Red Hat Ceph Storage

Past, present, and future

Neil Levine, Federico Lucifredi, Uday Boppana
Storage Product Management

Agenda

- Red Hat Storage Overview
- Red Hat Ceph Storage use-cases
- Technology roadmap

Red Hat Storage Overview

KEY MARKET TRENDS

NON TRADITIONAL STORAGE BUYERS

Developers, Cloud Admins, and DevOps are decision makers

CLOUD NATIVE IS THE NEW NORMAL

Driven by the rise of Kubernetes for infrastructure orchestration

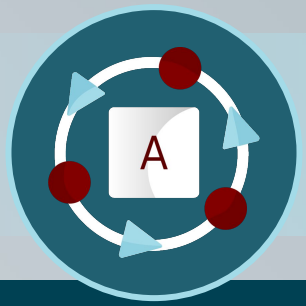

MODERN, STORAGE INTENSIVE APPLICATIONS

AI, ML and emerging workloads need scalable object storage

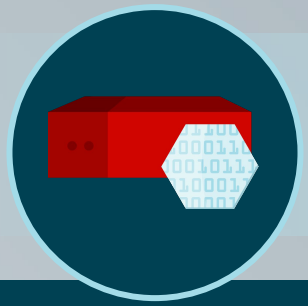

HYPERCONVERGENCE

Customers value simplicity and completeness of infrastructure

STORAGE FOR THE OPEN HYBRID CLOUD

OPENSHIFT CONTAINER STORAGE

COMPLETE DATA PORTABILITY

For OpenShift Across the Hybrid Cloud

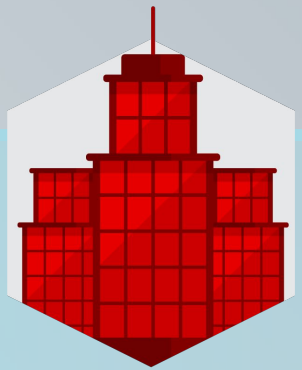

HYBRID CLOUD OBJECT STORAGE

MOST SCALABLE DATA PLATFORM

For Data Analytics, AI/ML, and emerging workloads

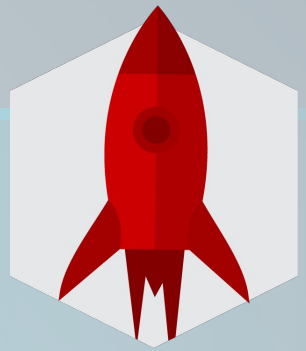

HYPERCONVERGENCE

ELASTIC INFRA ACROSS THE DATA CENTER AND EDGE

Built to enable flexibility, scale, and ease of use

Use-cases

Complete Storage for OpenStack

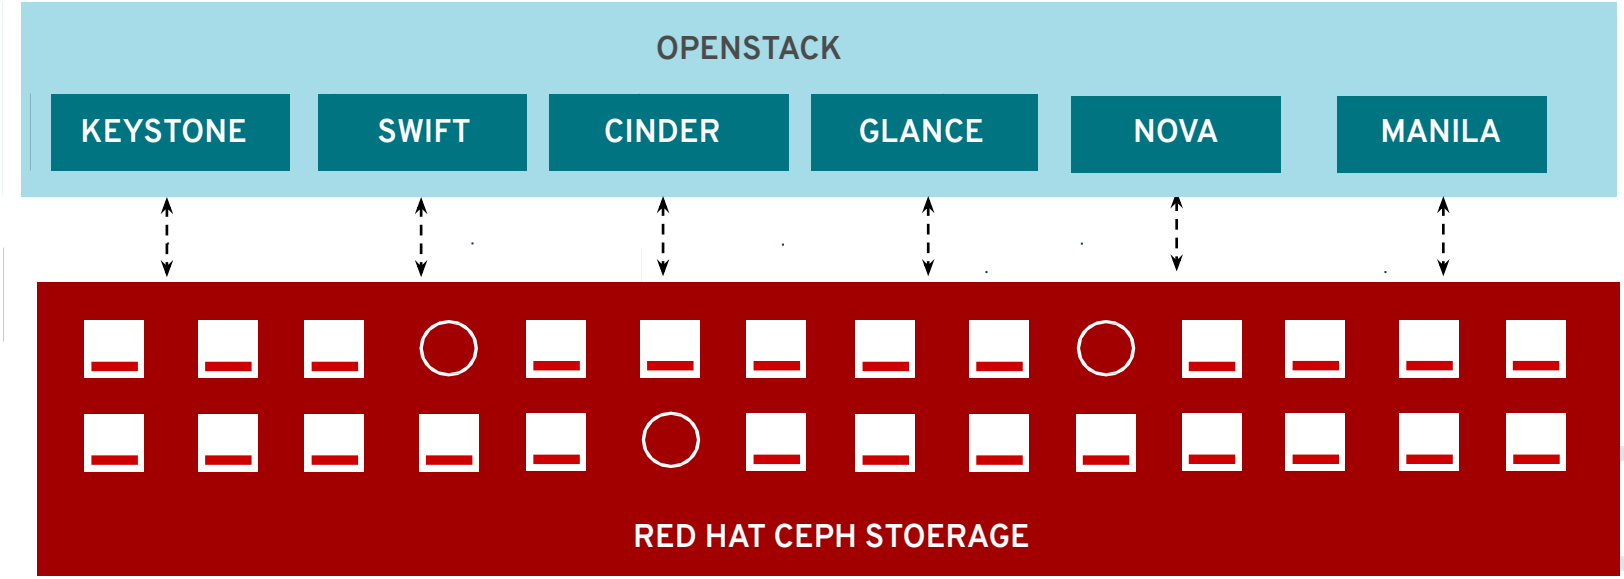

Object storage solutions

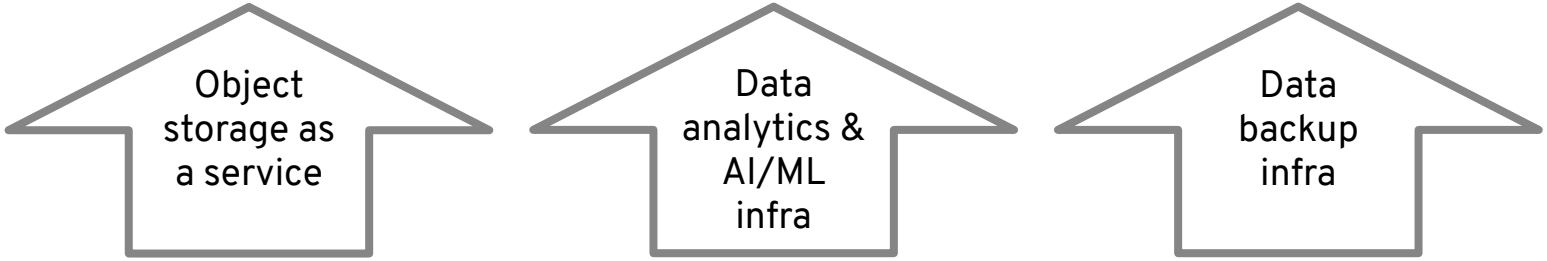

Expand workload support

ISV partnerships
&
certifications

Quick time-to-value

Hardware configurations
&
Reference architectures

Red Hat Ceph Storage features and functionality

Object storage - Data backup infrastructure

ISV partnerships for multi-workload S3/cloud backup

Multi-workload backup

VMware backup

OpenStack backup

Object storage - Red Hat data analytics infrastructure

Multi tenant workload isolation with Shared Data Context

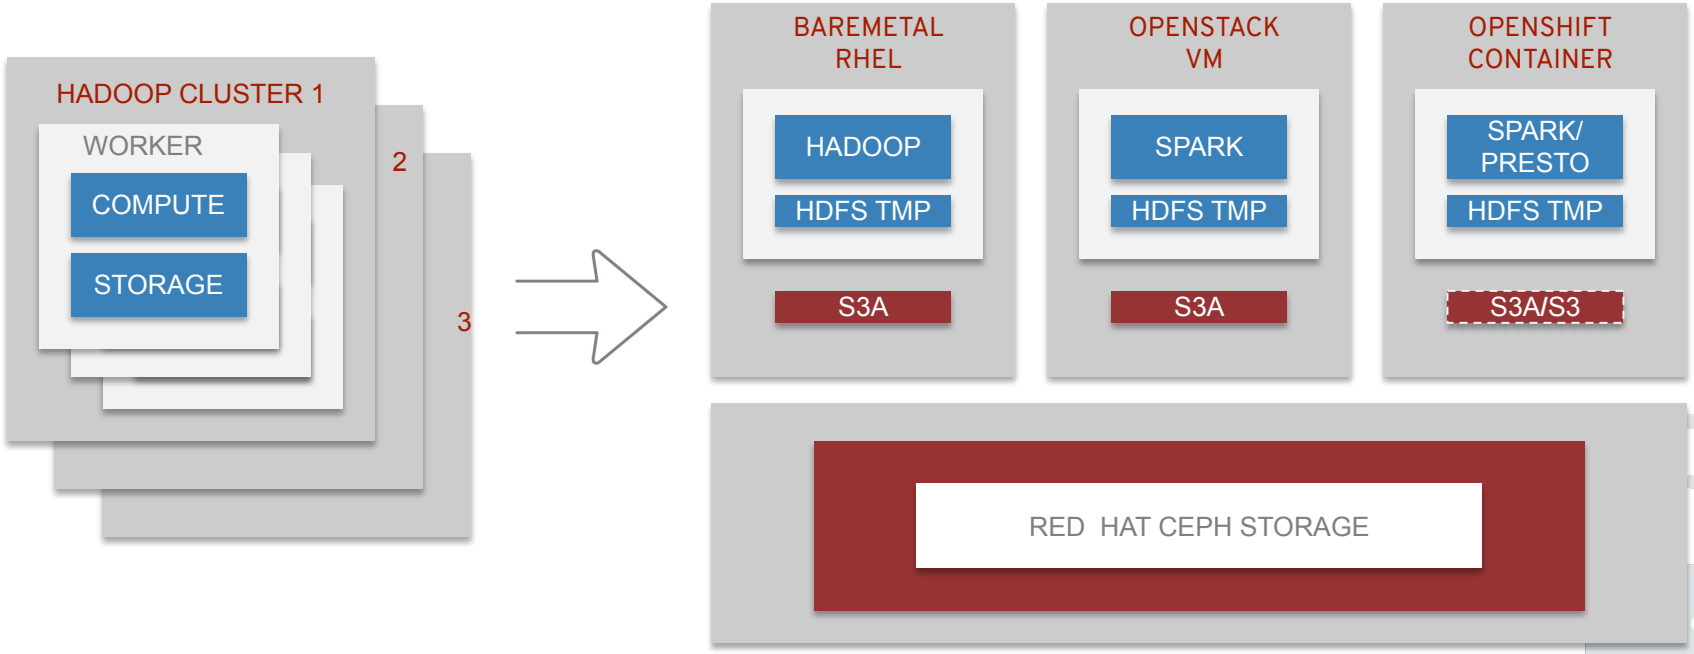

Object storage - Red Hat data analytics infrastructure

Extending analytics infrastructure solution to AI/ML-Available Now at OpenDataHub.io

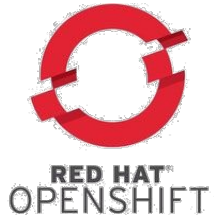

- Container platform
- Certified Kubernetes
- Hybrid cloud

- Unified, distributed storage
- RESTful gateway
- S3 and Swift compatible

- Unified analytics engine
- Large-scale data
- Runs on Kubernetes

- Multi-user Jupyter
- Used for data science and research

Ease of deployment

Faster time-to-value: Recommended and validated hardware configurations

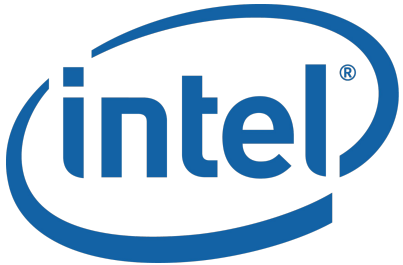

Themes and Features

Red Hat Ceph Storage versions

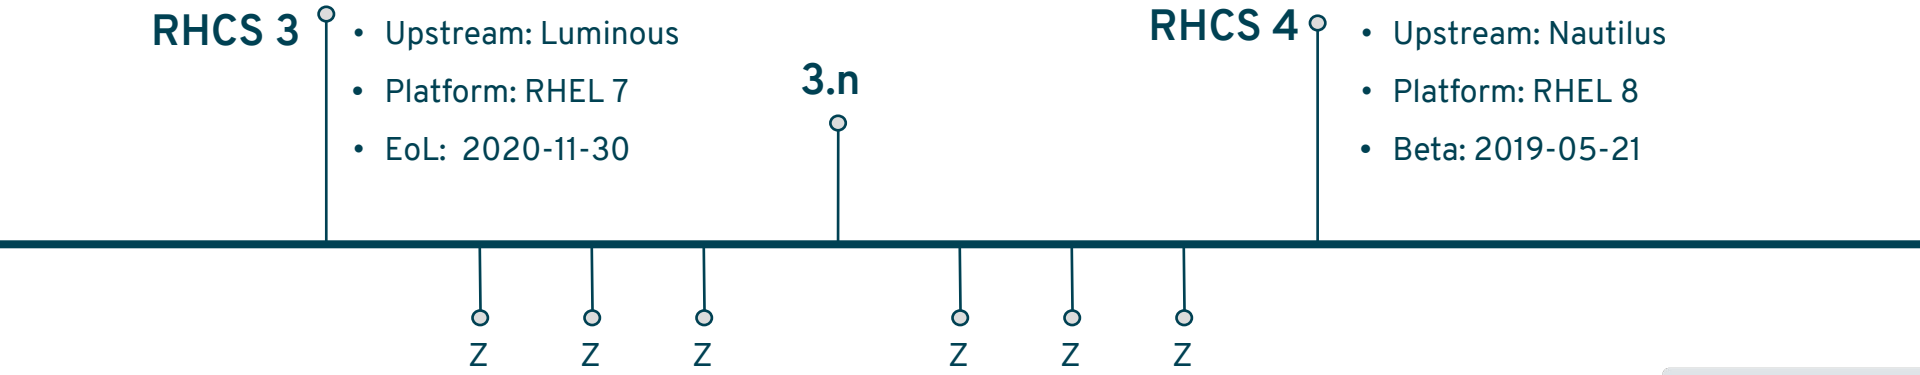

Regular 6 week Z-stream updates

Object Storage

Scale!

3.2

Dynamic Sharding of Bucket Indices

Per-object compression

Per-object encryption (SSE-C)

Next

New RGW web server (3.3)

Bucket Notifications (tech preview)

Future

AWS Secure Token Service

Public cloud tiering and archive

Data at rest encryption (SSE-S3 and SSE-KMS)

Usability

Increase TB/admin and make common tasks simpler to perform

3.2

RHCS Dashboard for monitoring

Cleaner CLI

Improved Logging

Prometheus support

Next

RHCS Dashboard for management (4.0)

Automatic Placement Group Management (4.0)

Access Insights (4.n)

Future

Extended RGW management

SAAS Monitoring

DATADOG

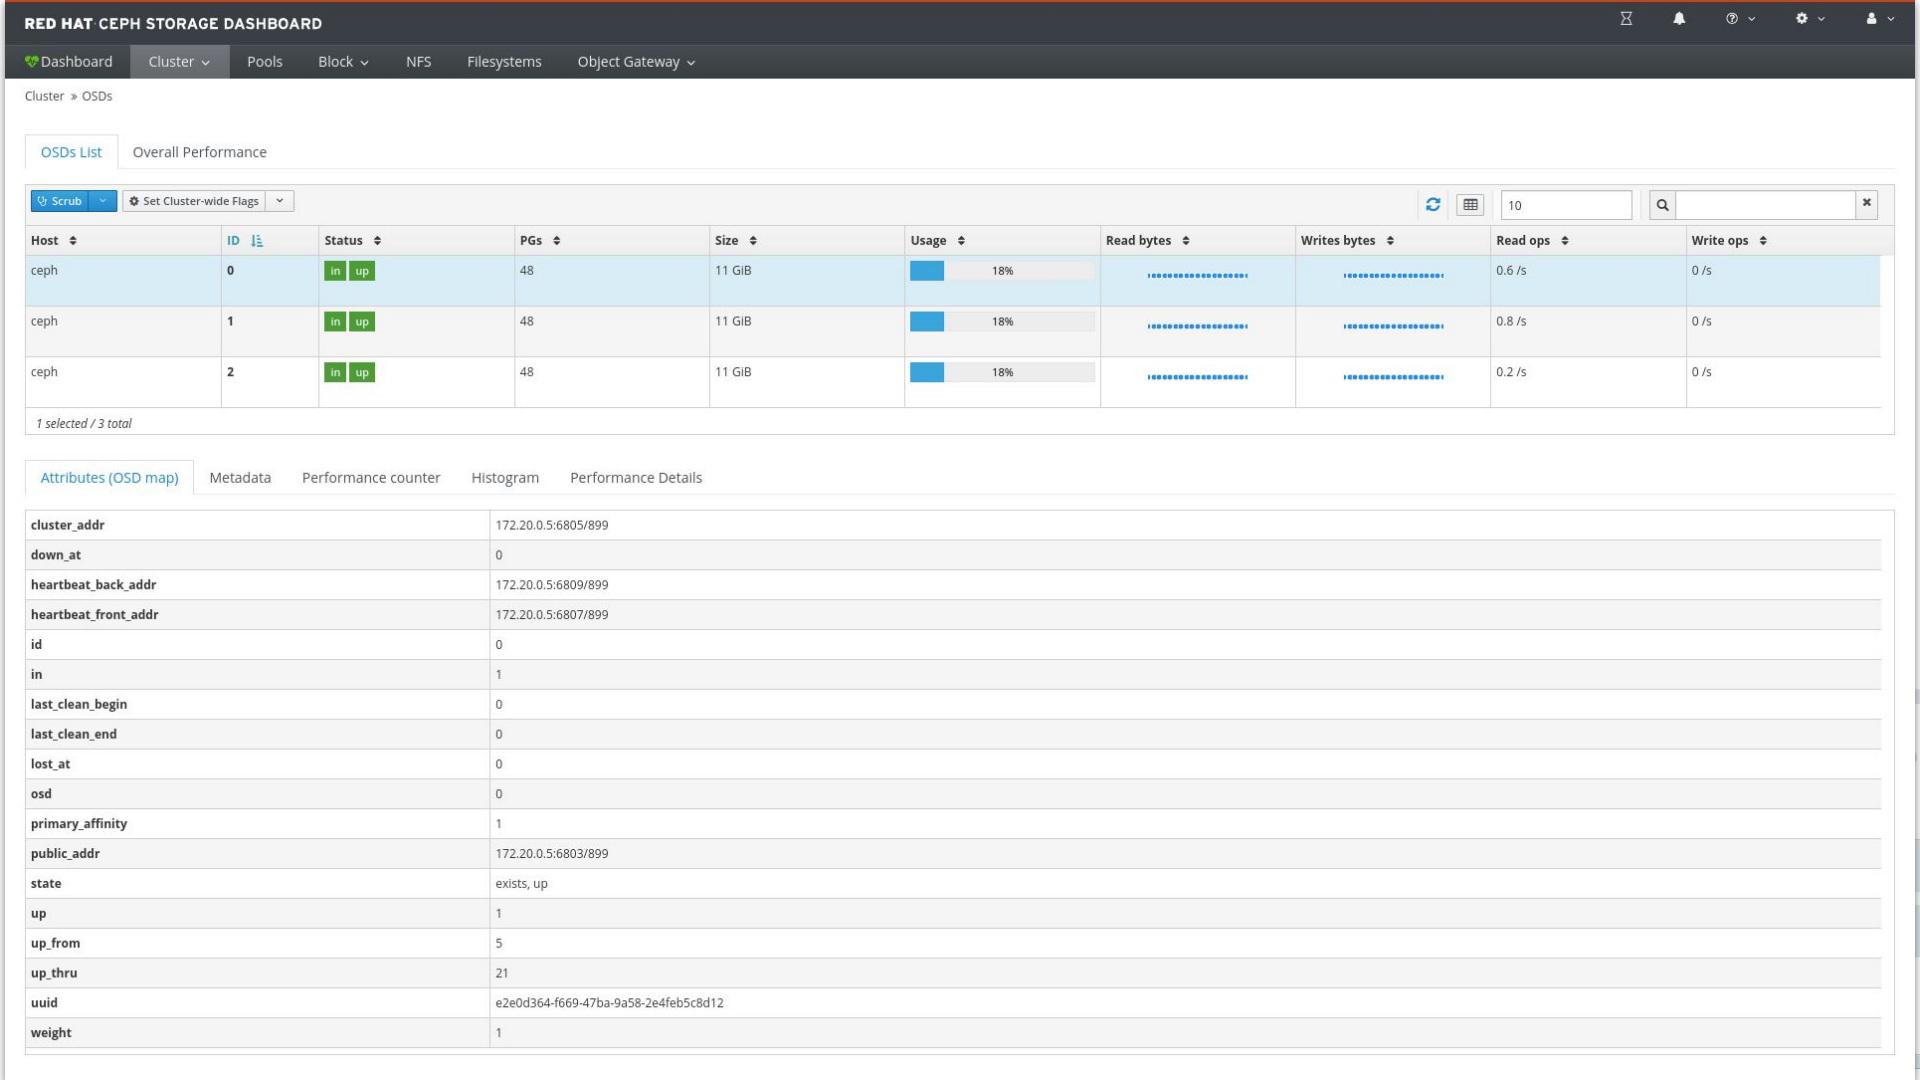
RED HAT CEPH STORAGE DASHBOARD
Dashboard
Cluster
Pools
Block
NFS
Filesystems
Object Gateway
Cluster » OSDs
OSDs List
Overall Performance
Scrub
Set Cluster-wide Flags
10
Host
ID
Status
PGs
Size
Usage
Read bytes
Writes bytes
Read ops
Write ops
ceph 0 in up 48 11 GiB 18% 0.6 /s 0 /s
ceph 1 in up 48 11 GiB 18% 0.8 /s 0 /s
ceph 2 in up 48 11 GiB 18% 0.2 /s 0 /s
1 selected / 3 total
Attributes (OSD map)
Metadata
Performance counter
Histogram
Performance Details
cluster_addr 172.20.0.5:6805/899
down_at 0
heartbeat_back_addr 172.20.0.5:6809/899
heartbeat_front_addr 172.20.0.5:6807/899
id 0
in 1
last_clean_begin 0
last_clean_end 0
lost_at 0
osd 0
primary_affinity 1
public_addr 172.20.0.5:6803/899
state exists, up
up 1
up_from 5
up_thru 21
uuid e2e0d364-f669-47ba-9a58-2e4feb5c8d12
weight 1

RED HAT CEPH STORAGE DASHBOARD

Dashboard | Cluster | Pools | Block | NFS | Filesystems | Object Gateway

Pools

Pools List | Overall Performance

Edit

10

Name	Type	Applications	PG Status	Replica Size	Last Change	Erasure Coded Profile	Crush Ruleset	Usage	Read bytes	Write bytes	Read ops	Write ops
.rgw.root	replicated	rgw	8 active+clean	3	17		replicated_rule	0%			0 /s	0 /s
cephfs.a.data	replicated	cephfs	8 active+clean	3	16		replicated_rule	0%			0 /s	0 /s
cephfs.a.meta	replicated	cephfs	8 active+clean	3	15		replicated_rule	0%			0 /s	0 /s
default.rgw.control	replicated	rgw	8 active+clean	3	19		replicated_rule	0%			0 /s	0 /s
default.rgw.log	replicated	rgw	8 active+clean	3	23		replicated_rule	0%			0 /s	0 /s
default.rgw.meta	replicated	rgw	8 active+clean	3	21		replicated_rule	0%			0 /s	0 /s

1 selected / 6 total

Details | Performance Details | Configuration

Grafana Time Picker: Last 1 hour (Default)

Pool '.rgw.root' Performance Details

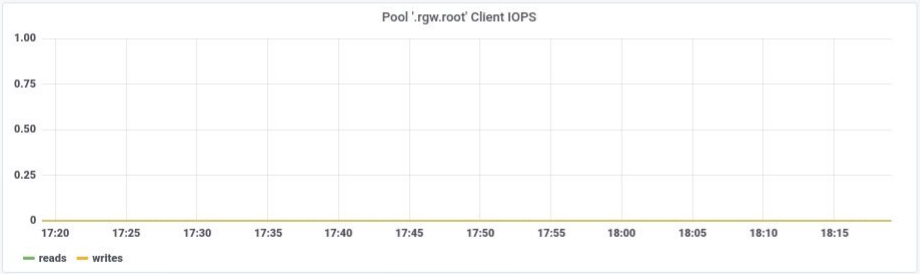

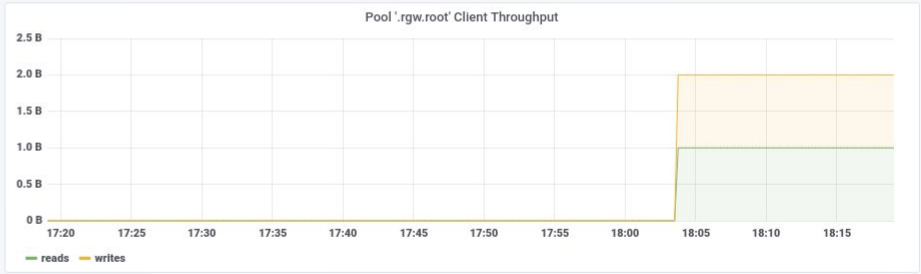

Images Trash Overall Performance

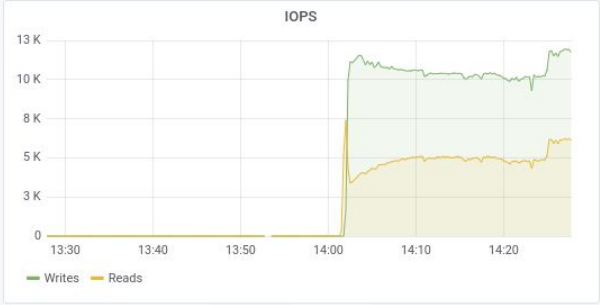

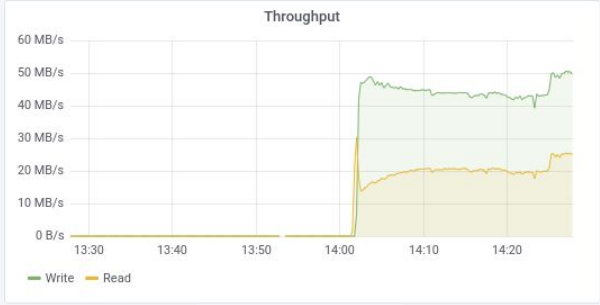

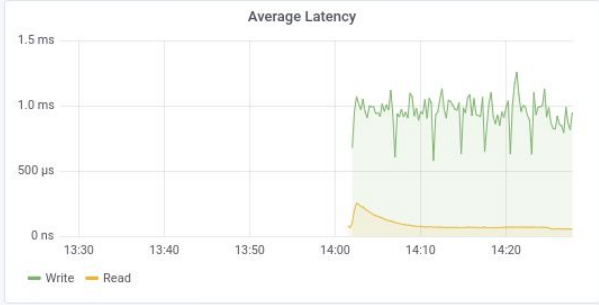

Highest IOPS

Image	Pool	IOPS
image2	rbd	11.76K iops
image1	rbd	3.06K iops
image3	rbd	3.06K iops
image4	rbd	0 iops
foo	rbd	0 iops

Highest Throughput

image	Pool	Throughput
image2	rbd	49.73 MB/s
image1	rbd	12.52 MB/s
image3	rbd	12.47 MB/s
foo	rbd	0 B/s
image4	rbd	0 B/s

Highest Latency

Image	Pool	Value
image2	rbd	947.14 µs
image1	rbd	55.31 µs
image3	rbd	53.94 µs
foo	rbd	0 ns
image4	rbd	0 ns

Performance and Scale

Reduce $/IOPS and $/Gb

BlueStore GA

OSD compression & dedup (4.n)

Consistent IO in recovery (4.n)

QoS

Bluestore

Tuning for all-flash configuration

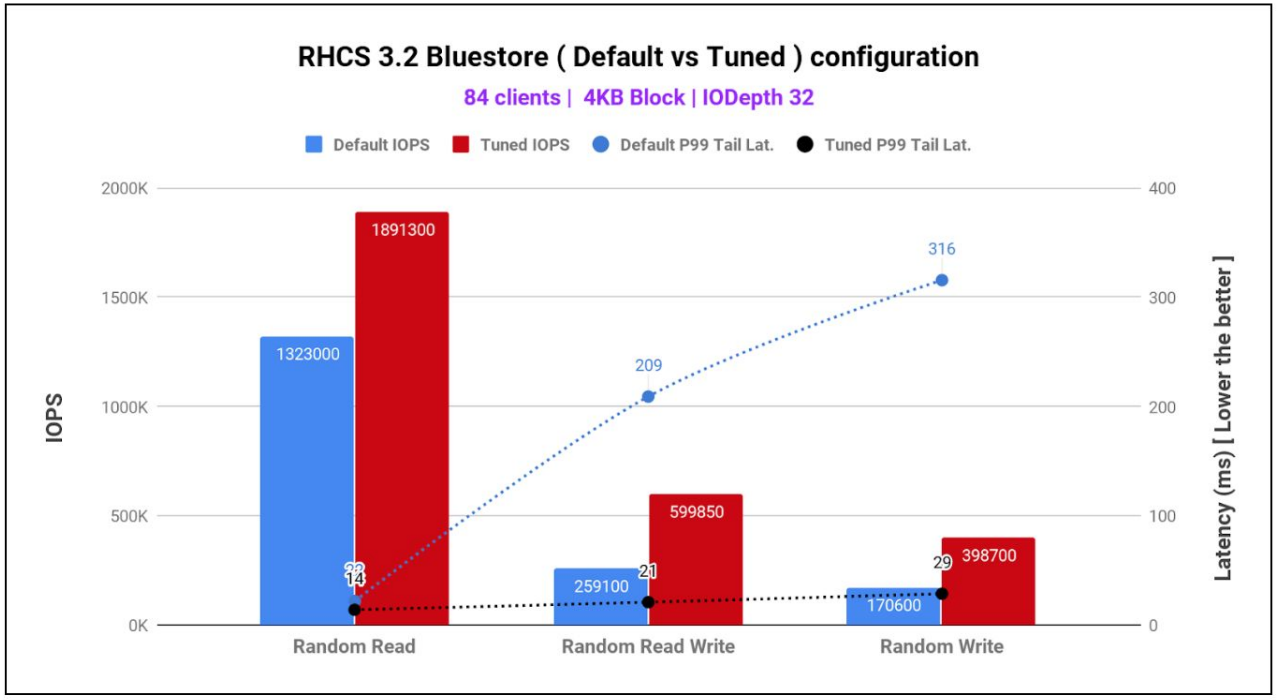

Security

Meet compliance requirements

3.2

Per-Object encryption

Security Guidebook

Next

FIPS Compliance (4.n)

RBD Namespaces (4.0)

On-The-Wire Encryption (4.n)

Future

Kerberos integration with RADOS

OpenStack

Complete and tightly integrated storage for OpenStack

3.2

Manila support for CephFS (OSP 13)

Cinder encryption with RBD (OSP 13)

Hyperconverged deployment with Director (OSP 13)

Next

Distributed compute with Director (OSP 16)

Glance Image clones (OSP 16)

RHEL 8 enablement (4.0)

Future

Cinder QoS with RBD

OpenShift Container Storage 4

Persistent Container Storage on OpenShift

Ceph sessions at summit 2019

Ceph sessions over the next three days

Title: Experiences from building production-ready, massive hybrid cloud infrastructures on Red Hat Ceph Storage
Type: Panel
Day/Time: Wednesday, May 8 10:30 AM-11:15 AM

Title: Transforming Cloudera analytics agility in a fast-moving world
Day/Time:Thursday, May 9, 11:00 a.m.-11:45 a.m.
Type: Breakout

Title: Next-generation data lake architecture with Red Hat
Day/Time: Thursday, May 9, 1:30 p.m.-1:50 p.m.
Type: Mini session

Title: Ceph and the future of storage
Day/Time: Thursday, May 9, 2:00 p.m.-2:45 p.m.
Type: Birds of a feather

FIND US AT RED HAT SUMMIT

redhat.com/storage

@redhatstorage

redhatstorage.redhat.com

- At the Storage lockers
- At the Red Hat booth
- At one of Storage dedicated sessions (red.ht/storageatsummit)
- At the Community Happy Hour (Tues 6:30, Harpoon Brewery)
- At the Hybrid Cloud Party (Wed, 7:30, “Committee” restaurant)

Red Hat OpenShift Container Storage
red.ht/videos-RHOCS

Red Hat data analytics infrastructure solution
red.ht/videos-RHDAIS

Red Hat Hyperconverged Infrastructure
red.ht/videos-RHHI

RED HAT SUMMIT

THANK YOU


linkedin.com/company/Red-Hat


facebook.com/RedHatinc

youtube.com/user/RedHatVideos


twitter.com/RedHat