



# CEPH DAYS

## SEATTLE 2026

OS NEWS CLUSO



# Today's Agenda



09:00 - 10:00	<i>Welcome, Check-in, Coffee, Pastries</i>
10:00 - 10:30	<b>Introduction / State of the Cephalopod 2026</b> – Steve (OSNexus) & Dan (CLYSO)
10:30 - 11:15	<b>Proxmox and Ceph: Ditch the old SAN</b> – Alex Gorbachev (Intelligent Systems Services)
11:15 - 11:30	<i>Coffee / Tea Break</i>
11:30 - 12:15	<b>Reclaim CephFS Capacity Through Pool Tiering + Migration</b> – Anthony D'Atri (Ceph Ambassador)
12:15 - 13:15	<i>Lunch</i>
13:15 - 14:00	<b>Ceph Bucket Replication</b> – Seth Cagampang + Sheetal Garde (OSNexus)
14:00 - 14:45	<b>Ceph Operations Toolkit: Ideas and Prototypes</b> – Dan van der Ster (CLYSO)
14:45 - 15:00	<i>Coffee / Tea Break</i>
15:00 - 15:45	<b>Ceph Cloud Storage Transitions and Beyond</b> – John Shubeck (IBM)
15:45 - 16:30	<b>Designing Ceph Clusters: What's the Right Durability?</b> – Steve Umbehocker (OSNexus)
16:30 - 17:00	<b>Speaker Panel Q&amp;A</b>
17:00 - 18:00	<i>Networking!</i>

# State of the Cephalopod - 2026

Dan van der Ster (CLYSO & Ceph Executive Council)

Ceph Days Seattle - May 28, 2026

# Outline

- What is Ceph?
- Peek into the Ceph Archives
- Ceph Project / Community Status
- Getting Started

# Storage Requirements

- **Multi-Protocol:** POSIX, Block, and Object from the one system
- **Use Commodity Hardware:** Avoid expensive hardware RAIDs
- **Fault Tolerant:** No Single Point of Failure
- **Horizontal Scalability:** Adding Capacity and IOPS should be trivial
- **Solution Maturity:** Avoid NIH-syndrome, leverage open source
- **Multiple Data Centres:** Scale globally

# Storage Requirements *(in 2013)*



- **Multi-Protocol:** POSIX, Block, and Object from the one system
- **Use Commodity Hardware:** Avoid expensive hardware RAIDs
- **Fault Tolerant:** No Single Point of Failure
- **Horizontal Scalability:** Adding Capacity and IOPS should be trivial
- **Solution Maturity:** Avoid NIH-syndrome, leverage open source
- **Multiple Data Centres:** Scale globally

# Storage Requirements *(in 2013)*



## **Proposal for a Petabyte-Scale Generic Storage Proof-of-Concept**

*Dan van der Ster, Arne Wiebalck / February 2013*

---

### **Executive Summary**

The medium-term needs related to (a) block storage for Agile Infrastructure VMs and (b) backend storage for AFS and NFS motivate a generic consolidated storage system for IT. Ceph is an appealing option because of its scalability/reliability design goals and relevant set of native access methods (object, block, and file). We therefore propose a proof-of-concept project to deploy a 1 PB (usable) Ceph evaluation cluster.



**Date:** Tue, 26 Feb 2013 15:37:05 +0100

**Subject:** subscribe

**From:** Dan van der Ster <dan@cern>

**To:** ceph-users-join@lists.ceph.com

# Storage Requirements *(in 2026)*



- **Multi-Protocol:** POSIX, Block, and Object from the one system
- **Use Commodity Hardware:** Avoid expensive hardware RAIDs
- **Fault Tolerant:** No Single Point of Failure
- **Horizontal Scalability:** Adding Capacity and IOPS should be trivial
- **Solution Maturity:** Avoid NIH-syndrome, leverage open source
- **Multiple Data Centres:** Scale globally



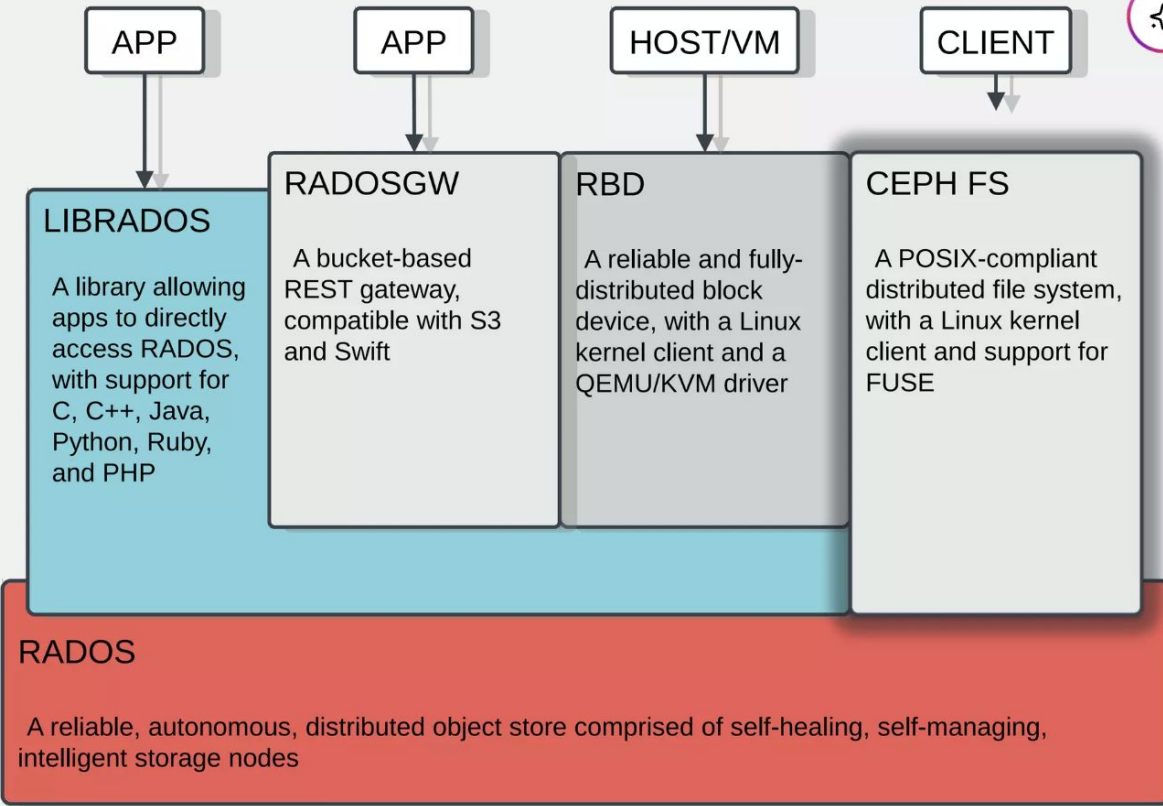
# Future of CephFS

Sage Weil

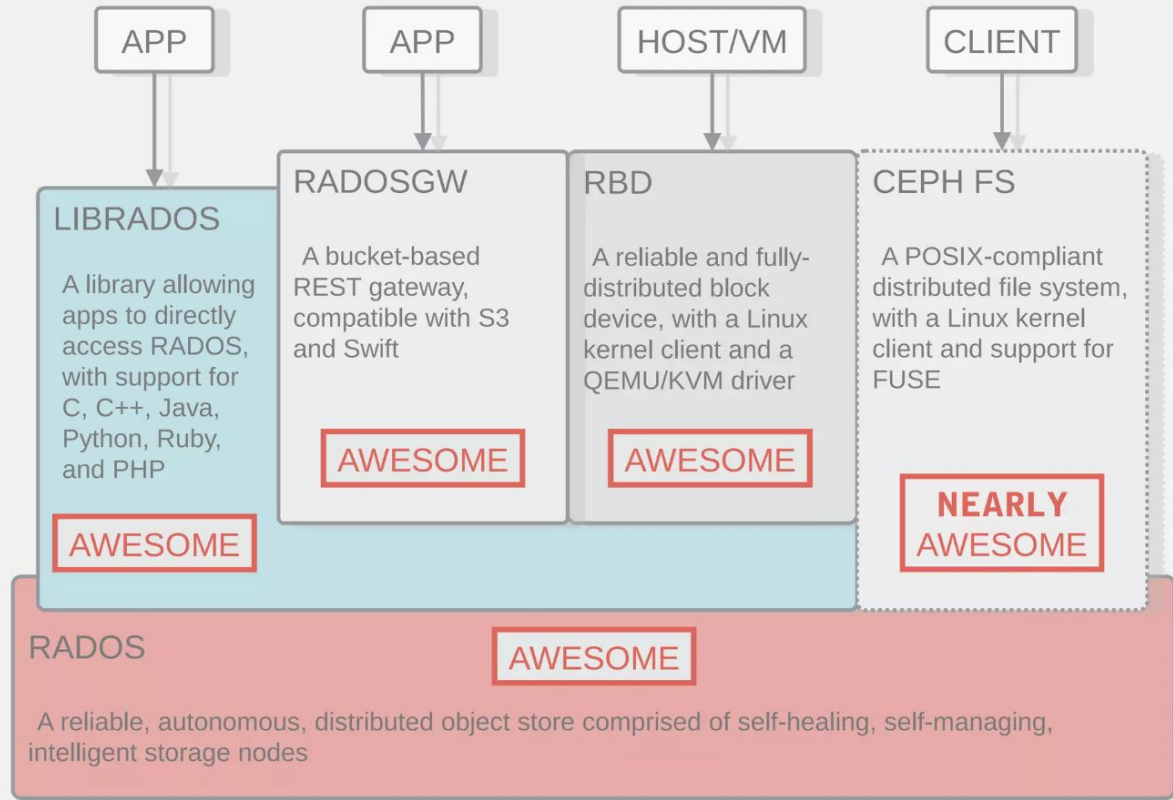
*Ceph Community Day London 2013*



**Sage Weil**  
Creator of Ceph



**Sage Weil**  
Creator of Ceph



**Sage Weil**  
Creator of Ceph

# How I got started with Ceph

## Proposal for a Petabyte-Scale Generic Storage Proof-of-Concept

*Dan van der Ster, Arne Wiebalck / February 2013*

---

### Executive Summary

The medium-term needs related to (a) block storage for Agile Infrastructure VMs and (b) backend storage for AFS and NFS motivate a generic consolidated storage system for IT. Ceph is an appealing option because of its scalability/reliability design goals and relevant set of native access methods (object, block, and file). We therefore propose a proof-of-concept project to deploy a 1 PB (usable) Ceph evaluation cluster.



**Date:** Tue, 26 Feb 2013 15:37:05 +0100

**Subject:** subscribe

**From:** Dan van der Ster <dan@cern>

**To:** ceph-users-join@lists.ceph.com

# How I got started with Ceph

## Proposal for a Petabyte-Scale Generic Storage Proof-of-Concept

*Dan van der Ster, Arne Wiebalck / February 2013*

---

### Executive Summary

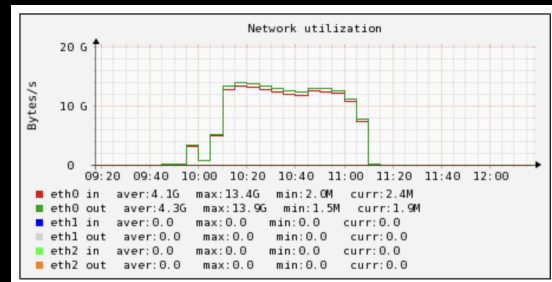
The medium-term needs related to (a) block storage for Agile Infrastructure VMs and (b) backend storage for AFS and NFS motivate a generic consolidated storage system for IT. Ceph is an appealing option because of its scalability/reliability design goals and relevant set of native access methods (object, block, and file). We therefore propose a proof-of-concept project to deploy a 1 PB (usable) Ceph evaluation cluster.

**Date:** Tue, 26 Feb 2013 15:37:05 +0100

**Subject:** subscribe

**From:** Dan van der Ster <dan@cern>

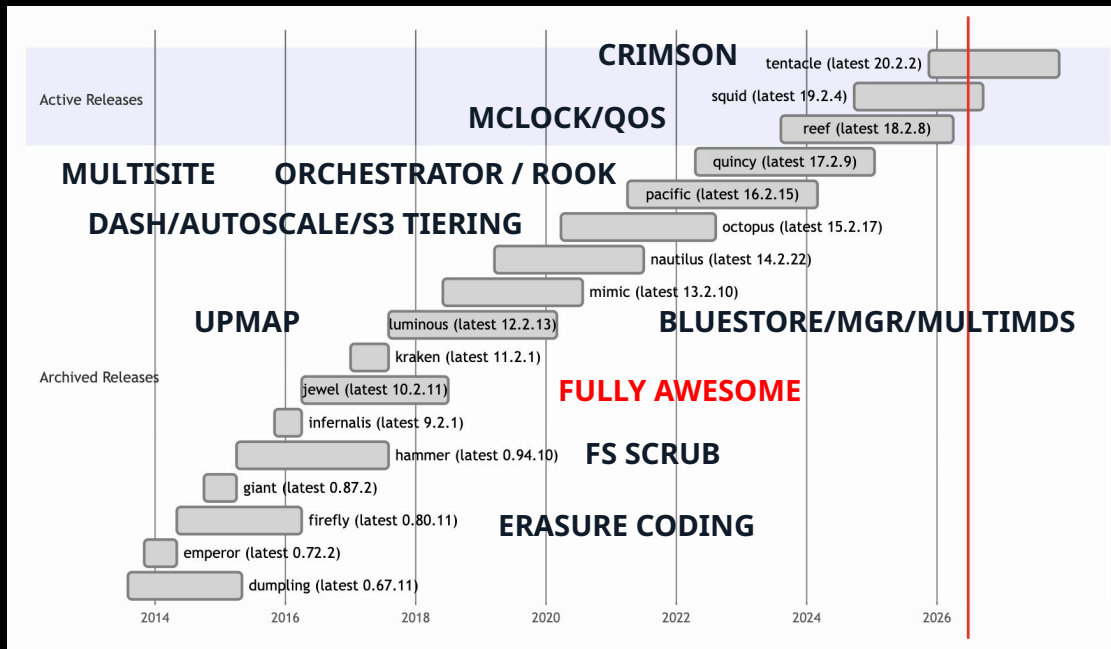
**To:** ceph-users-join@lists.ceph.com



# History of Ceph

- 2004: Research at UC Santa Cruz: ***Build a modern replacement for Lustre!***
- 2004: **First commit** to Ceph source code
- 2006: CRUSH research paper published
- 2010: RBD contributed by a community member
- 2010: OpenStack starts as a project of Rackspace and NASA
- 2011: Inktank Storage founded
- 2012: Argonaut, the first stable release of Ceph
- 2014: “Red Hat Buys Inktank”
- 2015-16: Development of Newstore/Bluestore.
- **2016: CephFS is All Ceph components are now “FULLY AWESOME”**
- 2015-18: Ceph Advisory Board formed
- 2018: Ceph joins the Linux Foundation - Our first Cephalocon (Beijing)

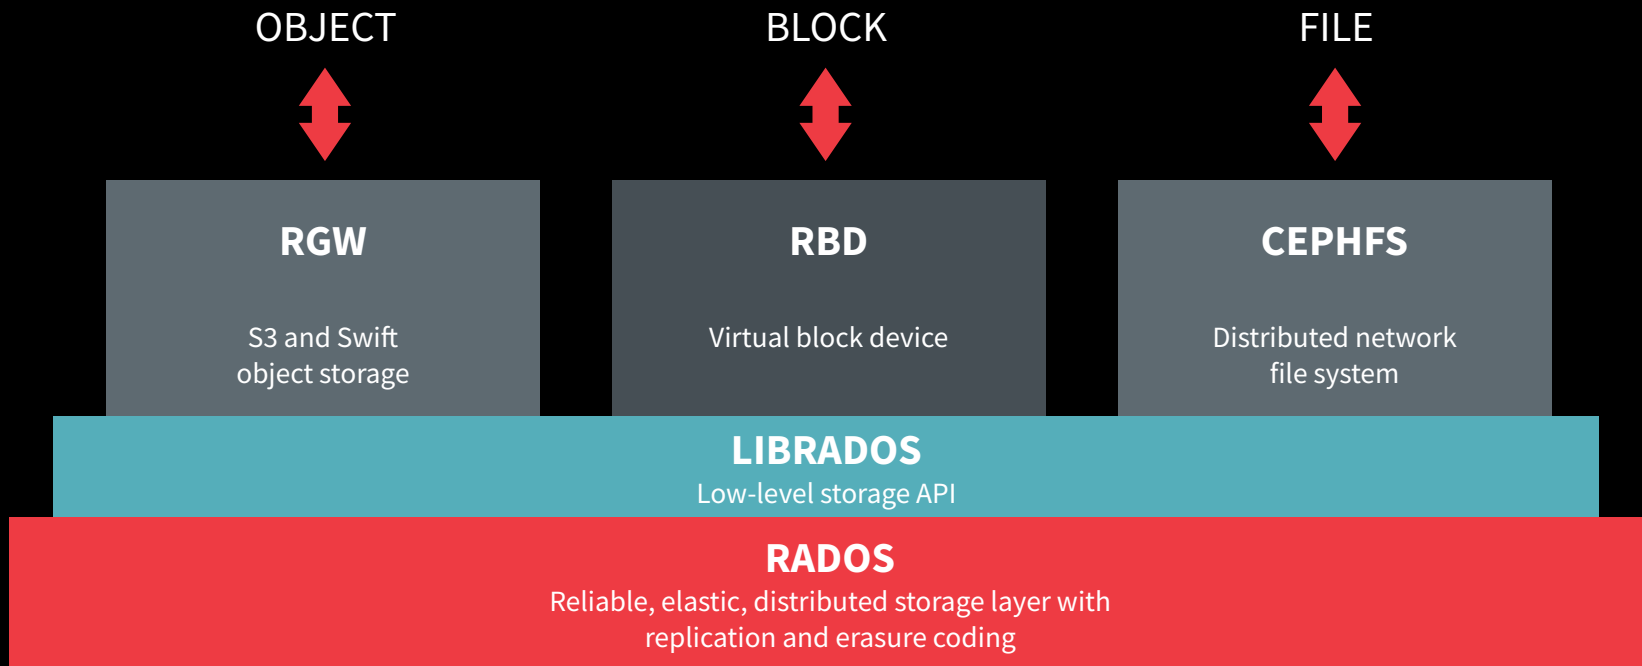
# HISTORY OF THE CEPH PROJECT ceph days



# Ceph: Unified, Distributed, Open Storage ceph days

- **Truly Unified Platform:** Block (RBD), object (RGW), and file (CephFS)
- **Massive Scale in Prod:** More than 2 Exabytes across ~4,000 active clusters.
- **Dominant Integration:** Powers over 90% of Open Infra clouds + K8s
- **Extreme Performance:** Real world 1 TiB/s and 25 Million+ IOPS.
- **Total Freedom:** Fully open source with zero hardware or licensing lock-in.

# Unified Storage with Ceph



Who uses Ceph? Here's a sample of speakers at our events

Flipkart



UK Research and Innovation



HUAWEI



Elettra Sincrotrone Trieste



ceph days

OVHcloud



engin

INDITEX



etraveli group

Switch

SPACEX

Proton

VULTR



LINE



Deutsche Telekom

SAMSUNG



NEURALINK

Your. Online



DigitalOcean




pawsey

Bloomberg

ETH zürich

Walmart



Sony Interactive Entertainment


workday

# Ceph Project Status - May 2026



- **Active Releases:**
  - The current active Ceph releases are **Squid (v19) and Tentacle (v20)**.
  - Real-world deployment trends can be seen on the [Ceph Telemetry Dashboard](#).
- **Tentacle (v20.2.1):**
  - Dramatically improves **erasure coding performance** for faster, more cost-efficient storage.
  - Enhanced **gateway support** with major upgrades to **Samba, NFS, and NVMe-oF**.
- **Ceph is *really* joining the Linux Foundation:**
  - The Ceph Foundation was formed in 2018 - Not-For-Profit in support of the Ceph Project
    - “Ceph Project” remained independent, community managed by RedHat
  - Last year we started the process of formally establishing Ceph in the Linux Foundation
    - Ceph Steering Committee signed a Technical Charter with the LF
- **Introducing our new dedicated Tech Writer Emmanuel Ameh:**
  - <https://ceph.io/en/news/blog/2026/from-laminated-pages-to-petabytes/>

# Ceph Release Schedule



- Stable, named release every 12 months
- Backports for 2 releases
  - 18.2.8 was the final Reef release
- Upgrade up to 2 releases at a time
  - Pacific → Reef, Quincy → Squid, Reef → **Tentacle**

# Ceph Project Health

- [LFX Insights for Ceph](#)

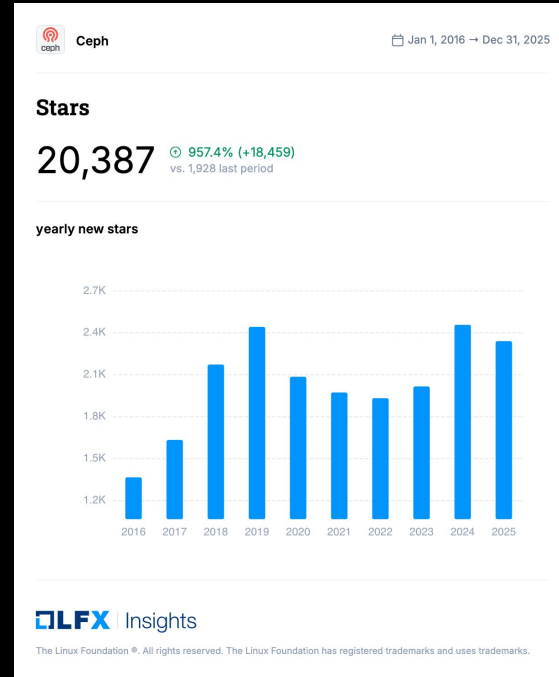


## ● Healthy

The Insights Health Score combines the four key areas to measure an open source project's overall trustworthiness. [Learn more](#)

# Ceph Project Health

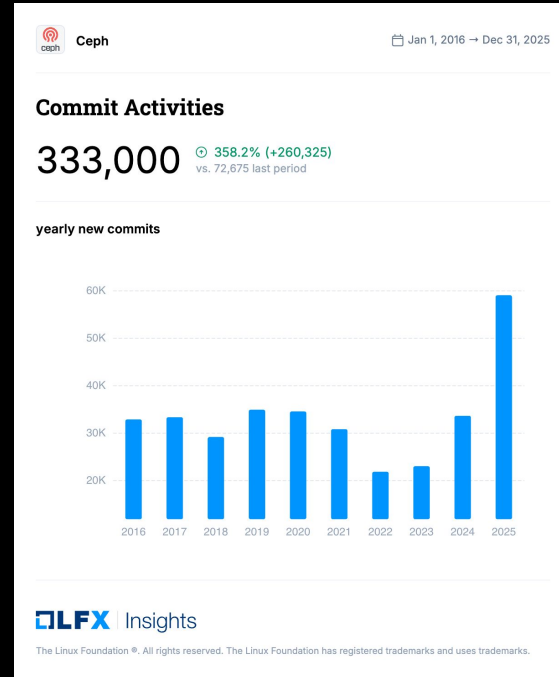
- LFX Insights for Ceph



# Ceph Project Health



- LFX Insights for Ceph



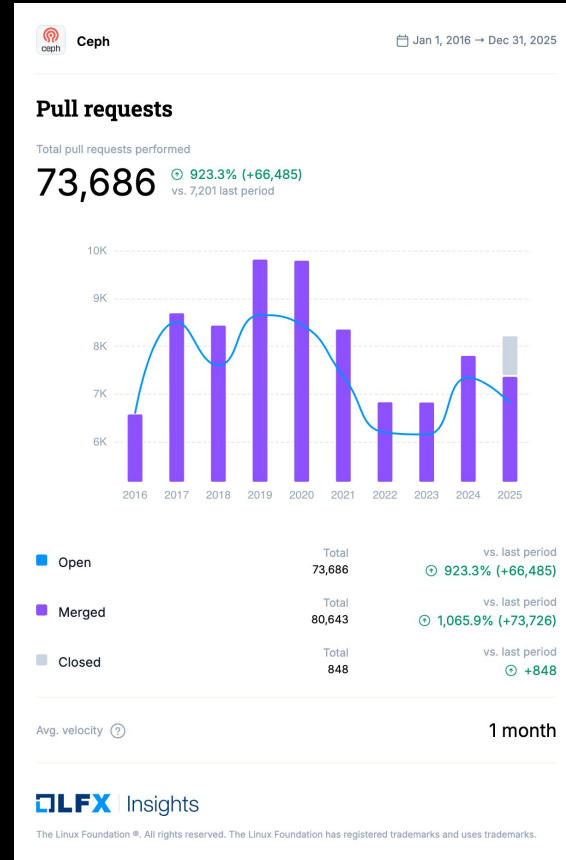
# Ceph Project Health

- LFX Insights for Ceph



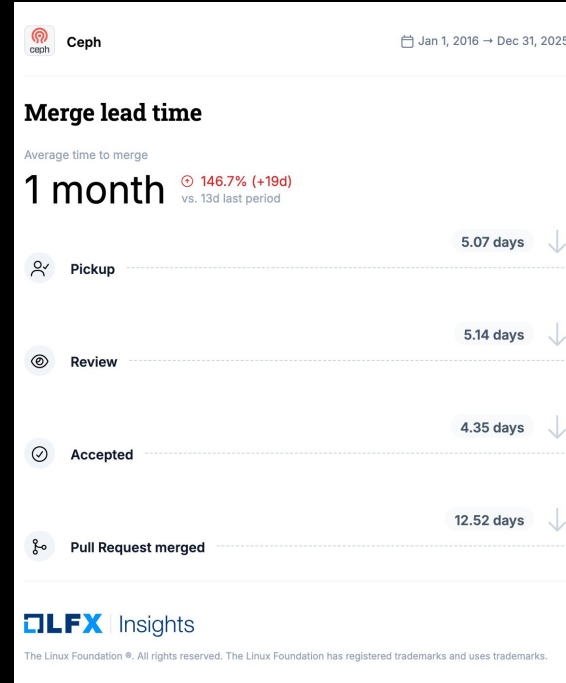
# Ceph Project Health

- LFX Insights for Ceph



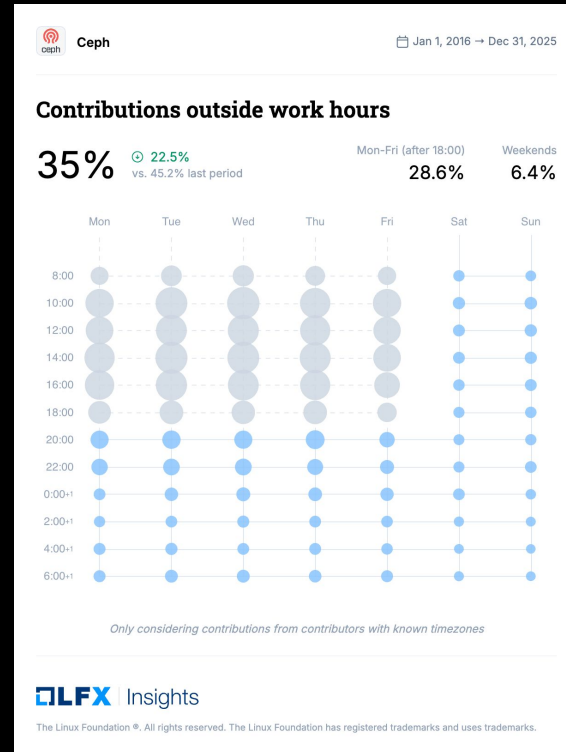
# Ceph Project Health

- LFX Insights for Ceph



# Ceph Project Health

- LFX Insights for Ceph



# Getting Started

- **Try It Out:** Install with cephadm:

```
curl -O https://download.ceph.com/rpm-tentacle/e19/noarch/cephadm  
chmod +x cephadm && sudo ./cephadm bootstrap
```

- **Learn More**

- Docs: [docs.ceph.com](https://docs.ceph.com)
- Slack and Mailing Lists: <https://ceph.io/en/community/connect/>

- **Get Involved**

- Join a [Ceph Day](#) or [Cephalocon](#)
- Contribute on [GitHub](#)
- Become a member of the [Ceph Foundation](#)

*Ceph is open, active, and ready — dive in and build something!*

## Coming up



### Ceph Days Seattle 2026

May 28, 2026

Bringing Ceph back to Seattle Join us for a full-day celebration of Ceph in Seattle, bringing...



### Ceph Days London 2026

Jun 17, 2026

Bringing Ceph to London yet again A full-day event dedicated to sharing Ceph's transformative power...



### Ceph Meetup Frankfurt June 2026

Jun 22, 2026

Ceph Meetup Frankfurt — June 2026 Join the Ceph community in Frankfurt for an evening of technical...



09:00 - 10:00	<i>Welcome, Check-in, Coffee, Pastries</i>
10:00 - 10:30	<b>Introduction to Ceph &amp; The State of the Cephalopod in 2025</b> – Steve (OSNexus) & Dan (CLYSO)
10:30 - 11:15	<b>Proxmox and Ceph: Ditch the old SAN</b> – Alex Gorbachev (Intelligent Systems Services)
11:15 - 11:30	<i>Coffee / Tea Break</i>
11:30 - 12:15	<b>Reclaim CephFS Capacity Through Pool Tiering + Migration</b> – Anthony D'Atri (Ceph Ambassador)
12:15 - 13:15	<i>Lunch</i>
13:15 - 14:00	<b>Ceph Bucket Replication</b> – Seth Cagampang + Sheetal Garde (OSNexus)
14:00 - 14:45	<b>Ceph Operations Toolkit: Ideas and Prototypes</b> – Dan van der Ster (CLYSO)
14:45 - 15:00	<i>Coffee / Tea Break</i>
15:00 - 15:45	<b>Ceph Cloud Storage Transitions and Beyond</b> – John Shubeck (IBM)
15:45 - 16:30	<b>Designing Ceph Clusters: What's the Right Durability?</b> – Steve Umbehoeker (OSNexus)
16:30 - 17:00	<b>Speaker Panel Q&amp;A</b>
17:00 - 18:00	<i>Networking!</i>