Real-World Challenges of S3 Compatibility

Ben McClelland
Sr Principal Software Engineer

What is S3

PutObject: Uploads an object to a bucket.

GetObject: Retrieves an object from a bucket.

DeleteObject: Deletes an object from a bucket.

HeadObject: Retrieves metadata about an object without returning the object itself.

ListObjects: Lists some or all of the objects in a bucket.

What is S3 - actually

AbortMultipartUpload
CompleteMultipartUpload
CopyObject
CreateBucket
CreateBucketMetadataConfiguration
CreateBucketMetadataTableConfiguration
CreateMultipartUpload
CreateSession
DeleteBucket
DeleteBucketAnalyticsConfiguration
DeleteBucketCors
DeleteBucketEncryption
DeleteBucketIntelligentTieringConfiguration
DeleteBucketInventoryConfiguration
DeleteBucketLifecycle
DeleteBucketMetadataConfiguration
DeleteBucketMetadataTableConfiguration
DeleteBucketMetricsConfiguration
DeleteBucketOwnershipControls
DeleteBucketPolicy
GetBucketLogging
GetBucketMetadataConfiguration
GetBucketMetadataTableConfiguration
GetBucketMetricsConfiguration
GetBucketNotification
GetBucketNotificationConfiguration
GetBucketOwnershipControls
GetBucketPolicy
GetBucketPolicyStatus
GetBucketReplication
GetBucketRequestPayment
GetBucketTagging
GetBucketVersioning
GetBucketWebsite
GetObject
GetObjectAcl
GetObjectAttributes
GetObjectLegalHold
GetObjectLockConfiguration
GetObjectRetention
GetObjectTagging
GetObjectTorrent
GetPublicAccessBlock
HeadBucket
HeadObject
ListBucketAnalyticsConfigurations
ListBucketIntelligentTieringConfigurations
ListBucketInventoryConfigurations
ListBucketMetricsConfigurations
ListBuckets
ListDirectoryBuckets
ListMultipartUploads
ListObjects
ListObjectsV2
ListObjectVersions
ListParts
PutBucketAccelerateConfiguration
PutBucketAcl
PutBucketAnalyticsConfiguration
PutBucketCors
PutBucketEncryption
PutBucketIntelligentTieringConfiguration
PutBucketInventoryConfiguration
PutBucketLifecycle
PutBucketLifecycleConfiguration
PutBucketLogging
PutBucketMetricsConfiguration
PutBucketNotification
PutBucketNotificationConfiguration
PutBucketOwnershipControls
PutBucketPolicy
PutBucketReplication
PutBucketRequestPayment
PutBucketTagging
PutBucketVersioning
PutBucketWebsite
PutObject
PutObjectAcl
PutObjectLegalHold
PutObjectLockConfiguration
PutObjectRetention
PutObjectTagging
PutPublicAccessBlock
RenameObject
RestoreObject
SelectObjectContent
UpdateBucketMetadataInventoryTableConfiguration
UpdateBucketMetadataJournalTableConfiguration
UploadPart
UploadPartCopy
WriteGetObjectResponse

SDK Compatibility

Service

S3 Standard – General purpose, low latency, high durability storage for frequently accessed data.

S3 Standard-IA (Infrequent Access) – Cheaper than Standard, meant for infrequently accessed data, with retrieval charges.

S3 One Zone-IA – Same as Standard-IA but stored in a single AZ (lower cost, less resilience).

S3 Glacier Instant Retrieval – Archive storage with milliseconds retrieval, for data that is rarely accessed but must be retrieved quickly.

S3 Glacier Flexible Retrieval – Low-cost archive with minutes–hours retrieval (formerly “Glacier”).

S3 Glacier Deep Archive – Lowest-cost storage, retrieval in hours, for long-term compliance or rarely accessed data.

S3 Express One Zone – High-performance, single-AZ storage with microsecond latency and very high throughput.

S3 on Outposts – Brings S3 storage APIs and features to AWS Outposts (on-premises racks).

Client

AWS SDK for **Java**

AWS SDK for **JavaScript**

AWS SDK for **Python** (Boto3)

AWS SDK for **PHP**

AWS SDK for **.NET** (C# etc.)

AWS SDK for **Ruby**

AWS SDK for **Go**

AWS SDK for **C++**

AWS SDK for **Kotlin**

AWS SDK for **Swift**

AWS SDK for **Rust**

AWS Tools for **PowerShell**

AWS CLI

SDK Compatibility

Even AWS has trouble with cross product compatibility

How is Amazon S3 on Outposts different from Amazon S3?

S3 on Outposts specifications

- The maximum Outposts bucket size is 50 TB.
- The maximum number of Outposts buckets is 100 per AWS account.
- Outposts buckets can be accessed only by using access points and endpoints.
- The maximum number of access points per Outposts bucket is 10.
- Access point policies are limited to 20 KB in size.
- The Outpost owner can manage access within your organization in AWS Organizations by using AWS Resource Access Manager. All accounts that need access to the Outpost must be within the same organization as the owner account in AWS Organizations.
- The S3 on Outposts bucket owner account is always the owner of all objects in the bucket.
- Only the S3 on Outposts bucket owner account can perform operations on the bucket.
- Object size limitations are consistent with Amazon S3.
- All objects stored on S3 on Outposts are stored in the OUTPOSTS storage class.
- By default, all objects stored in the OUTPOSTS storage class are stored by using server-side encryption with Amazon S3 managed encryption keys (SSE-S3). You can also explicitly choose to store objects by using server-side encryption with customer-provided encryption keys (SSE-C).
- If there is not enough space to store an object on your Outpost, the API returns an insufficient capacity exception (ICE).

S3's Split Personality

- **S3 for Storage Vendors:** Often seen as a simple protocol for put/get access. A basic checklist of features. This is a low bar for “compatibility."

- **S3 for Applications:** The true test. Applications may rely on a vast ecosystem of client-side libraries, features, and specific behaviors. This is where the real incompatibilities possibly arise.

Is this S3 Compatible?

Maybe? Depends on client assumptions/needs?

SDK Compatibility

The Moving Target: AWS constantly evolves the S3 API with new features. 3rd party S3 storage vendors are perpetually playing catch-up.

The Grey Area: The public S3 specification is incomplete. We often have to reverse-engineer behaviors from the official AWS SDKs/service to ensure compatibility.

De-facto Standard: Only real protocol spec is how AWS service currently behaves.

SDK Compatibility

Chasing AWS S3 API is a constantly moving target

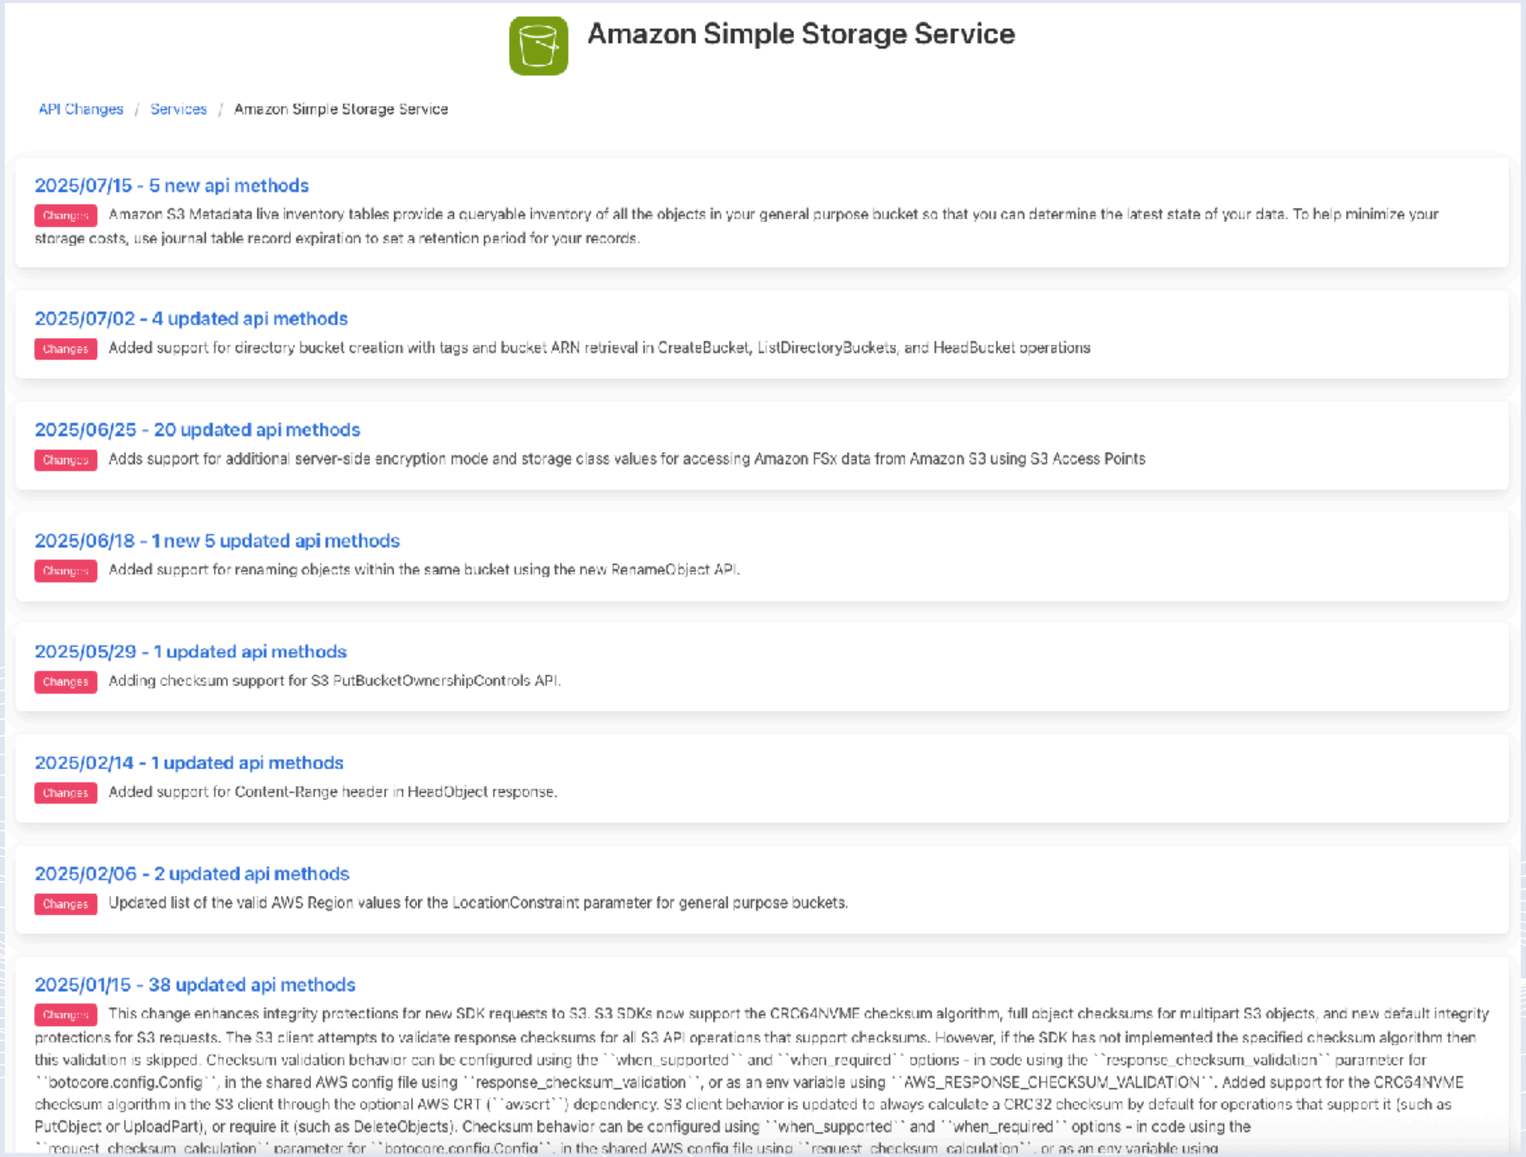

Amazon Simple Storage Service

API Changes / Services / Amazon Simple Storage Service

2025/07/15 - 5 new api methods

Changes: Amazon S3 Metadata live inventory tables provide a queryable inventory of all the objects in your general purpose bucket so that you can determine the latest state of your data. To help minimize your storage costs, use journal table record expiration to set a retention period for your records.

2025/07/02 - 4 updated api methods

Changes: Added support for directory bucket creation with tags and bucket ARN retrieval in CreateBucket, ListDirectoryBuckets, and HeadBucket operations

2025/06/25 - 20 updated api methods

Changes: Adds support for additional server-side encryption mode and storage class values for accessing Amazon FSx data from Amazon S3 using S3 Access Points

2025/06/18 - 1 new 5 updated api methods

Changes: Added support for renaming objects within the same bucket using the new RenameObject API.

2025/05/29 - 1 updated api methods

Changes: Adding checksum support for S3 PutBucketOwnershipControls API.

2025/02/14 - 1 updated api methods

Changes: Added support for Content-Range header in HeadObject response.

2025/02/06 - 2 updated api methods

Changes: Updated list of the valid AWS Region values for the LocationConstraint parameter for general purpose buckets.

2025/01/15 - 38 updated api methods

Changes: This change enhances integrity protections for new SDK requests to S3. S3 SDKs now support the CRC64NVME checksum algorithm, full object checksums for multipart S3 objects, and new default integrity protections for S3 requests. The S3 client attempts to validate response checksums for all S3 API operations that support checksums. However, if the SDK has not implemented the specified checksum algorithm then this validation is skipped. Checksum validation behavior can be configured using the ``when_supported`` and ``when_required`` options - in code using the ``response_checksum_validation`` parameter for ``botocore.config.Config``, in the shared AWS config file using ``response_checksum_validation``, or as an env variable using ``AWS_RESPONSE_CHECKSUM_VALIDATION``. Added support for the CRC64NVME checksum algorithm in the S3 client through the optional AWS CRT (``awscrt``) dependency. S3 client behavior is updated to always calculate a CRC32 checksum by default for operations that support it (such as PutObject or UploadPart), or require it (such as DeleteObjects). Checksum behavior can be configured using ``when_supported`` and ``when_required`` options - in code using the ``request_checksum_calculation`` parameter for ``botocore.config.Config``, in the shared AWS config file using ``request_checksum_calculation``, or as an env variable using

Versity

SDK Compatibility

AWS S3 server has to support older clients
AWS client SDK only needs to support current AWS servers

SDK Compatibility

Describe the bug

I read the announcement about the integrity check behaviour change, and tried to disable the integrity check by setting `request_checksum_calculation` and `response_checksum_validation` to `when_required`.

However, it seems that the integrity check cannot be disabled for multipart upload by `Uploader.Upload` function in `feature/s3/manager` package. I checked the HTTP header of the SDK's requests using `tcpdump` command. The results were as follows:

adev-code added **p3** **queued** and removed **investigating** on Feb 17

This is a minor priority issue

benmcclelland on Feb 18

Is there any possibility in raising the priority to something more than "minor" since this is currently breaking all 3rd party S3 servers that don't (yet) support the integrity checks? This is forcing clients to pin the following versions to be able to support non-AWS S3 endpoints:

```
require (
    github.com/aws/aws-sdk-go-v2/feature/s3/manager v1.17.49
    github.com/aws/aws-sdk-go-v2/service/s3 v1.72.3
)
```

6

[redacted] on Feb 21 Contributor

Compatibility with arbitrary 3rd-party S3 clones is not going to be a driving factor for prioritization.

1 9

Versity

Versity Gateway

S3 → POSIX
Enables S3 workloads to utilize POSIX filesystems and tape via ScoutAM

Modular
POSIX and ScoutAM backends supported, easy to add more

Scalable
Load balance across multiple instances

High Performance
Written from scratch, Go / GoFiber
Fast processing and response times

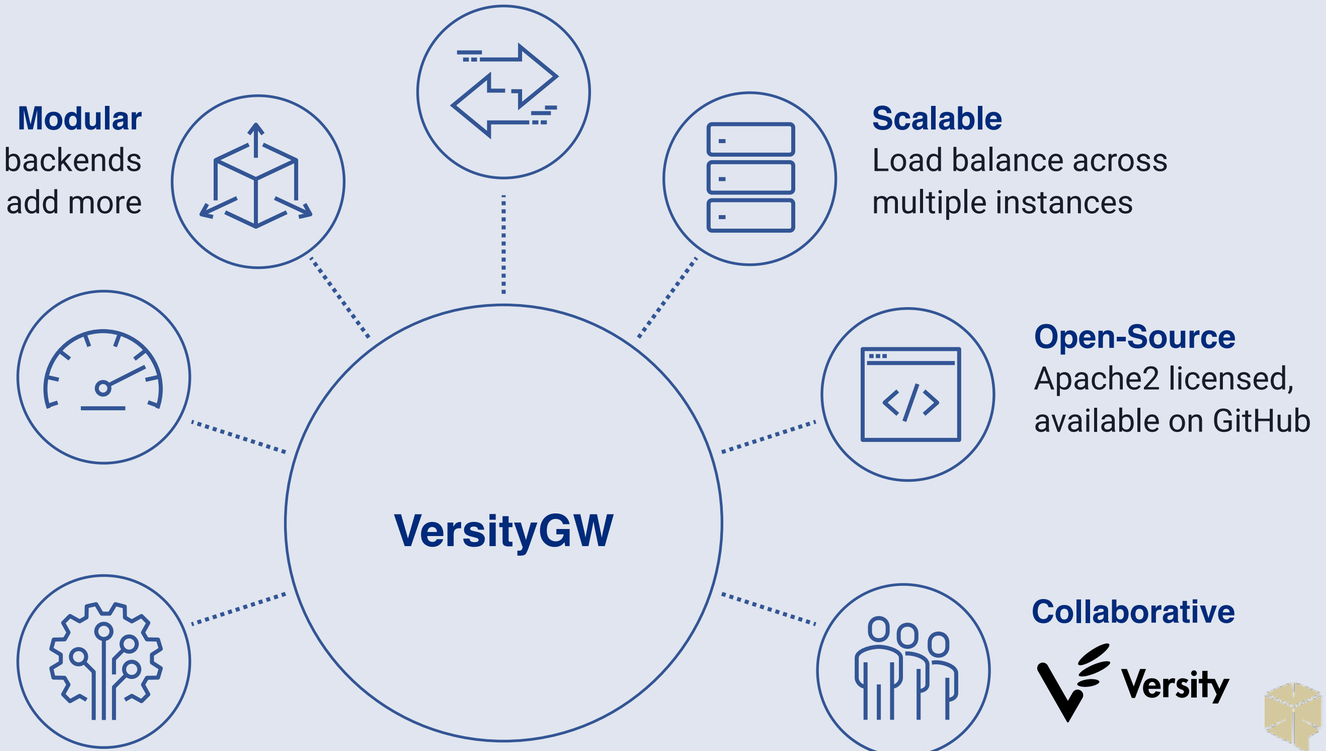

Open-Source
Apache2 licensed, available on GitHub

Seamless Integration
Familiar interfaces

Collaborative

Versity

pawsey

Los Alamos NATIONAL LABORATORY

https://github.com/versity/versitygw

Versity Gateway

April 2023 - Started in collaboration with LANL
June 2024 - First production release

Versity Gateway Design Requirements

- Scalable
- High Performance
- Stateless
- Modular Backend Support
- Flexible Open Source License

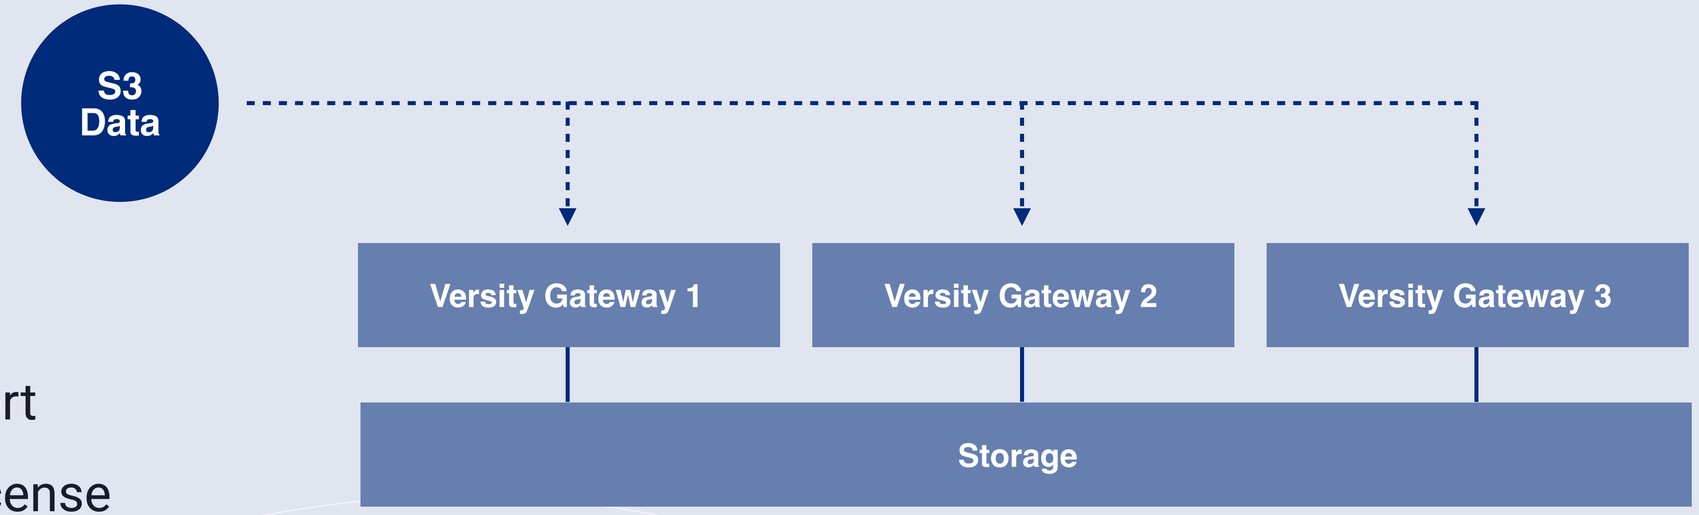

Versity

Versity Gateway

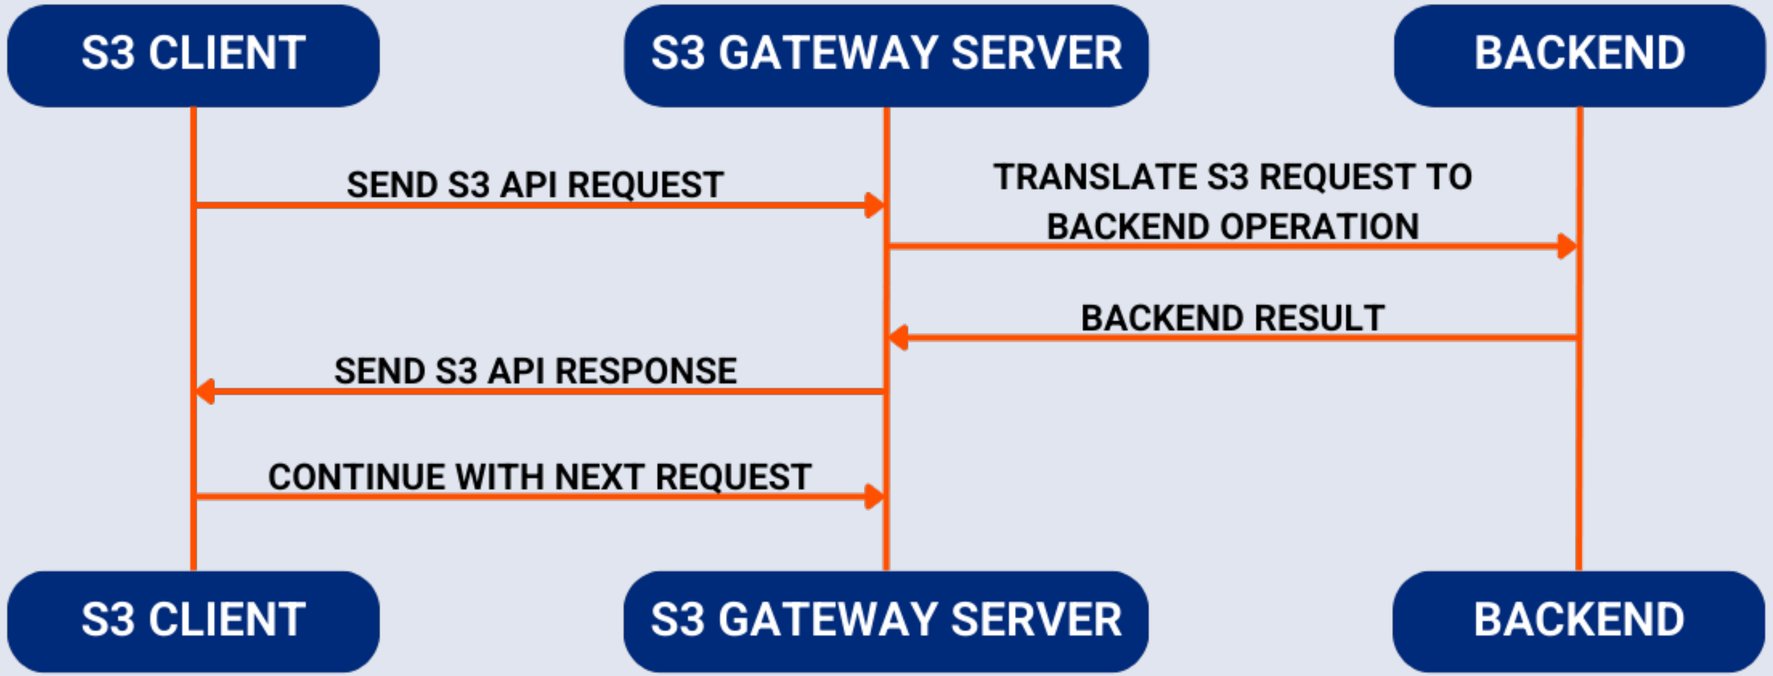

Versity Gateway

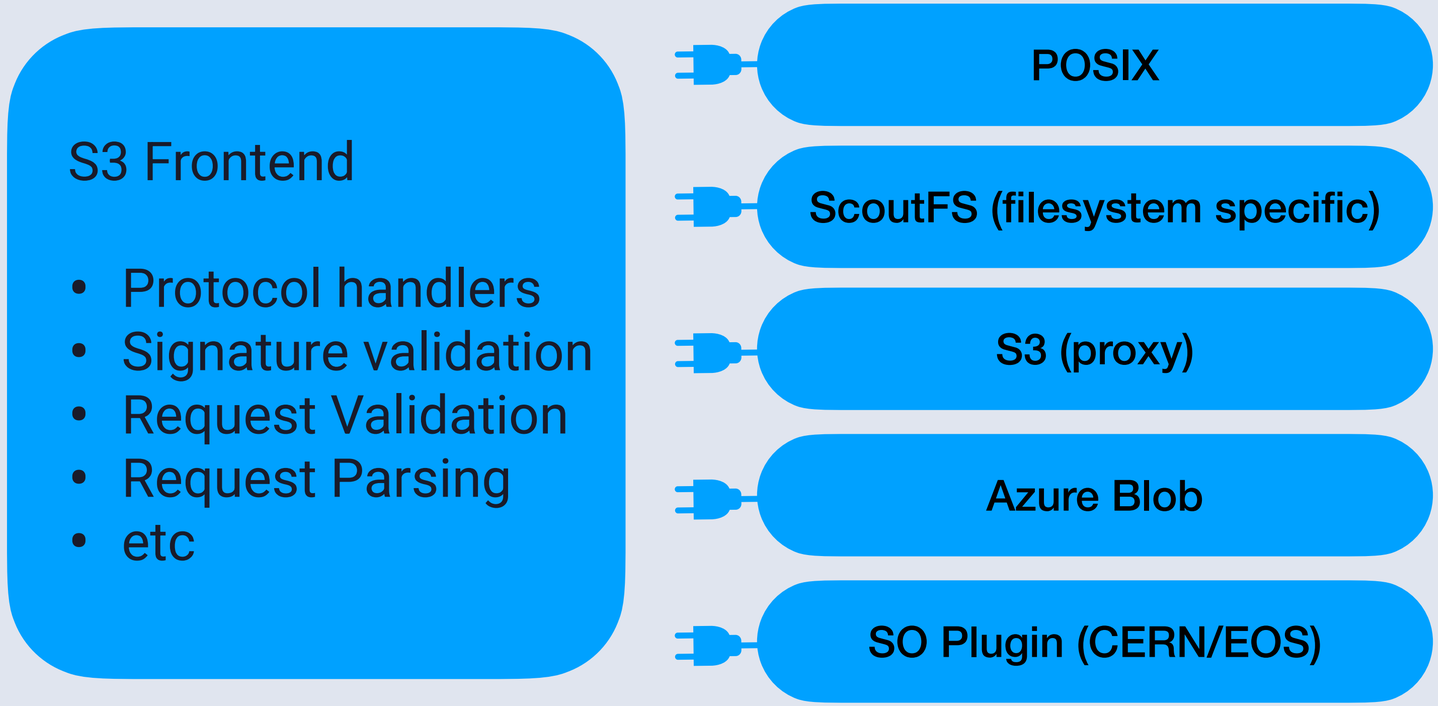

Versity Gateway - POSIX

S3 Frontend

- Protocol handlers
- Signature validation
- Request Validation

POSIX - split object names on “/“, translate to file

Example:
gateway root: /mnt/fs/gateway
bucket: mybucket
object: 2023/Jan/myobject

/mnt/fs/gateway/mybucket/2023/Jan/myobject

Versity Gateway - POSIX

Support for multipart uploads

Stores temporary part uploads in hidden directory

Hands off final file construction to filesystem with copy_file_range()

Versity Gateway - POSIX

Not strictly S3 compatible though

```
// Non-AWS errors
ErrExistingObjectIsDirectory
ErrObjectParentIsFile
ErrDirectoryObjectContainsData
ErrDirectoryNotEmpty
```

“good enough” for majority of workloads

Versity Gateway

S3 Frontend

- Protocol handlers
- Signature validation
- Request Validation

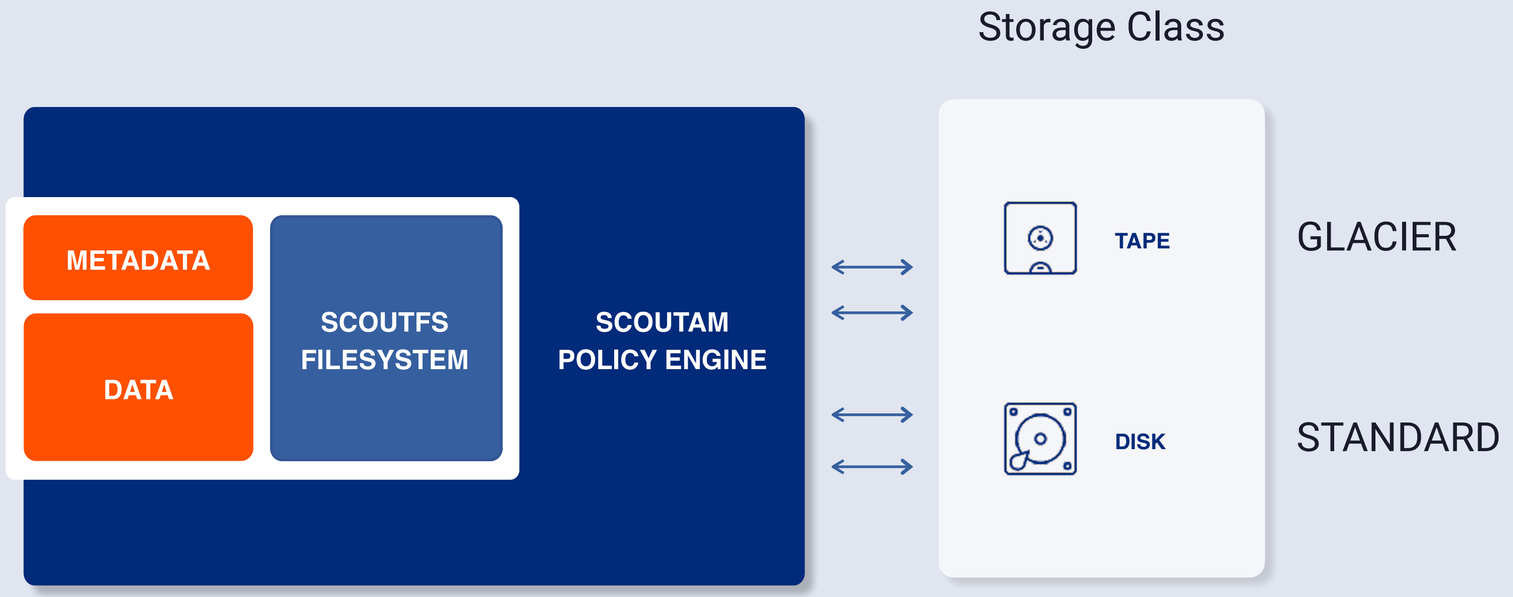

Versity Gateway

Client Glacier workflows expect to be able to manage storage class within the same storage system endpoints

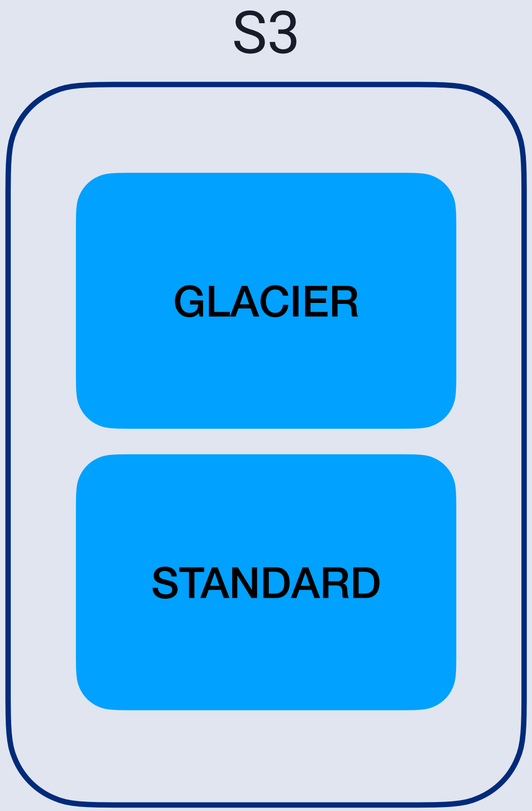

Test Strategy

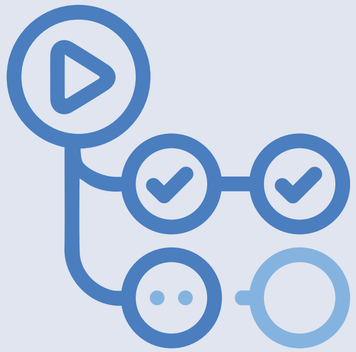

Automated tests pipeline for all pull requests
All tests are in GitHub repo

- Regression Tests - targeted for specific bug fixes
- System Tests - targeted at real clients
  - aws cli
  - s3cmd
  - mc
- System Tests REST/curl - targeted at protocol details

Cloud Object Storage Plugfest

https://www.snia.org/forums/csti/plugfests

Cloud Object Storage Plugfest

Versity

Thank You

sales@versity.com

@versitysoftware

Versity