

Object Storage Challenges

Ethan Miller

Storage Systems Research Center
University of California, Santa Cruz
elm@cs.ucsc.edu



UC Santa Cruz



Where are we now?

- ◆ We can
 - Build object storage disks (OSDs)
 - Hook them together to store lots of data
 - Find data we've stored in the system
- ◆ But...
 - We don't have much experience with this
 - We tend to map traditional file system concepts onto object-based storage systems
- ◆ We can build an OSD cluster today!
 - People on this panel are doing it
- ◆ Can we do it *well*?



Challenges for OSD systems

- ◆ Managing the data
 - What should an OSD file system look like?
 - Indexing and using all of that data
 - Metadata management (particularly scaling)
 - Safely sharing data between users (locking, etc.)
 - Can't require too much human intervention
- ◆ Scalability
 - Thousands of devices are great, but...
 - How do you actually find anything without consulting a central “oracle”?
 - What happens when you add new disks or retire old ones?
 - Simply putting new data on the new disks may not be a good idea...
- ◆ Wide-area access
 - Where does the data live?
 - How does it move from place to place?



Challenges for OSD systems

◆ Reliability

- Large systems will lose multiple disks per day
- Disk rebuild takes ~1 day (or more!)
- System can't drop in performance when a component fails
- Detecting a failure can be difficult

◆ Security

- Devices are on the Internet!
- Security is tougher in an open, shared environment
- Ensuring integrity will be a full-time job

◆ We've made some progress!

◆ There's lots left to do before object-based storage can realize its potential

