

# Software Defined Storage Overview

August 2019

Juan Jose Floristan  
Cloud Specialist Solution Architect

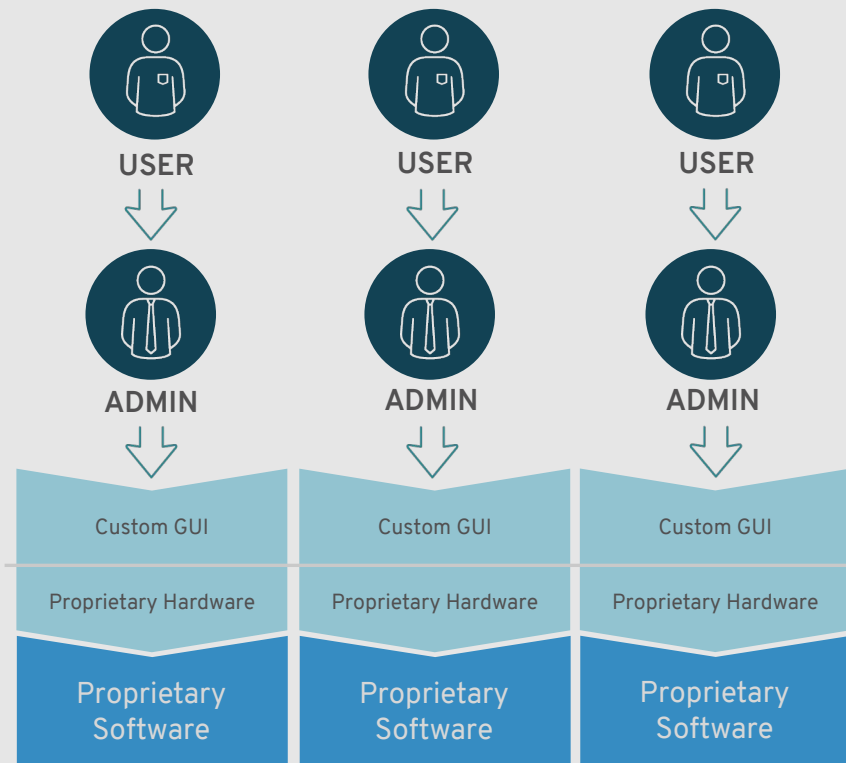
## AGENDA

1. Why Red Hat Storage?
2. Red Hat Ceph Storage
3. Red Hat Gluster Storage
4. Red Hat Openshift Container Storage

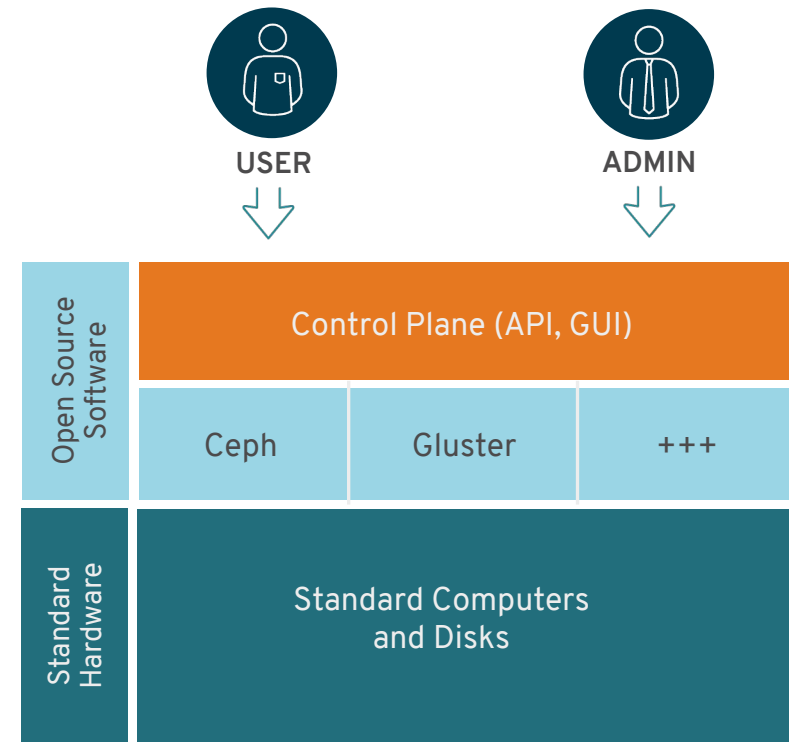
# Why Red Hat Storage?

# STORAGE IS EVOLVING

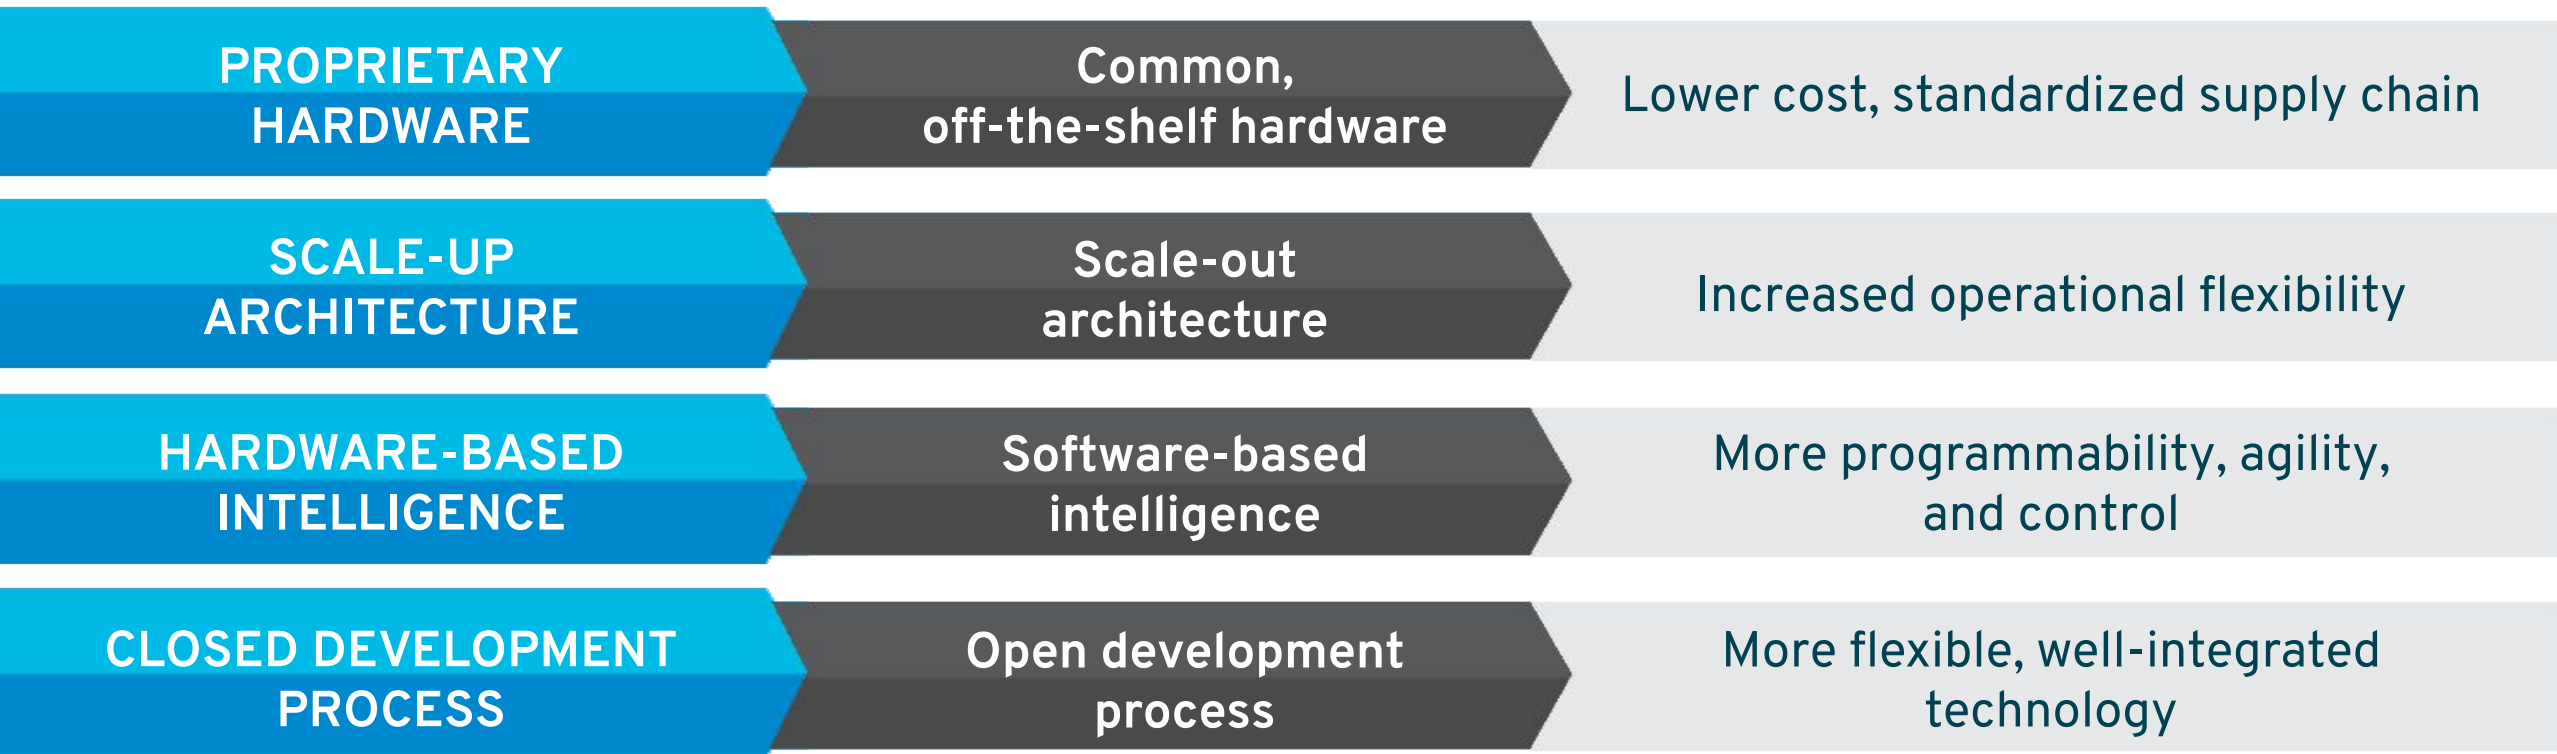
## TRADITIONAL STORAGE Complex proprietary silos



## OPEN, SOFTWARE-DEFINED STORAGE Standardized, unified, open platforms



# WHY THIS MATTERS



# A RISING TIDE

Software-Defined Storage is leading a shift in the global storage industry, with far-reaching effects.

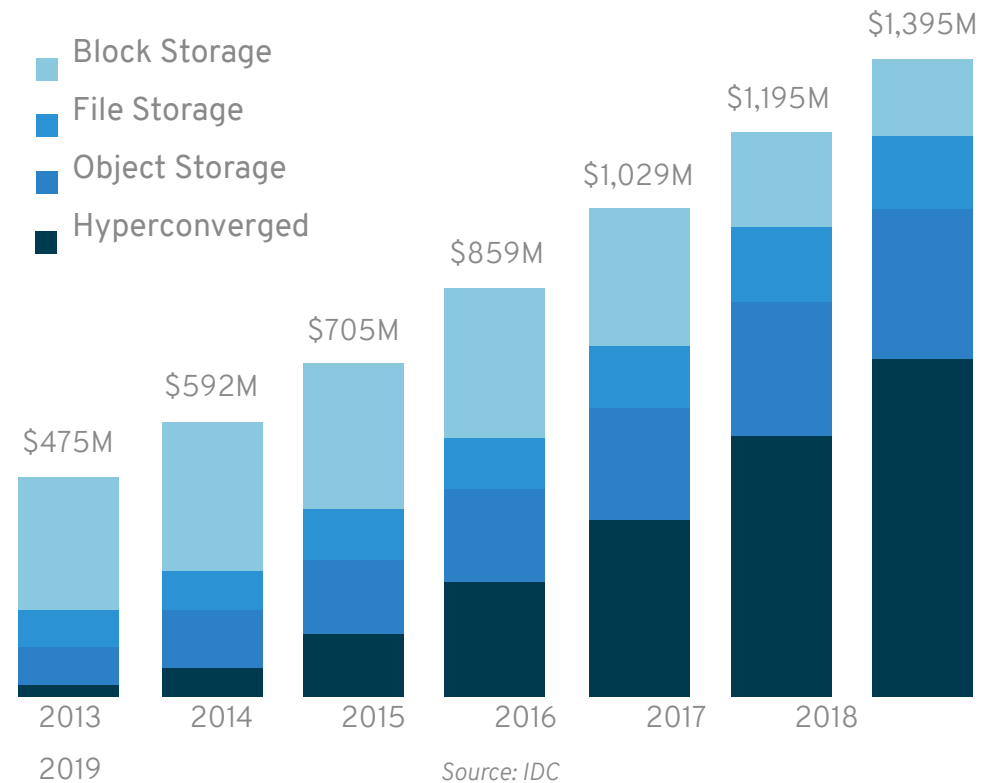
“By 2020, between 70%-80% of unstructured data will be held on lower-cost storage managed by SDS.”

Innovation Insight: Separating Hype From Hope for Software-Defined Storage

“By 2019, 70% of existing storage array products will also be available as software-only versions.”

Innovation Insight: Separating Hype From Hope for Software-Defined Storage

SDS-P MARKET SIZE BY SEGMENT



# THE RED HAT STORAGE MISSION

To offer a unified, open software-defined storage portfolio that delivers a range of data services for next generation workloads, thereby accelerating the transition to modern IT infrastructures.

# Red Hat Ceph Storage



# BUSINESS BENEFITS

## OPEN SOURCE

No proprietary lock-in, with a large commercial ecosystem and broad community

## PEACE OF MIND

Over a decade of active development, proven in production and backed by Red Hat

## LOWER COST

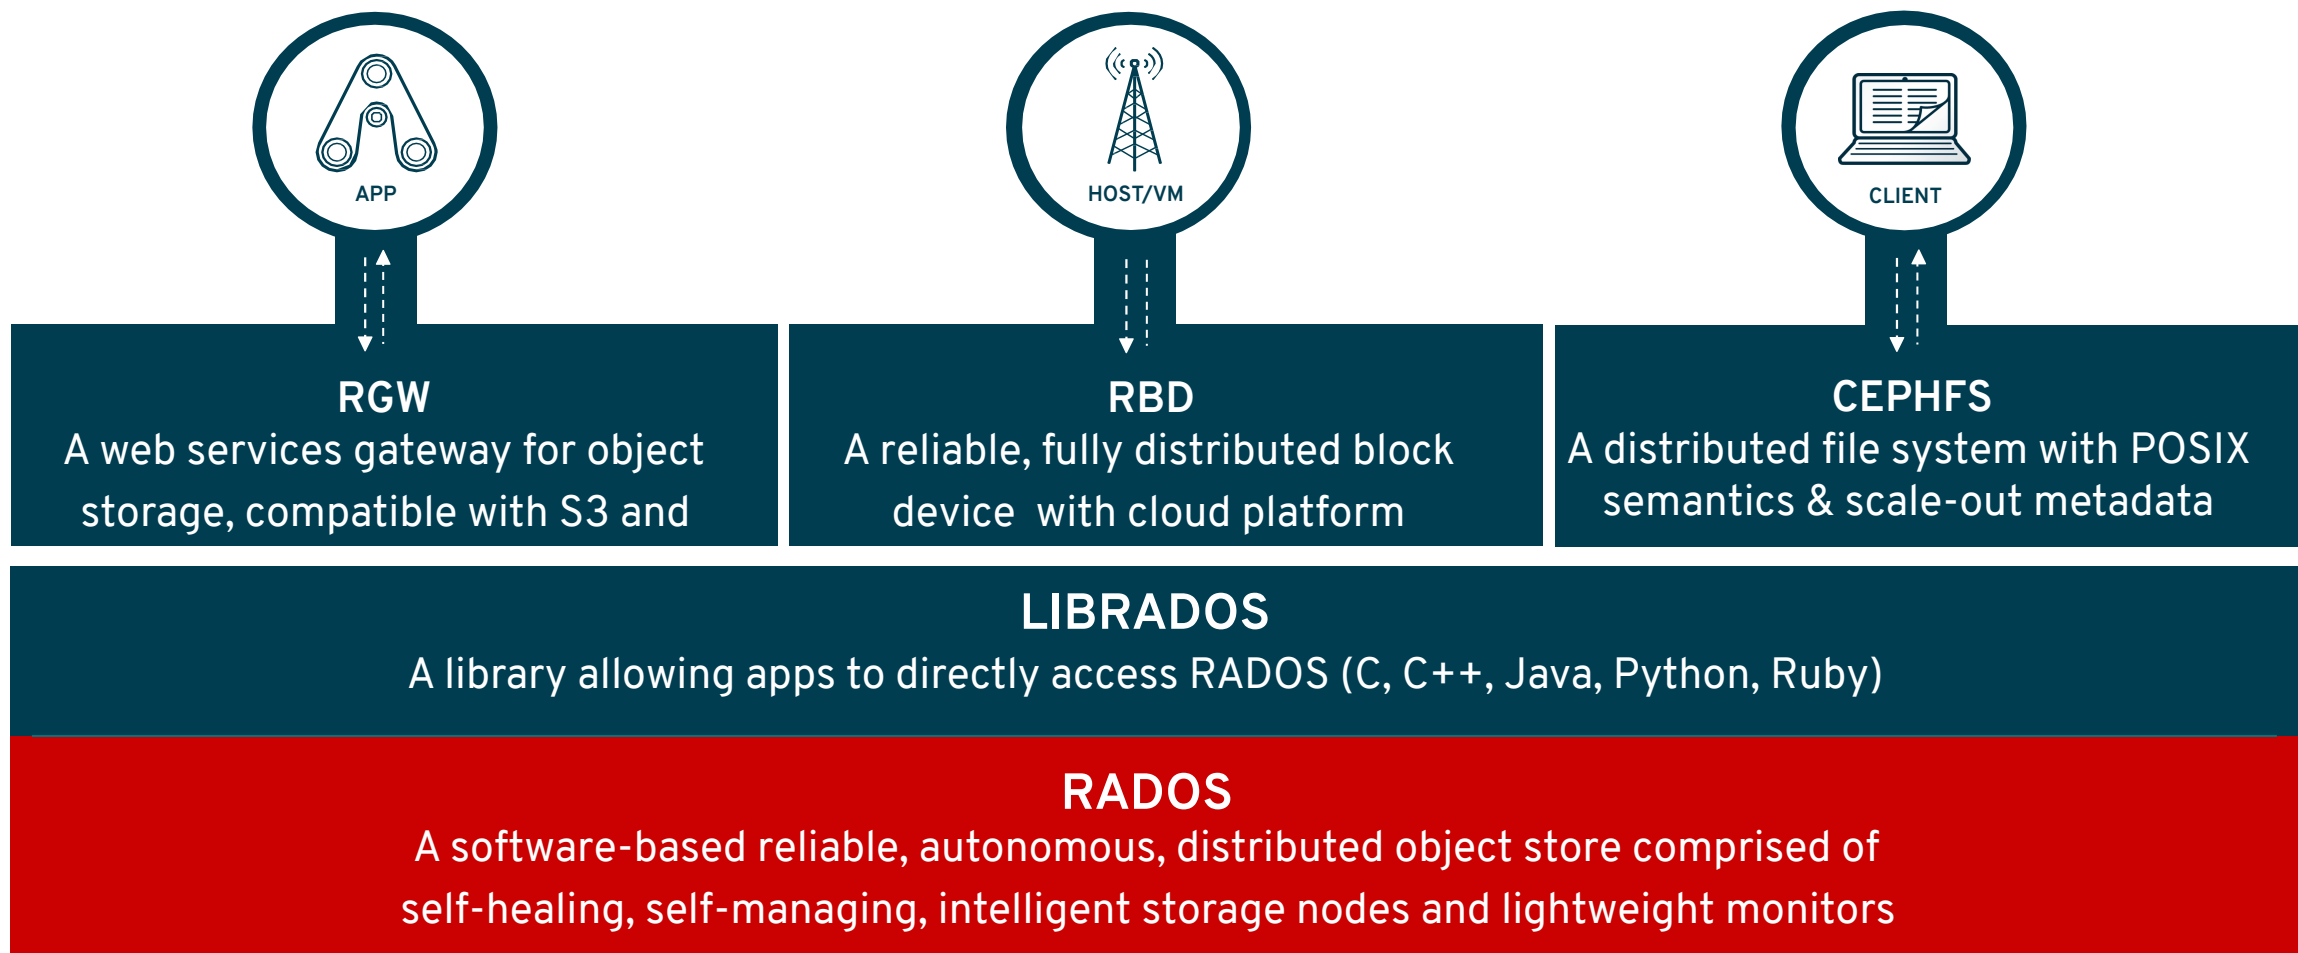
More economical than traditional NAS/SAN, particularly at petabyte scale

# RED HAT CEPH STORAGE

Distributed, enterprise-grade object storage, proven at web scale

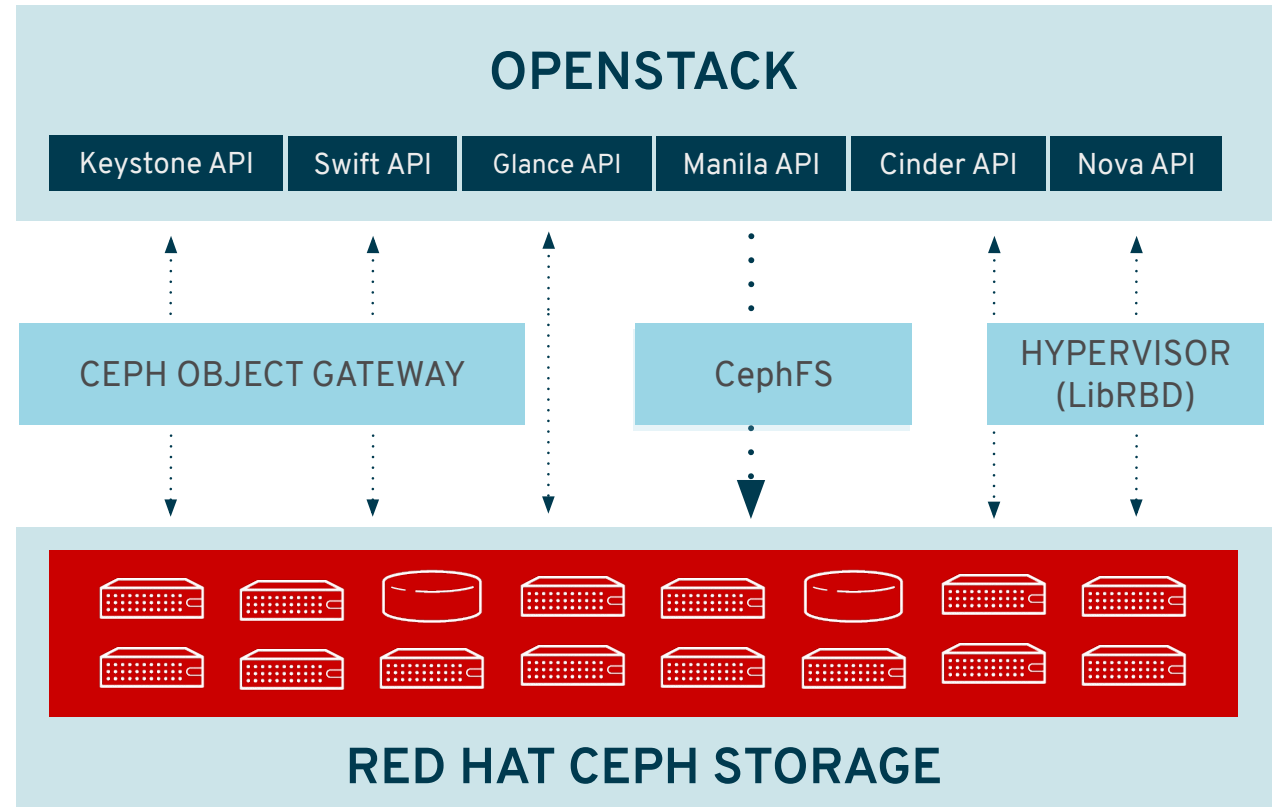
- Open, massively-scalable, software-defined
- Flexible, scale-out architecture on clustered standard hardware
- Single, efficient, unified storage platform (object, block, file)
- User-driven storage lifecycle management with 100% API coverage
  - S3 compatible object API
- Designed for modern workloads like cloud infrastructure and data lakes

# ARCHITECTURAL COMPONENTS



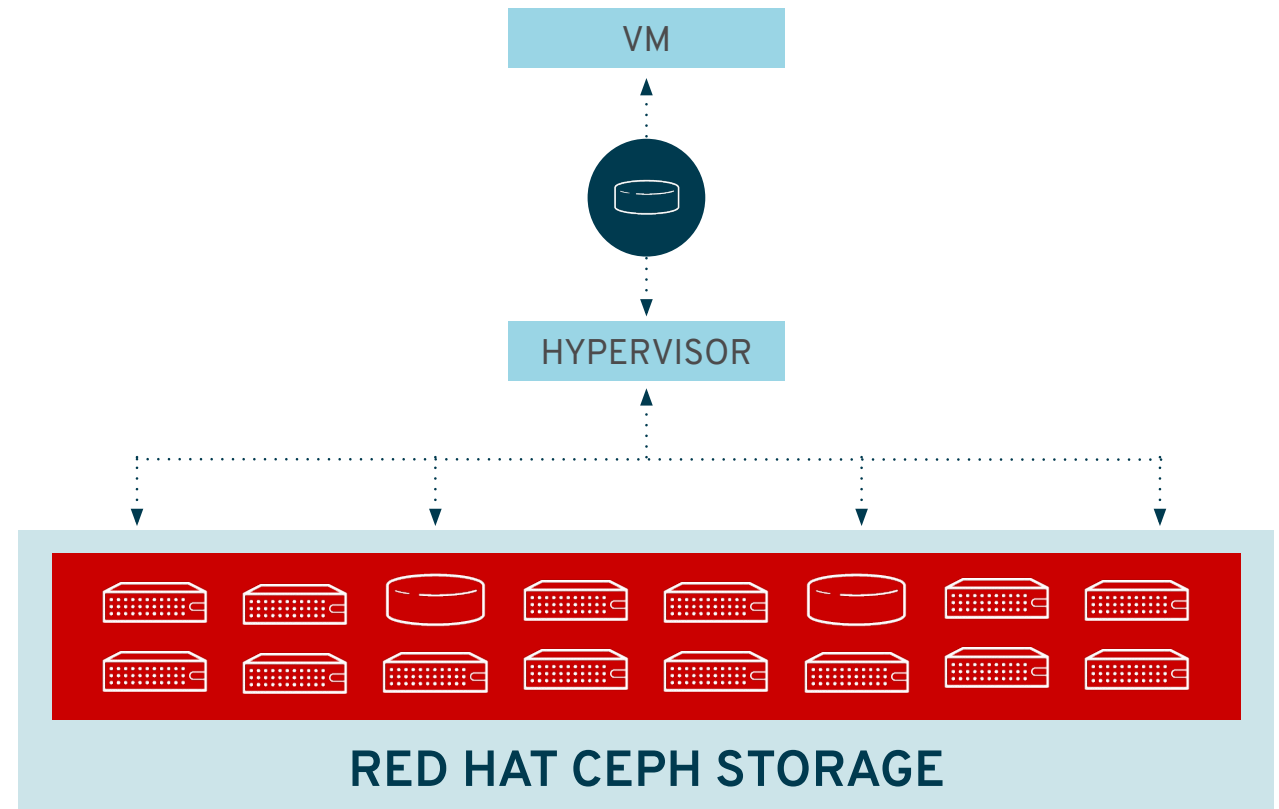
# COMPLETE OPENSTACK STORAGE

- Deeply integrated with modular architecture and components for ephemeral & persistent storage
  - Nova, Cinder, Manila, Glance, Keystone, Ceilometer, Swift, Gnocchi
- 64 TB bundled 'once' in Red Hat OpenStack Platform for evaluation



# ADVANTAGES FOR OPENSTACK USERS

- Instantaneous booting of 1 or 100s of VMs
- Instant backups via seamless data migration between Glance, Cinder, Nova
- Tiered I/O performance within single cluster
- Multi-site replication for disaster recovery or archiving

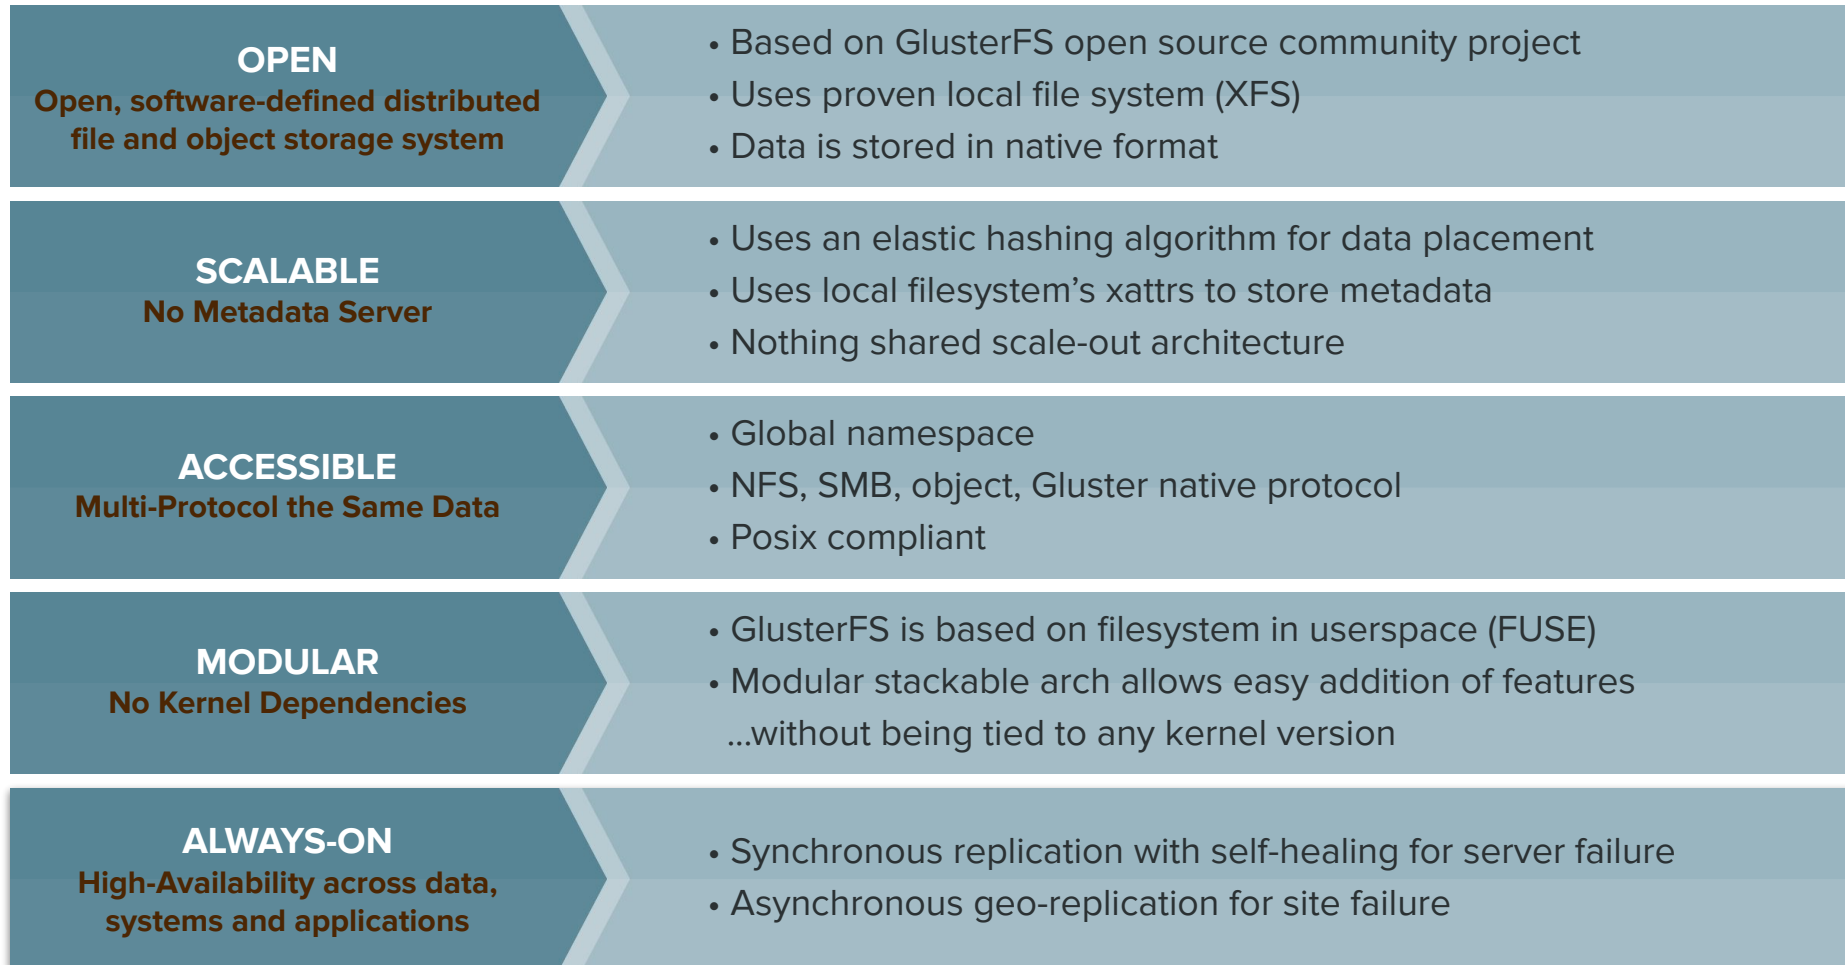


# TARGET USE CASES

- **Private Cloud** - enterprise deployments growing for test & dev and production application deployments. FSI, retail and technology sectors.
- **Archive & Backup**: object storage as a replacement for tape and expensive dedicated appliances. Hybrid cloud compatibility critical.
- **NFVi (new)** - OpenStack with Ceph dominant reference platform for next-generation telco networks. Global demand for Ceph now standalone and hyperconverged.
- **Enterprise Virtualization (new)**: legacy protocol support for legacy VM storage to be managed on same platform as modern, private cloud storage.
- **Big Data (new)** - object storage providing common, data lake for multiple analytics applications for greater efficiencies and better business insights

# Red Hat Gluster Storage

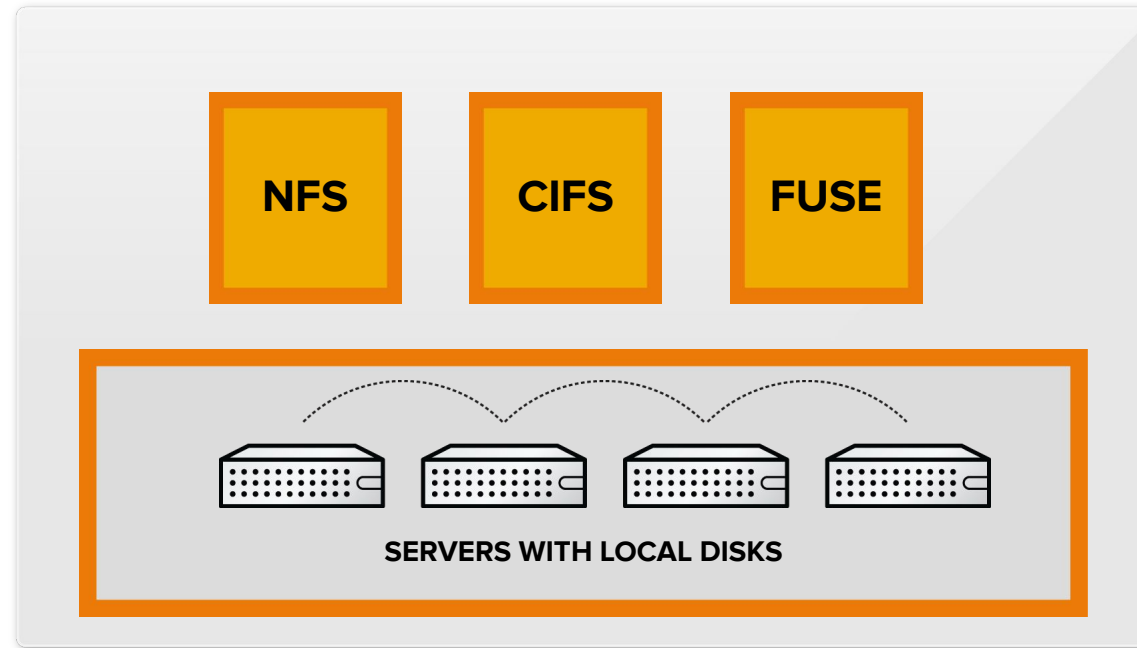
# RED HAT GLUSTER STORAGE ADVANTAGES





# GLUSTER ARCHITECTURE

Distributed scale out storage using industry standard hardware




Aggregates systems to one cohesive unit and presents using common protocols

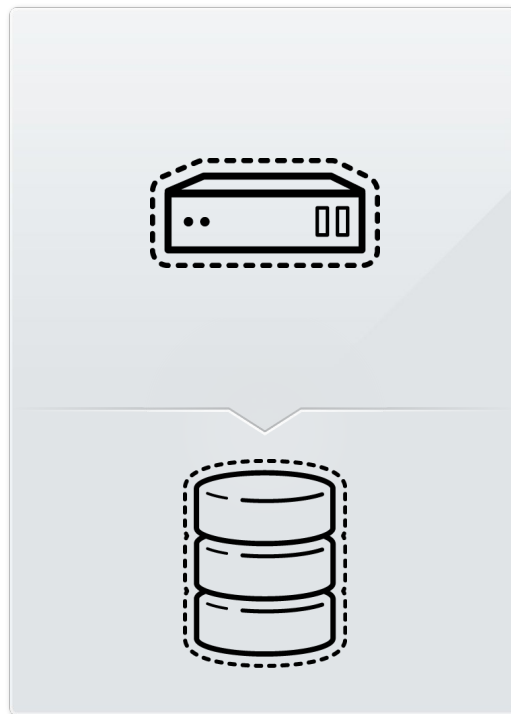
# WHAT IS A SYSTEM?

Can be physical, virtual or cloud

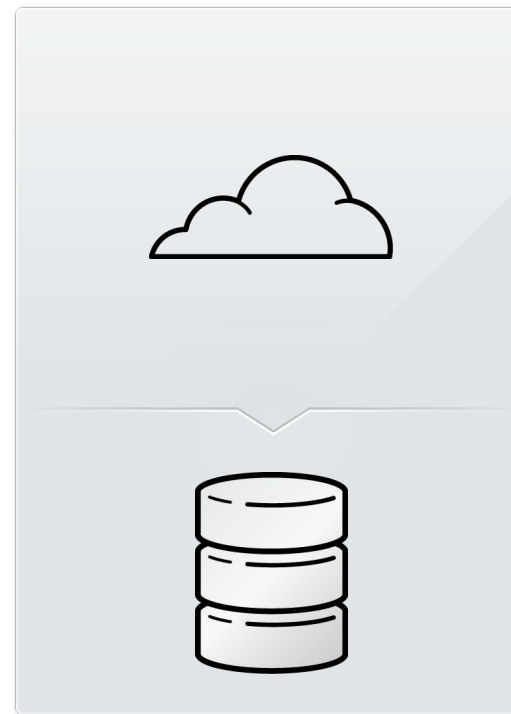
## PHYSICAL



## VIRTUAL



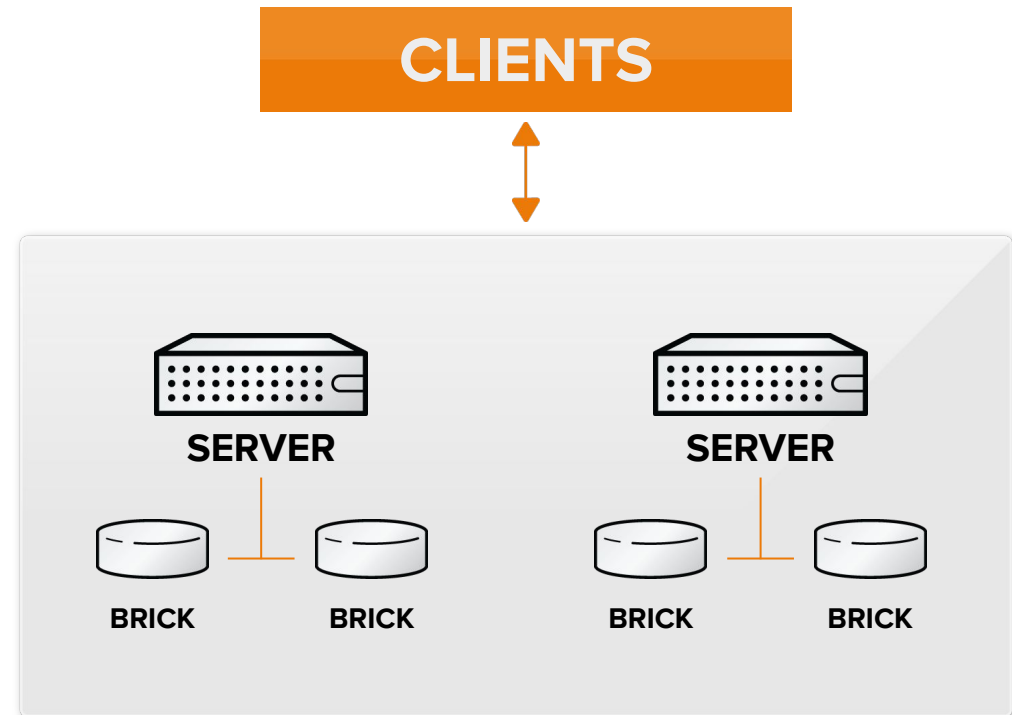
## CLOUD



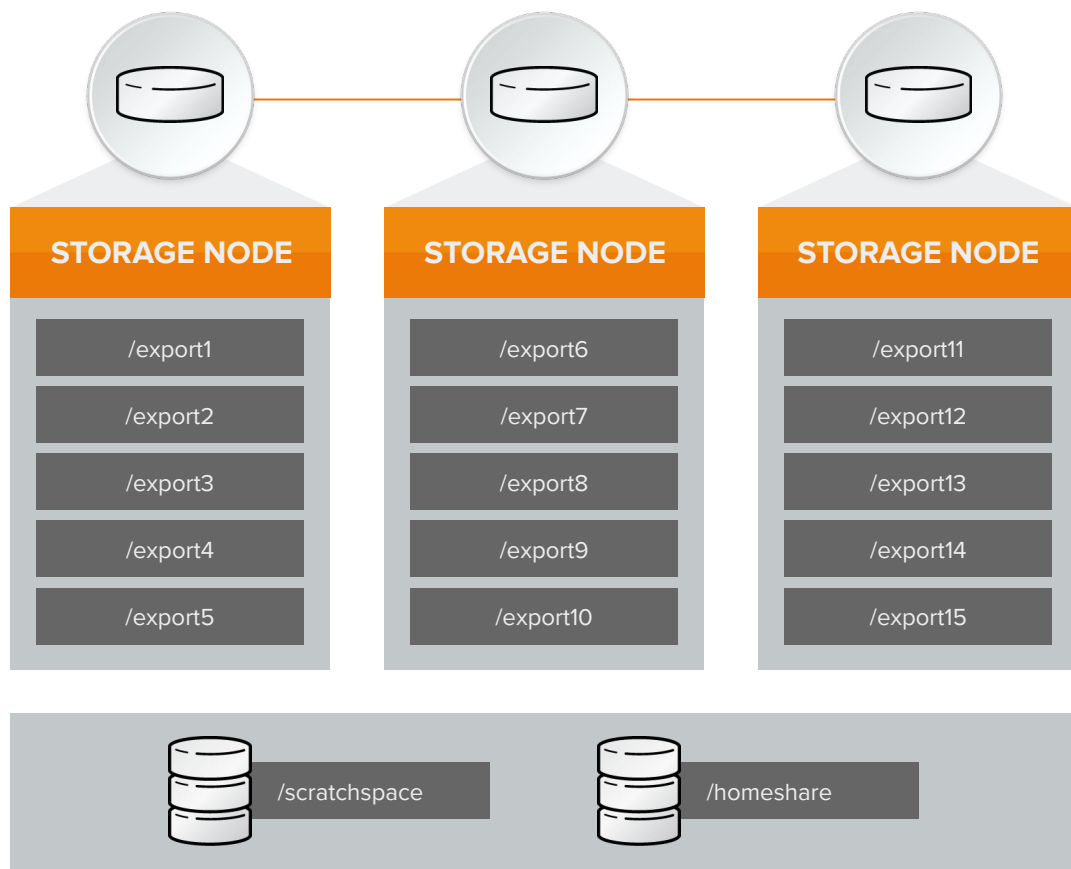
# VOLUMES

Bricks taken from multiple hosts become one addressable unit

- High availability as needed
- Load balanced data
- Managed by Gluster



# VOLUMES



**A Volume is some number of bricks >1,  
exported by RHGS**

- Volumes have administrators assigned names
- A brick can be a member of one volume
- Data in different volumes physically exists on different bricks
- Volumes can be mounted on clients

# NFS

## Accessibility from UNIX and Linux systems

- **STANDARD NFS** connects to NFS Ganesha process on storage node
- **MOUNT GlusterFS VOLUME** from any storage node
- **NFS GANESHA** includes network lock manager to synchronize locks
- **LOAD BALANCING** must be managed externally
- **STANDARD AUTOMOUNTER** is supported.
- **SUPPORTED FEATURES:** ACLs, NFSv4, Kerberos

**Better performance** reading many small files from a single client

# SAMBA/CIFS

## Accessibility from Windows systems

- **STORAGE NODE** uses Samba with winbind to connect with AD
- **SMB CLIENTS** can connect to any storage node running Samba
- **SMB VERSIONS** 2.0 & 3.0 supported
- **LOAD BALANCING** must be managed externally
- **CTDB** is required for Samba clustering

**Samba uses RHGS gfapi library** to communicate directly with GlusterFS server process without going through FUSE

# OBJECT ACCESS

## of GlusterFS Volumes

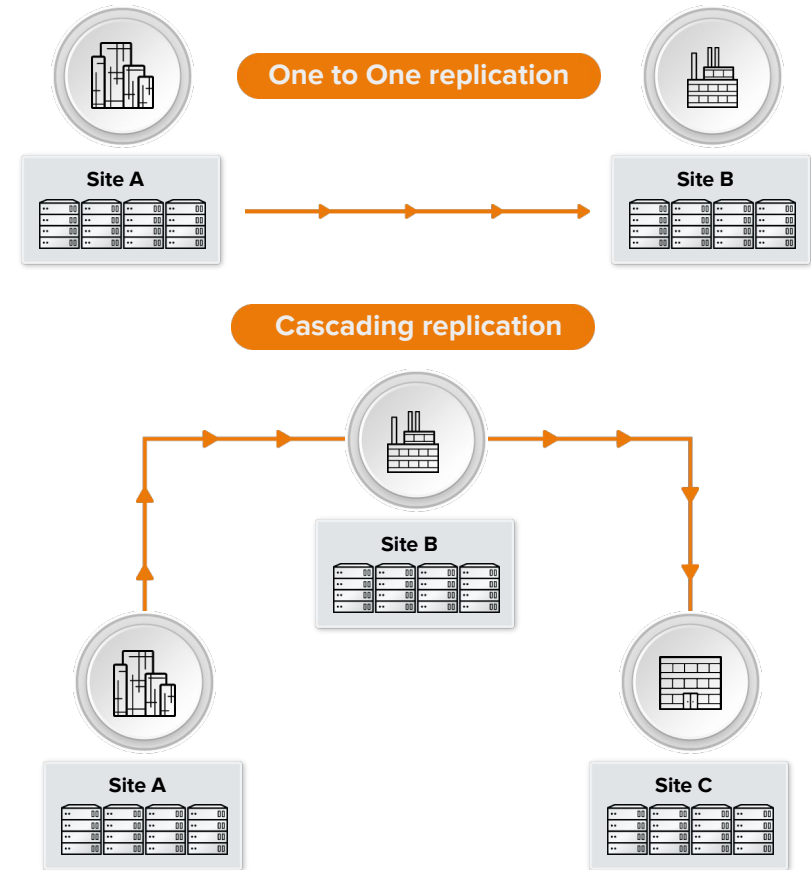
- **BUILT UPON** OpenStack's Swift object storage system
- **BACK-END FILE SYSTEM** for OpenStack Swift Accounts as GlusterFS volumes
- **STORE AND RETRIEVE** files using the REST interface
- **SUPPORT INTEGRATION** with SWAuth and Keystone authentication service

**Implements objects** as files and directories  
under the container ("Swift on File")

# GEO-REPLICATION

## Multi-site content distribution

- Asynchronous across LAN, WAN, or Internet
- Master-slave model, cascading possible
- Continuous and incremental
- Multiple configurations
  - One to one
  - One to many
  - Cascading





A vertical red graphic on the left side of the slide. It features various icons: a cloud with a keyhole, a database cylinder, a server rack, a window with a cursor, and several arrows pointing in different directions. There are also some 'X' and 'O' symbols scattered throughout the graphic.


# Red Hat OpenShift Container Storage

# WHAT STORAGE OPTIONS DO YOU HAVE FOR YOUR CONTAINERS?



## OUTDATED STORAGE ARRAYS AND APPLIANCES

- Scalability and high availability fall short of customer needs
- Vendor lock in and high TCO
- Monolithic appliance model



## SILOED OR POINT PLAY STORAGE SOLUTIONS

- No hybrid cloud support
- No unified control plane (K8s)
- Lack of automated or dynamic provisioning of storage



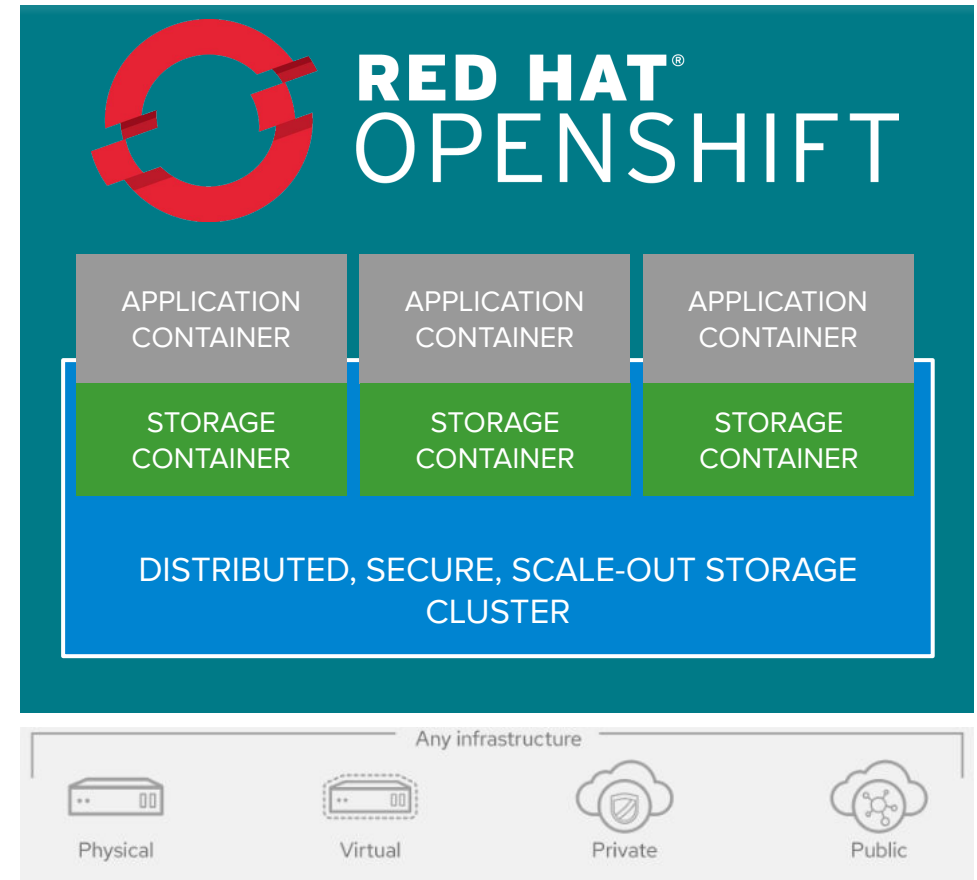
## PORTABLE STORAGE ACROSS ON-PREM / PUBLIC CLOUDS

### **RED HAT** OPENSIFT Container Storage

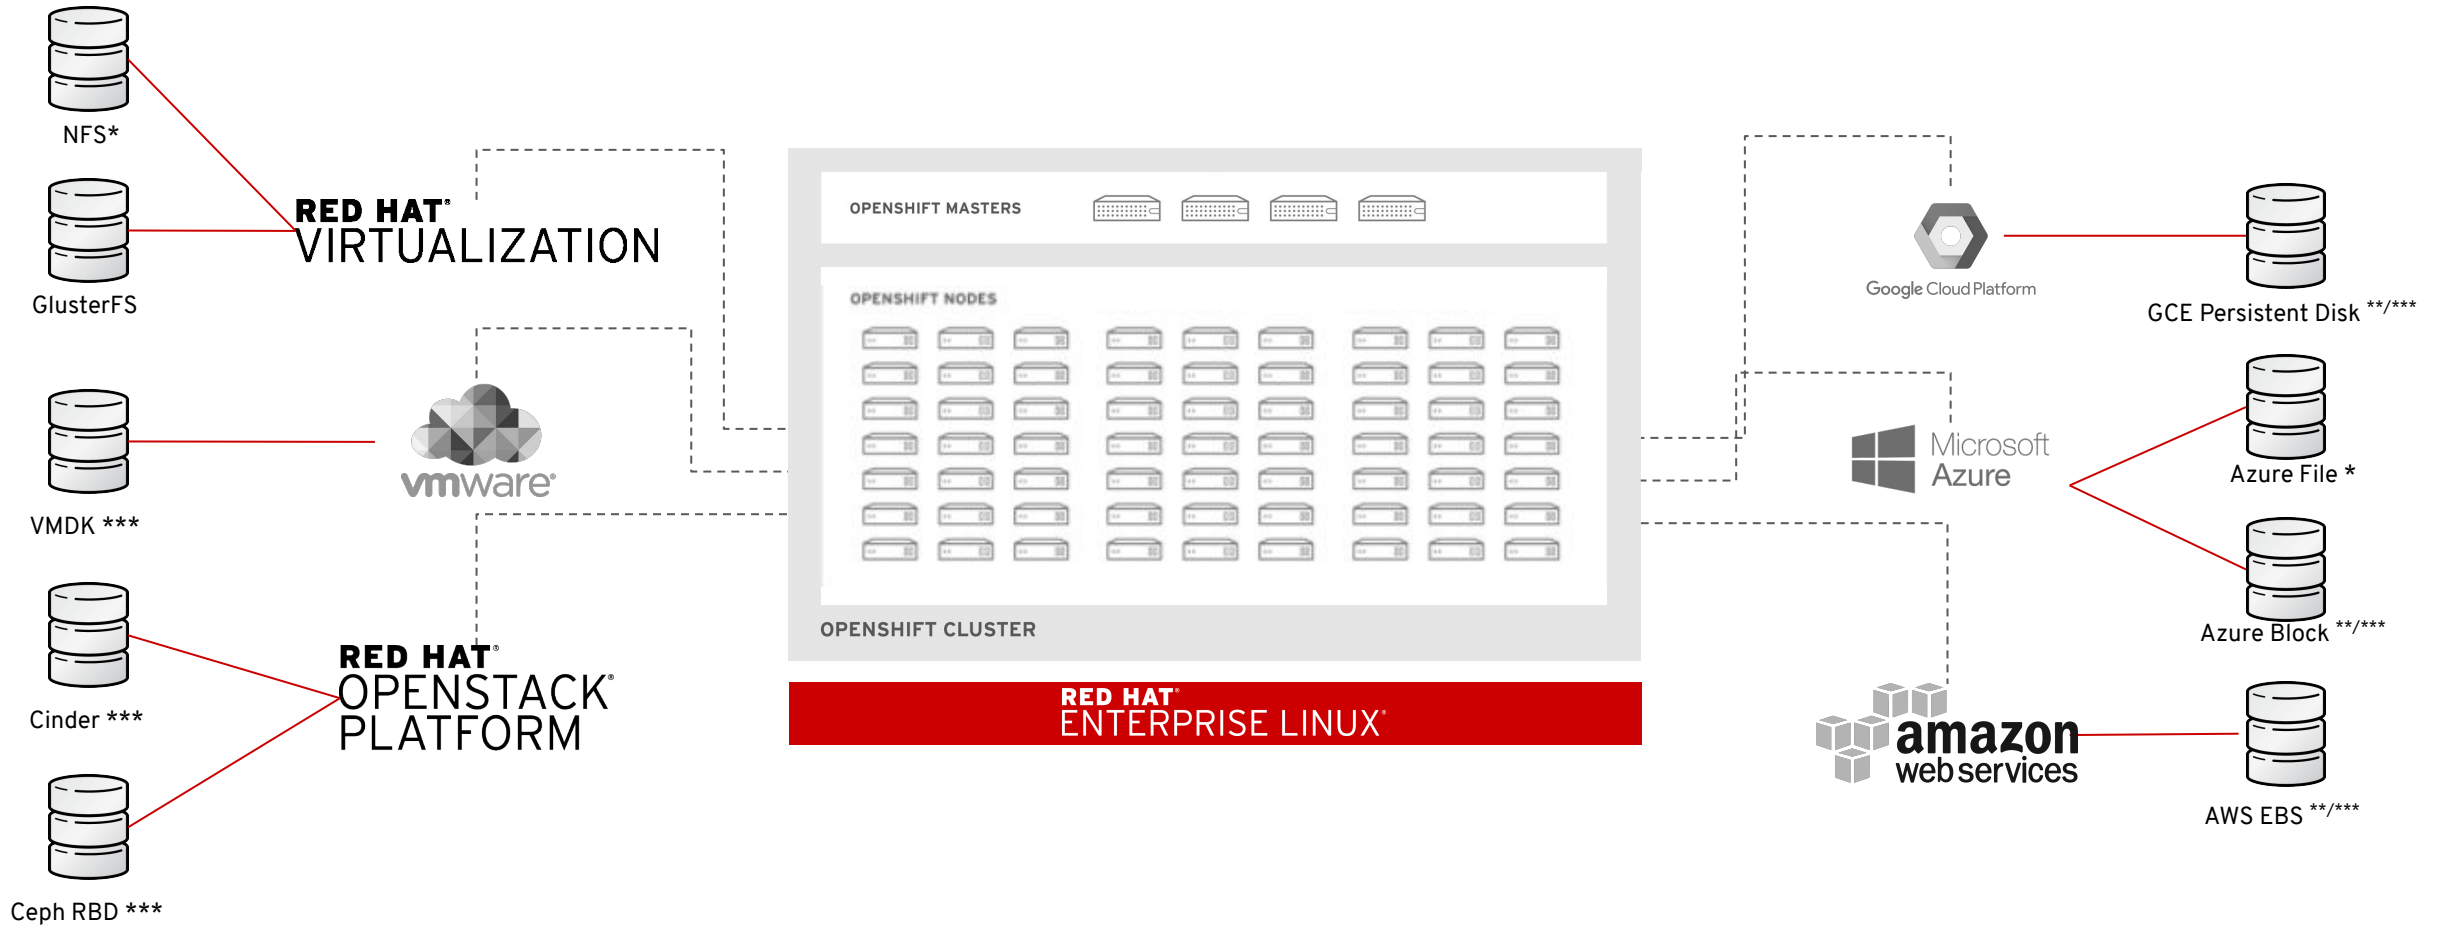
- Runs to all 3 public clouds
- Single Vendor Support model
- Seamless user experience for developers and DevOps

# WHAT IS OPENSIFT CONTAINER STORAGE?

- Highly scalable, enterprise-grade storage
- 2 Deployment options: independent and converge mode
- Native integration with OpenShift
- Unified Orchestration using Kubernetes for applications and storage
- Greater control & ease of use for developers
- Developed, maintained, and deployed in synchrony with Red Hat OpenShift releases
- Single vendor Support



# STORAGE SUPPORT BEFORE OCS



\* No dynamic provisioning

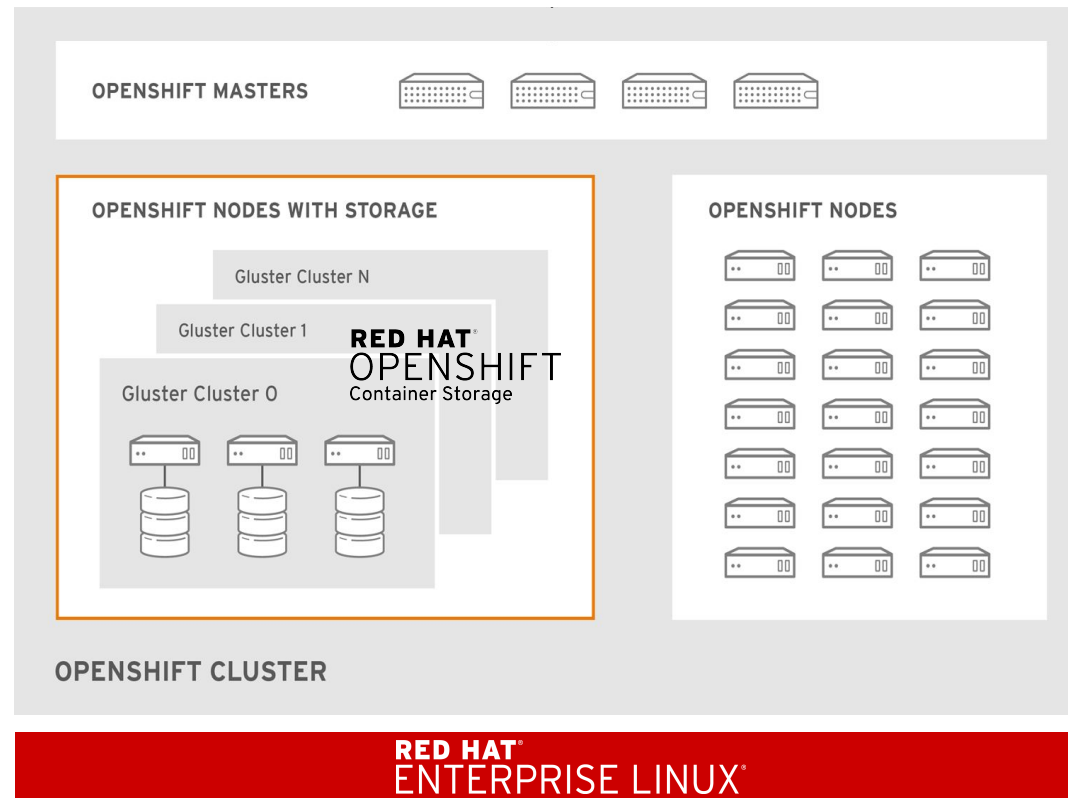
\*\* No Cross-Availability Zone Support

\*\*\* No Shared-Storage

# STORAGE SUPPORT AFTER OCS

- **Scalable**  
(1000+ volumes)
- **Highly-Available**  
(across availability zones)

*Persistent, resilient and elastic storage...*



- **Automated**  
(Dynamic Provisioning)
- **Integrated**  
(installs with / runs on OpenShift)

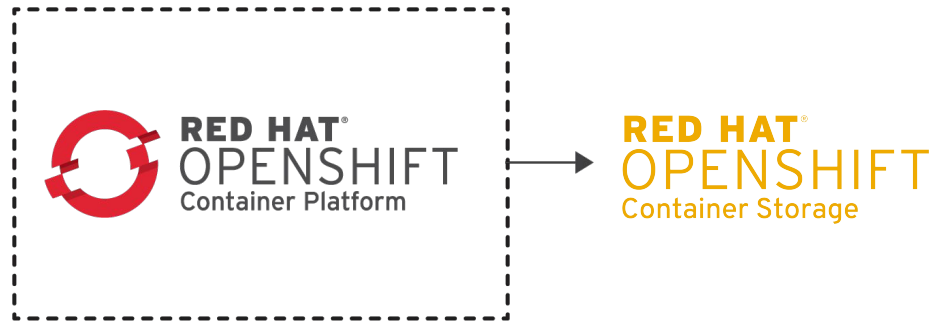
*... that travels with the platform.*

\* No dynamic provisioning

\*\* No Cross-Availability Zone Support


\*\*\* No Shared-Storage

# TWO FLAVORS OF CONTAINER STORAGE



## **OPENSIFT CONTAINER STORAGE INDEPENDENT MODE**

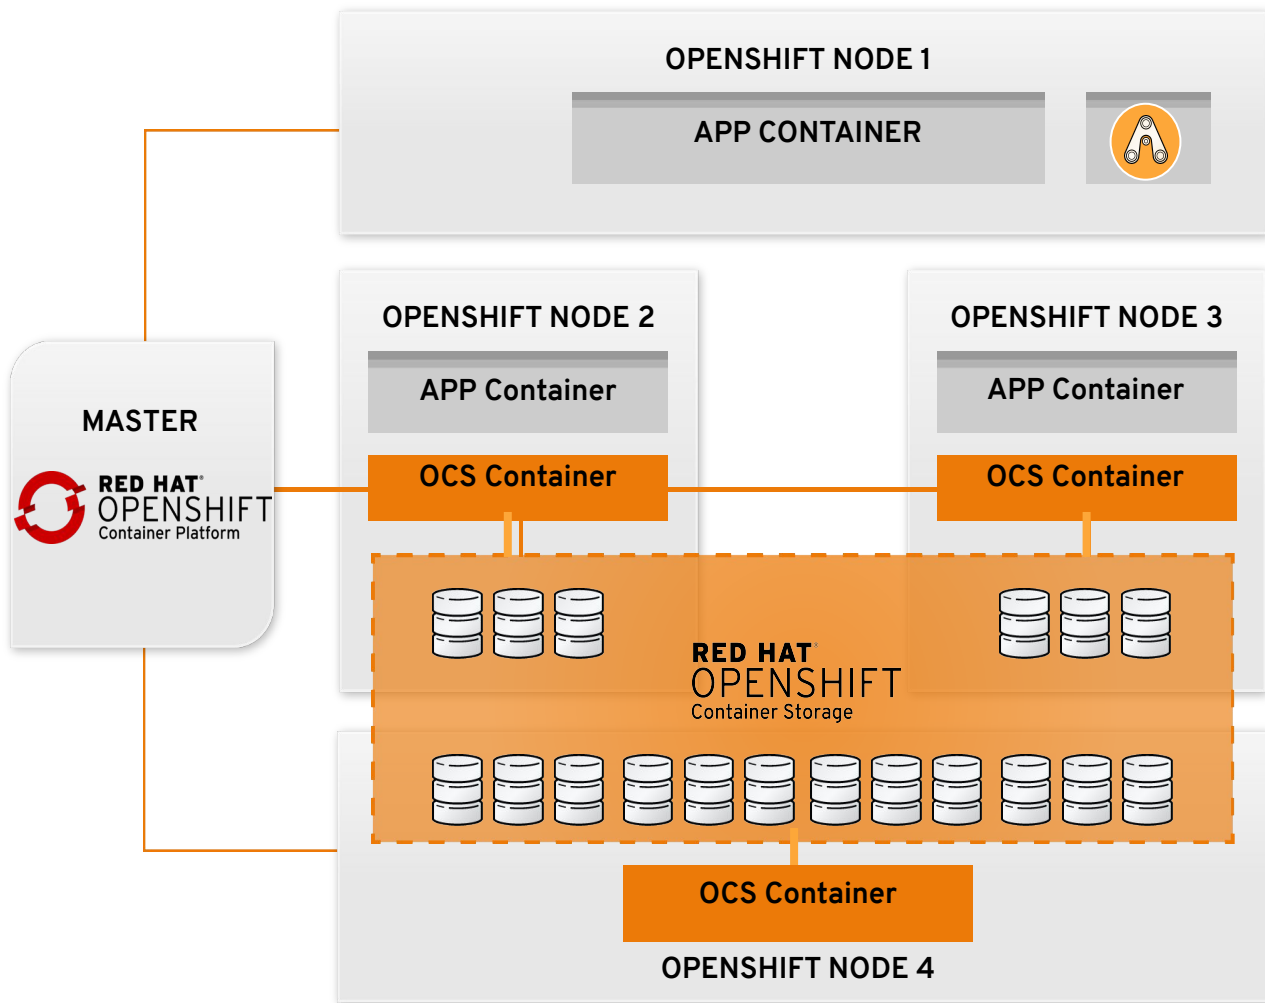
Use existing investment in traditional storage, managed by storage admin – attach to standalone storage



## **OPENSIFT CONTAINER STORAGE CONVERGED MODE**

Highly scalable, enterprise-grade storage, fully integrated into OpenShift Container Platform

# CONVERGED MODE



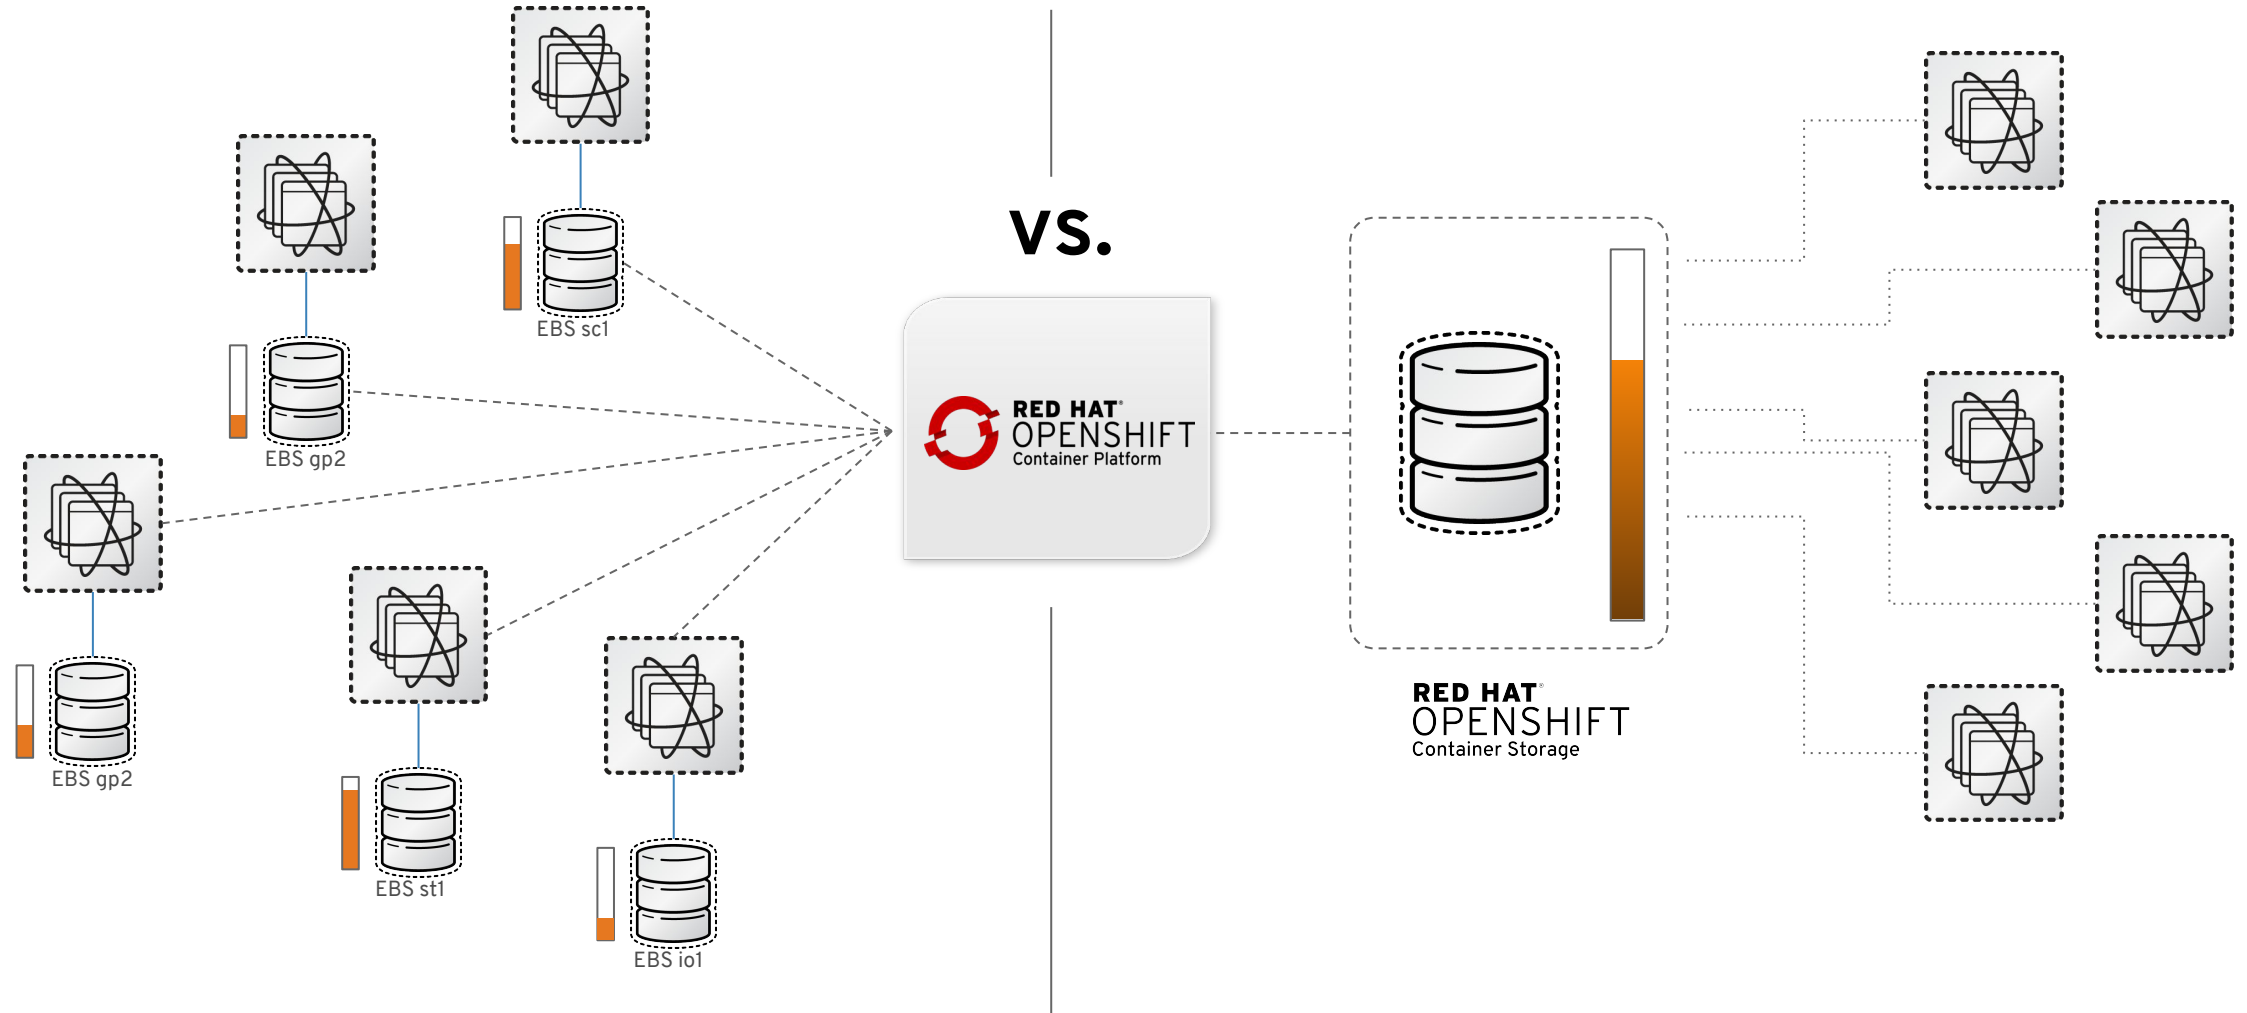
Co-Locate Storage and Apps

Dynamic Provisioning

Managed by OpenShift

Infrastructure-Agnostic

# STORAGE CONSOLIDATION IN THE CLOUD





# STORAGE CONSOLIDATION ON PREM

PERSISTENT VOLUMES PROVIDED BY OPENSIFT CONTAINER STORAGE

**RED HAT**  
**OPENSIFT**  
Container Storage



OPENSIFT  
CONTAINER  
STORAGE  
POD



BLOCK  
DEVICE



OPENSIFT  
CONTAINER  
STORAGE  
POD



BLOCK  
DEVICE



OPENSIFT  
CONTAINER  
STORAGE  
POD



BLOCK  
DEVICE



OPENSIFT  
CONTAINER  
STORAGE  
POD



BLOCK  
DEVICE




OPENSIFT  
CONTAINER  
STORAGE  
POD




BLOCK  
DEVICE



FIBRE-CHANNEL  
ARRAY




ISCSI SAN




SHARED SAS

# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 [linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)

 [facebook.com/redhatinc](https://www.facebook.com/redhatinc)

 [youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)

 [twitter.com/RedHat](https://twitter.com/RedHat)

