ceph

CEPH DAYS CHICAGO · 16 NOVEMBER 2026

State of the Cephalopod

Where Ceph is, where it is going, and who is taking it there

Neha Ojha · Dan van der Ster



Photo: NOAA Photo Library, CC BY 2.0

NEW TO CEPH?

What is Ceph?

Open source storage: block, file and object from one self-healing cluster

RBD	NVMe/TCP	CephFS	SMB ·NFS	RGW
block for VMs and Kubernetes	block for any host	POSIX file system	file shares	S3 and Swift objects

RADOS

reliable, autonomic, distributed object store: replication and erasure coding

Commodity hardware ·no single point of failure ·scales out by adding nodes ·built into Linux, OpenStack, Proxmox VE and Kubernetes (Rook)

NEW TO CEPH?

Ceph in the real world

Some of what the community has publicly shared

2.38

EiB reported by telemetry, November 2026

4.2k

clusters opted in to telemetry

100+

PB at CERN, across several clusters

21

stable releases since 2012

From home labs to enterprises, supercomputing centres, clouds and flight-test data.

Sources: telemetry-public.ceph.com, November 2026 (opt-in clusters only, so the real total is far larger); CERN IT at Cephalocon 2024, "Beyond Particle Physics"; Argonaut v0.48, June 2012; case studies at ceph.io/en/discover/case-studies.

THE COMMUNITY

Twelve months of Ceph

November 2025 to October 2026, in the main branch alone

6,947
commits merged

3,226
pull requests merged

~320
people wrote code

169
of them for the first time

About 19 commits and 9 merged pull requests, every day of the year.

Method: git log of ceph/ceph main, 1 Nov 2025 to 31 Oct 2026; commits exclude merges; contributors counted by author email (324 addresses).

RELEASES

A stable release every year

Squid reaches end of life with v19.2.8: target date 2 December

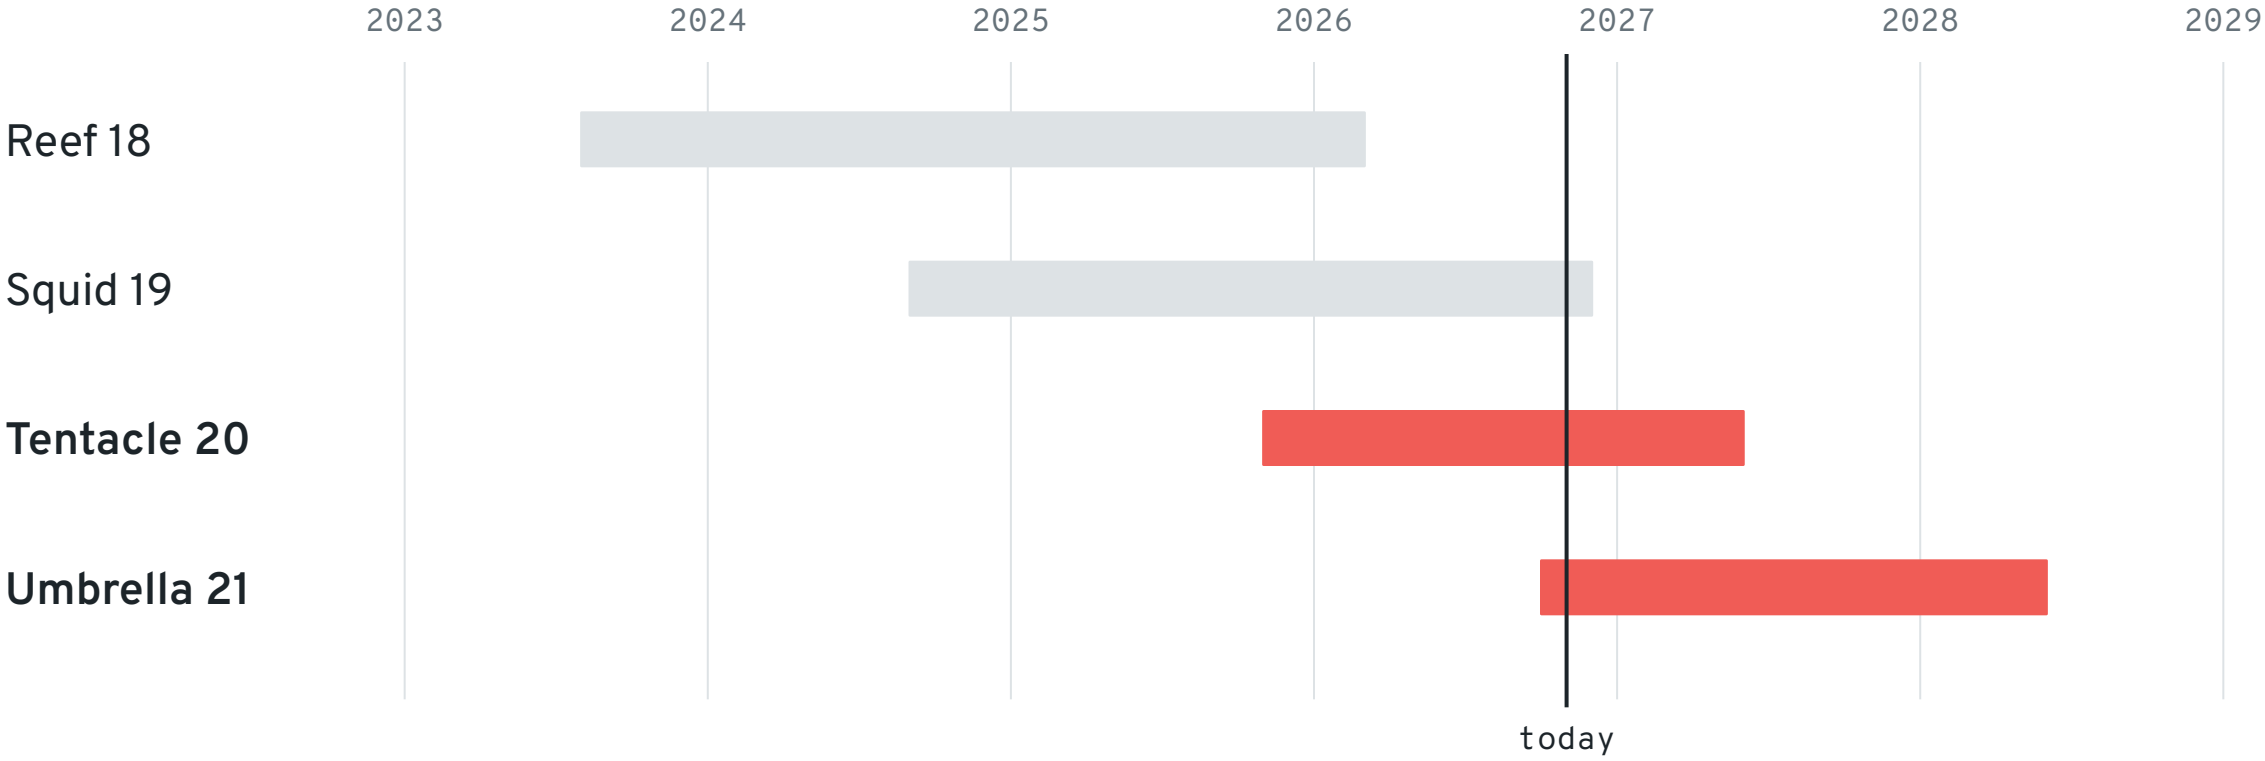

Bars run from the first stable release (x.2.0) to the actual or target end of life. Source: doc/releases/releases.yml.

Umbrella, released 28 October

The 21st stable release of Ceph

Core

- Fast EC: direct reads, OMAP, object classes
- CephX: AES-256 keys, online rotation
- Crimson: EC, recovery, PG splitting
- SeaStore: +30% IOPS, 86% lower P99 than Classic

Interfaces

- CephFS: fscrypt, faster mirroring
- RGW: dedup preview, bucket logging, cloud restore
- NVMe/TCP and SMB grow up

Operations

- Upgrades by CRUSH bucket, with ok-to-upgrade
- Monitor backup and restore
- Certificates managed by cephadm
- Manager cache follows map updates

Upgrading to Umbrella

What to check before you start

Before

- Container images need x86-64-v3 CPUs
- Python 3.9 or newer on every host for cephadm
- Small-memory or split-SSD OSDs: consider pinning the old PG limits

During and after

- Autoscaler now aims for 200 PGs per OSD, so pools may split
- pg_autoscale_during_upgrade paces the splitting
- Upgrade OSDs a CRUSH bucket at a time with ok-to-upgrade

Heading to Vampire (22)

Already merged on main for the next stable release

New

- Pool Migration
- EC pools in stretch clusters
- CephFS: faster lookups in large directories
- RGW: returns free memory to the OS
- RBD: weekly mirror schedules

Safer

- MDS failover avoids the failed host
- cephadm checks CPU support before upgrading
- Stricter bilog trimming in multisite

Retiring

- PMEM backends (Optane is gone)
- S3 Select, off by default
- Elasticsearch sync module

TODAY

Today at Ceph Days Chicago

The roadmap, in the words of the people building it

9:30 **Ceph on a Shoestring**
Francois Baillargeon

10:00 **BaraCUDA, Riding the Torpedo**
Kyle Bader

10:45 **Kubernetes-Native Identity for Ceph Object Storage**
Billy Ridgway

11:00 **Superconducting Ceph in the Age of AI**
Dan van der Ster

11:30 **Preparing Ceph for the SMR Era**
Dennis Maisenbacher

1:00 **Releasing the Kraken (Responsibly)**
Patrick Donnelly

1:15 **Minimizing Data Movement on CRUSH Failure Domain Changes**
Frédéric Nass

1:45 **The Blue Screen of Ceph**
Goutham Pacha Ravi

2:15 **Powered by People, Part 2: Where You Fit in Ceph**
Anthony Middleton

2:45 **The Power of RADOS: Making fsync (Almost) Free**
Gregory Farnum

3:15 **The Best of Both Worlds**
Steven Umbehocker

JOIN US

Everyone can contribute

You do not need to write code to make Ceph better

Use it and tell us

- Ask and answer on ceph-users@ceph.io
- Report bugs at tracker.ceph.com
- Opt in to telemetry: ceph telemetry on
- Give a talk at the next Ceph Day

Build it with us

- Code and reviews at github.com/ceph/ceph
- Improve the docs at docs.ceph.com
- Talk development on dev@ceph.io and Slack
- Join the community meetings and tech talks

ceph

Thank you

to the more than 2,000 people who have written Ceph, and to everyone here today

ceph.io · github.com/ceph

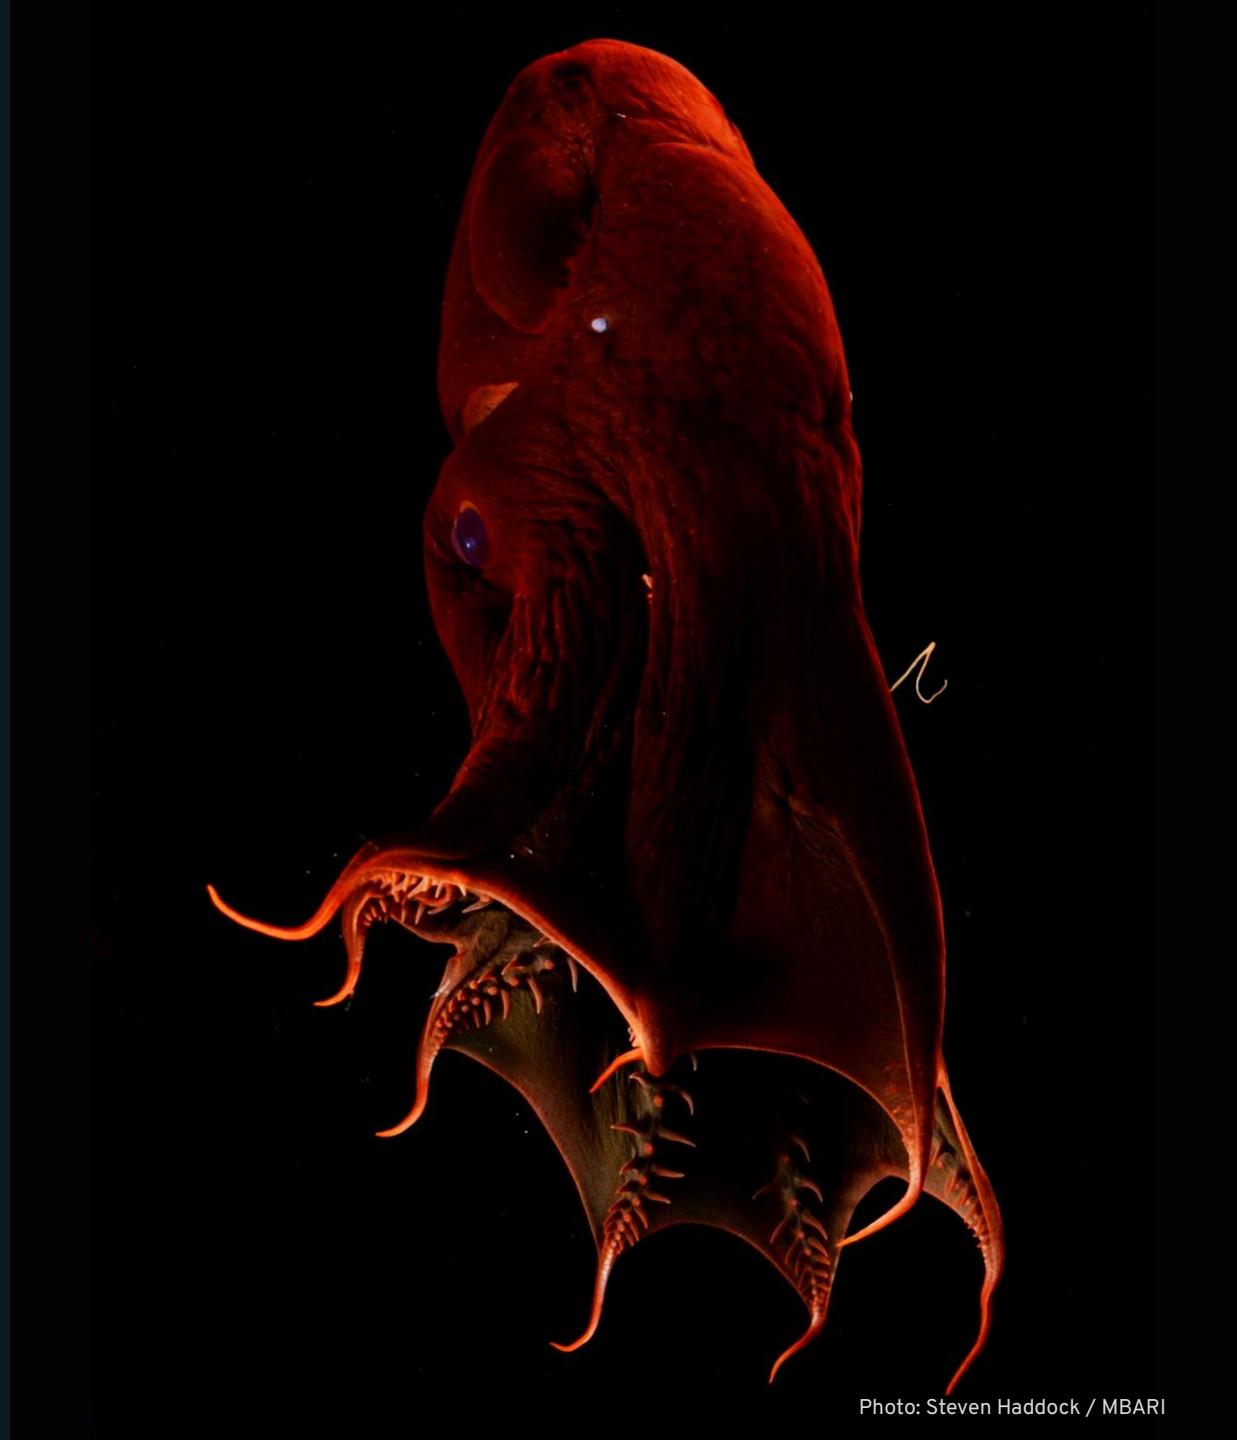

Photo: Steven Haddock / MBARI