

OS NEXUS



CL'USO

Today's Agenda



| | |
|---------------|---|
| 12:00 - 13:00 | <i>Welcome, Check-in, Coffee, Pastries</i> |
| 13:00 - 13:30 | Introduction to Ceph: The State of the Cephalopod in 2025 – Dan van der Ster (CLYSO) |
| 13:30 - 14:00 | Choosing the Right Data Protection Strategies For Your Ceph Deployments – Anthony D'Atri (IBM) |
| 14:00 - 14:30 | Ceph Solution Design Tool – Steven Umbehocker (OSNEXUS) |
| 14:30 - 15:00 | <i>Coffee / Tea break</i> |
| 15:00 - 15:30 | Ceph in Proxmox VE – Alex Gorbachev (ISS) |
| 15:30 - 16:00 | NVMe over TCP and Block Performance – Mike Burkhart (IBM) |
| 16:00 - 16:30 | Ceph Object Storage - Keycloak ID Broker With Azure – Seth Cagampang (OSNEXUS) |
| 16:30 - 17:00 | Ceph Durability: How Safe Is My Data? – Dan van der Ster (CLYSO) |
| 17:00 - 17:30 | Optimizing Ceph RGW for Specific Workloads Including AI – Steven Umbehocker (OSNEXUS) |
| 17:30 - 19:00 | <i>Dinner & Drinks & Networking oh my!</i> |

Introduction to Ceph

State of the Cephalopod in 2025

Dan van der Ster (CLY'SO & Ceph Executive Council)

Ceph Days Seattle - May 15, 2025



Outline

- What is Ceph?
- Why Ceph? Who uses it?
- Project Status & Roadmap
- Getting Started & Community

What is Ceph?

- **Open Source Software-Defined Storage**
- **Unified:** Object, Block, and File in one system
- **Scalable:** From a few nodes to exabyte-scale environments
- **Resilient:** No single point of failure, self-healing, self-managing
- **Community-Driven:** Built by and for operators, researchers, enterprises
- **Proven:** Runs at scale at CERN, Bloomberg, cloud providers, and beyond

Ceph lets you build your own cloud-grade storage, with commodity hardware and full control.

Ceph is the *Linux* of Storage

- **Free & Open Source**
 - No licenses. No vendor lock-in. Total control.
- **Community-Powered**
 - Built and improved by a global ecosystem of users, developers, and companies.
- **Runs Everywhere**
 - From homelabs to hyperscale datacenters. Bare metal, VMs, Kubernetes, cloud.
- **Modular & Flexible**
 - Ceph is pluggable, extensible, adaptable to any workload.

If Linux gave us freedom to compute, Ceph gives us freedom to store.

Unified Storage with Ceph



OBJECT



RGW

S3 and Swift
object storage

BLOCK



RBD

Virtual block device

FILE



CEPHFS

Distributed network
file system

LIBRADOS

Low-level storage API

RADOS

Reliable, elastic, distributed storage layer with
replication and erasure coding

Ultra-Reliable Storage from Unreliable Parts



- **Ceph vs. Hardware**

- Exposed HDDs lying about flush — disabled volatile caches.
- Caught NVMe and CPU firmware bugs causing crashes after long uptimes.
- Detected unsafe power-loss behavior → drove improvements in vendor QA.

- **Ceph vs. Network**

- Detects and survives silent packet corruption, flaky links, and bad routing.
- Often the first system to notice when the network breaks.

- **Ceph vs. Linux**

- Found bugs in memory allocators, block layers, and LZ4 compression.
- BTRFS and XFS too unpredictable → built BlueStore for performance and reliability.

Ceph isn't just resilient — it makes the whole stack better.

Who uses Ceph?

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



DigitalOcean



pawsey


Bloomberg

ETH zürich



Your.
Online

Walmart



Sony
Interactive
Entertainment

workday

Who uses Ceph?




Speakers at recent Ceph events:

UK Research and Innovation | Elettra Sincrotrone Trieste
OVHcloud | Engin IT | INDITEX | etraveli | SWITCH | SpaceX |
Vultr | MLB | LINE | Deutsche Telekom | Samsung | Workday
Neuralink | Your.Online | Digital Ocean | Huawei | CERN
Pawsey Supercomputing Research Centre | Bloomberg | Flipkart
ETH Zurich | Walmart | Sony Interactive Entertainment | Proton

Nearly 4000 clusters reporting to our public telemetry.

Who uses Ceph?



Source: <https://www.openstack.org/analytics>

What do people use Ceph for?

- **Cloud Infrastructure**
 - Object, block, and file storage for OpenStack, Kubernetes, and sovereign clouds
- **High-Performance Computing**
 - Shared filesystems (CephFS) and fast object stores for scientific data & AI training
- **Enterprise Backup & Archiving**
 - Scalable, durable object storage for long-term data retention
- **AI / ML Pipelines**
 - High-throughput block and object storage for model training and inference
- **Media & Content Delivery**
 - Ceph RGW used behind CDNs and media platforms for video, image, and file delivery
- **Internal Developer Platforms**
 - Backing internal services, CI pipelines, artifact storage, and observability data

Ceph adapts to your needs — from pet projects to petabytes in production.

Ultra High Performance Ceph



Full Cluster Msgr Thread Scaling - FIO 4MB Throughput

LibRBD, 3X Rep, 256K PGs, 8 Shards, 2 Threads/Shard, 504 Client Procs, Reef RocksDB Tuning



Full Cluster Msgr Thread Scaling - FIO 4KB IOPS

LibRBD, 3X Rep, 256K PGs, 8 Shards, 2 Threads/Shard, 504 Client Procs, Reef RocksDB Tuning



- **Myth:** Ceph is too slow for modern workloads
- **Reality:** Ceph exceeds 1 TiB/s and 25 million IOPS in real-world tests
 - <https://ceph.io/en/news/blog/2024/ceph-a-journey-to-1tibps/>

Ceph isn't slow. It's just waiting for the right tuning.

Innovation through Collaboration



- Ceph thrives because of contributions from users, researchers, vendors, and integrators around the world.
- A few notable highlights:
 - RBD was contributed by a user in 2010!
 - BlueStore had major input from Red Hat, SanDisk, SUSE, ZTE, XSKY, and others
 - Samsung is leading cluster-wide deduplication efforts
 - Crimson is backed by IBM, Intel, Samsung, Qi An Xin, and more



Ceph Reef Credits, 2023

Ceph isn't built by one company — it's built by all of us.

Ceph Project Status - May 2025



- **Active Releases:**
 - The current active Ceph releases are **Reef (18.2.x)** and **Squid (19.2.x)**.
 - Real-world deployment trends can be seen on the [Ceph Telemetry Dashboard](#).
- **Tentacle (20.x):**
 - Development for the next major release has just entered feature freeze.
 - The release is expected later in 2025.
- **What's Coming in Tentacle:**
 - Dramatically improved **erasure coding performance** for faster, more cost-efficient storage.
 - Enhanced **gateway support** with major upgrades to **Samba**, **NFS**, and **NVMe-oF**.

Ceph 20 brings practical improvements in performance, cost-efficiency, and protocol support.

Getting Started with Ceph

- **Try It Out:** Install with cephadm (Squid release):

```
curl -O https://download.ceph.com/rpm-squid/el9/noarch/cephadm  
chmod +x cephadm && sudo ./cephadm bootstrap
```

- **Learn More**

- Docs: 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!



cephalocon

2025

"Pssst ... Did you hear? " ...

"What? Tell me!" ...

"Cephalocon 2025 is on!" ...

"When?Where?" ...

**"You didn't hear it from me: It looks like it
might be at end of October in Vancouver" ...**

"Can't wait for the formal announcement" ...

"Yes, going to belepica!"

More Info Coming Soon



| | |
|---------------|---|
| 12:00 - 13:00 | <i>Welcome, Check-in, Coffee, Pastries</i> |
| 13:00 - 13:30 | Introduction to Ceph: The State of the Cephalopod in 2025 – Dan van der Ster (CLYSO) |
| 13:30 - 14:00 | Choosing the Right Data Protection Strategies For Your Ceph Deployments – Anthony D'Atri (IBM) |
| 14:00 - 14:30 | Ceph Solution Design Tool – Steven Umbehocker (OSNEXUS) |
| 14:30 - 15:00 | <i>Coffee / Tea break</i> |
| 15:00 - 15:30 | Ceph in Proxmox VE – Alex Gorbachev (ISS) |
| 15:30 - 16:00 | NVMe over TCP and Block Performance – Mike Burkhart (IBM) |
| 16:00 - 16:30 | Ceph Object Storage - Keycloak ID Broker With Azure – Seth Cagampang (OSNEXUS) |
| 16:30 - 17:00 | Ceph Durability: How Safe Is My Data? – Dan van der Ster (CLYSO) |
| 17:00 - 17:30 | Optimizing Ceph RGW for Specific Workloads Including AI – Steven Umbehocker (OSNEXUS) |
| 17:30 - 19:00 | <i>Dinner & Drinks & Networking oh my!</i> |