Red Hat Gluster Storage Roadmap - Past, Present & Future

Red Hat Gluster Storage Roadmap

Sayan Saha
Director, Product Management, Storage Business

Vijay Bellur
Senior Principal Software Engineer

02-May-2017

Why?

THE RED HAT STORAGE ECOSYSTEM

Standard interfaces and full APIs ease integration with applications and systems

Self-managing and self-healing software provides durability and adapts to changes

Industry-standard hardware provides choice and can be tailored for specific workloads

RED HAT STORAGE IS DEEPLY INTEGRATED

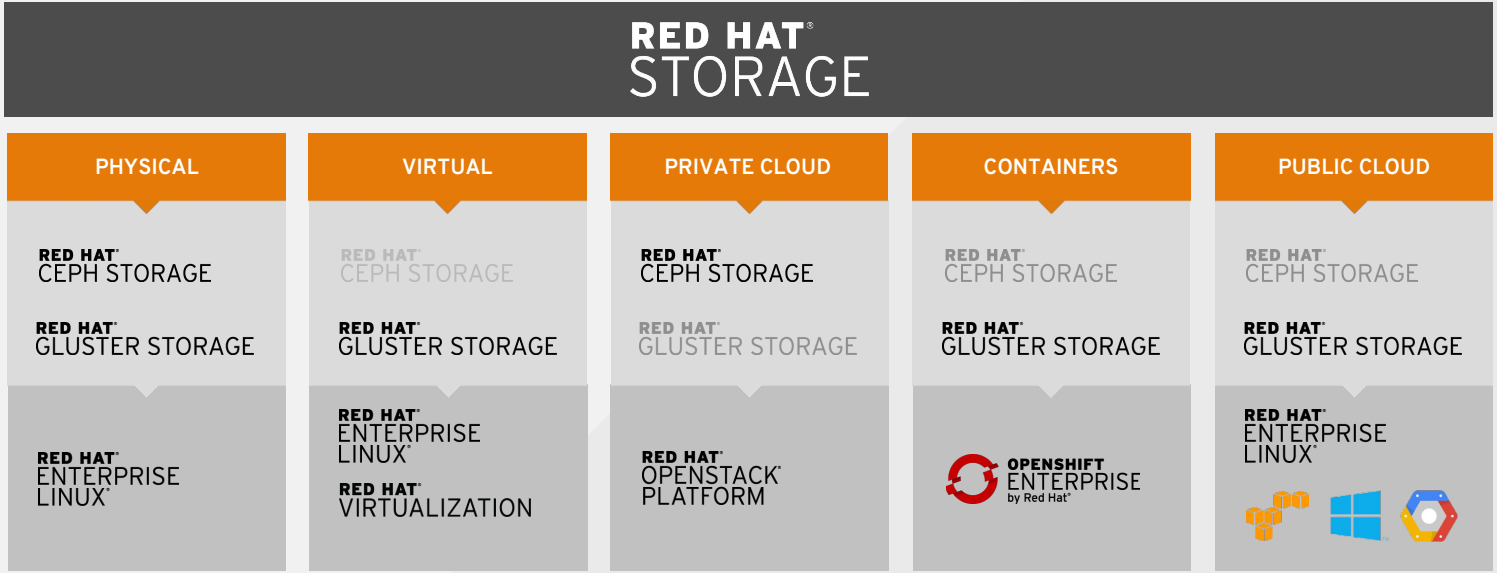

Red Hat Gluster Storage

RED HAT® GLUSTER STORAGE

TARGET USE CASES

Container-Native Storage
- Persistent storage
- Containerized storage

Enterprise File Sharing
- Media streaming
- Active Archives

Enterprise Virtualization

Rich Media and Archival

Shared-nothing, scale out file storage

- Purpose-built as a distributed file store with a straightforward architecture suitable for public, private, and hybrid cloud
- Simple to install and configure, with a minimal hardware footprint
- Offers mature NFS, SMB and native interfaces for enterprise use

Thriving Community

Hundreds of customers running in production

GLUSTER ARCHITECTURE

Distributed scale out storage using industry standard hardware

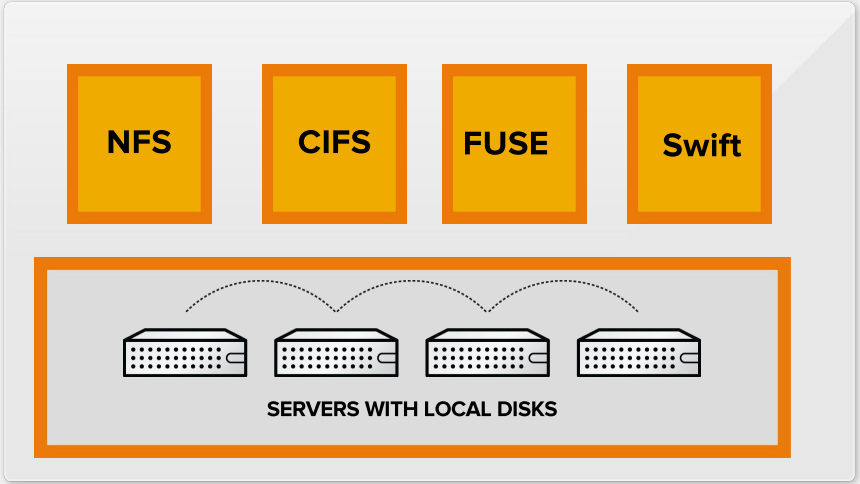

Aggregates systems to one cohesive unit and presents using common protocols

PAST

Gluster's journey @ Red Hat

Stability & Performance
2011–2013

Enterprise-Class Features
2014–Present

Enabling container-native storage & hyper-convergence
Now

A STEADY FLOW OF INNOVATION

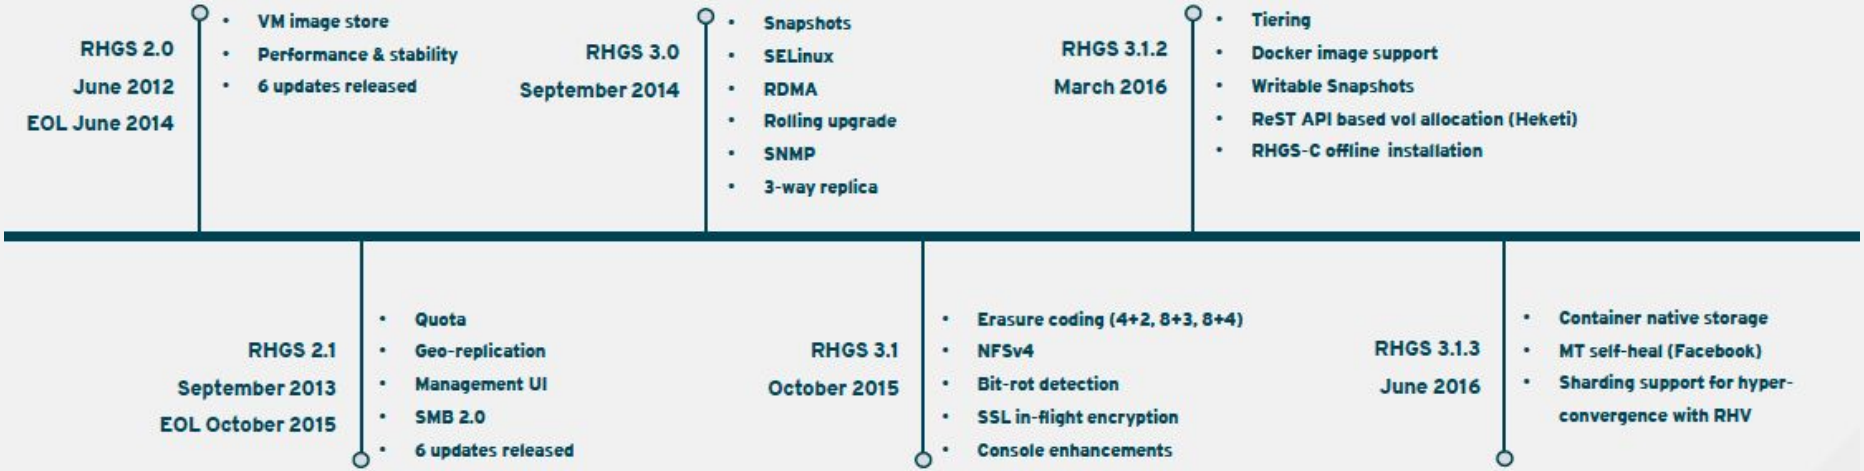

Key Features: Red Hat Gluster Storage

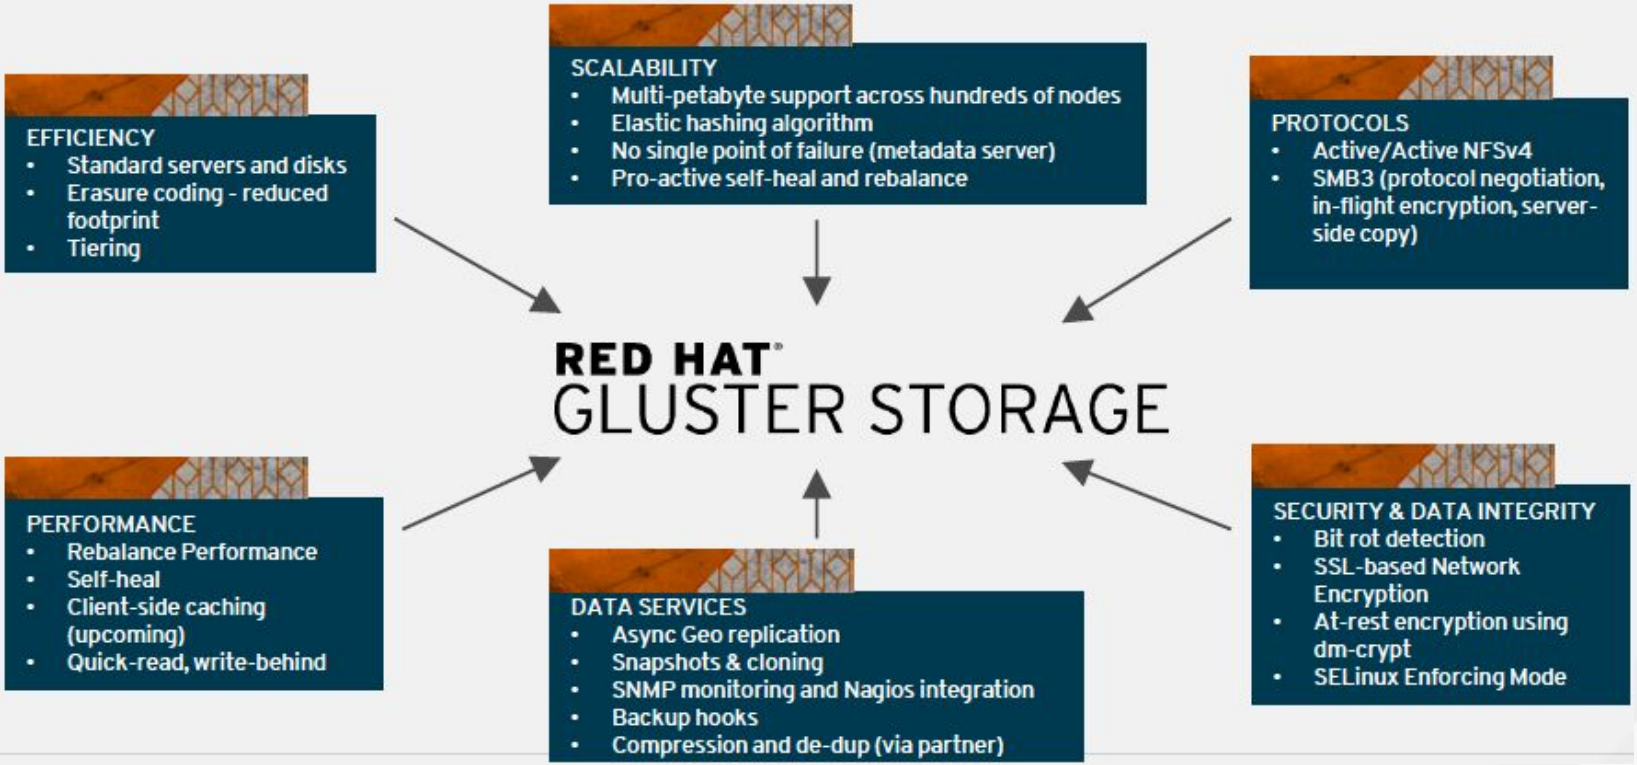

Flexible Deployment Options

- Same software bits across on-prem, VMs, containers and all three public clouds
- Applications can be ported across deployments without expensive re-writes
- Close integration with RHEL, RHV, and OpenShift

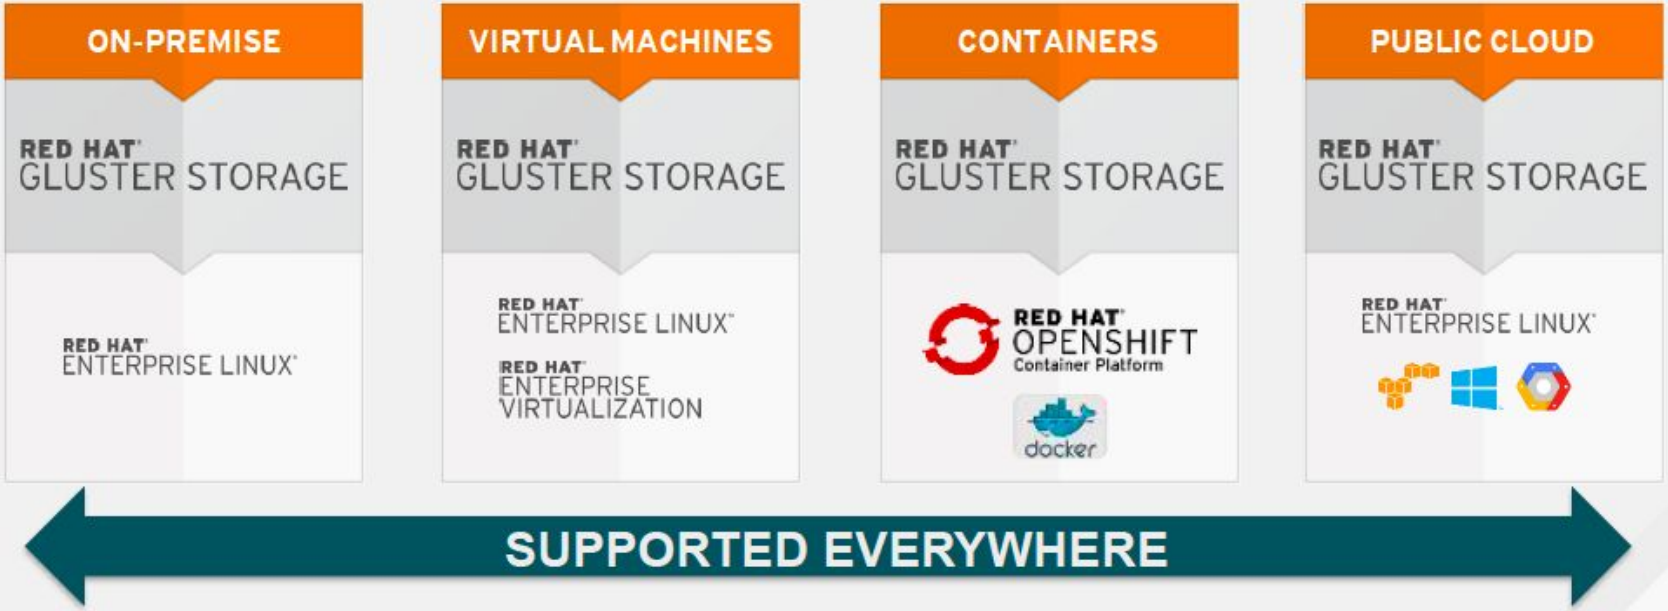

RED HAT RECOGNIZED AS A STORAGE VISIONARY

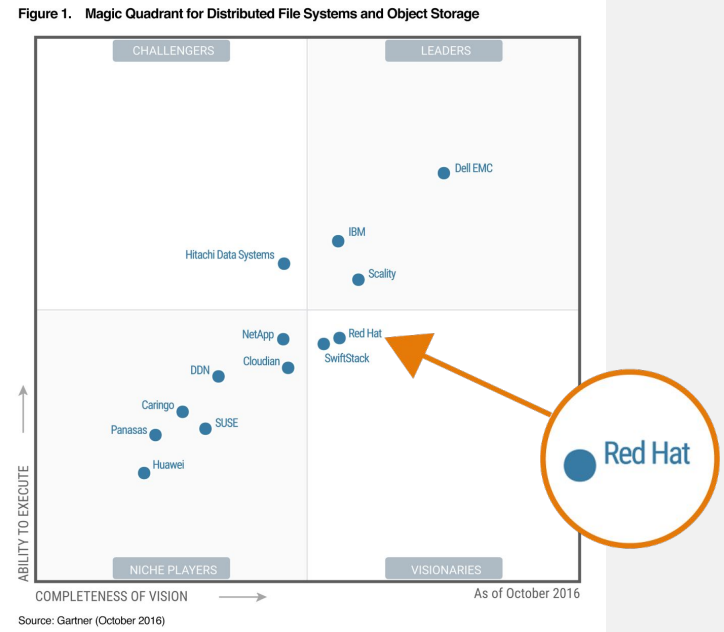

Red Hat Storage recognized as a Visionary
by Gartner in their first Magic Quadrant for Distributed File Systems and Object Storage.

This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request at https://engage.redhat.com/gartnermagic-quadrant-storage-s-201610121525

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

Customer Momentum

"Red Hat Gluster Storage met our requirements in terms of being a **scalable**, standard-based and **future-secure solution** based on proven technology from a trusted vendor, and adhering to our **high demands on accessibility and security**."
~ Anders Hellquist, Infrastructure Architect

"Our costs, including various **procurement costs and operating fees**, fell to less than half of what we had been before implementing Red Hat Gluster Storage. The solution's **flexibility** enabled us to build a storage environment using commodity servers and its **ease of operational control** was also a major advantage."
~ Kazuyasu Yamazaki, Group Mgr, IT Infrastructure Group

Customer Momentum

Across industry verticals, geographies, and use cases

BRINKER INTERNATIONAL

intuit

FINRA

COX AUTOMOTIVE™

Autotrader

verizon

MACQUARIE BANK

BRITISH COLUMBIA

OSLO BØRS

PRESENT & FUTURE

Strong interest and sales in

- Software-defined alternative for File/NAS
- Container Storage across hybrid cloud
- Hyper-converged storage
- POSIX compatible Storage in Public Cloud

Focus Area

Software-defined alternative for File/NAS

Container Storage across Hybrid Cloud

Hyper-converged storage

POSIX compatible Storage in Public Cloud

Software-defined alternative for File/NAS

Why Red Hat Gluster Storage?

Superior price-performance
80% features, 30% cost
Migrate away from dead-end NAS appliances
Comprehensive public cloud support!

Use-Case Focus

MEDIA REPOSITORY

like photos and videos, at enormously large scale

BACKUP

where elasticity and recovery time is critical

SURVEILLANCE

storage for customers with 1000+ cameras

Roadmap - RHGS offering for File/NAS

Red Hat Gluster Storage (2017)

RHGS 3.2 (Mar 2017)

- Arbiter Volumes
- Metadata & Small files perf
- New capabilities in container
- Faster self-heal for EC & sharded volumes

Roadmap - RHGS offering for File/NAS

Red Hat Gluster Storage (2017)

RHGS 3.2 (Mar 2017)

- Arbiter Volumes
- Metadata & Small files perf
- New capabilities in container
- Faster self-heal for EC & sharded volumes

RHGS 3.3 (Summer 2017)

- SMB performance enhancements (reduce lookups - negative lookup cache)
- Expanded EC configurations (8+2, 16+4 in addition to 4+2, 8+3 and 8+4)

Roadmap - RHGS offering for File/NAS

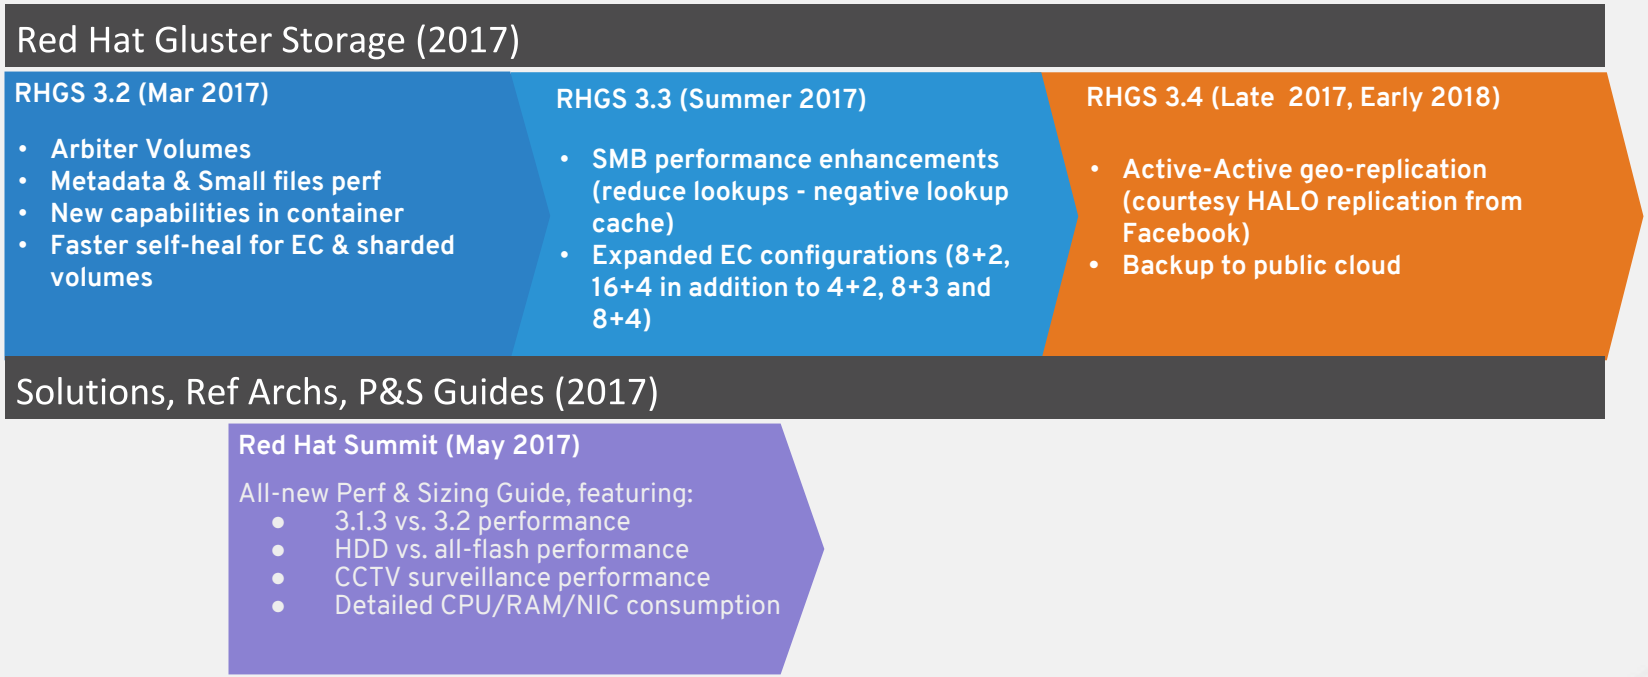

Let’s zoom in on some of the key upcoming features

Multi-master Replication (courtesy Facebook)

- Why?
  - Enable data federation across multiple-sites
- What?
  - Single Gluster volume spans multiple sites
  - Write/update a single Gluster volume from multiple geographical sites
  - Policy-based reconciliation of writes (last writer wins, one site always wins)
  - All reads are local, writes are propagated using self-heal to remote sites
- When?
  - End of CY 2017

GF Proxy (courtesy Facebook)

- Why?
  - Less frequent and minimally disruptive client upgrades for applications
  - Reduced churn as GlusterFS client changes less frequently
- What it is?
  - Thin, volume-topology aware client stack
  - Reduced footprint on the client
  - Distribution, Replication and Erasure Coding handled on the servers
- When?
  - CY 2018 with RHGS 4.0

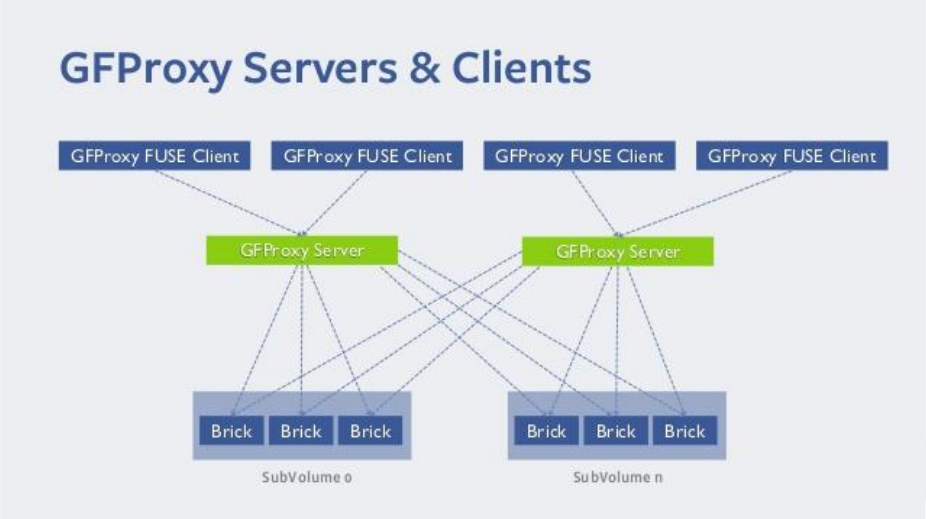

GlusterD 2.0 - Next Gen Management Plane

- Why?
  - Expand the number of nodes per cluster
  - Streamlined cluster membership (reduce chatter)
  - Strongly consistent state management
- What?
  - Highly scalable management plane
  - Autonomous management - low admin intervention for DevOps, containerized workloads (“give me 32 TBs of erasure coded storage”)
- When?
  - CY 2018 with RHGS 4.0

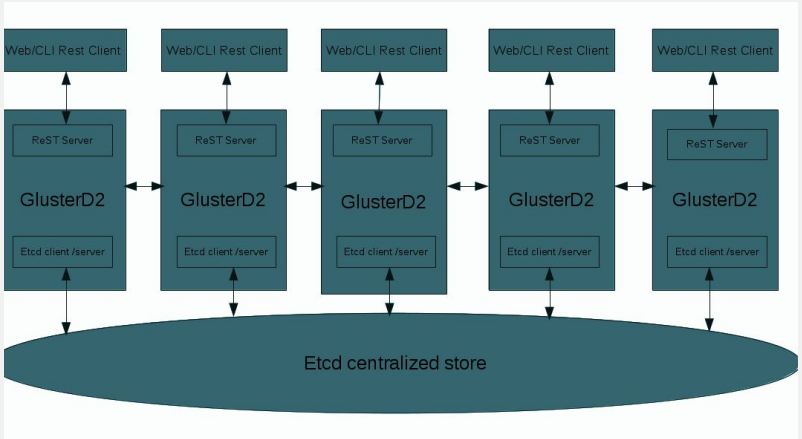

Backup data to object stores in Public Cloud

- Why?
  - Data management across hybrid cloud
- What it is?
  - Backup data in Gluster volumes to Amazon S3, Google Cloud Storage, Azure Storage etc.
  - Leverages changelog infrastructure used by geo-replication to synchronize data
  - HTTPs based data transfer vs ssh in geo-replication.
- When?
  - Late CY 2017/ Early CY 2018

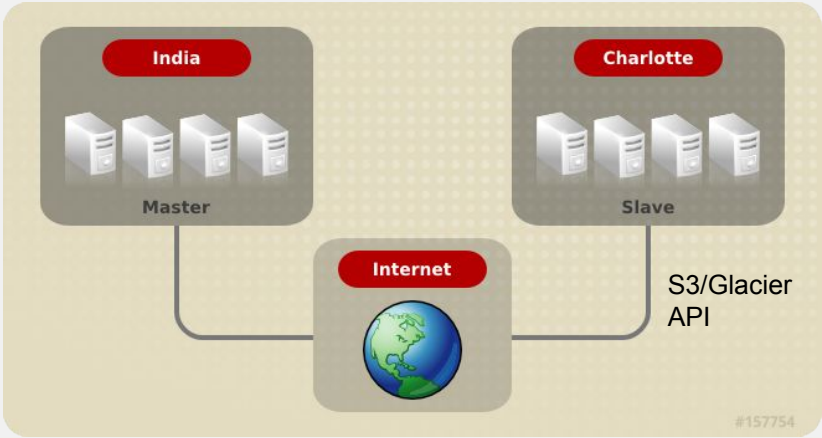

Focus Area

Software-defined alternative for File/NAS

Container Storage across hybrid Cloud

Hyper-converged storage

POSIX compatible Storage in Public Cloud

Container Storage across Hybrid Cloud

Why Red Hat Gluster Storage?

- RHGS provides consistent data management across the hybrid cloud for OpenShift - runs everywhere!
- Tight integration with RHGS running containerized inside OpenShift as kubernetes pods
- No need to throw away your existing SAN/NAS

Container Storage – Use Cases

Persistent Storage for containerized apps

OpenShift Infra (Registry)

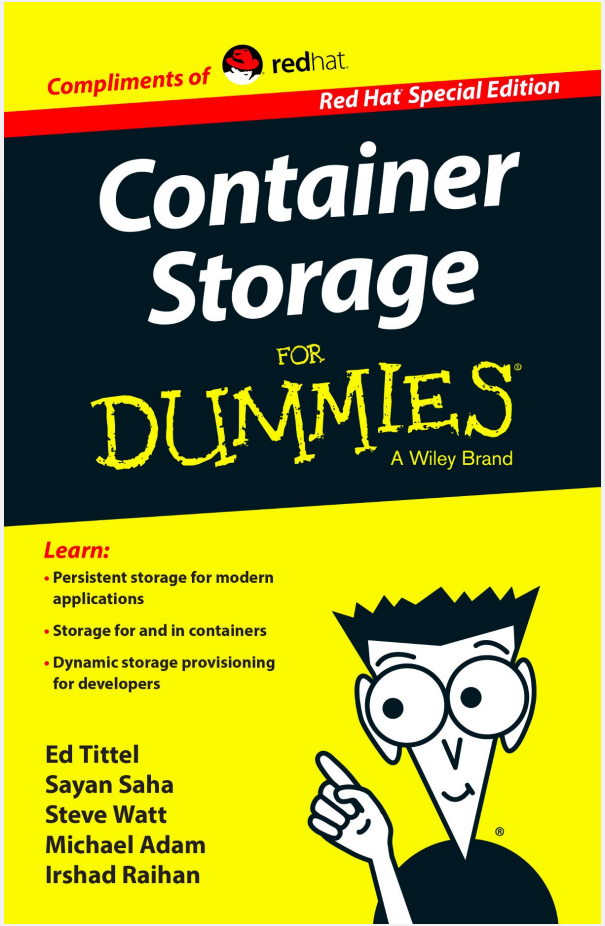
Compliments of redhat
Red Hat Special Edition
Container Storage
FOR
DUMMIES
A Wiley Brand
Learn:
• Persistent storage for modern applications
• Storage for and in containers
• Dynamic storage provisioning for developers
Ed Tittel
Sayan Saha
Steve Watt
Michael Adam
Irshad Raihan

Deployment Option-1: Container Ready Storage (CRS)

- Storage provided by a dedicated Red Hat Gluster Storage Cluster over the network outside OpenShift
- Red Hat Gluster Storage can run bare metal or VMs
- Storage and compute can scale independently

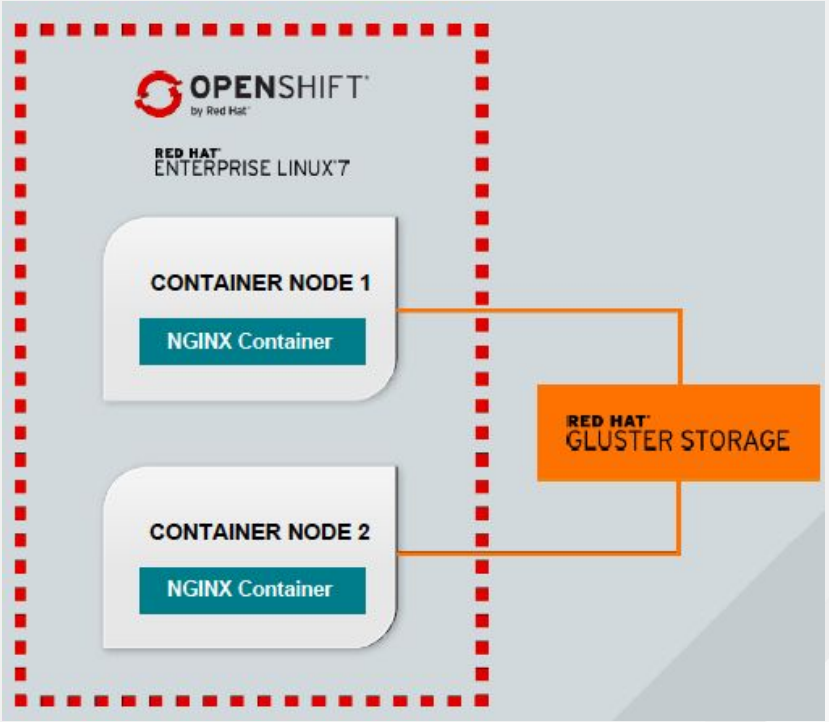

Deployment Option-2: Container-Native Storage (CNS)

- Fully featured scale-out storage platform (RHGS) runs inside OpenShift in containers (kubernetes pods)
- Application and storage containers can be co-located
- Storage services managed, scaled and upgraded like app containers

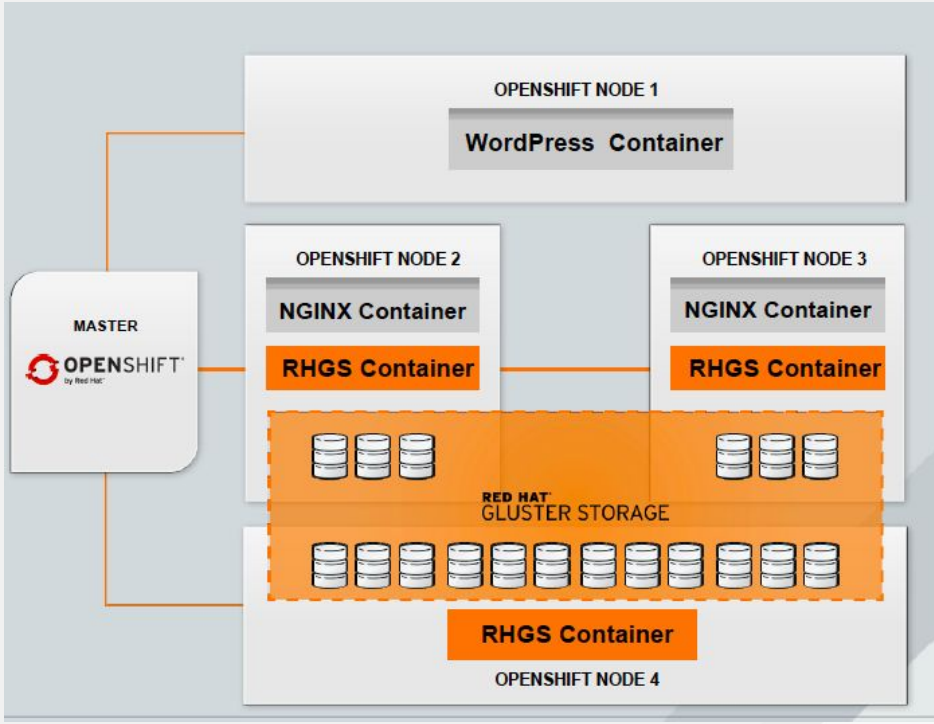

Zooming In on Deployment Options

- **Container Ready Storage (CRS) serving out storage to OpenShift with:**
  - RHGS in stand-alone bare-metal storage clusters
  - RHGS inside VMs fronting Enterprise Storage Arrays
  - RHGS in AWS, Azure and Google Cloud
- **Container-Native Storage (CNS)**
  - RHGS runs containerized inside OpenShift Container Platform (always with Heketi)
  - CNS runs anywhere OCP runs!

Container Storage Roadmap

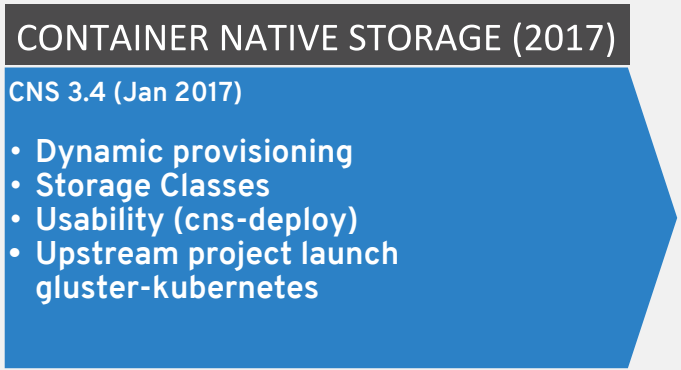

Container Storage Roadmap

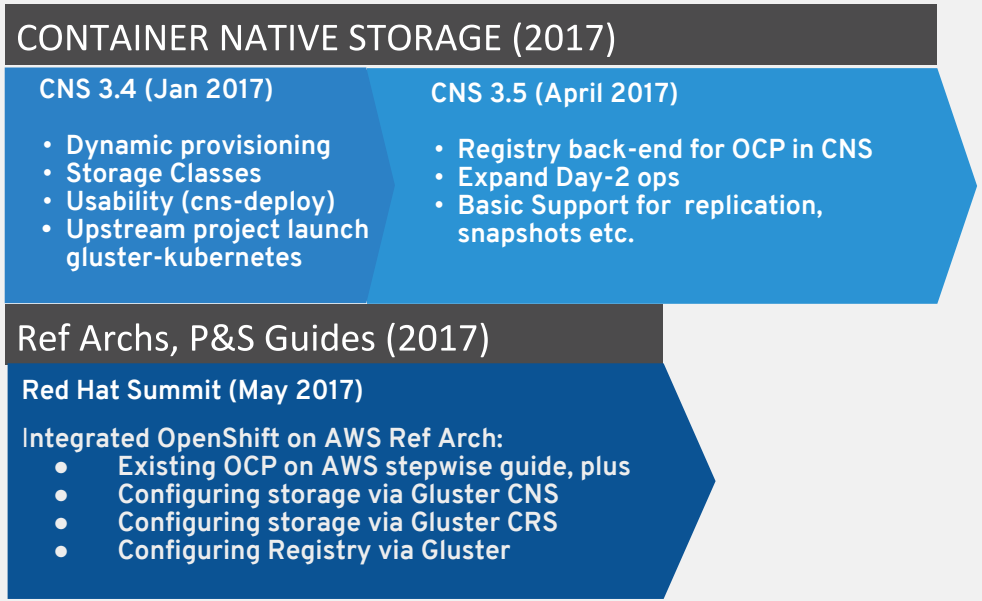

Container Storage Roadmap

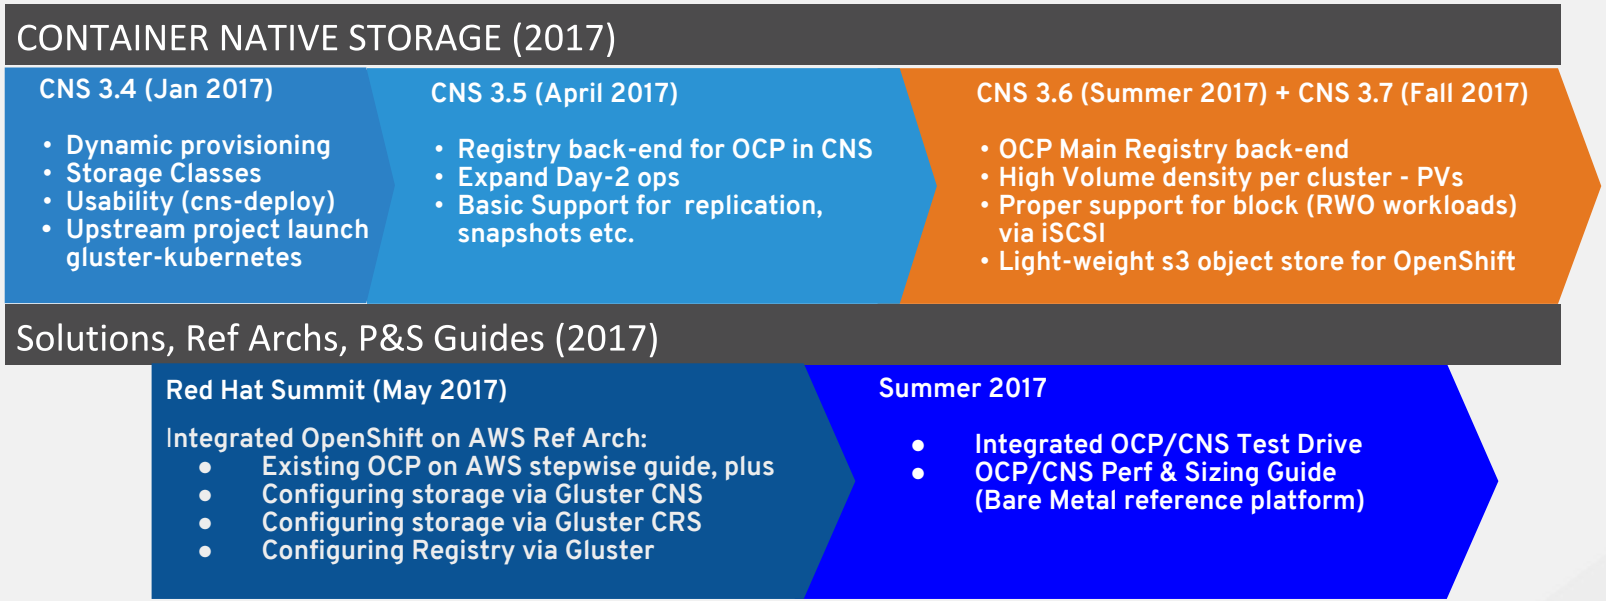

Red Hat is making strong investments in RHGS to make it the default storage platform for OpenShift

Brick Multiplexing

- Why?
  - Radically lower CPU, memory and port consumption/usage
  - Order of magnitude increase in the number of volumes per cluster especially for OpenShift storage
- What it is?
  - Move from 1 brick per process to multiple bricks per process
  - Many bricks can consume **one** port, **one** set of global data structures, and **one** pool of global threads
  - Configurable number of bricks per Server Process
  - Provides more flexibility for I/O scheduling across volumes
- When?
  - Summer 2017 for container-storage use-cases

Gluster-Block support for OpenShift/Container Storage

- Why
  - Proper and low latency block support for containerized workloads needing RWO access mode as defined by kubernetes
- What
  - iSCSI access to volumes with tcmu-runner and libgfapi
  - Orchestration of block devices managed by gluster-block
  - New gluster-block provisioner for kubernetes
- When
  - Summer CY 2017

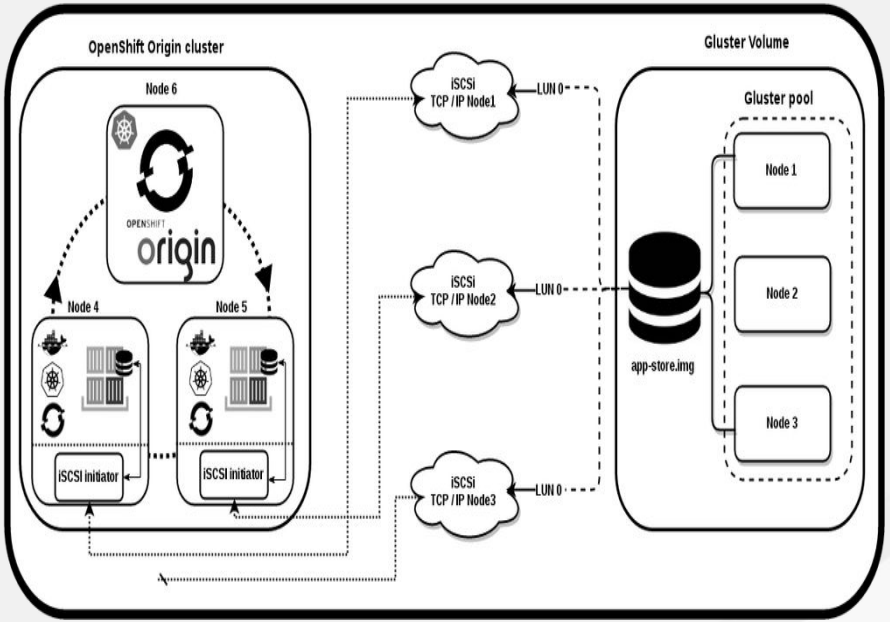

S3 support for OpenShift

- Why?
  - s3 support for developers and application teams building cloud-native apps
- What it is?
  - Gluster-s3 container provides a S3 interface to a gluster volume
  - Embedded as a micro-service inside OpenShift
  - Targeted for binary build archival, backups and registry storage in OpenShift
- When?
  - 2-H CY 2017

RHGS based Container Storage – Differentiators

- Storage can run as a set of micro-services inside OpenShift - no external appliance needed
- Consistent storage platform across the hybrid cloud
- Multi purpose (RWO, RWX, object), flexible & versatile

Summit Announce: OpenShift and Container Native Storage (CNS) on AWS

- **What Did We Announce?**
  - Reference Architecture for deploying OpenShift with CNS persistence layer on AWS
  - Ansible quickstart playbooks for automating deployments
- **Why:**
  - Provides OpenShift deployments with common storage abstraction layer across private and public cloud infrastructures.
  - Enhances application portability through common persistence layer semantics.
- **Examples:**
  - Multiple NGINX web server instances writing concurrently to a common Persistent Volume.
  - Single MySQL database instance writing to a fast StorageClass Persistent Volume.

Focus Area

Software-defined alternative for File/NAS

Container Storage across Hybrid Cloud

Hyper-converged storage with RHV

POSIX compatible Storage in Public Cloud

Hyper-converged storage with RHV

Why?

- Customers are looking for true open source alternatives to proprietary HCI solutions
- Price-performance
- Pricing & hardware selection flexibility

Hyperconverged VM store with RHV + RHGS

- World's first HCI solution with true open source components (oVirt + Gluster)
- Single SKU product with unified management & life-cycle
- Initial Use-Case: Remote-Office-Branch-Office (ROBO)/edge computing
- Arbiter Volumes Supported - Low footprint (full 3-way replication not needed!)
- Ansible based wizard driven deployment

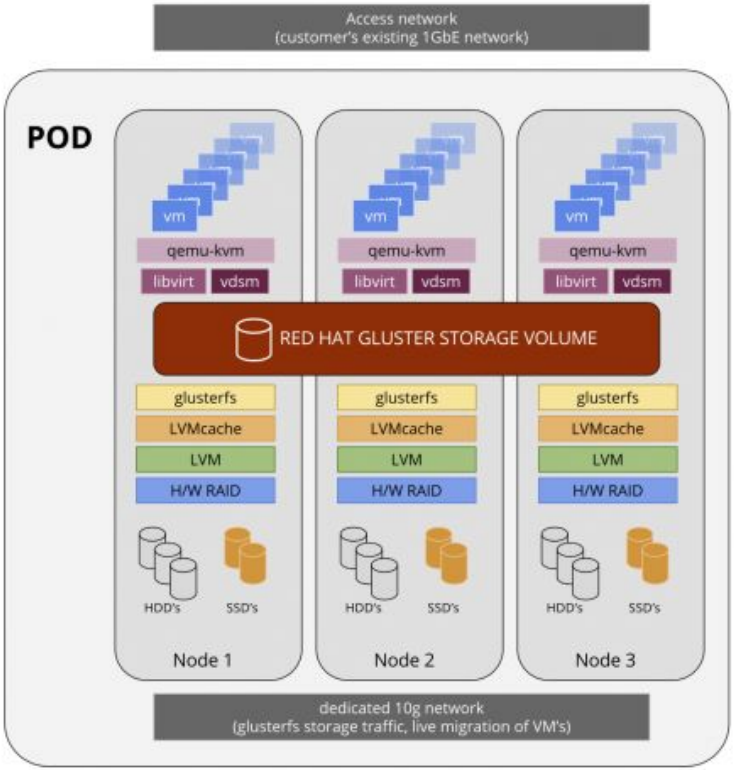

6 steps to nirvana

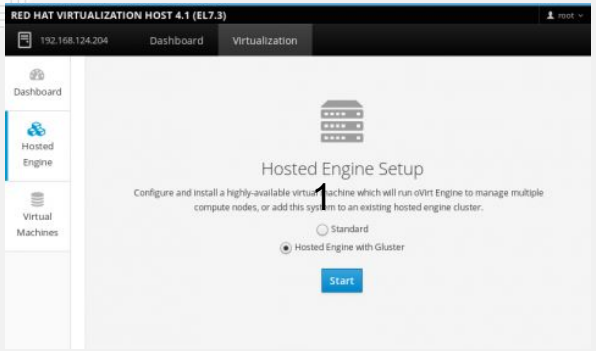

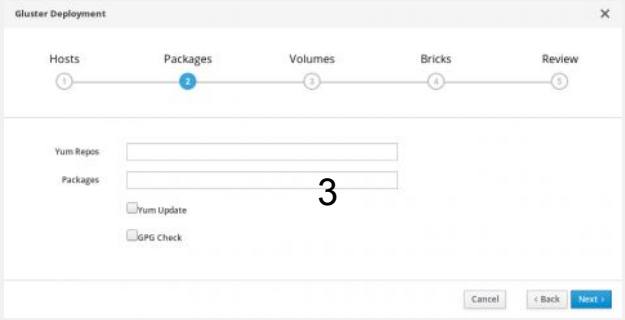

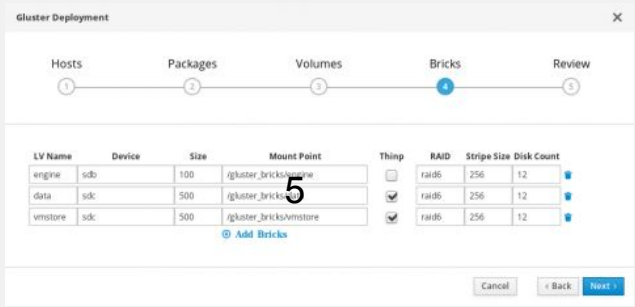

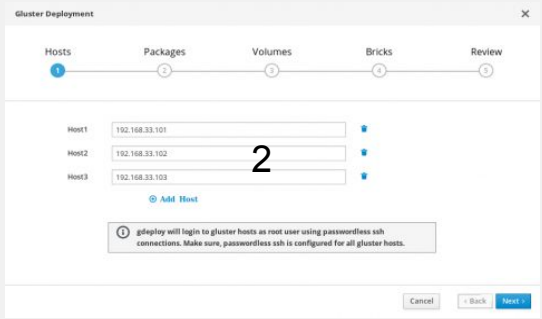

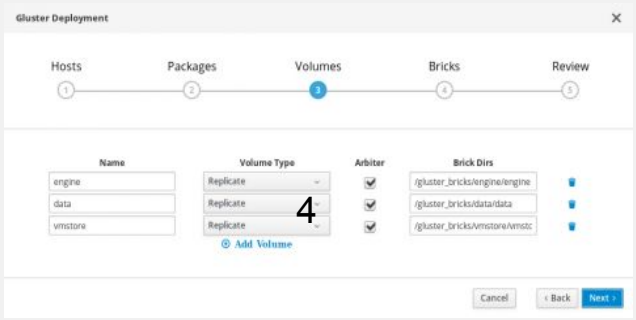

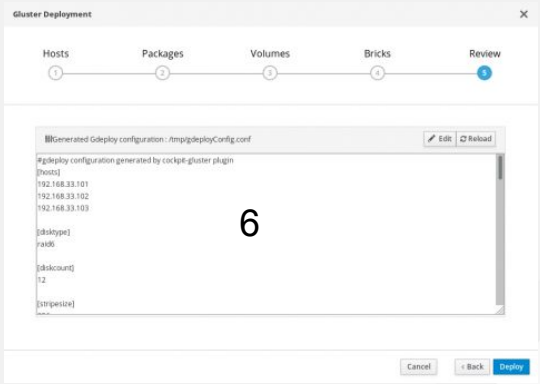

redhat.

Roadmap

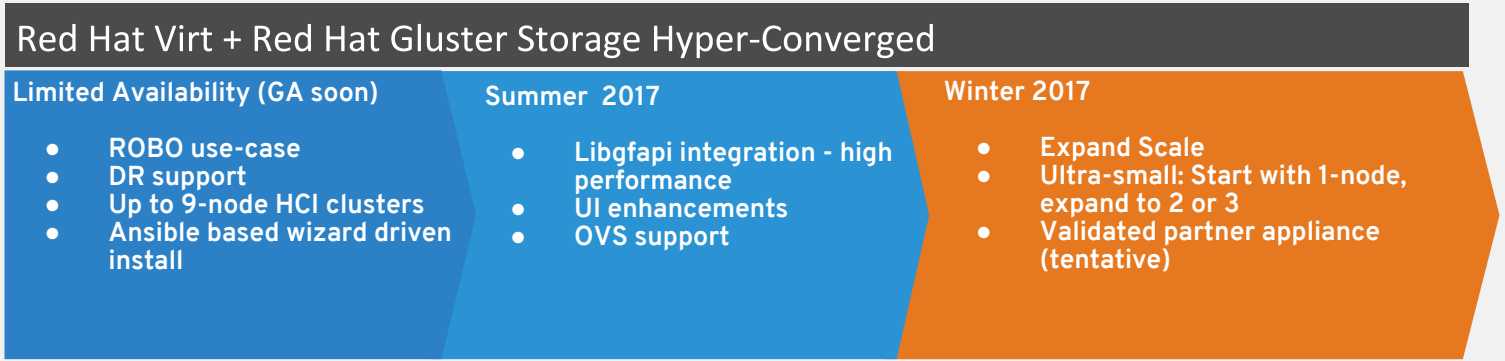
Red Hat Virt + Red Hat Gluster Storage Hyper-Converged
Limited Availability (GA soon)
ROBO use-case
DR support
Up to 9-node HCI clusters
Ansible based wizard driven install
Summer 2017
Libgfapi integration - high performance
UI enhancements
OVS support
Winter 2017
Expand Scale
Ultra-small: Start with 1-node, expand to 2 or 3
Validated partner appliance (tentative)

Sharding

- Why?
  - Better back-end for large files like Virtual Machine disks.
- What it is?
  - Chunk VM image files into smaller fragments across the storage cluster
  - Efficient healing and parallelized i/o
  - File sizes are not bound by disk or brick limits
- When?
  - Shipping now

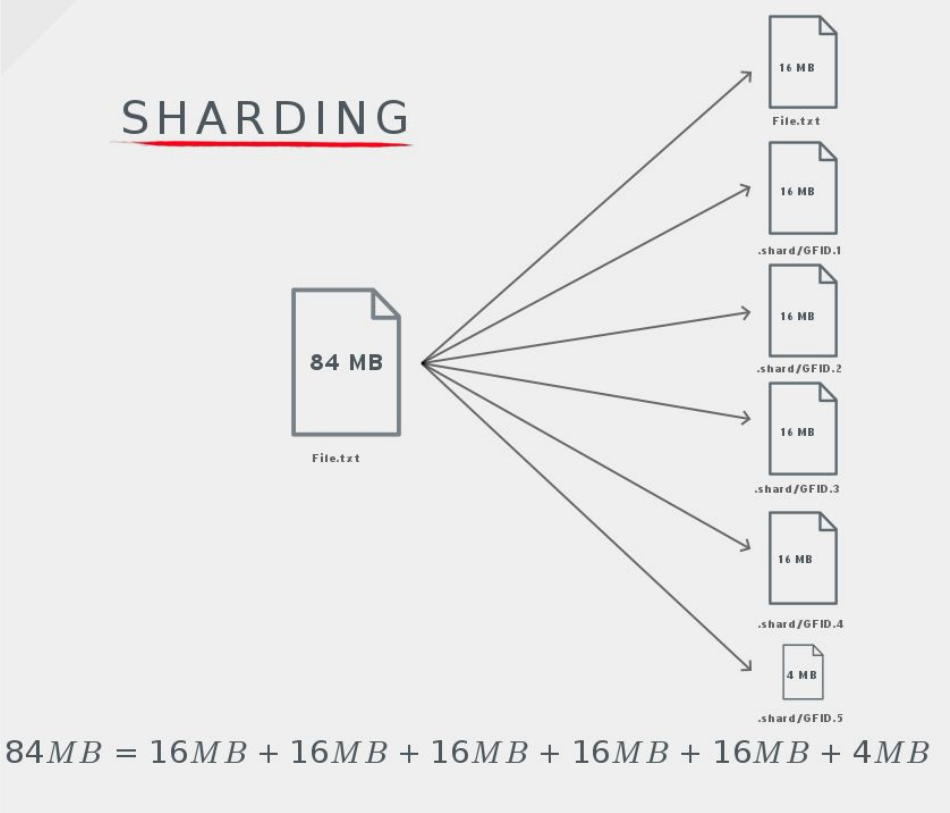

Focus Area

Software-defined alternative for File/NAS

Container Storage across Hybrid Cloud

Hyper-converged storage with RHV

POSIX compatible Storage in Public Cloud

POSIX compatible Storage in Public Cloud

Why?

- Public clouds seen as alternative to data center!
- Customers want to run the same apps, up there!
- Spotty support for POSIX compatible storage in public cloud

“Lift & Shift” to Public Cloud

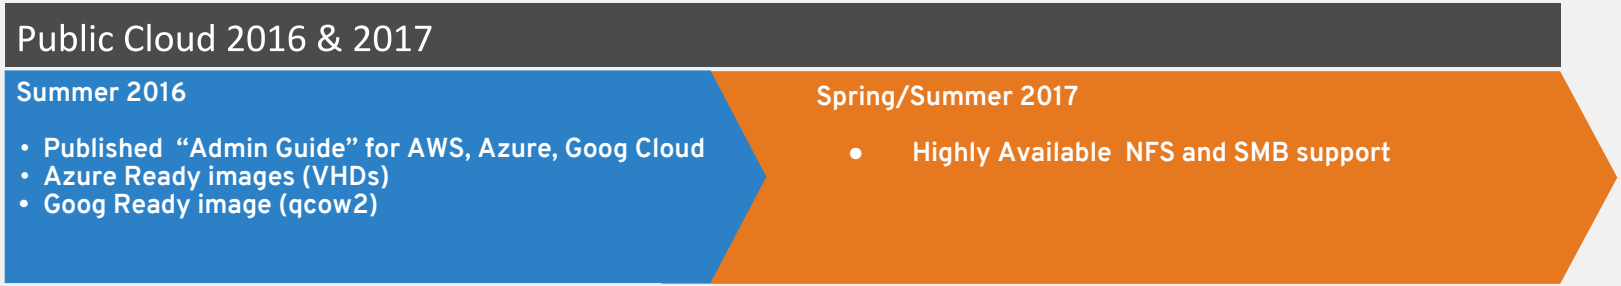

RED HAT SUMMIT

Container-Native Storage for Modern Applications with OpenShift and Red Hat Gluster Storage

Thursday, May 4th, 11:30 PM - 12:15 PM, Room # 157A

Sayan Saha
Director, Product Management, Storage Business, Red Hat

Michael Adam
Engineering Manager, Container-Native Storage, Red Hat

Annette Clewett
Senior Storage Architect, Storage Business, Red Hat

Daniel Messer
Technical Marketing, Storage Business, Red Hat

MODERNIZING STORAGE INFRASTRUCTURE WITH OPEN SCALE-OUT STORAGE

Wolfgang Schulze, Storage Practice Lead

Wednesday, May 3, 12:30 PM - 1:15 PM

Located at the Consulting Discovery Zone at the Services Showcase in the Partner Pavilion

To learn more, **visit red.ht/discoveryzone**

THANK YOU


plus.google.com/+RedHat


facebook.com/redhatinc


linkedin.com/company/red-hat


twitter.com/RedHatNews

youtube.com/user/RedHatVideos

LEARN. NETWORK.
EXPERIENCE
OPEN SOURCE.