

Object Storage...

An update in brief.

Christopher MacGown,
Co-Founder and Chief Scientist, Piston Cloud



- What is Object Storage
- Community Update
- Features Update
- Challenges



What is Object Storage?

- Storage without file-system semantics
 - Files accessed over simple web API
 - Often provide versioned writes
 - Simple ACLs
- Scalable and durable by replication and distribution
 - Allows for dynamic capacity adjustment
 - Some leverage distribution to allow for highly concurrent reads and writes...
- How large technology companies store multiple hundreds of petabytes of data that needs to be accessed at scale quickly



Community

- Presenting update on two projects...
 - Ceph: Strongly Consistent Block and Object
 - Swift: Eventually Consistent Object
- Others exist:
 - RiakCS
 - Amazon S3
 - Azure
 - Facebook Haystack



Community: Ceph

- Two major releases since last September
 - Emperor (November 2013)
 - Firefly (May 2014 LTS)
- Primary commercial developer was acquired by Red Hat in April
 - Red Hat now coupling Ceph with Gluster
- Commercial adoption has drastically increased
 - All major OpenStack vendors support
 - New storage appliances based on ceph announced this month at IDF
- About 175 contributors



Community: OpenStack Swift

- Two major OpenStack releases
 - Havana (October 2013 - security supported)
 - Icehouse (April 2014 - current stable)
 - Six releases of OpenStack swift
- Continued commercial adoption
 - Primarily by large cloud operators
 - New storage appliances announced this year by Seagate
- Almost two hundred contributors



Features Update: Ceph

- Multiple Data Center replication
- Support for iSCSI block export
- Erasure coding
- Cache tiering
- Affinity skew to improve read throughput
- Experimental support for key-value backends...



Features Update: Swift

- Multi-realm storage replication
- Policy based storage tiering
- Replicator speed and reliability improvements
- Better object proxy middleware¹
- Early quorum responses²
- ZeroVM...



Challenges...

- Between the two projects there are 2000 developers and 200 companies
 - Lots of competing interests
 - Mostly commercial
 - Academic focus is primarily computer science
 - Features driven by interest
 - Scalability
 - Mobility
 - Live analysis



Thank you.

