

RED HAT  
**SUMMIT**

# Red Hat Ceph Storage Past, Present, and Future

Neil Levine, Federico Lucifredi, Uday Boppana  
Product Management, Storage  
2018-05



RED HAT  
**SUMMIT**

# DON'T FORGET YOUR RAFFLE CARD!

Collect raffle cards at all Storage sessions, workshops, and booths. Cash them in for prizes at the Storage Launchpad by event registration and storage lockers. The raffle cards also enter you into a drawing for one of two Star Wars drones.

#UnStorage



# AGENDA

- **Red Hat Storage Overview**
- **Ceph Overview and Product Versions**
- **Roadmap Themes**

# RED HAT STORAGE OVERVIEW

# RED HAT STORAGE IS DEEPLY INTEGRATED

## RED HAT<sup>®</sup> STORAGE

### PHYSICAL

RED HAT<sup>®</sup>  
CEPH STORAGE

RED HAT<sup>®</sup>  
GLUSTER STORAGE

RED HAT<sup>®</sup>  
ENTERPRISE  
LINUX<sup>®</sup>

### VIRTUAL

RED HAT<sup>®</sup>  
CEPH STORAGE

RED HAT<sup>®</sup>  
GLUSTER STORAGE

RED HAT<sup>®</sup>  
ENTERPRISE  
LINUX<sup>®</sup>

RED HAT<sup>®</sup>  
VIRTUALIZATION

### PRIVATE CLOUD

RED HAT<sup>®</sup>  
CEPH STORAGE

RED HAT<sup>®</sup>  
GLUSTER STORAGE

RED HAT<sup>®</sup>  
OPENSTACK  
PLATFORM

### CONTAINERS

RED HAT<sup>®</sup>  
CEPH STORAGE

RED HAT<sup>®</sup>  
GLUSTER STORAGE

 **OPENSIFT**  
ENTERPRISE  
by Red Hat<sup>®</sup>

### PUBLIC CLOUD

RED HAT<sup>®</sup>  
CEPH STORAGE

RED HAT<sup>®</sup>  
GLUSTER STORAGE

RED HAT<sup>®</sup>  
ENTERPRISE  
LINUX<sup>®</sup>



# RED HAT - A STORAGE VISIONARY



Source: Gartner (October 2017)

Red Hat Storage recognized as a Visionary by Gartner in the Magic Quadrant for Distributed File Systems and Object Storage. Access the report [here](#).

Red Hat Storage ranked as a Visionary two years in a row

Red Hat Storage positioned farthest and highest in both Completeness of Vision and Ability to Execute in the Visionaries quadrant.

## Highlights

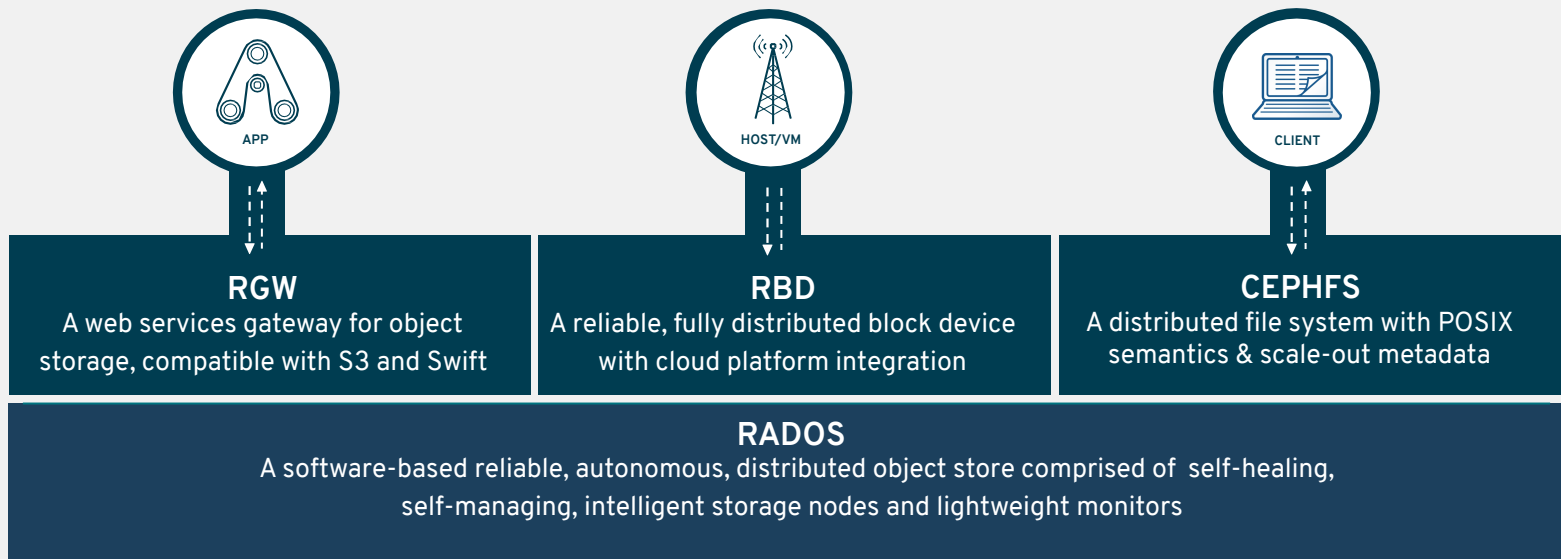
- Strong customer traction across geographies/verticals
- Solid product strategy and road map
- Leadership in container-native storage and storage for the hybrid cloud, Private Cloud
- Object Storage, Hyperconverged Infrastructure

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.

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 from <https://engage.redhat.com/gartnermagic-quadrant-storage-s-201610121525>.

# CEPH, VERSIONS AND TIMELINE

# Ceph Overview





# RED HAT CEPH STORAGE VERSIONS

**RHCS 3.0**  
**Nov 2017**

- Upstream: Luminous
- Platform: RHEL 7
- EoL: 2020-11-30

**3.n**

**RHCS 4.n**  
**2019**

- Upstream: Nautilus
- Platform: RHEL 7+



*Regular 6 week Z-stream updates*

# ROADMAP THEMES

# THEMES

- Usability
- Performance and Scale
- Security
- OpenStack
- Containers
- Object Storage
- Protocols

# THEME: USABILITY

*Increase TB/admin and make common tasks simpler to perform*



**RHCS Dashboard  
for monitoring**

**Cleaner CLI**

**Improved Logging**



**RHCS Dashboard  
for management**

**Prometheus  
support (3.1)**

**Access Insights  
(4.0)**



**Automatic  
placement group  
management**

Cluster Name ceph Percentile 95

Shortcuts

Dashboard Row

Cluster [click to go to: Ceph Pools](#) OSDs S3/Swift OSD Hosts Network

Active Alerts  
No alerts

At a Glance

<b>UP</b> 100% 100/100 100/100	<b>100%</b> 100/100 100/100 100/100	<b>100%</b> 100/100 100/100 100/100	<b>OSD Hosts</b> total - 23 up - 22 down - 1	<b>300%</b> 100/100 100/100 100/100	<b>100%</b> 100/100 100/100 100/100	<b>Capacity Utilization</b> 46%	<b>7d Growth Rate</b> 17.4 TB	<b>Weeks Remaining</b> 13	<b>Placement Group Status</b> No data points
---	--	--	---	--	--	------------------------------------	----------------------------------	------------------------------	---

Performance


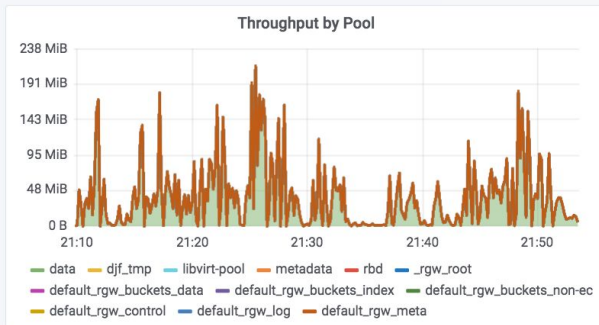
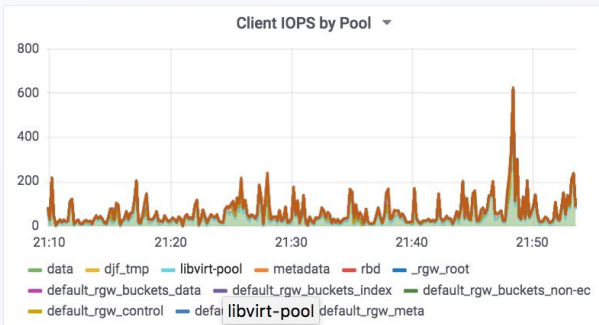
<b>Scrub</b> ACTIVE	<b>Recovery</b> 0 B/s	<b>Client IOPS</b> 102	<b>Client Throughput</b> 11.4 MB/s	<b>Pools</b> 11	<b>RBDs</b> 3	<b>Client Read/Write Ratio</b> No data points	<b>OSD Apply vs Commit Latency</b> Apply Latency Commit Latency
------------------------	--------------------------	---------------------------	---------------------------------------	--------------------	------------------	--	--

Dashboard Row

<b>OSD Hosts CPU Busy</b> 13%	<b>Disk IOPS</b> 1680	<b>Disk Throughput</b> 234.7 MB/s	<b>Nearly Full Disks</b> 0	<b>RAM UTIL</b> 76%	<b>Disk Latency</b> 29 ms	<b>Disk Utilization</b> average %util %util @ 95%ile
----------------------------------	--------------------------	--------------------------------------	-------------------------------	------------------------	------------------------------	---

Cluster Name **ceph** Pool Name **All**

Pool Overview : All



Top 5

Top 5 Pools by Client IOPS	
Pool Name	IOPS
data	77
metadata	11
libvirt-pool	0
djf_tmp	0
default_rgw_log	0

Top 5 Pools by Throughput	
Pool Name	Throughput
data	5.33 MiB
metadata	180.68 KiB
libvirt-pool	336.00 B
default_rgw_log	252.00 B
djf_tmp	0 B

Top 5 Pools by Capacity Used	
Pool Name	Capacity Used
data	61.87%
default_rgw_buckets_data	2.07%
libvirt-pool	0.01%
metadata	0%
djf_tmp	0%

# THEME: PERFORMANCE AND SCALE

*Reduce \$/IOPS and \$/Gb*



**BlueStore  
(Tech Preview)**



**BlueStore (3.n)**

**Erasure Code for  
RBD (3.n)**

**OSD compression  
& dedup (3.n)**

**Consistent IO in  
recovery (4.n)**



**QoS**

# BLUESTORE

## *Micron test results*

### **4k Random Block**

Write:

18% higher IOPs

5% lower average latency

Up to 70%+ reduced 99.99% latency

70% read / 30% write:

14% higher IOPs

80%+ lower read tail latency

70%+ lower write tail latency

### **4M Object**

Write:

88% increase in throughput

47% decrease in average latency

70% read / 30% write:

64% increase in throughput

40% decrease in average latency





# THEME: SECURITY


*Meet compliance requirements*



Per-Object  
encryption



Security  
Guidebook (3.n)




Kerberos  
integration with  
RADOS

On-The-Wire  
Encryption

# THEME: OPENSTACK

*Complete and tightly integrated storage for OpenStack*



# THEME: OPENSTACK

*Complete and tightly integrated storage for OpenStack*



**Manila support for  
CephFS (OSP 13)**

**Cinder encryption  
with RBD (OSP 13)**

**Hyperconverged  
deployment with  
Director (OSP 13)**



**Persistent  
volumes via  
Cinder/Manila**

**Distributed  
compute with  
Director (OSP 14)**




**Cinder QoS with  
RBD**

# USE CASE: HYPERCONVERGED

*Common technology for core and edge deployments*

- Edge sites hosting content for low-latency delivery
- Remote POPs running virtual network functions
- Micro data centers capturing IOT telemetry for real-time processing



# THEME: CONTAINERS

*Simplify lifecycle operations and provide more flexible topologies*



**Containerized  
Service Daemons**



**Persistent  
Volumes via  
Cinder/Manila  
(OSP 14)**

**Kubernetes  
support (4.n)**



**Dynamic load  
balancing**

# THEME: OBJECT STORAGE

*Scale!*

3.0

Dynamic Sharding  
of Bucket Indices

Per-object  
compression

Per-object  
encryption  
(SSE-C)

Next

Ongoing S3  
compatibility  
enhancements  
(3.n)

New RGW web  
server (4.n)

Future

Cloud Sync

AWS Secure  
Token Service

# THEME: OBJECT STORAGE

*Modern backup infrastructure*

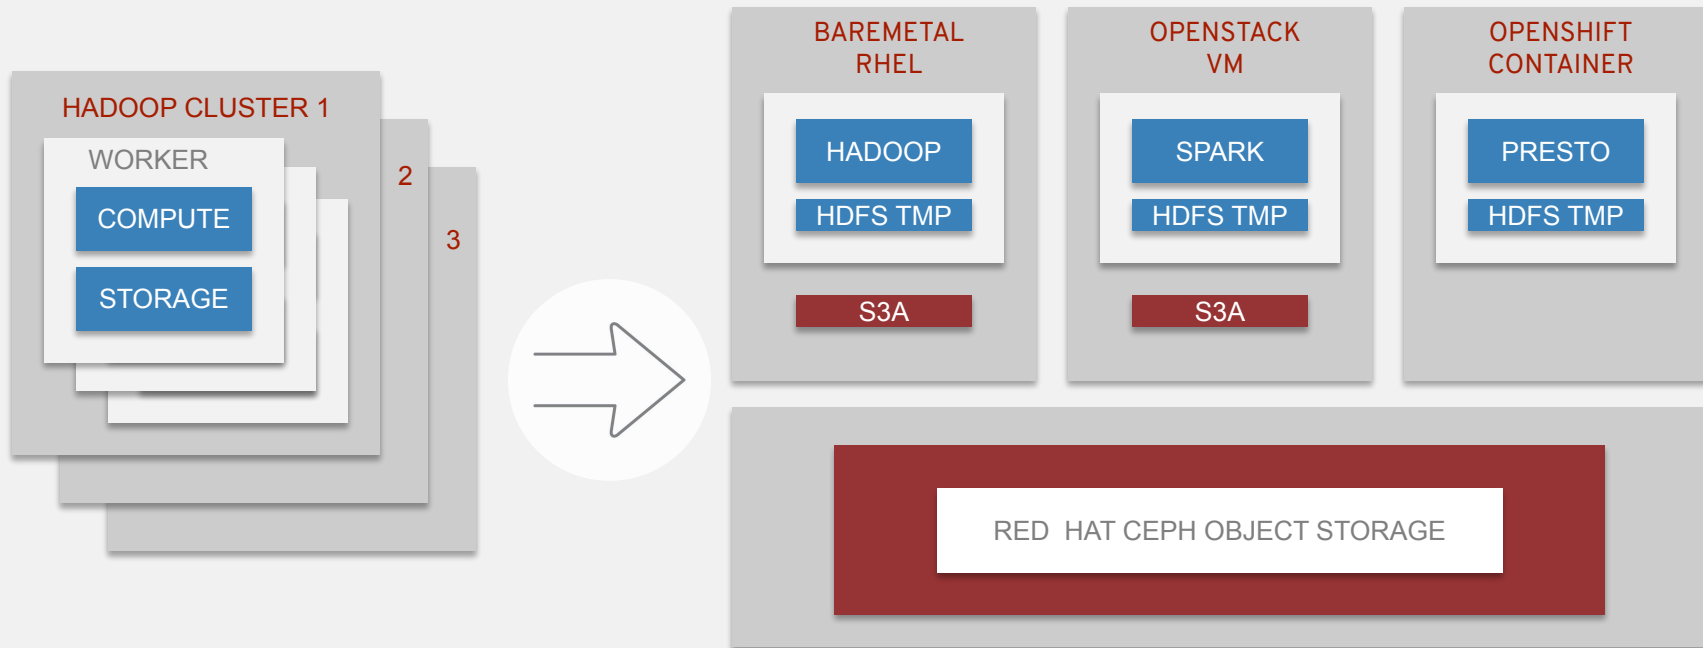
TRILI 

VERITAS™

 rubrik

# THEME: OBJECT STORAGE

*Data lakes with Ceph*





# THEME: PROTOCOLS

*Broaden workload options*



iSCSI

CephFS

NFS v3/v4  
Gateway for RGW



NFS for CephFS  
(4.n)



Ongoing  
enhancements

# SUMMARY

- Ongoing focus on usability
- Enhancing Red Hat Ceph Storage's ability to scale
- Deepening the integration with Red Hat OpenStack Platform

# SUMMIT TALKS

# Storage Sessions cont'd

Date	Time	Session Title	Speaker	Room
Wednesday, May 9	10:30 AM - 11:15 AM	<u>Red Hat Ceph Storage roadmap: Past, present, and future</u>	Neil Levine, Red Hat	2024
Wednesday, May 9	11:45 AM - 12:30 PM	<u>Automate complex software-defined storage in minutes with Ansible</u>	Dustin Black, Red Hat	2020
Wednesday, May 9	11:45 AM - 12:30 PM	<u>Code your infrastructure: ELO success story</u>	Flavio Andrade, Red Hat; Anderson Agapito, ELO Cartões	2022
Wednesday, May 9	11:45 AM - 12:30 PM	<u>Optimize Ceph object storage for production in multisite clouds</u>	Michael Hackett, Red Hat; John Wilkins, Red Hat	208

# Storage Sessions cont'd

Date	Time	Session Title	Speaker	Room
Wednesday, May 9	3:30 PM - 4:15 PM	<u>How the future of storage is shaping the Ceph roadmap</u>	Sage Weil, Red Hat	2020
Wednesday, May 9	4:30PM - 5:15PM	<u>DevSecOps with disconnected Red Hat OpenShift</u>	Mike Battles, Red Hat; Chase Barrette, MITRE Corporation; Stuart Bain, Red Hat; Jeremy Sontag, Red Hat	207
Wednesday, May 9	5:30 PM - 6:30 PM	<u>Bring your performance and technical problems to the experts</u>	Joe Mario, Red Hat	2016
Thursday, May 10	11:15 AM - 11:35 AM	<u>Using container-native storage: Myths and best practices</u>	Jose Rivera, Red Hat	2010
Thursday, May 10	1:00 PM - 1:45 PM	<u>Above the clouds with container-native storage</u>	Ryan Cook, Davis Phillips, Annette Clewett, Glenn West	207
Thursday, May 10	1:00 PM - 1:45 PM	<u>How Red Hat helps Vorwerk drive innovation with IoT</u>	Peter Mumenthaler, Red Hat; Michael Hosse, Vorwerk International Strecker & Co.; Sandro Koechli, Adfinis SyGroup AG	2009

# Storage Sessions cont'd

Date	Time	Session Title	Speaker	Room
Thursday, May 10	1:00 PM - 1:45 PM	<u>Scalable application platform on Ceph, OpenStack and Ansible</u>	Keith Hagberg, Fidelity; Senthivelrajan Lakshmanan, Fidelity Investments; Michael Pagan, Red hat; Sacha Dubois, Red Hat; Alexander Brovman, Solera Holdings Inc	2007
Thursday, May 10	2:00 PM - 2:45 PM	<u>Ceph Object Storage for Apache Spark data platforms</u>	Kyle Bader, Red Hat Meng Meng Lui	207
Thursday, May 10	2:00 PM - 2:45 PM	<u>OpenShift storage automation with container-native storage</u>	Annette Clewett, Red Hat Inc; Roya Vafadar, County of Los Angeles	215
Thursday, May 10	3:00 PM - 3:20 PM	<u>Ceph BlueStore performance on latest Intel Server Platforms</u>	Orlando Moreno, Intel	210

RED HAT  
**SUMMIT**

# THANK YOU



[plus.google.com/+RedHat](https://plus.google.com/+RedHat)



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



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



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



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