


S3 Object Deduplication in Ceph

Gabriel BenHanokh (gbenhano@ibm.com)
Yuval Lifshitz (ylifshit@ibm.com)



The Question:

- Full S3 Object dedup or Chunk level dedup?
 - Full S3 Object dedup is far more efficient
 - Consumes less resources, lower latency
 - Easier and faster to implement
 - Chunk level dedup yields higher dedup ratio
 - Does it worth it?
- Answering this requires access to real world user data

Help Wanted!

The Estimation Tool

- Easy to use and fast (up to 100K objects/sec) tool to estimate deduplication levels
- Based on eTags that are already available in S3 (not ideal for multipart uploads)
- Low impact on the cluster – the user controls the thread count
- No dependencies (container based)
- Open-source tool available at:
 - https://github.com/benhanokh/s3_dedup_estimate/tree/main

Full S3 Object: real-world data

- Cloud Computing system with huge objects
 - 114 buckets / 21,439 Objects / 49.3TB
 - dedup ratio 1.24
- "images" zone:
 - 651,770 Objects / 13.5TB (20MB avg)
 - dedup-ratio of 1.07 (1.35TB -> 12.6TB)
- Standard zone:
 - 4,872,996 Objects / 2.75TB (560KB avg)
 - dedup-ratio of 1.22
- Archived zone:
 - 4,872,996 Objects / 3TB
 - dedup-ratio of 1.27

IBM