

Patterns in S3 Data Access

Protecting and enhancing access to data banks, lakes, and bases

Josh Snyder @ fwd:CloudSec, 2023

Synopsis


1. Preamble: AWS request signing
2. Problem: Complex-pattern data in S3 (e.g. data lakes)
3. Common solutions
4. Just-in-time access

Go watch this: https://www.youtube.com/watch?v=BOz2_hgoob4

STG328

Solving large-scale data access challenges with Amazon S3


Becky Weiss
Senior Principal Engineer
Amazon Web Services

 © 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Disclaimer!

Do not go implement anything from this talk without first watching that talk!



Preamble: AWS request signing

Motivation: Josh's Photo Sharing Service

- Users (millions) upload their photos (billions)
- Photos are 1MB up to 1GB
- Photos are private by default
 - original uploader can always see their photo
- Photos can be added to albums (1->many relationship)
- Albums can be shared with other users (1->many relationship)
- Access to an album can be revoked at any time
- No photo needs to be super-popular (no CDN)

Signed URLs

https://artifacts-911be34e61abfc8d.s3.us-east-1.amazonaws.com/serverless/certs/dev/1608193965988-2020-12-17T08%3A32%3A45.988Z/certs.zip?response-content-disposition=attachment&X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20230607T083045Z&X-Amz-SignedHeaders=host&X-Amz-Expires=300&X-Amz-Credential=ASIA3LROMZGCSJI7XRXE%2F20230607%2Fus-east-1%2Fs3%2Faws4_request&X-Amz-Signature=57121f2315ae365434e3b70b54e2e72d3b4a9bf93fa0e864c3a960da4457d9c4&X-Amz-Security-Token=IQoJb3JpZ2luX2VjEjEaCjVzLXdlc3QtMSJHMEUCIAQWLF5nbq4T4PIfFoCc26X1aDgnLw%2BtR9flpwL9VFc5AiEAYFwpzVUB7%2BftMY%2BXVqzy3LZ2k%2FOMmaUs7H9bgZW%2FdSwqygMIWhACGgw3ODA3MDc2MTMwNjEiDJMAzsoSh3PhmuvJrCqnA0QsG0upWxiYAGjfflqVXEL14cW8sIX8qRFoeUUyVUUpAV2RoZACSI%2BJ8MgbYYmDpTDVe0Uc3FxFrFloyEA3pDuzXIVFnN9x9tdw5p2QQms1MCKkdQpY6QoNvnD8z3vh4oK7IT0WFMcto2Hw6LCJllruL2FfYgnGvN5bFBZzcgbTpxZCUI0bExKdZcXZLpkE7Lu%2Bq4IAooox1luqNVxdw00IGgeODpauSznSXLadQrlv4CC1Gw%2BMYuRMtfsB4P%2F0Rz5wpysfLSUcRsKRfwxhh8jv9%2FtD4ECHtlrADGhvMVdmcYBtgFAhG9YWjAy78MnfjmyoanHmdsMQKRXpsqQI%2Fzh%2BH%2FJX%2FnirVyGbmMHczfvemRV8WA3OP2gc98qhLkgsPGay%2FowpWdnXctHHL3ulKDh3bq%2FX6hn6Vwcl8lShztu67J1375Tu%2B4BdcBAM2yn8CcmKwyYlmLgFKNMVp%2BJk30V4zicbH9Ug9RABvQNb52aLNFsp%2FIQ11OgDgvSYDqR9Y3cKW%2F3HZLxgxtsyerqZbJG87dp8X1MzBKFxs6PKZGV91FI1%2FITDy%2F4CkBJqUAm8M60Cyphslo2CF7i3sDKE%2FwBiEPKqKf%2Frrj8mDTwnCAAp3eS0OjD2ph8Q3NDto0QiAeCh8DIYXgsNpjACeFANnqHCgldB2%2BunKXps9%2FZg7mdaoVyEfUn2X874QtZ%2B62kus9%2Ff8%2BkroF92lip%2B77%2BANKJ6DE%2BiETcf7F21D%2FRosmDfMfZLBeclJanxzdN5jllk64e6kkJA0O77LC1hpzVotmKACNAfwJJCfNUsNPkYjMeKzF

Signed URLs expire (good!)

```
-<Error>  
  <Code>AccessDenied</Code>  
  <Message>Request has expired</Message>  
  <X-Amz-Expires>300</X-Amz-Expires>  
  <Expires>2023-06-07T08:35:45Z</Expires>  
  <ServerTime>2023-06-07T08:54:53Z</ServerTime>  
  <RequestId>WSAXPXJKFTQFYDWP</RequestId>  
-<HostId>  
  +EGS5rRm7enuhyfw0/Osnq2t+JZslwVAt7fpaW4pwjpH6WepthJR2UJVv8dySiOEVuc4nhKNDwE=  
  </HostId>  
</Error>
```


What are signed URLs?

- Portable **capabilities** (antique term, from the ~1970s and 80s)
- Cryptographically authenticated!
- Useful for **browsers**
- Not useful for API-based clients

☰ Capability-based security 🌐 6 languages ▾

Article [Talk](#)

[Tools](#) ▾

From Wikipedia, the free encyclopedia


Capability-based security is a concept in the design of [secure computing](#) systems, one of the existing [security models](#). A **capability** (known in some systems as a **key**) is a communicable, unforgeable [token](#) of authority.

Portable Capabilities!?!

- the whole point of capabilities is to flow through a system
- this freaks security engineers out (pictured at right)
- docs used to say this (below); it was removed sometime late September–early October 2022

Important

If you make a request in which all parameters are included in the query string, the resulting URL represents an AWS action that is already authenticated. Therefore, treat the resulting URL with as much caution as you would treat your actual credentials. We recommend you specify a short expiration time for the request with the `X-Amz-Expires` parameter.



Making Signed URLs safer

- Options:
 - IP-binding
 - source-VPC binding
 - cookie-binding



<https://github.com/hashbrowncipher/safer-signed-urls>

Every request in AWS is a portable capability!



Cryptography of AWS Request Signing (dangerously oversimplified)

1. Authentication

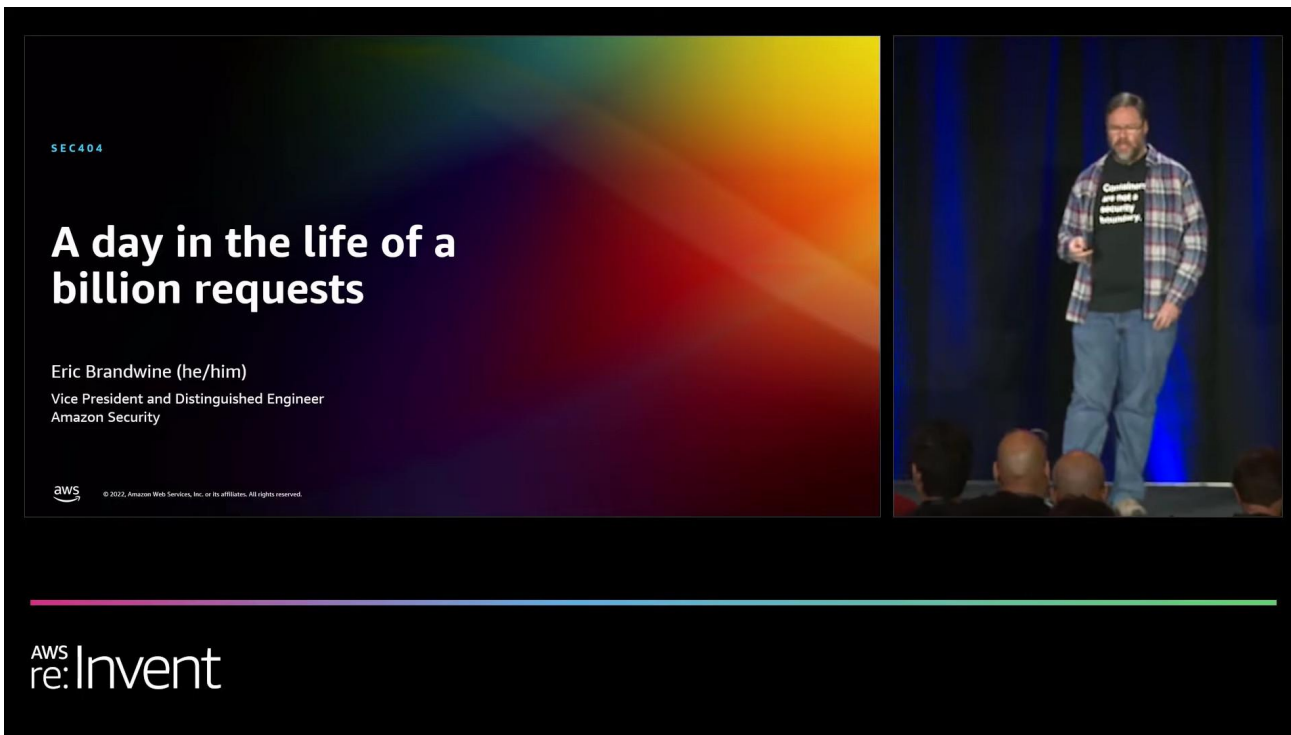
- a. I send: `HASH(AWS_SECRET_ACCESS_KEY || My Request) || My Request`
- b. Amazon computes: `HASH(AWS_SECRET_ACCESS_KEY || My Request)`
- c. If not equal: send HTTP 403 Forbidden

2. Authorization

3. Do the thing

4. Send the result

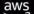
For more on SigV4: <https://www.youtube.com/watch?v=tPr1AgGkvc4>



SEC404

A day in the life of a billion requests

Eric Brandwine (he/him)
Vice President and Distinguished Engineer
Amazon Security

 © 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS
re:Invent

The slide features a vibrant, multi-colored background transitioning from purple to yellow. On the right side, there is a video thumbnail showing a man on stage wearing a plaid shirt over a black t-shirt with the text 'Contributors are not a security boundary.' The audience's heads are visible in the foreground of the video.



AWS Identity and Access Management

User Guide



- ▶ What is IAM?
- ▶ Getting set up
 - Getting started
- ▶ Tutorials
- ▶ Identities
- ▶ Access management
- ▶ Code examples
- ▶ Security
- ▶ IAM Access Analyzer
- ▶ Troubleshooting IAM
- ▶ Reference
 - Amazon Resource Names

Create a signed AWS API request

[PDF](#) | [RSS](#)

The following is an overview of the process to create a signed request. For more information, see the [code examples in the AWS SDKs](#).

Contents

- [Step 1: Create a canonical request](#)
- [Step 2: Create a hash of the canonical request](#)
- [Step 3: Create a string to sign](#)
- [Step 4: Calculate the signature](#)
- [Step 5: Add the signature to the request](#)
- [Temporary security credentials](#)
- [Code examples in the AWS SDKs](#)



Hopefully the non-AWS users haven't left yet

This talk *should be** compatible with any service that uses SigV4 authentication, including:

- Google Cloud Storage
- Cloudflare R2
- Backblaze B2
- <your favorite here>

Problem: complex-pattern data in S3

Motivation: Josh's Photo Sharing Service

- Users (**millions**) upload their photos (billions)
- Photos are 1MB up to 1GB
- Photos are private by default
 - original uploader can always see their photo
- Photos can be added to albums (**1->many relationship**)
- Albums can be shared with other users (**1->many relationship**)
- Access to an album can be granted and revoked **at any time**
- No photo needs to be super-popular (no CDN)

What *isn't* complex-pattern data?

- Lots of data ("billions of photos")
- Prefix-based layout by uploader ID
 - Very simple if no sharing is needed
 - `/<user 1>/<photo 1>.jpg`

What is complex-pattern data? (in other words)

- There are too many of them! (users, relationships)
- IAM doesn't know about them!
- Policy changes too fast!

Pivoting to data lakes

	Photo Sharing	Data Lake
Who?	End-users	Jobs (seconds -> days)
What code?	Browsers	Arbitrary code
Authentication	username/password	SSO portal
Cookie	HTTP cookie	X.509 certificate (or) JWT
Access method	Browser (HTTP GET)	AWS SDK
Authentication to S3	Query string	Authorization header

Similarities

- The callers run code we do not control
 - We cannot grant them unfettered access to the bucket
- We cannot reshape the data to fit the access problem (e.g. into nice prefixes)
 - and even if we could, shifting PBs of data doesn't sound fun
- We have a database of grants, relating callers to objects
 - the database can change rapidly
- There aren't enough IAM roles to model our callers

**How do I secure my data lake?
(and some slightly easier problems, too)**

Common solutions
(consider these first)

Overview






- Plain IAM (Weiss @ 5:12)
 - Tagged IAM principals (Weiss @ 19:42)
 - **Permissions Boundaries**
- Access Points (Weiss @ 25:02)
- STS Session Broker (Weiss @ 40:28)


Overview

See Weiss @ 3:35



Data access concepts in this session

	IAM session broker pattern 	 Structured data: Lake Formation
	S3 access points 	
	IAM techniques using S3 prefixes 	
	IAM basics	

 © 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

"Plain" IAM + S3 Bucket Policy

- IAM is a slow-moving control plane service
- very low rate limit (< 1 mutation / second)
- about 1000 roles, unless you go by account (Weiss @ 19:20)
 - inline policy: max 10 kB
 - attached policy: 6 kB (* 20 policies)
- S3 Bucket Policy (20 kB limit)
 - fits about 30 prefixes (Weiss @ 15:25)

Permissions Boundaries

- Gives an IAM role a way to create new IAM roles with fewer permissions
- Good for a group of services under common ownership
 - example: every Elasticsearch cluster owned by the ES team
- Useful for delegation
 - you delegate a set of permissions to the ES team
 - the ES team creates and manages policies under that boundary, without review by security team
 - the roles can be used directly as IAM instance profiles

Create an S3 Access Point per caller

Weiss @ 34:55

- Multiply out your bucket policy
 - by 10,000 per account-region
- Also useful for delegation
- Doesn't fit dynamic authorization



Do S3 access points fit your use case?

Best-fit use cases:

- Large number of static access patterns (up to 10,000 by default)
- Delegation-with-guardrails: Separate policies of access points and S3 buckets

Considerations:

- Discovery mechanism: How to find the right access point?
- Rate of change required for access patterns

... What to do if they're dynamic?

AWS
© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Let's build a proxy!

Weiss @ 36:05

- You can do anything!
 - Build your own authorization layer from scratch
- Spoiler alert: scalability
- S3 can do terabits / second
- Don't do it!



STS Session Broker

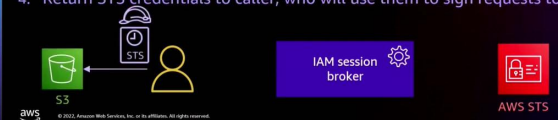
Weiss @ 40:28

- `sts:AssumeRole` creates temporary credentials.
 - can pass `Policy` and `PolicyArn` to downscope the granted role.
- STS is a data plane service



How to implement an IAM session broker

1. Authenticate caller
2. Make access decision: Should the caller be able to access the S3 location requested (e.g., `s3://example-bucket/folder/`)?
3. If yes, call STS AssumeRole API to create a session from a central IAM role with a superset access policy. Specify a session policy that grants access to the requested location.
4. Return STS credentials to caller, who will use them to sign requests to S3



STS Session Broker (drawbacks)

- Limit of 2048 characters in session policy
 - `AWS_SESSION_TOKEN`
- Account-wide ratelimit
 - Availability concern

Just-in-time authorization with signing

A basic signer

github.com/hashbrowncipher/fwdcloudsec-signers



```
$ curl -s -u josh:password \
-d '{
  "url":"https://s3.amazonaws.com/permanent/mykey",
  "method": "GET",
  "headers": {}
}' \
http://127.0.0.1:8000 | jq "."
{
  "url": "https://permanent-quoic7ui7jhvtjt6.s3.us-west-2.amazonaws.com/josh/mykey",
  "headers": {
    "x-amz-content-sha256": "e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855",
    "x-amz-expected-bucket-owner": "111111111111",
    "X-Amz-Date": "20230610T124405Z",
    "Authorization": "AWS4-HMAC-SHA256
Credential=AKIA3LROMZGCV47J47W6/20230610/us-west-2/s3/aws4_request,
SignedHeaders=host;x-amz-content-sha256;x-amz-date;x-amz-expected-bucket-owner,
Signature=7415160f39168fdb4fa47f67702c578ea3bcc97471743ff69933d518103fe473"
  }
}
```

A basic signer

github.com/hashbrowncipher/fwdcloudsec-signers

```
$ curl -s -u josh:password \
-d '{
  "url":"https://s3.amazonaws.com/permanent/mykey",
  "method": "GET",
  "headers": {}
}' \
http://127.0.0.1:8000 | jq ".
{
  "url": "https://permanent-quoic7ui7jhvtjt6.s3.us-west-2.amazonaws.com/josh/mykey",
  "headers": {
    "x-amz-content-sha256": "e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855",
    "x-amz-expected-bucket-owner": "111111111111",
    "X-Amz-Date": "20230610T124405Z",
    "Authorization": "AWS4-HMAC-SHA256
Credential=AKIA3LROMZGCV47J47W6/20230610/us-west-2/s3/aws4_request,
SignedHeaders=host;x-amz-content-sha256;x-amz-date;x-amz-expected-bucket-owner,
Signature=7415160f39168fdb4fa47f67702c578ea3bcc97471743ff69933d518103fe473"
  }
}
```



Home directories


```
$ export USER=josh
$ echo $USER > myname; aws s3 cp myname s3://permanent/myname
upload: ./myname to s3://permanent/myname
$ export USER=alice
$ echo $USER > myname; aws s3 cp myname s3://permanent/myname
upload: ./myname to s3://permanent/myname
$ AWS_SKIP_SIGNER=1 aws s3 ls --recursive permanent-quoic7ui7jhvtjt6
2023-06-03 06:05:29      6 alice/myname
2023-06-03 06:05:20      5 josh/myname
$ USER=alice aws s3 cp s3://permanent/myname -
alice
$ USER=josh aws s3 cp s3://permanent/myname -
josh
```

awscli plugin

Use our own credentials (e.g. JWT or x.509) to talk to signer

Use awscli's existing event hooks feature

```
1 import os, requests
2 from botocore.auth import AUTH_TYPE_MAPS, BaseSigner
3 from requests.auth import HTTPBasicAuth
4
5 def make_signing_request(request):
6     basic = HTTPBasicAuth(os.environ["USER"], "password")
7     resp = requests.post("http://localhost:8000", json=request, auth=basic)
8     return resp.json()
9
10 class ExternalSigner(BaseSigner):
11     REQUIRES_REGION = False
12
13     def __init__(self, credentials):
14         pass
15
16     def add_auth(self, request):
17         ea_request = dict(
18             method=request.method,
19             url=request.url,
20             params=request.params,
21             headers=dict(request.headers),
22         )
23
24         response = make_signing_request(ea_request)
25         request.url = response["url"]
26         request.headers = response["headers"]
27
28     def choose_signer(*args, **kwargs):
29         return "mysigner"
30
31     def awscli_initialize(event_hooks):
32         if os.environ.get("AWS_SKIP_SIGNER", "0") != "0":
33             return
34
35         AUTH_TYPE_MAPS["mysigner"] = ExternalSigner
36         event_hooks.register("choose-signer.s3", choose_signer)
```




Can I change responses?

- Like all things IAM, we can only affect requests, not responses
- Without a proxy, we can't change a response

Home directories: handling ListObjects(V2)


```
$ USER=josh aws s3 ls s3://permanent
2023-06-03 06:05:20          5 myname
$ USER=alice aws s3 ls s3://permanent
2023-06-03 06:05:29          6 myname
```

```
C: Hey 127.0.0.1:8000, could you sign this
ListObjectsV2 call please?
S: Sure! Your signed request is:
  GET /list?...
  Host: 127.0.0.1:8000
  Authorization: Basic am9zaDpwYXNzd29yZA==
C: <sends that request>
S: Here's a ListObjectsV2 response that
makes it look like your home directory is
the only one in the bucket.
```

Can I change responses?

- Like all things IAM, we can only affect requests, not responses
- Without a proxy, we can't change a response
- **BUT! We can selectively proxy by rewriting requests:**
 - to ourselves (e.g. previous slide)
 - to an S3 object lambda access point
 - to any other HTTP(S) service



How do I integrate this with my data lake?

Assume we're just signing, not modifying the requested bucket or key.

1. (in the caller) add the signer as a shim into your *AWS* client library
2. (in the signer) authenticate the caller
3. work backwards from key requested to data entity (e.g. table or column)
4. determine whether the caller should have access.
5. sign or reject the request



Does this work with other AWS services?

e.g. DynamoDB

Yes!

...but you might prefer to just
use a proxy



Does it multipart?

Yes, most definitely

And you can cache!



**Can I run the signer
as a lambda?**

Yep!



**Do I have to feel
comfortable running
the signer as a
service in prod?**

Yes!



Should I worry about latency?

Three components:

- Network
- Authorization
 - Caching!
- Signing

Scenario: My bucket is in the "wrong" account

Migration Solution:

1. Reconfigure the app to send all requests to the signer.
2. The signer routes writes into the new bucket.
3. For reads, perform a HEAD on the new bucket
 - a. If present, sign a request for the new bucket
 - b. If absent, sign a request for the old bucket
4. Copy the data in the background

Features that S3 doesn't have, but signers can simulate

- Atomic rename of one file
- Atomic rename of entire directories
- Symlinking (i.e. alias the most recent version of a blob)
- Hard linking (i.e. instant copy)
- Transparent sharding of "hot" keys
- Time travel queries
- Customized data positioning (for regulatory or lifecycle reason)



Thank you!