

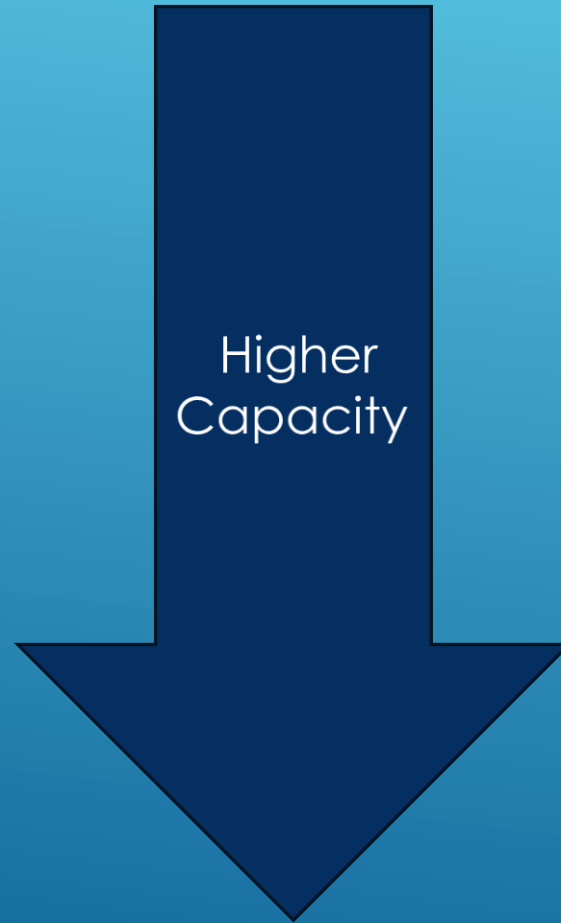

SCD CEPH SERVICES

The image features a solid blue background with a gradient from light blue at the top to a darker blue at the bottom. In the bottom right corner, there are several white, parallel diagonal lines that create a sense of motion or a stylized graphic element.

OUR CEPH SERVICES, SUMMARISED

Arided

Echo



Sirius

Deneb



SIRIUS

- ▶ Tell me about the cluster
 - ▶ 230TB usable NVMe flash (700TB raw)
 - ▶ 3*Replicated
- ▶ What is the access method?
 - ▶ RADOS Block Device
- ▶ What it is for?
 - ▶ Image storage for the SCD Cloud
- ▶ What is it not for?
 - ▶ Anything else
- ▶ Can I use it?
 - ▶ No

“I am the SCD
Cloud and I
need an image
store”

ARIDED

- Tell me about the cluster
 - 400TB usable SATA SSDs (1200TB raw)
 - 3*Replicated
- What is the access method?
 - Provisioned CephFS mounts via Manila
- What it is for?
 - A shared workspace for processing experimental data
 - Users can self-provision native CephFS shares and mount them from Cloud VMs
 - Diverse user requirements and workloads (3*replication is better for small files than EC)
 - Use cases that require a full-fat POSIX filesystem
- What is it not for?
 - Bulk data storage
 - Heavy parallel access
- Can I use it?
 - Yes, in principle, but it's not intended for SCD

**“I need a shared,
high-
performance
workspace”**

DENEB

- Tell me about the cluster
 - 5.4PB non-capacity HDDs (7.5PB raw)
 - 8+3 EC
- What is the access method?
 - Generic CephFS mounts
- What it is for?
 - Data sharing that needs a POSIX interface
 - Larger files
 - Bulk use cases
- What is it not for?
 - Long term archival
 - Heavy parallel access
- Can I use it?
 - Yes – ask the Ceph team

“I have a lot of data to share internally, and my users need a POSIX interface”

ECHO (S3/SWIFT)

- Tell me about the cluster
 - 100PB capacity HDDs (7.5PB raw)
 - 8+3 EC
- What is the access method?
 - Web-standard S3 and SWIFT protocols on discrete gateway hosts
- What it is for?
 - Access external to RAL
 - Bulk storage
 - Parallel access
- What is it not for?
 - Low-latency use cases
 - Direct use by individuals
- Can I use it?
 - Yes – ask the Ceph team

“I have some data that I want to share and I’m fine with a basic interface”

ECHO (XROOTD)

- Tell me about the cluster
 - 100PB capacity HDDs (7.5PB raw)
 - 8+3 EC
- What is the access method?
 - XrootD gateway processes on batch workers and discrete gateway hosts
- What it is for?
 - High-throughput grid computing
- What is it not for?
 - Normal people
 - Low-latency use cases
- Can I use it?
 - Almost certainly not

“We are a large collaboration, and want to do grid computing using the particle physics toolset”