

Ceph: An Open-Source Petabyte-Scale File System

Storage Systems Research Center
University of California, Santa Cruz



What is Ceph?

- ✦ Petabyte-scale file system using object-based storage
 - Based on research at UC Santa Cruz and elsewhere
 - Scalable: grows gracefully from terabytes to petabytes
- ✦ Goals
 - Capacity: 10s of petabytes
 - Bandwidth: 1+ TB/second
 - High I/O (metadata) rate
- ✦ Features
 - Decentralized data distribution
 - Scalable metadata service
 - Replication and fast recovery for reliability
 - High-performance local object-based file system
- ✦ Open source!

Ceph status

- ◆ Open source project hosted at SourceForge:
<http://ceph.sourceforge.net/>
 - Download source code
 - View documentation on setting up and running Ceph
 - Current active development areas
 - In-kernel client (default client is a user library)
 - Snapshots
- ◆ Contribute: Ceph is looking for developers
 - Still alpha-quality code
- ◆ Full list of Ceph-related publications at
<http://www.ssrc.ucsc.edu/proj/ceph.html>