

Ceph Storage Roadmap

Enterprise meets Community

Federico Lucifredi

Product Management Director, Ceph Storage Platform, *IBM | Red Hat*

Neha Ojha, Program Director, IBM Storage Ceph, *IBM*

History



SLIDES

CEPH STORAGE TECHNOLOGY ROADMAP

Manageability

Year/Quarter	Release Name	Key Features/Updates
2015	RHCS 1.2	• ceph-deploy
2015	RHCS 1.3	• Major version Upgrades • director intg. (puppet-ceph)
2016	RHCS 2.0	
2016	OSP 7	
2017	RHCS 3.0	• Director intg. (ceph-ansible)
2017	OSP 11	• Dashboard v.1 (Ceph Metrics)
2018	OSP 13	• ceph-ansible
2018	OSP 13	• Hyperconverged Ceph + OpenStack
2020 Q1	RHCS 4.0	• Rook • "Opinionated" design
2020 Q1	OCS 4.2	
2020 Q1	OCS 4.5	
2020 Q2	RHCS 5.0	• Stable mgmt API • Dashboard v.3 • RGW multisite • replacing OSDs • user mgmt • Cephadm • Independent mode
2020 Q2	OSP 16.1	• Edge OpenStack
2020 Q2	RHCS 5.0	• Dashboard v.2 (MGR) • Install UI • Bluestore migration

7 STRATEGIC ROADMAP – SUBJECT TO CHANGE

HEAR MY
TALK AT

 **OpenInfra**
SUMMIT > VANCOUVER '23

SPOTLIGHT
SPONSOR




Red Hat

Ceph: State of the Cephalopod



JOSH DURGIN



FEDERICO LUCIFREDI



JUNE 13-15 2023 · VANCOUVER CONVENTION CENTRE

[REGISTER NOW!](#)

Platform: roadmap and commit criteria

Responsible PM: Federico Lucifredi (Platform)

- Roadmap commitment
 - 70% jointly and severally
- Release Headliner (PO) commitment
 - 90% jointly and severally
 - A headliner should only ever fail if we are somehow “surprised”
- Schedule is our ultimate commitment
 - Quality that does not ship has no impact
 - Schedule changes have unacceptable consequences on marketing

FOR REFERENCE

- Community roadmap commitment
 - 50% jointly and severally

Platform: product dependencies

Responsible PM: Federico Lucifredi (Platform)

Ceph Platform 9.1

- Red Hat Ceph Storage
 - Simultaneous GA (9.1)
- IBM Storage Ceph
 - Simultaneous GA (new versioning: 9.9.1.0) – [ISCE 2150](#) – Blue Ceph: new versioning scheme
- OpenShift Data Foundation, Fusion Data Foundation, IBM Fusion HCI
 - 4.22 - June/July 2026 (*in planning)
- IBM Cloud – Bare metal (NVMeoF/RBD)
 - GA (phase 2) Q3 2026 (*in planning)
- Flashsystem 9000 (FCM)
 - TBD (*under review)

-
- IBM Cloud – Block (RBD)
 - IBM Cloud – File (NFS/CephFS)
 - Ceph Ready Nodes, next-generation Ceph Ready Nodes
 - Ceph as a Service
 - Nokia (NCP, NCS, CBIS)
 - Red Hat Upgrade Infrastructure (RHUI)

IBM Storage Ceph 2026 Roadmap

	IBM Storage Ceph 9.9.0 (1Q26)	IBM Storage Ceph 9.9.1 (2Q26)	IBM Storage Ceph 9.9.2 (3Q26)
Object Storage Analytics/ WatsonX	<ul style="list-style-type: none"> Performance Object Data-Lake Write-Through Caching D4N [GA] Bucket Inventory Bucket Index Performance Improvements Iceberg Rest Catalog Auth / Authz Integration S3 Access Logging 	<ul style="list-style-type: none"> the Bucket Level Security. Add the S3Control API for Account-Level Public Access Block Multitenancy. IAM API. Customer Managed Policies 	<ul style="list-style-type: none"> Sharding Ceph S3 Vector Buckets: API, Indexing, and Metadata Filtering [TP] Quota Enhancements(Per Storage Class)
Block Storage	<ul style="list-style-type: none"> NVMe/TCP DH HMAC CHAP NVMe/TCP BCDR Support via Active / Standby NVMe/TCP CSI-Driver (TP) 	<ul style="list-style-type: none"> Encrypted Namespaces and BYOK KMS management Automated service configuration management Automated listener management 	<ul style="list-style-type: none"> NVMeoF CSI driver Stretch Cluster Support Simplified Lifecycle Management Workflows Enhanced Realtime Metrics in Dashboard
File Storage	<ul style="list-style-type: none"> SMB IO-Rate Limiting NFS TLS-Support 	<ul style="list-style-type: none"> SMB Services Separation SMB IP-based access control SMB BYOK for IBM Cloud SMB per share quotas and capability reporting SMB Cluster-wide IO-rate limiting NFS over TLS [GA], Cluster-wide IO-rate limiting CephFS: MDS performance analysis 	<ul style="list-style-type: none"> SMB services on top of S3/RGW [TP] SMB Persistent File Handles [TP] CephFS: Improved Snapshot-based async file replication CephFS: MDS improved observability CephFS: optimized MDS defaults
Operational Ease/ Deployment	<ul style="list-style-type: none"> Zero Touch Provisioning [TP] Dashboard: Multi-Tenancy Enhancements Multi-Cluster Management [TP] Multi-Namespace workflows for NVMe/TCP in Ceph Dashboard and vSphere plugin 	<ul style="list-style-type: none"> Multi-Cluster Manager [GA] Ceph Provisioning Manager [GA] FCM Deployment Support & Dashboard validation readiness 	<ul style="list-style-type: none"> S3-Browser Ceph Smart Assistant [DP] Refinements for Dashboard, MCM and Ceph Provisioning
Ceph Core	<ul style="list-style-type: none"> Crimson OSD [TP] OSD Scaling improvements (>5k) Fast Erasure Coding Implementation ARM Cluster support 	<ul style="list-style-type: none"> FCM Support Defender Integration for RBD volumes 	<ul style="list-style-type: none"> RocksDB Performance counters

30

Squid(merged)

17

scrub: improving cluster-wide replica reservation for scrubs

core: integrate OpTracker into ceph-mgr

end2end tracing (RGW-OSD) using open-telemetry and jaeger

bluestore: remove zoned namespaces

bluestore: optimize Blob memory usage

QoS: snap trimming tuning

Tentacle(merged)

29

osd: ec partial stripe writes

osd: ec parity-delta-writes

EC reads could avoid reading padding

Fix the stretch modem encoding in pg_pool_t

bluestore: compression improvements

mon/mgr: track availability over time

Umbrella (merged)

5

osd: ec direct reads

mgr: autotune mgr_stats_period

mgr: perf counter support for modules

Enhance mgr cache

mgr: better module loading logic

RADOS

46

crimson

osd: omap in an ec pool

Support for pool migration

QoS Phase 2: client vs client foundation

Switch PGLog to DENC to reduce memory fragmentation

osd: ec efficient sparse objects

osd: ec distributed parity-delta-write

BlueStor

bluestor

bluestor improve

bluestor RocksDE

bluestor strategie

Improve

bluestor per retur

bluestor

Technology Outlook

Schedule Assumptions

DELIVERY

2026
H1

- 1/ IBMC 9.9 ✓
- 1/ RHCS 9 ✓
- 3/ ODF 4.21 ✓
- 6/ IBMSC 9.9.1

2026
H2

- 7/ ODF 4.22
- 9/ IBMSC 9.9.2, RHCS 9.2
- 11/ ODF 4.23
- 11/ Ceph “Umbrella”

2027

- 1/ IBMC 9.10 (Umbrella)
- 6/ IBMC 9.10.1
- 8/ Ceph “Vampire”
- 10/ IBMC 9.11 (Vampire)

2026 H1

- IBMC 9.9.1 ✓
- IBMC 9.9.2 ✓
- OSP 19

2026 H2

- IBMC 9.9.10
- CDS “V” (Future 9.11)

2027

- IBMC 9.10.1

PLANNING

Tech Roadmap: RADOS

2026 Q1
Ceph 9.0

- Fast EC
- ARM
- 10k OSD (control path)

2026 H2
Ceph 9.2

- Direct Read

2027 H2
Ceph 11

- Object packing
- Quantum safe
- Seastore w/Crimson full preview (TP)

2026 Q2
Ceph 9.1

- FCM5 support (Atollon)

2027 H1
Ceph 10

- Stretch EC
- Pool migration
- RGW EC on Crimson (TP)

2028+

- Dmcclock user QoS
- Crimson GA

Tech Roadmap: AI

2026 H1
Ceph 9.0


- Kv cache

2026 Q3
Ceph 9.2

- S3 Vectors

2027 H1
Ceph 10



- S3 over RDMA (cuObject / GPUdirect)
- Smart Assist
- MCP server



KV Caching with vLLM, LMCache, and Ceph

Kyle Bader (IBM), Tushar Gohad (Intel) → ceph.io

Community Roadmap



Major Technical Accomplishments in Tentacle

RADOS

Fast EC availability & performance improvements for block, file, and object. Various improvements to compression

NVMe-oF

General Availability of NVMe-oF Gateways with high availability support.

RGW

AWS IAM accounts for fine-grained control and self-service APIs.

Dashboard

IAM management, SMB support, and NVMe workflow additions.

CephFS

Case-insensitive directory trees for better interoperability (NFS, Samba).

Cephadm

Certificate management and critical scalability enhancements.

RBD

Live migration and namespace remapping during mirroring.

Crimson

SeaStore Tech Preview included in the Tentacle release.

Major Upcoming Roadmap Items

RADOS

Direct Reads & EC support for OMAP

Migration

Pool Migration tools and streamlined workflows

Robustness

Ceph-mgr cache implementation for system stability

Stretch Clusters

Full EC support for Stretch configurations

Efficiency & TCO

Small object packing for improved performance and lower total cost of ownership

Crimson & SeaStore

Block, object, and file workload support in Ceph's next-gen storage backend (TP)

Management

Improved upgrade experience and orchestration with cephadm enhancements

Thank You
