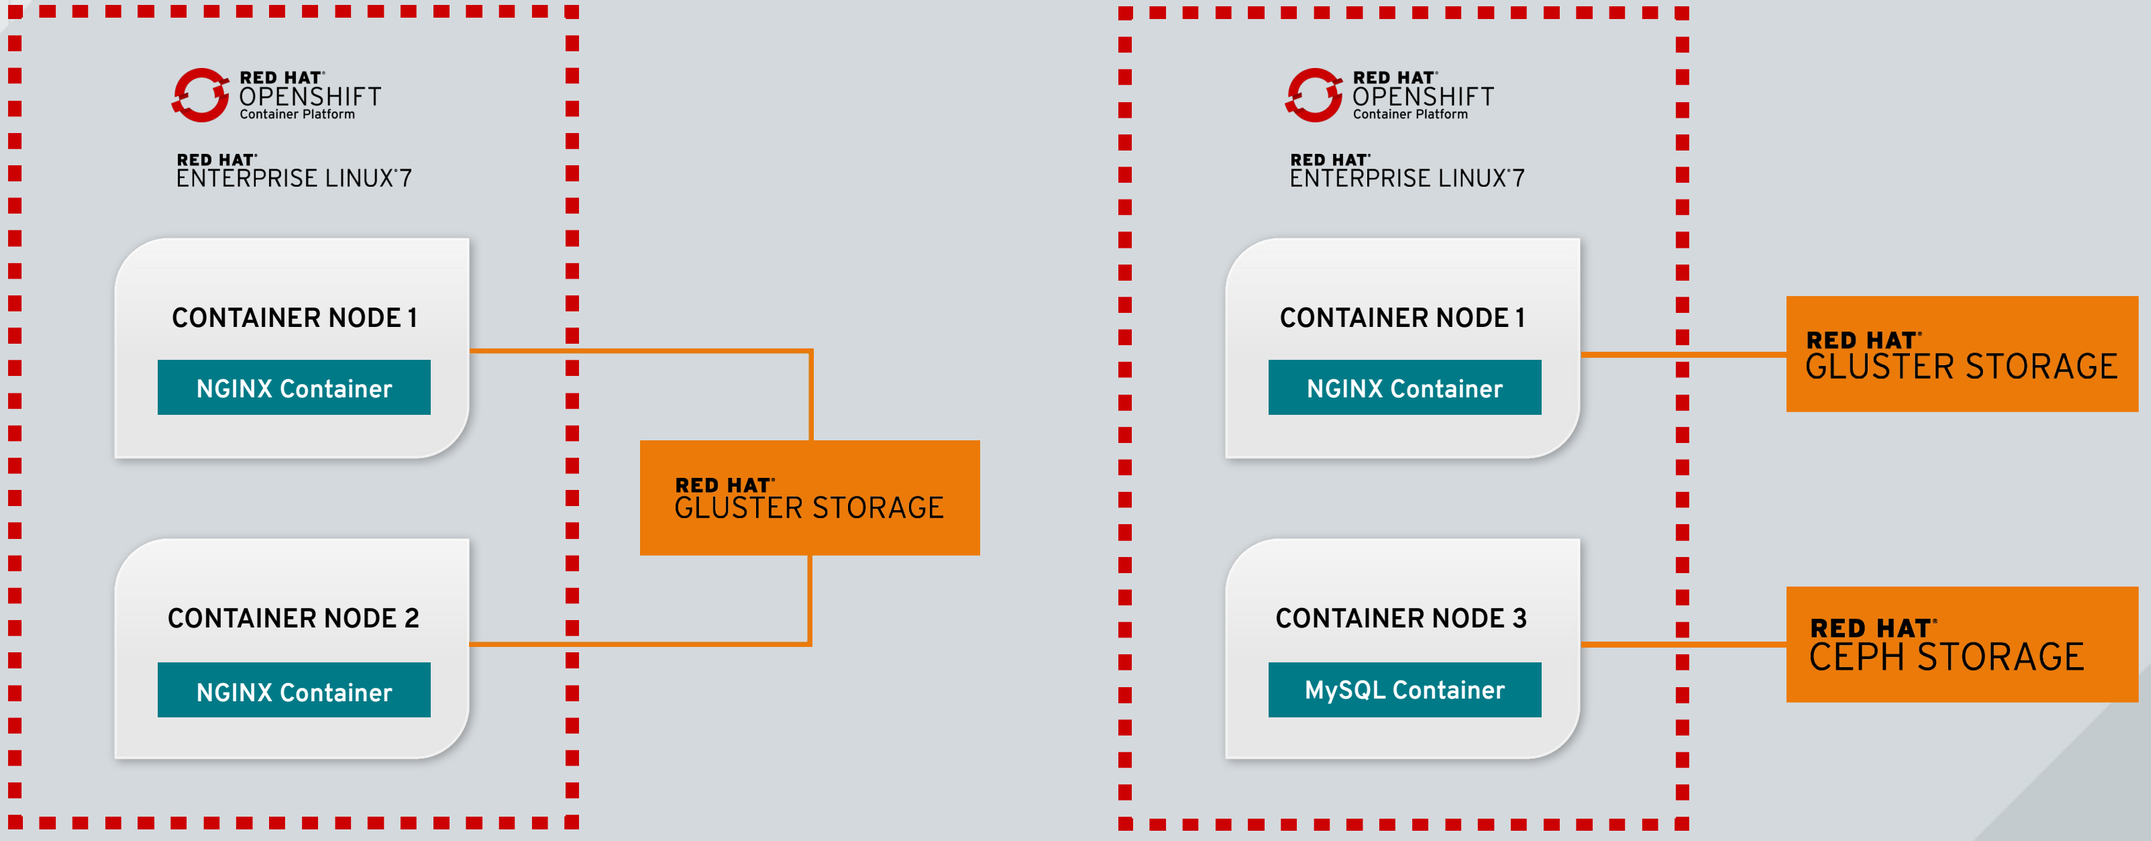

CONTAINERIZED RED HAT GLUSTER STORAGE

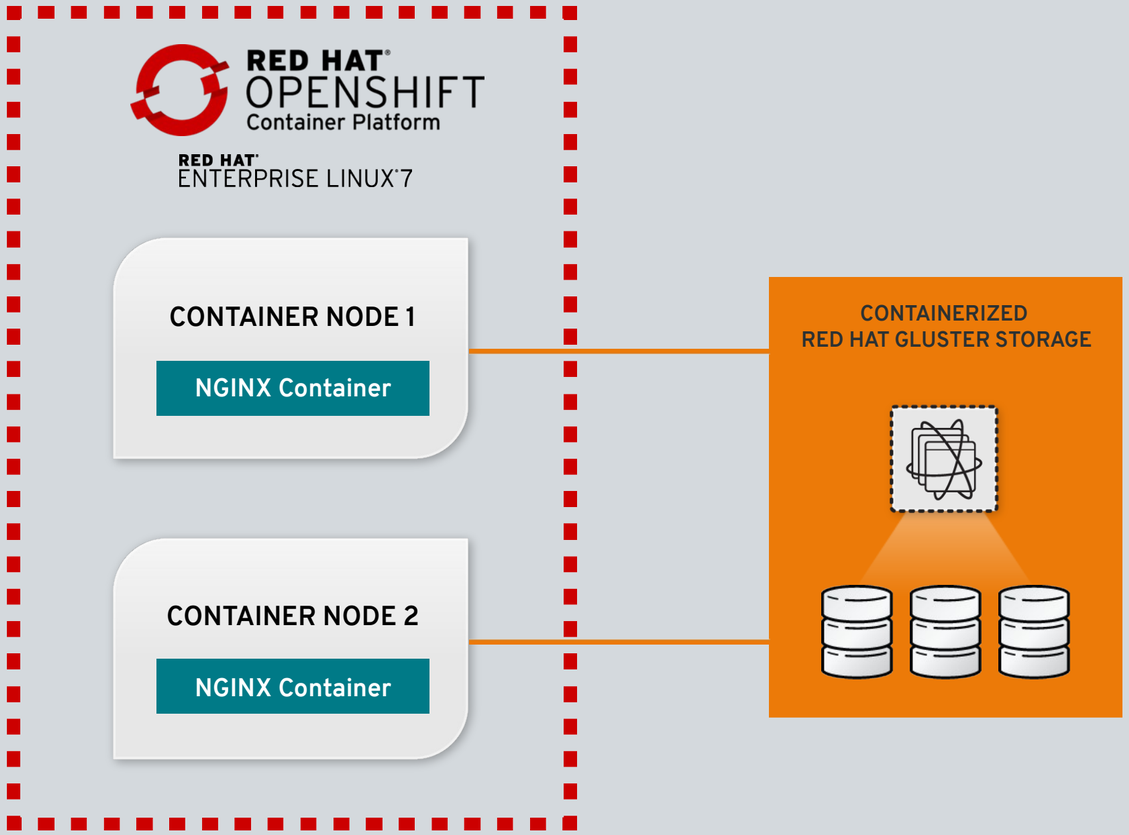

CONTAINER-NATIVE STORAGE

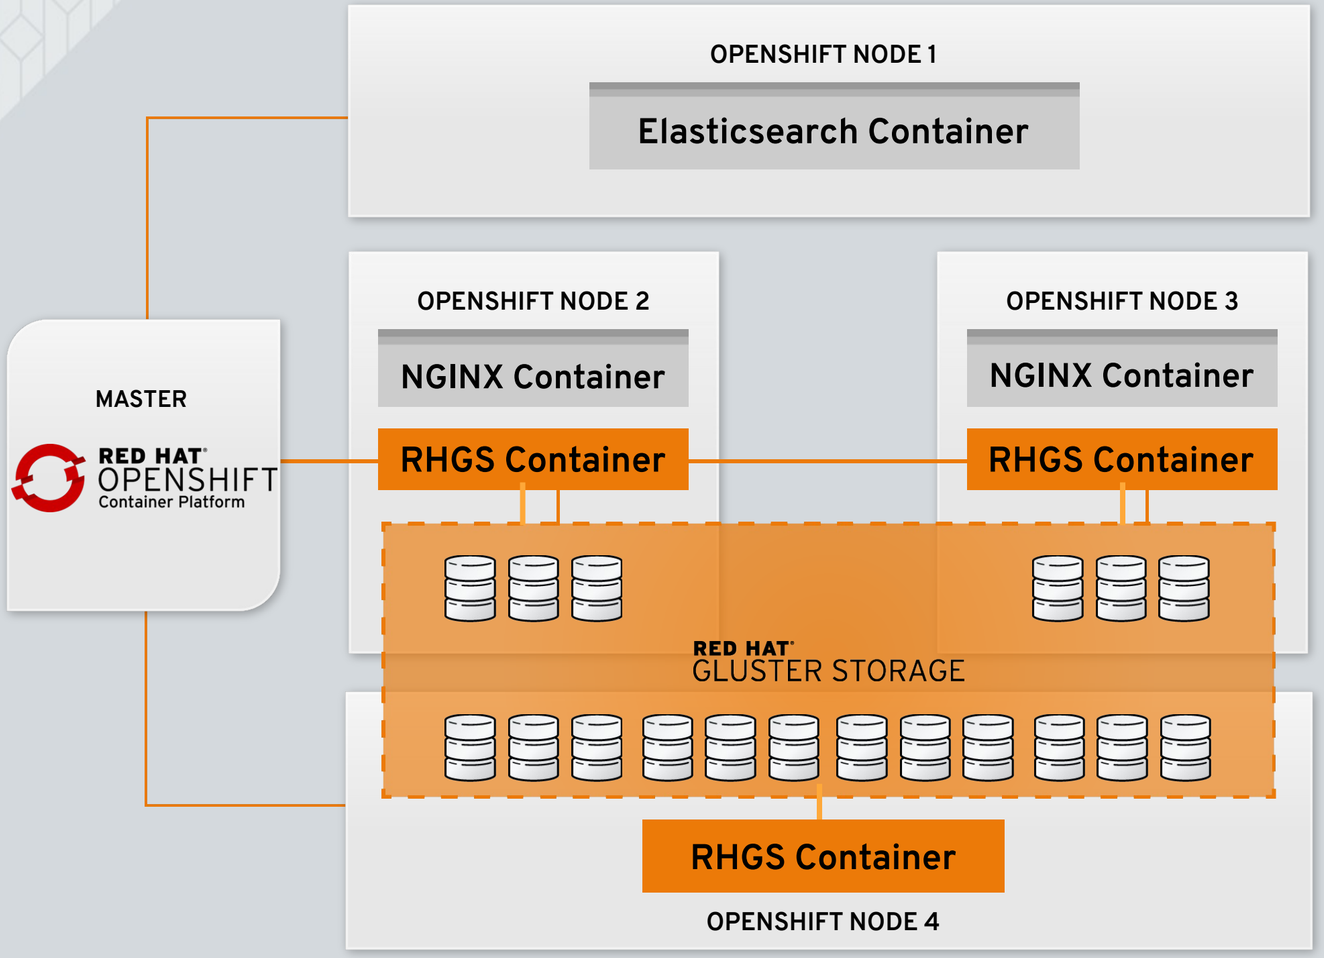

- Lower TCO
- Unified Orchestration
- Ease of Use
- Greater control

redhat.

THE RED HAT STACK – FROM PAAS TO STORAGE

OPENSHIFT marketplace

RED HAT® JBOSS® MIDDLEWARE

RED HAT® SATELLITE

xPaaS

RED HAT® STORAGE

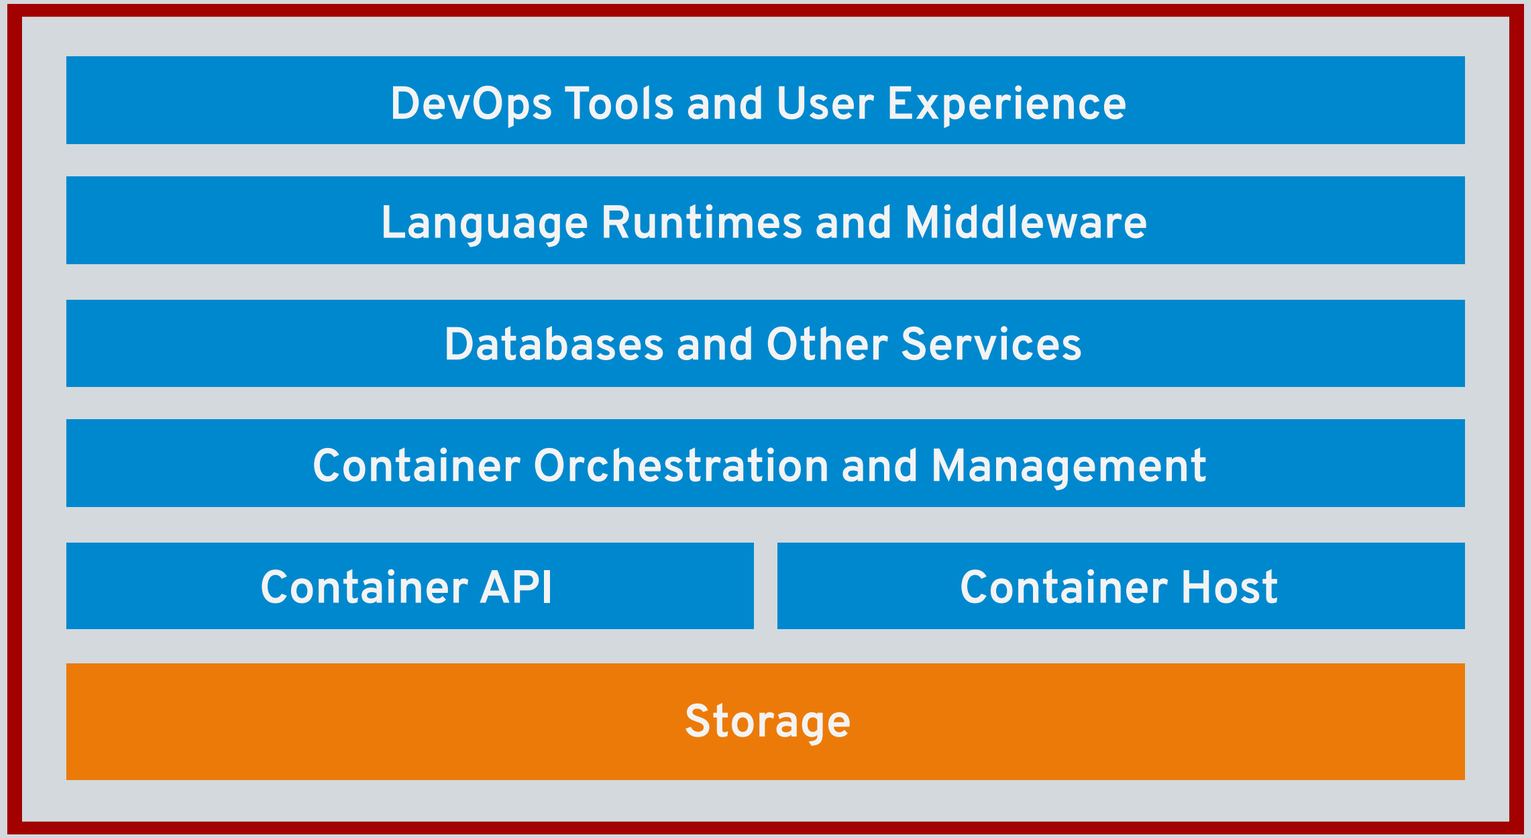

RED HAT® OPENSHIFT Container Platform

kubernetes by Google

docker

redhat. LINUX

DRIVING THE FUTURE OF STORAGE

CONTAINER READY STORAGE

RED HAT® GLUSTER STORAGE

Nov 2015

- Dedicated storage cluster for containerized and PaaS environments
- Supported for OpenShift Enterprise

CONTAINERIZED RHGS

RED HAT® GLUSTER STORAGE

Mar 2016

- Containerized Red Hat Gluster Storage serving storage from a dedicated storage cluster
- Optimized for applications running on RHEL 7, OpenShift Enterprise, and RHEL Container Host

CONTAINER-NATIVE STORAGE

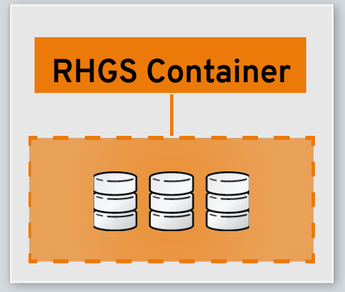

Summer 2016

- Containerized Red Hat Gluster Storage inside OpenShift Container Platform hyper converged with application containers
- Red Hat Gluster Storage cluster comprised of disks from multiple container cluster nodes

CONVERGENCE OF STORAGE AND COMPUTE

redhat.

RED HAT GLUSTER STORAGE ADVANTAGES

OPEN **Open, software-defined distributed file and object storage system**	• Based on GlusterFS open source community project • Uses proven local file system (XFS) • Data is stored in native format
SCALABLE **No Metadata Server**	• Uses an elastic hashing algorithm for data placement • Uses local filesystem's xattrs to store metadata • Nothing shared scale-out architecture
ACCESSIBLE **Multi-Protocol the Same Data**	• Global name space • NFS, SMB, object, HDFS, Gluster native protocol • Posix compliant
MODULAR **No Kernel Dependencies**	• GlusterFS is based on filesystem in userspace (FUSE) • Modular stackable arch allows easy addition of features ...without being tied to any kernel versioniant
ALWAYS-ON **High-Availability across data, systems and applications**	• Synchronous replication with self-healing for server failure • Asynchronous geo-replication for site failure

HOW IS GLUSTER DEPLOYED?

Red Hat Gluster Storage

PHYSICAL	VIRTUAL	CONTAINERS	CLOUD
RED HAT® GLUSTER STORAGE	RED HAT® GLUSTER STORAGE	RED HAT® GLUSTER STORAGE	RED HAT® GLUSTER STORAGE
RED HAT® ENTERPRISE LINUX®	RED HAT® ENTERPRISE LINUX® RED HAT® ENTERPRISE VIRTUALIZATION	RED HAT® ENTERPRISE LINUX® ATOMIC HOST OPENSHIFT ENTERPRISE by Red Hat®	RED HAT® ENTERPRISE LINUX® Google Cloud Platform amazon web services™

redhat.

DATA PLACEMENT BEST PRACTICE

Distributed-Replicated Volume

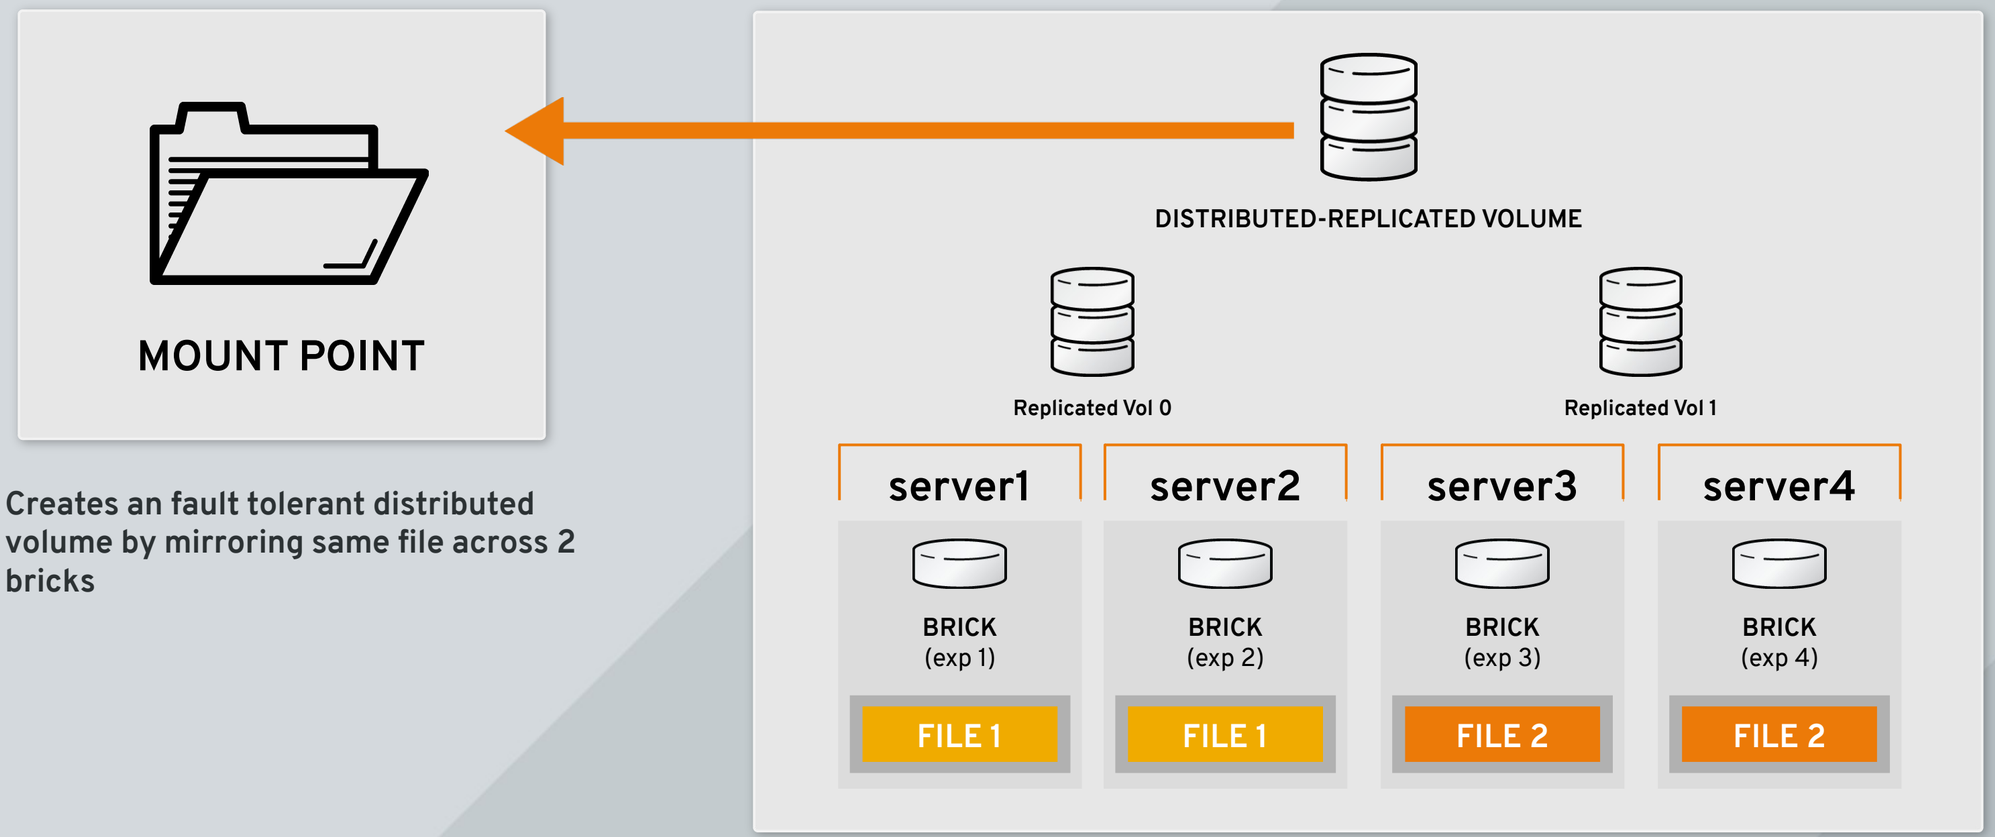

Creates an fault tolerant distributed volume by mirroring same file across 2 bricks

DATA PLACEMENT FOR CONTAINER NATIVE

Replicated Volume

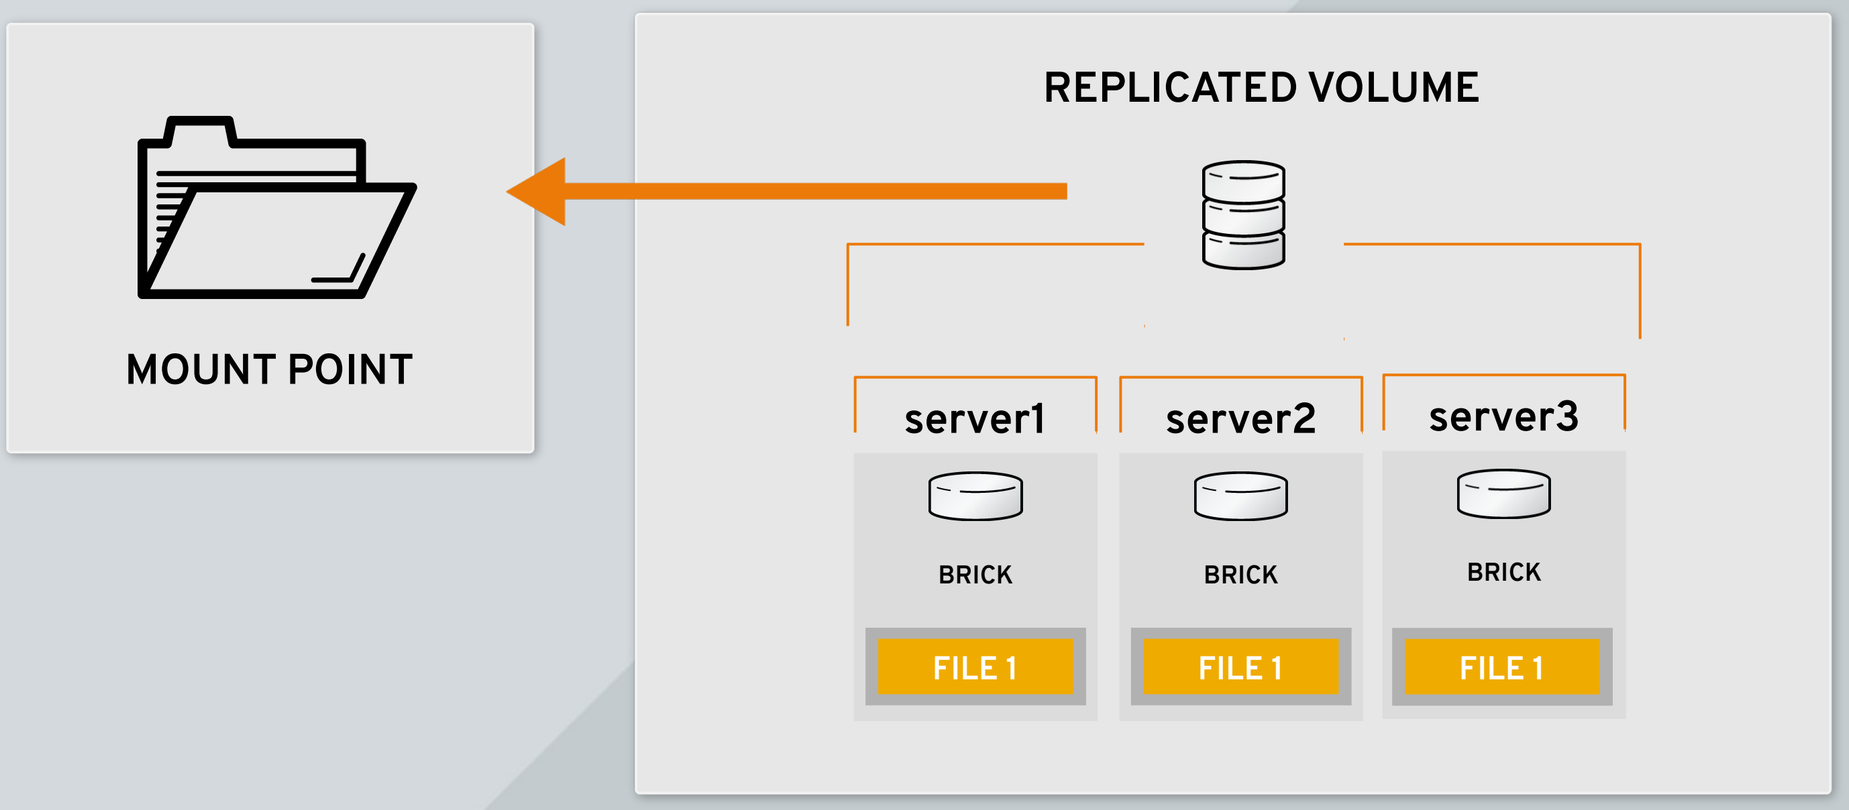

GlusterFS NATIVE CLIENT

- **BASED ON FUSE KERNEL MODULE,** which allows the file system to operate entirely in userspace
- **SPECIFY MOUNT** to any GlusterFS server
- **NATIVE CLIENT** fetches volfile from mount server, then communicates directly with all nodes to access data

Load inherently balanced across distributed volumes
Recommended for high concurrency & high write performance

A PEEK OVER THE HORIZON

Community Innovation

Greater Integration

OPENSHIFT origin

Storage as a Microservice

Storage Communities

ceph

GLUSTER

redhattechnicalseries.com/storage

redhat.

THANK YOU

plus.google.com/+RedHat

facebook.com/redhatinc

linkedin.com/company/red-hat

twitter.com/RedHatNews

youtube.com/user/RedHatVideos