SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Navigating the Complexities of Object Storage Compatibility

Live Webinar

January 9, 2024

10:00 am PT / 1:00 pm ET

Today's Presenters

Michael Hoard
Chair, SNIA Cloud Storage Technologies Initiative
Intel

Gregory Touretsky
Chair, SNIA Cloud Storage Technical Work Group
Seagate

Arnold Jones
SNIA Technical Director

The SNIA Community

200
Corporations, universities, startups, and individuals

2,500
Active contributing members

50,000
Worldwide IT end users and professionals

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

What We Do

Educate vendors and users on cloud storage, data services and orchestration

Support & promote business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage

Understand Hyperscaler requirements Incorporate them into standards and programs

Collaborate with other industry associations

SNIA Legal Notice

- The material contained in this presentation is copyrighted by SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Agenda

- Object Storage Implementations Landscape
- Incompatibility Examples
- Interoperability Plugfest and Call for Action

Object Storage Implementations Landscape

Gregory Touretsky

Object Storage Protocols and Implementations - Server

- Simple Storage Service (S3)
  - Amazon, Ceph-RGW, Minio, Cloudian, Wasabi, Backblaze, Seagate Lyve, IBM COS, DellEMC, NetApp, Hitachi, Oracle Cloud Storage, Scality, Google Cloud Storage (*), CloudFlare, Akamai, DigitalOcean, …
- Google Cloud Storage
- Azure Blob Storage
- OpenStack Swift
- SNIA Cloud Data Management Interface (CDMI™)
  - Several open source and commercial (NetApp StorageGrid, Scality Ring, etc) implementations

Object Storage Client Implementations

Not a comprehensive list

SDKs

Go
Python
Java
Ruby
CPP
.Net
Rust
…

Go
Python
Java
CPP
.Net
Rust
…

Poll Question #1

- What is your company's role in the object space?
  - A) Customer of one or more object storage platforms
  - B) Developer of one or more object storage platforms
  - C) Service provider offering object storage solutions
  - D) Researcher/Academic in the field of data storage
  - E) Other

Poll Question #2

- If you are a customer of an object storage solution, what is your primary production implementation?
  - A) Open source (ex: Ceph, Minio)
  - B) Commercial on-prem solution (ex: Dell ECS, IBM COS, Cloudian, etc)
  - C) Managed cloud service offering (Amazon S3, Azure Blob, Seagate Lyve, Wasabi, etc)
  - D) Other
  - E) None (not using any object storage solution)

Poll Question #3

- What is the total object storage capacity used by your business across all locations and implementations?
  - Below 1 TB
  - 1 TB – 1 PB
  - 1 PB – 10 PB
  - Over 10 PB
  - None

Incompatibilities? What incompatibilities?...

Some real world examples

As an example - Amazon S3

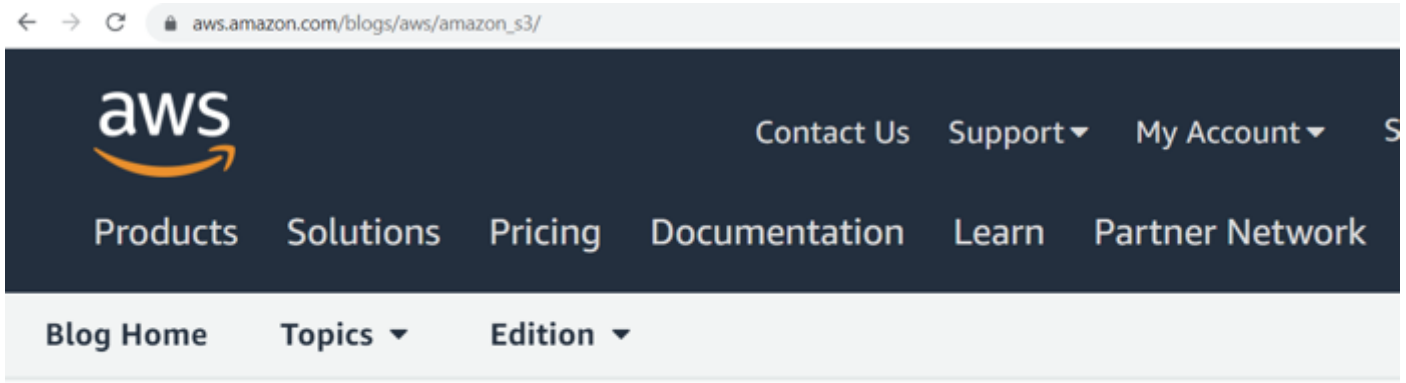

AWS News Blog

Amazon S3

by Jeff Barr | on 14 MAR 2006 | Permalink | Share

Earlier today we rolled out Amazon S3, our reliable, highly scalable, low-latency data storage service.

Using SOAP and REST interfaces, developers can easily store any number of blocks of data in S3. Each block can be up to 5 GB in length, and is associated with a user-defined key and additional key/value metadata pairs.

- >350 Trillion objects
- >100M requests per second
- Traffic peaks at >1PB/sec
- 4B checksum calculations per second
- Millions of hard drives, Exabytes of data

Amazon S3 API

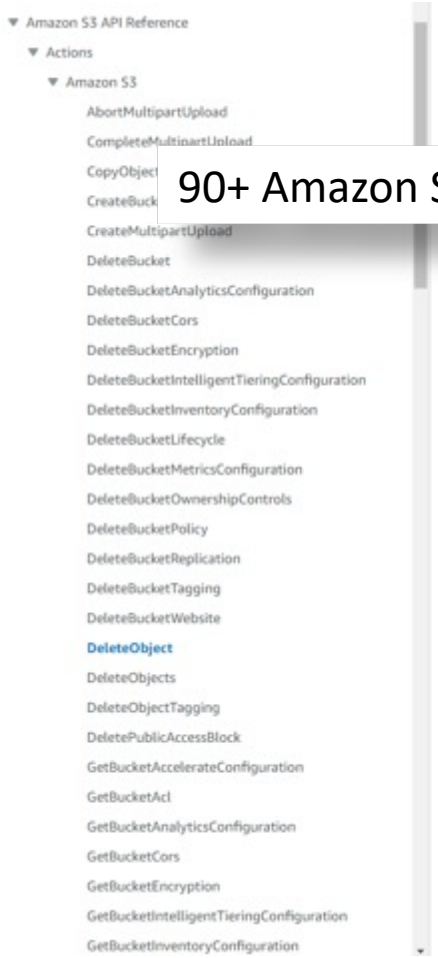

90+ Amazon S3 Actions

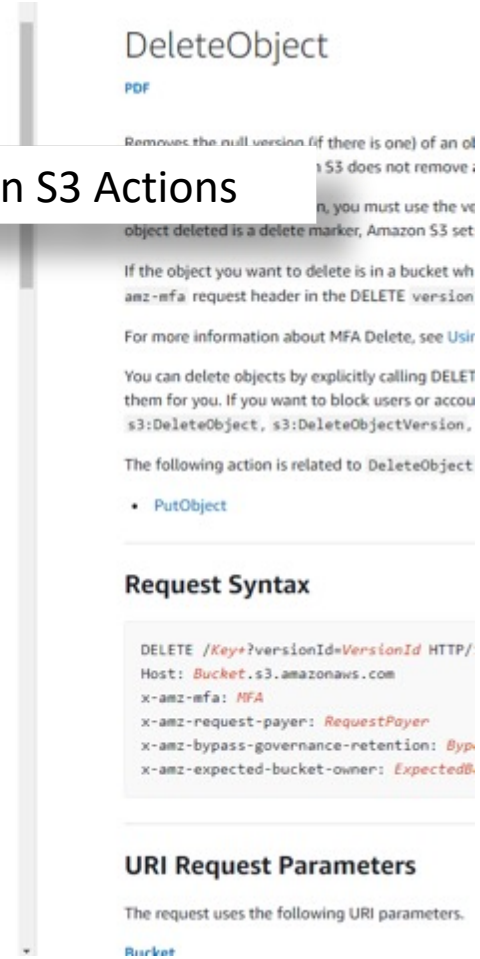

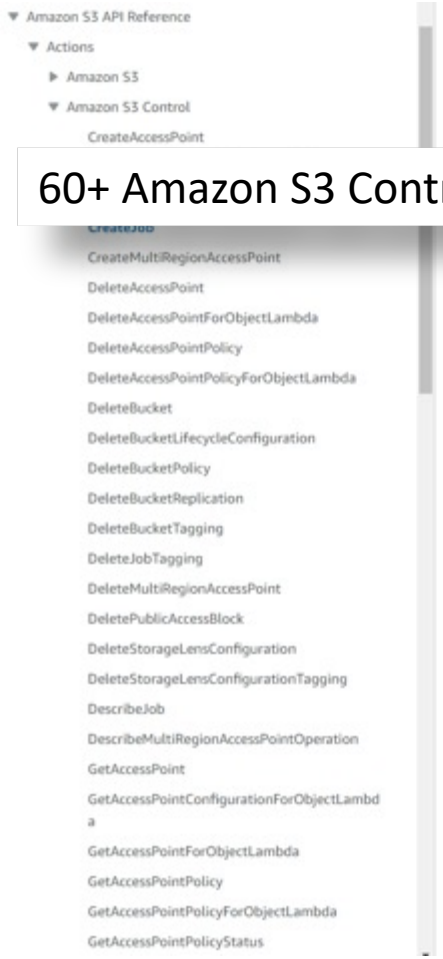

60+ Amazon S3 Control Actions

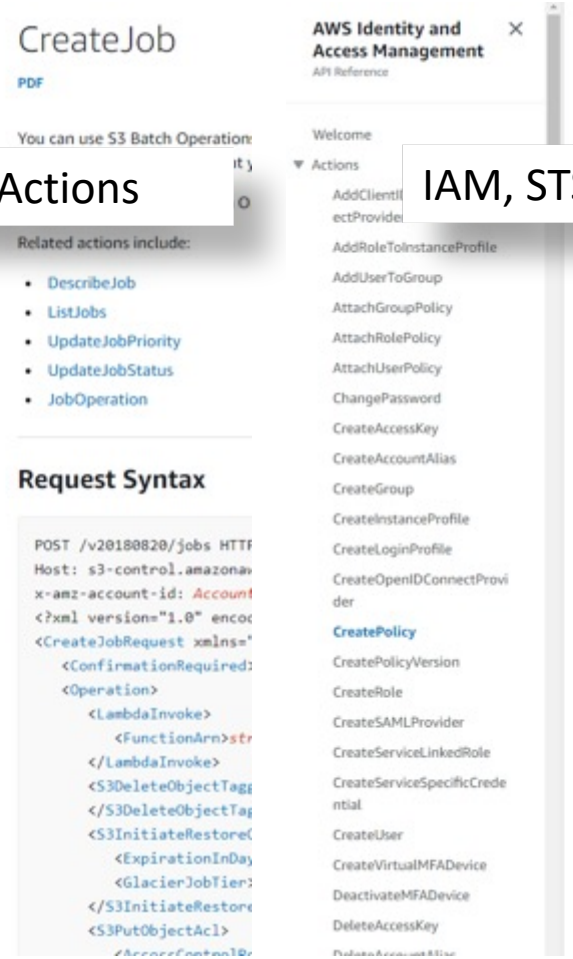

IAM, STS Actions

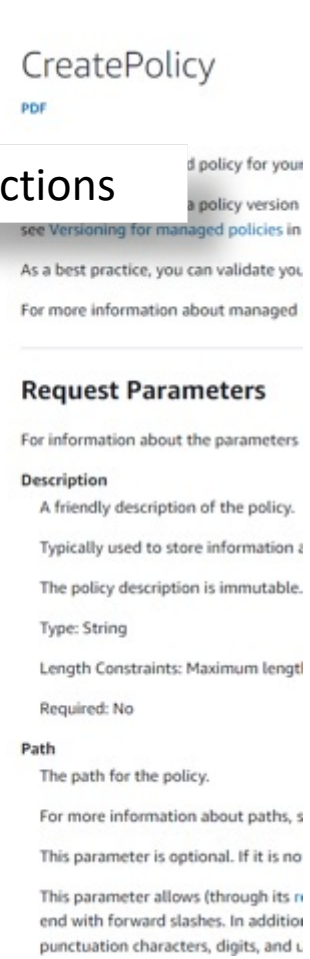

Documented Incompatibilities

ibm.com/docs/en/coss/3.10.2?topic=development-rest-api-reference

IBM Documentation | Search in IBM Cloud Object Storage System 3.10.2

IBM Cloud Object Storage System

List Parts (GET) Returns a list of parts that are associated wi

The following operations are not supported in the IBM COS implementation of the S3 API:

- DELETE Bucket analytics
- DELETE Bucket inventory
- DELETE Bucket lifecycle
- DELETE Bucket metrics
- DELETE Bucket policy
- DELETE Bucket replication
- DELETE Bucket website
- GET Bucket accelerate

Show full table of contents

Filter on titles

System Cloud Storage Object API

ECS 3.6.2 Data Access Guide | Do

dell.com/support/manuals/en-us/ecs-appliance-software-with-encr...

Unsupported S3 APIs

Table 3. Unsupported S3 APIs

FEATURE	NOTES
DELETE Bucket tagging	-
DELETE Bucket website	-
GET Bucket location	ECS is only aware of a single Virtual Data Center (VDC).

is only defined for

github.com/scality/cloudserver/blob/development/8.7/constants.js#L98

```
98      unsupportedQueries: [
99          'accelerate',
100         'analytics',
101         'inventory',
102         'logging',
103         'metrics',
104         'policyStatus',
105         'publicAccessBlock',
106         'requestPayment',
107         'torrent',
108     ],
109
110     // Headers supported by AWS that we do not currently support.
111     unsupportedHeaders: [
112         'x-amz-server-side-encryption-customer-algorithm',
113         'x-amz-server-side-encryption-context',
114         'x-amz-server-side-encryption-customer-key',
115         'x-amz-server-side-encryption-customer-key-md5',
116     ],
```

Use metadata search queries

The metadata search feature provides a rich query language that enables objects that have indexed metadata to be searched.

Table 21. API Syntax

API Syntax	Response Body
`GET /(bucket)/?query=(expression)&attributes=(fieldname,…)&sorted=(selector)&include_older_version=(true|false)&max-keys=(num_keys)&marker=(marker value)`	`<BucketQueryResult xmlns:ns2="http://s3.amazonaws.com/doc/2006-03-01/"><Name>mybucket</Name><Marker/><IsTruncated>false</IsTruncated><MaxKeys>0</MaxKeys><ObjectMatches><object><objectName>file4</objectName><objectId>09998027b1b7fbb21f50e13fabb481a237ba2f60f352d437c8da3c7c1c8d7589</objectId><versionId>0</versionId><queryMds><type>SYSMD</type><mdMap>`

NOTE: Prefix capability is added to the metadata search. See Prefix capability in metadata search.

min.io/docs/minio/linux/operations/checklists/thresholds.html

Unsupported S3 Bucket APIs

MinIO does not support the following API calls available in S3. These APIs are either redundant or only provide functionality within AWS S3.

- `BucketACL`, `ObjectACL` (use Policies)
- `BucketCORS` (CORS enabled by default on all buckets for all HTTP verbs)
- `BucketWebsite` (use `caddy` or `nginx`)
- `BucketAnalytics`, `BucketMetrics`, `BucketLogging` (use Bucket Notifications)
- `BucketRequestPayment`

Protocol Compliance Tools

- Home-grown
- 3rd party applications
- https://github.com/ceph/s3-tests
- https://github.com/splunk/s3-tests
- https://github.com/open-io/ceph-s3-tests
- https://github.com/minio/mint

Example: Running a particular test suite against different implementations

Solution	Failed Mint tests
Minio	0
Backblaze B2	14
Google Cloud Storage – S3	15
AWS S3	12

Object Key (Path)

- Up to 1,024 bytes long

- The following objects can coexist in a bucket:
  - mybucket/myfolder/obj
  - mybucket/myfolder/obj/
  - mybucket/myfolder/obj//
  - mybucket/myfolder/../obj/
  - mybucket/myfolder/

Solution 1

Solution 2

- Up to 1,024 bytes long
- "/" is interpreted as a directory
- Directory segments are limited to 255 bytes

- "//", ".", ".." are not allowed
- mybucket/myfolder and mybucket/myfolder/obj objects can't co-exist

```
$ aws --profile minio s3api put-object --bucket
bucket1 --key NameWith//Inside --body ~/empty

An error occurred (XMinioInvalidObjectName) when
calling the PutObject operation: Object name
contains unsupported characters
```


Complete-Multipart-Upload Response Caching

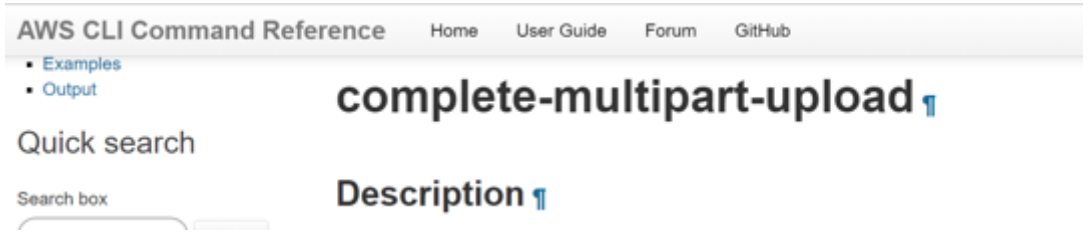

Complete Multipart Upload is an idempotent operation. After your first successful complete multipart upload, if you call the operation again within a short period, the operation will succeed.

```
$ aws s3api complete-multipart-upload --bucket A
--key B --uploadId XX --multipart-upload file
POST 200 None
$ aws s3api complete-multipart-upload --bucket A
--key B --uploadId XX --multipart-upload file
POST 200 None
```

Solution 1

```
$ aws s3api complete-multipart-upload --bucket A
--key B --uploadId XX --multipart-upload file
POST 200 325
$ aws s3api complete-multipart-upload --bucket A
--key B --uploadId XX --multipart-upload file
POST 404 449
```


CompleteMultipartUpload: Undocumented Implementation

- Amazon S3 API documentation:

Request Syntax

```
POST /Key+?uploadId=UploadId HTTP/1.1
Host: Bucket.s3.amazonaws.com
          * * *
<CompleteMultipartUpload xmlns="http://s3.amazonaws.com/doc/2006-03-01/">
   <Part>
      <ChecksumCRC32>string</ChecksumCRC32>
      <ChecksumCRC32C>string</ChecksumCRC32C>
      <ChecksumSHA1>string</ChecksumSHA1>
      <ChecksumSHA256>string</ChecksumSHA256>
      <ETag>string</ETag>
      <PartNumber>integer</PartNumber>
   </Part>
   ...
</CompleteMultipartUpload>
```

https://aws.amazon.com/sdk-for-javascript/ (prior to v 3.35.0):

```
<CompletedMultipartUpload xmlns="http://s3.amazonaws.com/doc/2006-03-
01/">
   <Part>
      <ChecksumCRC32>string</ChecksumCRC32>
      <ChecksumCRC32C>string</ChecksumCRC32C>
      <ChecksumSHA1>string</ChecksumSHA1>
      <ChecksumSHA256>string</ChecksumSHA256>
      <ETag>string</ETag>
      <PartNumber>integer</PartNumber>
   </Part>
   ...
</CompletedMultipartUpload>
```

- Worked with Amazon S3
- Failed with some other implementations
- Fixed https://github.com/aws/aws-sdk-js-v3/issues/1814

Get-object-attributes


```
$ aws s3api get-object-attributes --bucket gt-test-006 --key awscliv2.zip --object-attributes "ETag"

{
    "LastModified": "2023-04-20T01:01:27+00:00",
    "VersionId": "2m1EHmC6ALD_FVJ6HZBJ9znvBEbi6pNa",
    "ETag": "75c77163c337dfd5bb5a5f9f7a6473dd-1"
}
```

Solution 1

```
$ aws s3api get-object-attributes --bucket gt-test-006 --key awscliv2.zip --object-attributes "ETag"


Unable to parse response (not well-formed (invalid token): line 1, column 2), invalid XML received. Further retries may succeed:
b'PK\x03\x04\x14\x00\x00\x00\x00\x00F\x8f\xf4T\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x04\x00\x00\x00aws/PK\x03\x04\x14\x00\x00\x00\x00\x00E\x8f\xf4T\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00t\x00\x00\x00aws/dist/PK\x03\x04\x14\x00\x00\x00\x08\x00\x1a\x8e\xf4TW\x92\xbe\n<\x02\x00\x00\xb
```

Solution 3

Head-bucket


```
$ aws s3api head-bucket --bucket gt-test-006 --expected-bucket-owner wronguser --debug
|& grep HEAD
urllib3.connectionpool - DEBUG - https://gt-test-006.s3.us-east-1.amazonaws.com:443
"HEAD / HTTP/1.1" 400 0
```

Solution 1

Solution 4

```
$ aws s3api head-bucket --bucket gt-test-006 --expected-bucket-owner wronguser --debug
|& grep HEAD

urllib3.connectionpool - DEBUG - http://10.0.0.83:9000 "HEAD /gt-test-006 HTTP/1.1" 200
0
```


Summary

Complexity + nuance

Incompatibilities

Deep understanding of the API

Thorough testing

Compatibility = customer adoption

Poll Question #4

- Have you encountered compatibility issues between various object storage implementations?
  - Yes
  - No

Cloud Object Storage Plugfest

Arnold Jones

SNIA Cloud Object Storage Plugfest

The purpose of the Plugfest is for vendors to bring their Object Storge implementations to test, identify, and fix bugs in a collaborative setting with the goal of providing a forum in which companies can develop interoperable products.

Plugfest Operational Overview

- Similar format to the SMB3 Interoperability Lab held at the Storage Developer Conference (SDC) every year.
  - https://www.snia.org/SMB3IOlab
- SNIA provides and supports the network and infrastructure for the Plugfest, creating a collaborative framework for testing.
- The participants of the Plugfest work together to define the testing process, assuring that objectives are accomplished.
  - Planning calls held prior to the event to work out the details
- There is a fee to participate which includes a full SDC conference pass for each participating engineer
- All Plugfest participants are under an NDA to protect disclosure of issues found

Plugfest Plan

- Participation from Protocol Support experts onsite
  - Working on getting an Object Storage Protocol Support experts from Amazon, Google, and Microsoft onsite during the Plugfest
- Availability of protocol test suites
  - Working on getting companies with protocol test suites and equipment involved
  - Participants are encouraged to bring and share test software

Poll Question #5

- Would you be interested to learn more about the Plugfest and/or participate in it?
  - Yes
  - No

Call for Participation

Driving Compatibility and Collaboration in Cloud Object Storage

- Foster Ecosystem Collaboration
  - Facilitate collaboration and knowledge sharing among Object Storage developers by establishing a platform for discussions, forums, and workshops
- Host an Object Storage Interoperability Plugfest
  - Facilitated by SNIA and co-located at Storage Developer Conference (SDC) 2024
  - September 16-18, 2024, Santa Clara, CA
- Interested in more details?
  - Email us @ cloudtwgchair@snia.org
  - Fill out Plugfest survey at https://www.surveymonkey.com/r/objectplugfest

Thanks for Viewing this Webinar

- Please rate this presentation and provide us with feedback
- This webinar and a copy of the slides are available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webinar will be posted to the SNIA Cloud blog: www.sniacloud.com/
- Follow us @SNIACloud

Thank You!

