According to AWS...

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the AWS CLI to validate the files in the location where CloudTrail delivered them.

... while this statement is technically true, when rubber hits the road, their implementation leaves much room for improvement.

## Log File Validation

In order to prevent this type of log tampering, AWS provides a feature called "log file validation", which creates a hash of the log files CloudTrail creates every hour. Administrators can then verify that there logs still match the hash via the various AWS interfaces to the AWS control plane.

We can see if log file validation is enabled by viewing the information relating to CloudTrail:

**Terminal**

```
root@ip-10-0-1-251:~# aws cloudtrail describe-trails --region us-east-2
{
    "trailList": [
        {
            "LogFileValidationEnabled": false,
            "HomeRegion": "us-east-2",
            "IncludeGlobalServiceEvents": true,
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
            "IsMultiRegionTrail": true,
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ"
        }
    ]
}
```

You can see in this output that the CloudTrail AWS service logging is enabled across multi regions and that the log file validation is not currently enabled.

## Exercise

The Plan:

- Explore how CloudTrail logs actions

-- Write a Lambda function to clean a specific IP address from the logs stored within the S3 bucket

## Enable Log File Validation

We can enable log file validation via the following command:

**Terminal**

```
root@ip-10-0-1-251:~# aws cloudtrail update-trail --name "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ" --enable-log-file-validation --r
{
    "IsMultiRegionTrail": true,
    "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
    "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
    "IncludeGlobalServiceEvents": true,
    "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
    "LogFileValidationEnabled": true
}
root@ip-10-0-1-170:~# aws cloudtrail describe-trails --region us-east-2
{
    "trailList": [
        {
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "HomeRegion": "us-east-2",
            "IncludeGlobalServiceEvents": true,
            "LogFileValidationEnabled": true,
            "IsMultiRegionTrail": true,
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy"
        }
    ]
}
```

Now if we browse over to the S3 Bucket, we should see a new folder called "CloudTrail-Digest":

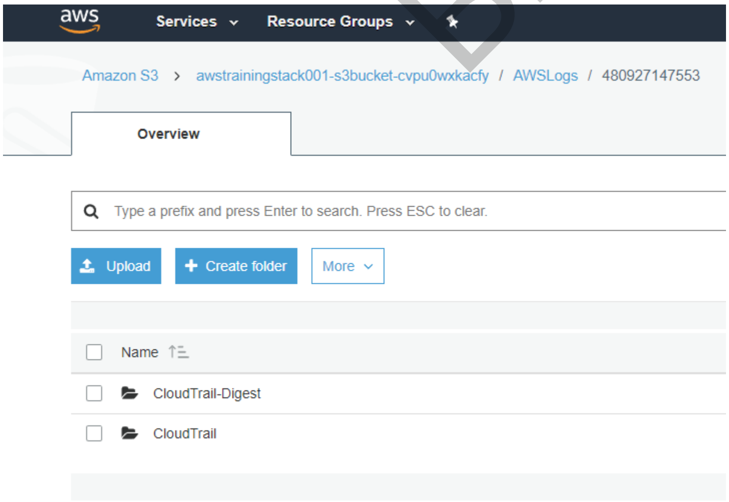
aws
Services
Resource Groups
Amazon S3 › awstrainingstack001-s3bucket-cvpu0wxkacfy / AWSLogs / 480927147553
Overview
Type a prefix and press Enter to search. Press ESC to clear.
Upload
Create folder
More
Name
CloudTrail-Digest
CloudTrail

Within this folder, hashes for each log file will be created every hour.

If we run the following command we can validate that log files have not been tampered with:

**Terminal**

```
root@ip-10-0-1-251:~# aws cloudtrail validate-logs --trail-arn "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ" --start-time 2018-01-01T12:31:41Z --region us-east-2
Validating log files for trail arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ between 2018-01-01T12:31:41Z and 2018-07-10T02:07:47Z

Results requested for 2018-01-01T12:31:41Z to 2018-07-10T02:07:47Z
Results found for 2018-07-10T01:01:08Z to 2018-07-10T02:01:08Z:

1/1 digest files valid
```

And within the AWS Console, under the CloudTrail service we should see:

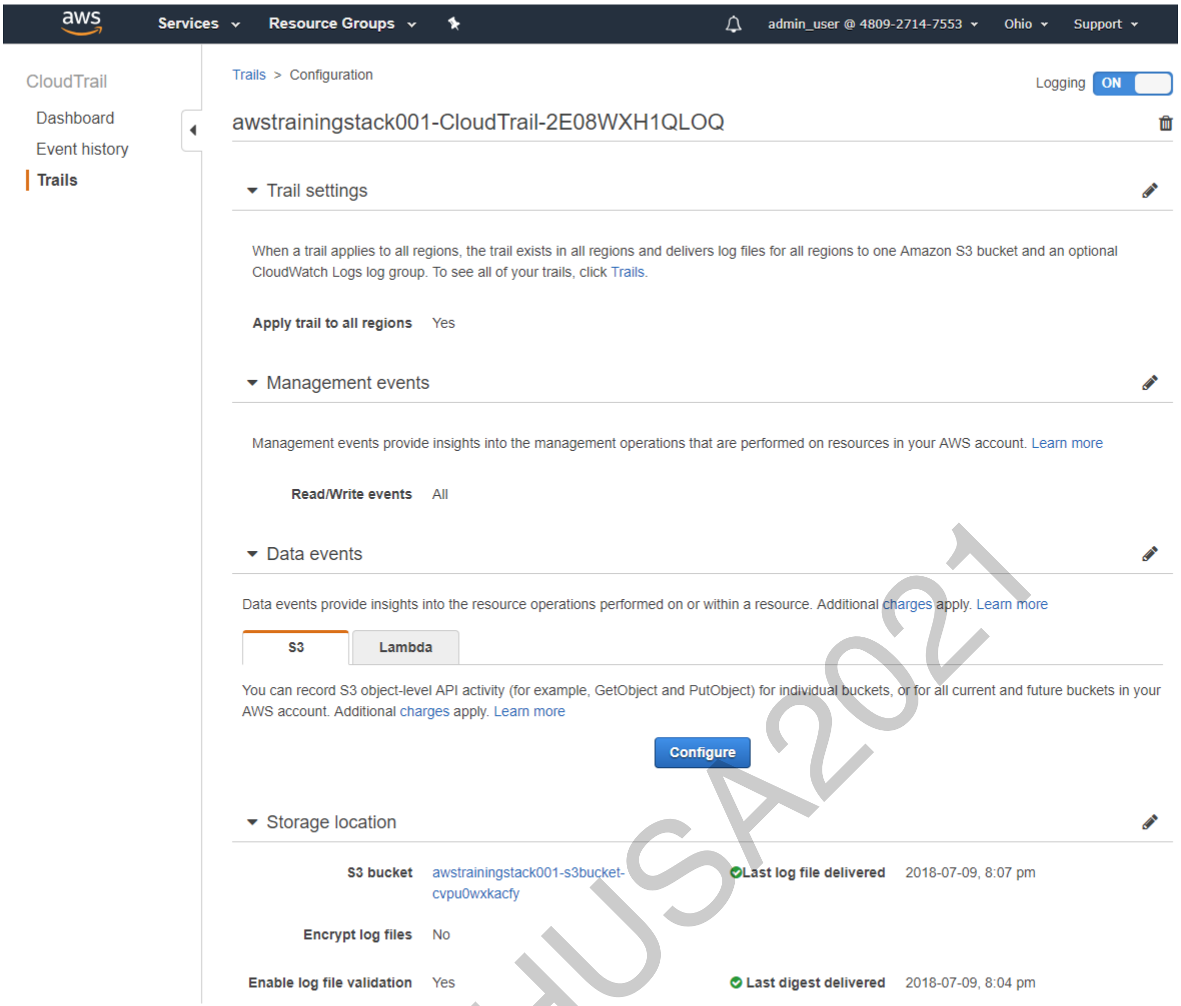

awstrainingstack001-CloudTrail-2E08WXH1QLOQ

Trail settings

When a trail applies to all regions, the trail exists in all regions and delivers log files for all regions to one Amazon S3 bucket and an optional CloudWatch Logs log group. To see all of your trails, click Trails.

**Apply trail to all regions** Yes

Management events

Management events provide insights into the management operations that are performed on resources in your AWS account. Learn more

**Read/Write events** All

Data events

Data events provide insights into the resource operations performed on or within a resource. Additional charges apply. Learn more

You can record S3 object-level API activity (for example, GetObject and PutObject) for individual buckets, or for all current and future buckets in your AWS account. Additional charges apply. Learn more

Storage location

**S3 bucket** awstrainingstack001-s3bucket-cvpu0wxkacfy

✓ **Last log file delivered** 2018-07-09, 8:07 pm

**Encrypt log files** No

**Enable log file validation** Yes

✓ **Last digest delivered** 2018-07-09, 8:04 pm

That log file validation is enabled and that the last digest delivered has been validated.

Log verification is a process which Cloudtrail uses to prove the log files have not been changed or deleted.

It does this by continuously creating "digest files" which contain hashes used to verify the log files.

Every digest points to the previous digest before it in the chain, forming a chain into the past.

S3 protects digest files from deletion where a subsequent digest file is pointing to it.

If we temporary disable log verification and then enable it again the chain of digests will be started again as if from the beginning, effectively breaking the previous chain of digests.

The previous digest will not have another digest pointing to it, and hence can now be deleted from the S3 bucket.

Hence if we stop and start log verification we can delete previous digests deleting our way from the newest digests to the oldest digests up the chain.

If digest are deleted, then modified log files will no longer fail validation.

Let's try this out by browsing to the Lambda service within AWS and clicking the "Create a function" button.

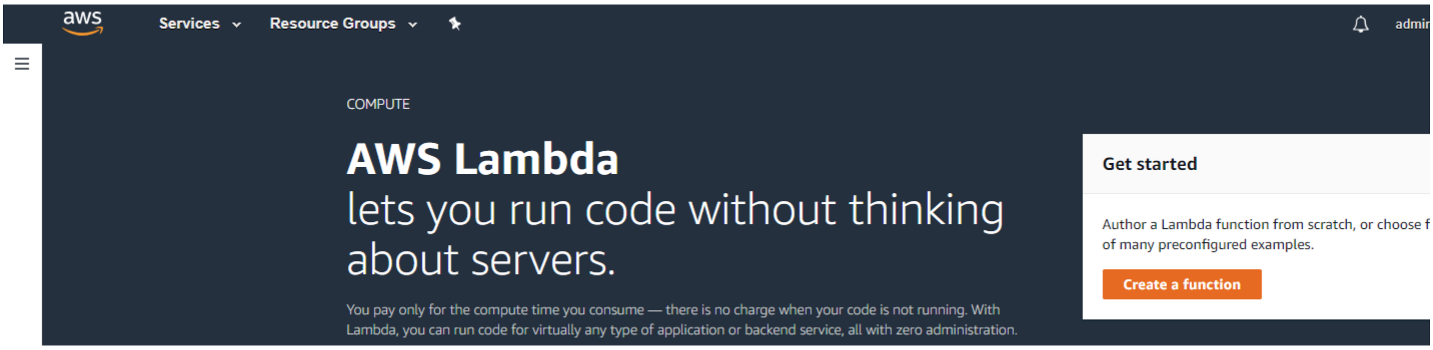

**AWS Lambda**
lets you run code without thinking about servers.

You pay only for the compute time you consume — there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or backend service, all with zero administration.

**Get started**

Author a Lambda function from scratch, or choose f of many preconfigured examples.

Let's create a Lambda function with the following settings:

- "Author from scratch"
- Name: myValidationFunction001
- Runtime: Node.js 10
- Role: Choose an existing role
- Existing role: lambda_clean_execution_001

And then clicking the "Create function" button.

Next let's find our IP address and start creating some logs to clean within a terminal:

Terminal

```
root@ip-10-0-1-251:~# curl ipcurl.net/n
18.216.151.35

root@ip-10-0-1-251:~# watch -n 5 -d aws cloudtrail describe-trails --region us-east-2
Every 5.0s: aws cloudtrail describe-trails --region us-east-2                                                                  Tue Jul 10 02:36:08 2018

{
    "trailList": [
        {
            "IsMultiRegionTrail": true,
            "IncludeGlobalServiceEvents": true,
            "HomeRegion": "us-east-2",
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "LogFileValidationEnabled": true,
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy"
        }
    ]
}
```

Now let's continue creating a Lambda function by downloading the following code: https://www.dropbox.com/s/runqcp4cy9fdro2/disable-cloudtrail-lambda-master.zip?dl=0

This Lambda function will automatically delete digest files and filter out logs containing a specified IP address.

Unzip the file and open up the "hiddenIPs.json" file within a text editor and add the IP address we want to filter:

**hiddenIPs.json**

```
["18.216.151.35"]
```

Now create a new zip file from the folder you extracted with the edited hiddenIPs.json file inside (depending on your operating system you can probably right click the folder and click compress).

And now within your Lambda Management Console for the function you just created upload the zip you just created.

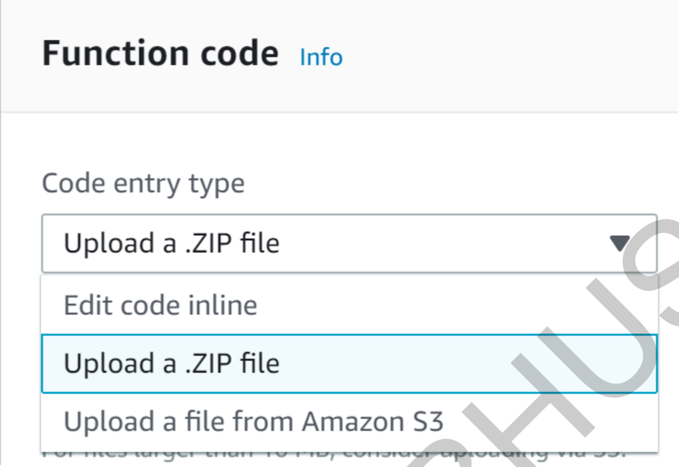

Code entry type: Upload a .ZIP file

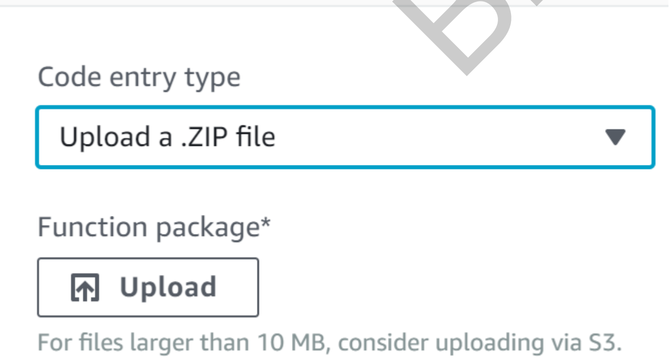

Code entry type: Upload a .ZIP file

For files larger than 10 MB, consider uploading via S3.

Now your Function code environment should look similar to this:

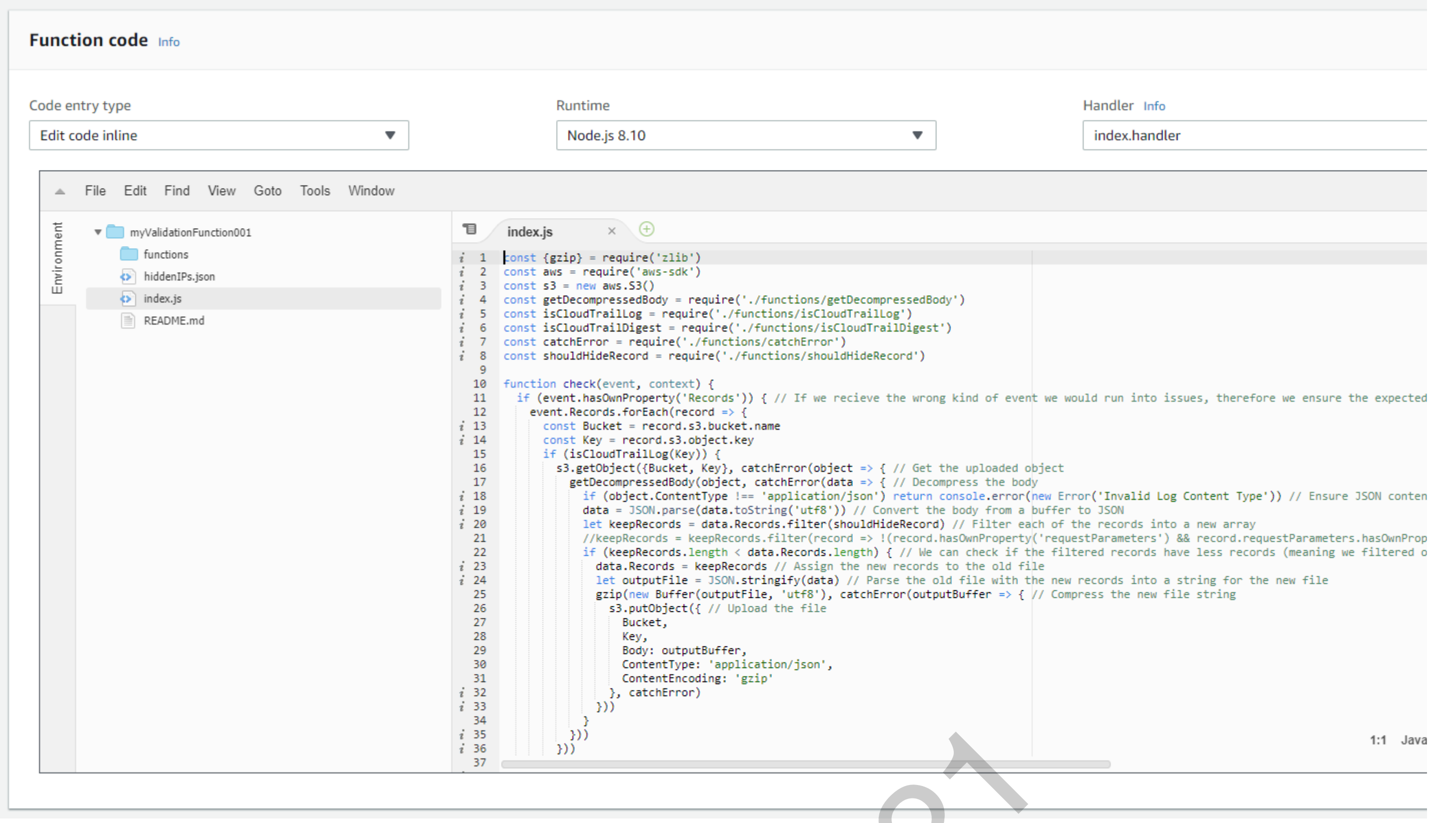

**Function code** Info

Code entry type: Edit code inline

Runtime: Node.js 8.10

Handler Info: index.handler

```
const {gzip} = require('zlib')
const aws = require('aws-sdk')
const s3 = new aws.S3()
const getDecompressedBody = require('./functions/getDecompressedBody')
const isCloudTrailLog = require('./functions/isCloudTrailLog')
const isCloudTrailDigest = require('./functions/isCloudTrailDigest')
const catchError = require('./functions/catchError')
const shouldHideRecord = require('./functions/shouldHideRecord')

function check(event, context) {
  if (event.hasOwnProperty('Records')) { // If we recieve the wrong kind of event we would run into issues, therefore we ensure the expected
    event.Records.forEach(record => {
      const Bucket = record.s3.bucket.name
      const Key = record.s3.object.key
      if (isCloudTrailLog(Key)) {
        s3.getObject({Bucket, Key}, catchError(object => { // Get the uploaded object
          getDecompressedBody(object, catchError(data => { // Decompress the body
            if (object.ContentType !== 'application/json') return console.error(new Error('Invalid Log Content Type')) // Ensure JSON conten
            data = JSON.parse(data.toString('utf8')) // Convert the body from a buffer to JSON
            let keepRecords = data.Records.filter(shouldHideRecord) // Filter each of the records into a new array
            //keepRecords = keepRecords.filter(record => !(record.hasOwnProperty('requestParameters') && record.requestParameters.hasOwnProp
            if (keepRecords.length < data.Records.length) { // We can check if the filtered records have less records (meaning we filtered o
              data.Records = keepRecords // Assign the new records to the old file
              let outputFile = JSON.stringify(data) // Parse the old file with the new records into a string for the new file
              gzip(new Buffer(outputFile, 'utf8'), catchError(outputBuffer => { // Compress the new file string
                s3.putObject({ // Upload the file
                  Bucket,
                  Key,
                  Body: outputBuffer,
                  ContentType: 'application/json',
                  ContentEncoding: 'gzip'
                }, catchError)
              }))
            }
          }))
        }))

```

We can then add the trigger for this function by clicking "S3" on the left hand-side of the interface:

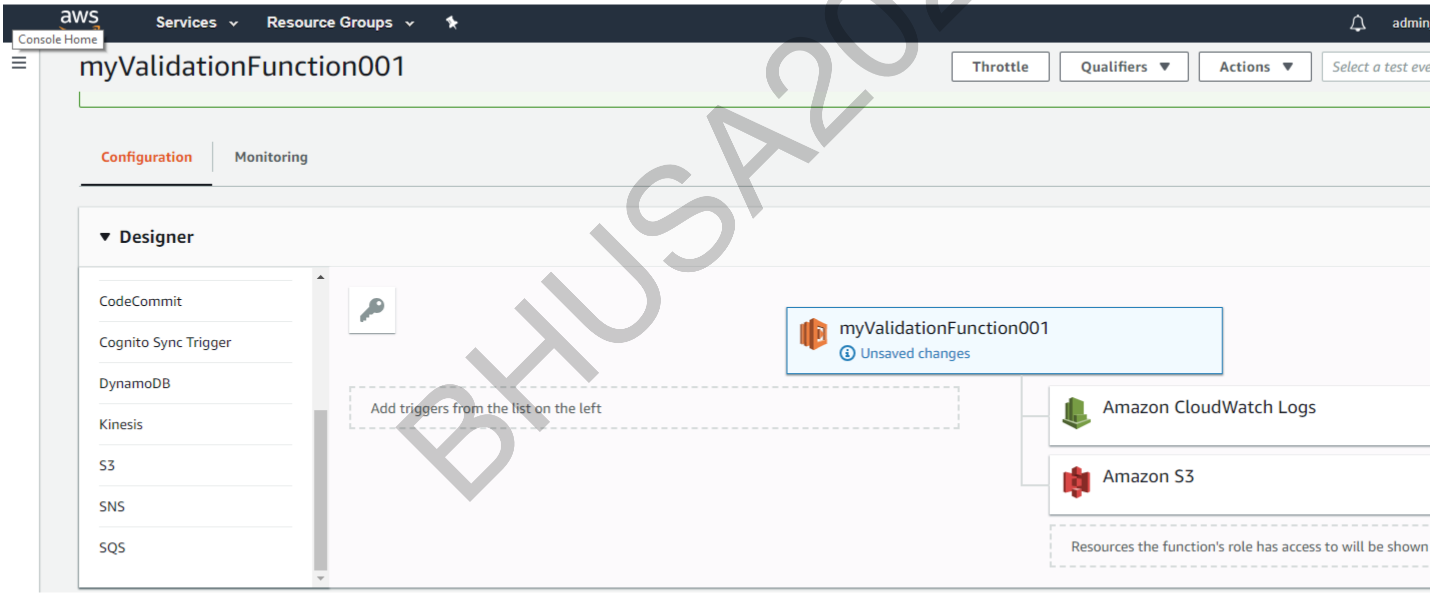

myValidationFunction001

**Designer**

Add triggers from the list on the left

myValidationFunction001
Unsaved changes

Resources the function's role has access to will be shown

We then can select the S3 bucket with the logs and the click the "Add" button.

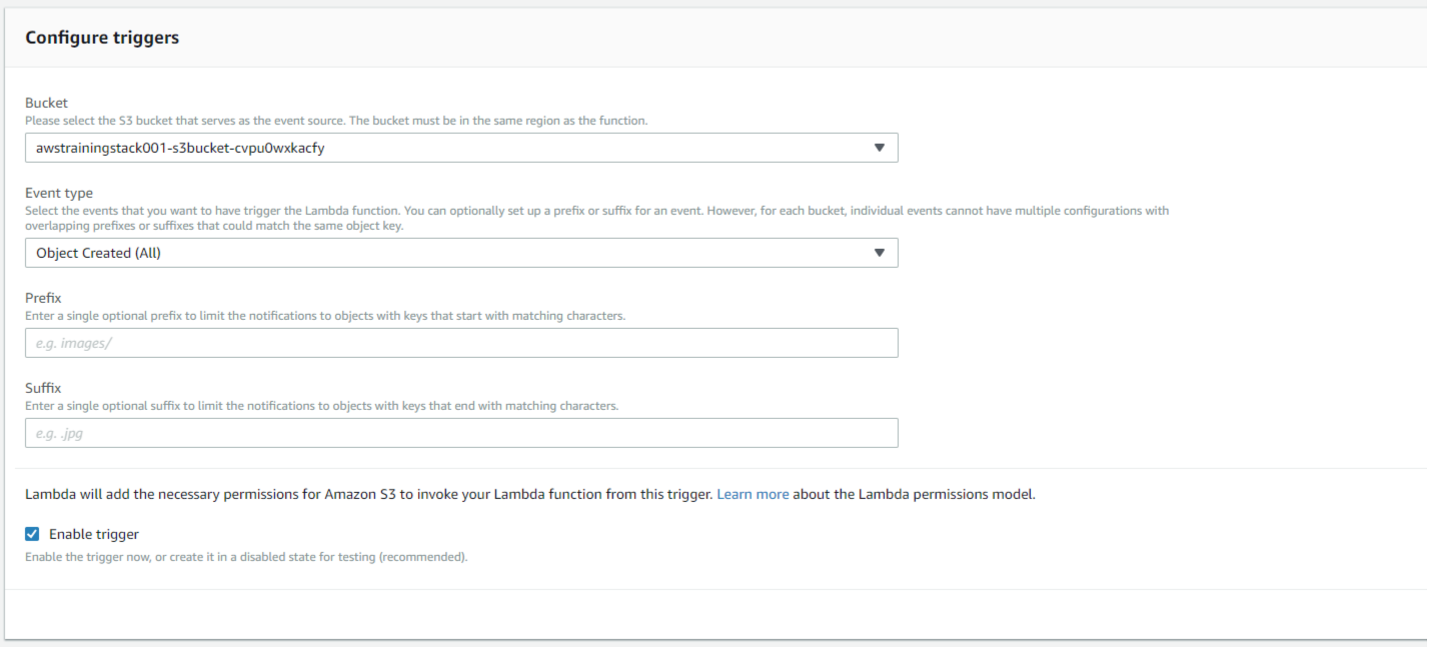

**Configure triggers**

Bucket
Please select the S3 bucket that serves as the event source. The bucket must be in the same region as the function.
awstrainingstack001-s3bucket-cvpu0wxkacfy

Event type
Select the events that you want to have trigger the Lambda function. You can optionally set up a prefix or suffix for an event. However, for each bucket, individual events cannot have multiple configurations with overlapping prefixes or suffixes that could match the same object key.
Object Created (All)

Prefix
Enter a single optional prefix to limit the notifications to objects with keys that start with matching characters.
*e.g. images/*

Suffix
Enter a single optional suffix to limit the notifications to objects with keys that end with matching characters.
*e.g. .jpg*

Lambda will add the necessary permissions for Amazon S3 to invoke your Lambda function from this trigger. Learn more about the Lambda permissions model.

☑ Enable trigger
Enable the trigger now, or create it in a disabled state for testing (recommended).

We should then see something similar to this:

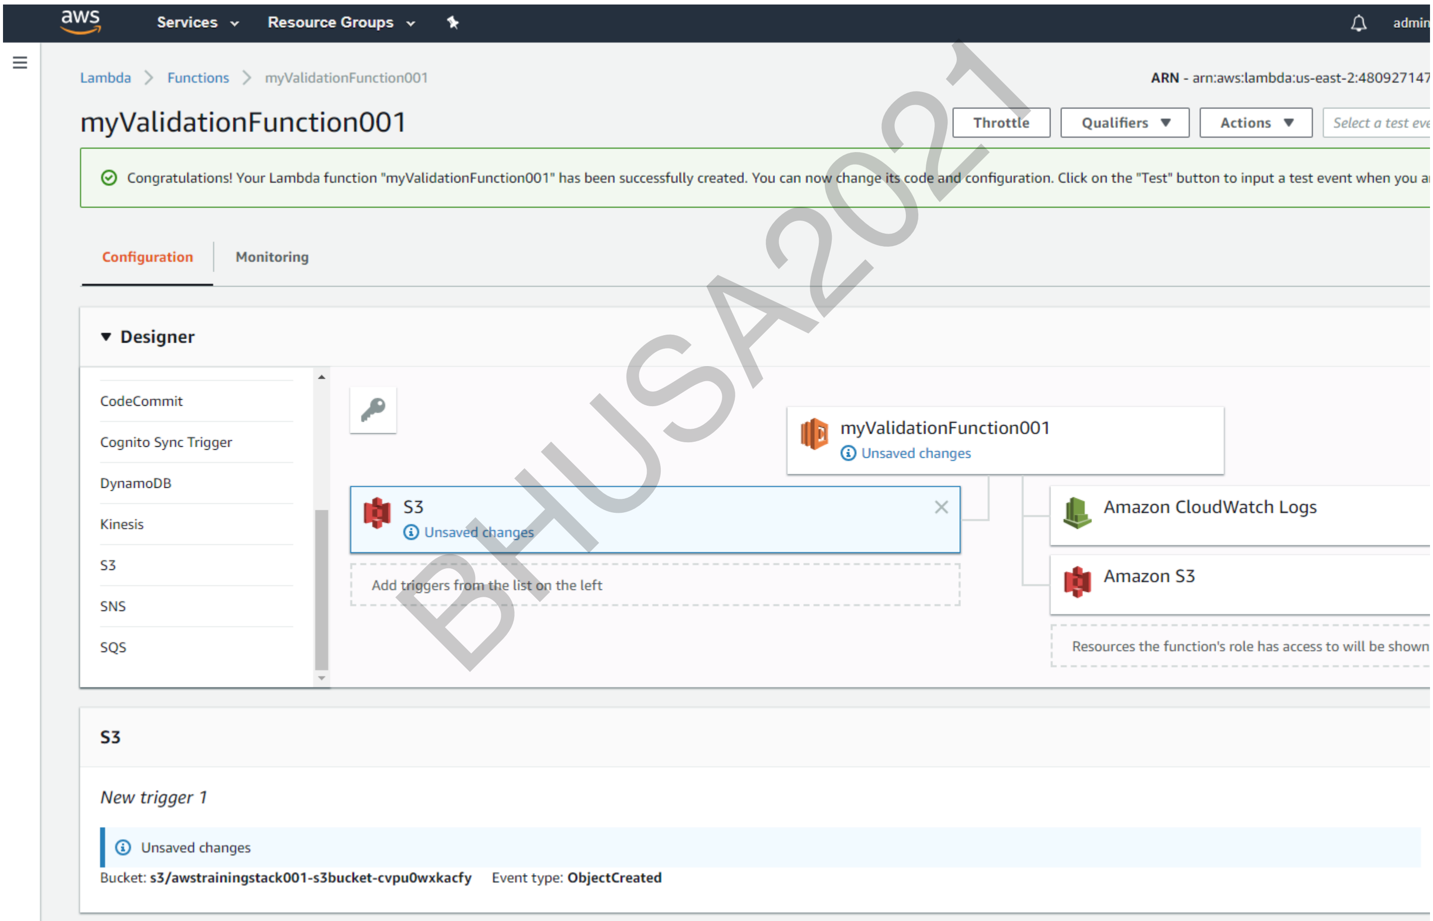

ARN - arn:aws:lambda:us-east-2:480927147

# myValidationFunction001

Congratulations! Your Lambda function "myValidationFunction001" has been successfully created. You can now change its code and configuration. Click on the "Test" button to input a test event when you a

**S3**

*New trigger 1*

Unsaved changes

Bucket: **s3/awstrainingstack001-s3bucket-cvpu0wxkacfy** Event type: **ObjectCreated**

Then we click the "Save" button in the upper right-hand corner to save & start running the Lambda function and we should see output similar to this:

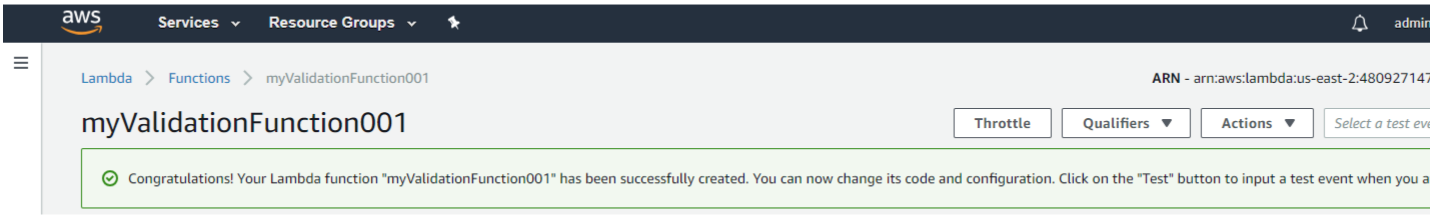

ARN - arn:aws:lambda:us-east-2:480927147

# myValidationFunction001

Congratulations! Your Lambda function "myValidationFunction001" has been successfully created. You can now change its code and configuration. Click on the "Test" button to input a test event when you a

We now disable and then enable the log file verification via the following commands:

Terminal

```
root@ip-10-0-1-251:~# aws cloudtrail update-trail --name "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ" --no-enable-log-file-validation
{
    "IsMultiRegionTrail": true,
    "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
    "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
    "LogFileValidationEnabled": false,
    "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
    "IncludeGlobalServiceEvents": true
}

root@ip-10-0-1-251:~# aws cloudtrail describe-trails --region us-east-2
{
    "trailList": [
        {
            "IncludeGlobalServiceEvents": true,
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
            "HomeRegion": "us-east-2",
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "IsMultiRegionTrail": true,
            "LogFileValidationEnabled": false,
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ"
        }
    ]
}

root@ip-10-0-1-251:~# aws cloudtrail update-trail --name "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ" --enable-log-file-validation --r
{
    "IsMultiRegionTrail": true,
    "IncludeGlobalServiceEvents": true,
    "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
    "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
    "LogFileValidationEnabled": true,
    "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ"
}

root@ip-10-0-1-251:~# aws cloudtrail describe-trails --region us-east-2
{
    "trailList": [
        {
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "HomeRegion": "us-east-2",
            "LogFileValidationEnabled": true,
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy",
            "IsMultiRegionTrail": true,
            "IncludeGlobalServiceEvents": true,
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ"
        }
    ]
}
```

If we give it a while (e.g. over an hour) we can then click the "Monitoring" link to see this function get fired at the same time that CloudTrail writes new digests into the S3 bucket, approximately every 60 minutes:

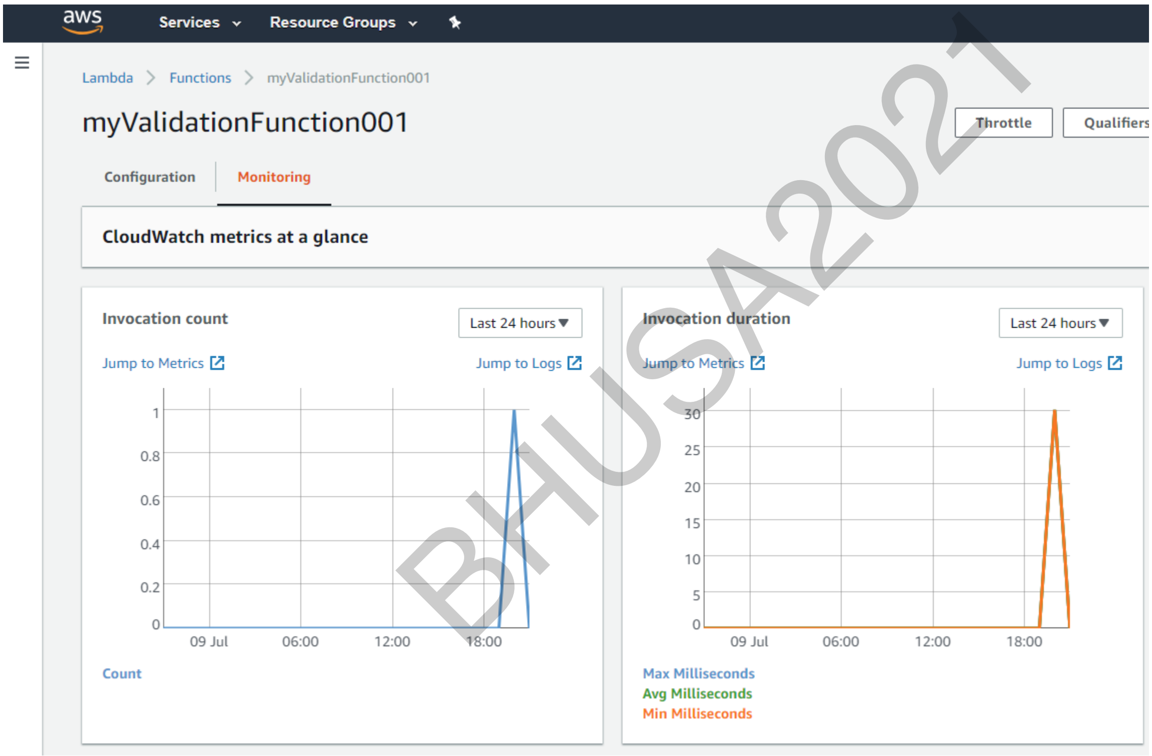

If we run the following command we can validate that log files have not been tampered with:

Terminal

```
root@ip-10-0-1-251:~# aws cloudtrail describe-trails --region us-east-2
{
    "trailList": [
        {
            "TrailARN": "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "LogFileValidationEnabled": true,
            "Name": "awstrainingstack001-CloudTrail-2E08WXH1QLOQ",
            "HomeRegion": "us-east-2",
            "IsMultiRegionTrail": true,
            "IncludeGlobalServiceEvents": true,
            "S3BucketName": "awstrainingstack001-s3bucket-cvpu0wxkacfy"
        }
    ]
}

root@ip-10-0-1-251:~# aws cloudtrail validate-logs --trail-arn "arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ" --start-time 2018-01-01T12
Validating log files for trail arn:aws:cloudtrail:us-east-2:480927147553:trail/awstrainingstack001-CloudTrail-2E08WXH1QLOQ between 2018-01-01T12:31:41Z and 2018-07-10T02:07:47Z

Results requested for 2018-01-01T12:31:41Z to 2018-07-10T02:07:47Z
Results found for 2018-07-10T01:01:08Z to 2018-07-10T02:01:08Z:

1/1 digest files valid
```

So we can see that log file validation is enabled and that the previous log files still validate but that no new digest files are now appearing, because they are automatically being deleted via the lambda function.

We can see within the web console that log validation is enabled:

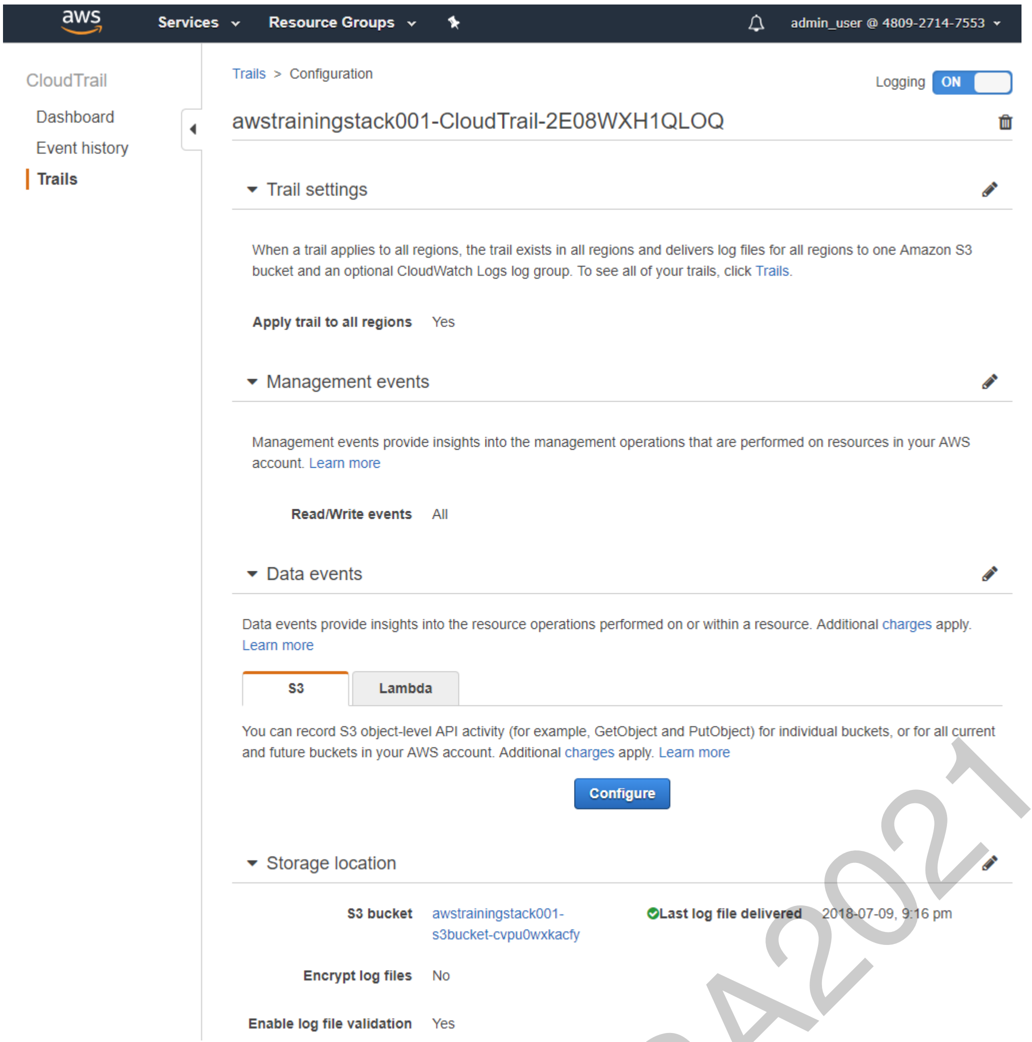

awstrainingstack001-CloudTrail-2E08WXH1QLOQ

Trail settings

When a trail applies to all regions, the trail exists in all regions and delivers log files for all regions to one Amazon S3 bucket and an optional CloudWatch Logs log group. To see all of your trails, click Trails.

**Apply trail to all regions** Yes

Management events

Management events provide insights into the management operations that are performed on resources in your AWS account. Learn more

**Read/Write events** All

Data events

Data events provide insights into the resource operations performed on or within a resource. Additional charges apply. Learn more

You can record S3 object-level API activity (for example, GetObject and PutObject) for individual buckets, or for all current and future buckets in your AWS account. Additional charges apply. Learn more

Storage location

**S3 bucket** awstrainingstack001-s3bucket-cvpu0wxkacfy

**Last log file delivered** 2018-07-09, 9:16 pm

**Encrypt log files** No

**Enable log file validation** Yes

But we can also see that the last digest delivered is no longer listed, because the digests are being automatically delete via the Lambda function.

## Exercise

The Plan:

```
- Explore how digest work
-- Create a Lambda function to remove digests on the fly and in real-time
```

## References

Check out the following references for more information:

- https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html
- Disrupting AWS logging - https://danielgrzelak.com/disrupting-aws-logging-a42e437d6594
- Disrupting AWS S3 Logging - http://blog.thinkst.com/2017/08/disrupting-aws-s3-logging.html
- Validating CloudTrail Log File Integrity - https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html