Demo of CloudTrail using the CLI

1. Create an S3bucket:

```
aws s3 mb s3://trailbucketdemo2015
```

2. Now run the command to create a trail

```
aws cloudtrail create-trail
--name mytrail20254679
--s3-bucket-name trailbucketdemo2015
```

You will notice that you get this response:

```
An error occurred (InsufficientS3BucketPolicyException) when
calling the CreateTrail operation: Incorrect S3 bucket policy
is detected for bucket: trailbucketdemo2015
```

What is going on? It looks like we need a bucket policy on our bucket that allows CloudTrail to dump logs here.

This is something that is sorted out for you when you use the AWS Management Console. So, I will create a trail using the AWS Management Console, then I will look at the S3 bucket to see the kind of policy that has been automatically added to it.

Well, I have just used the Management Console to create a Trail and it adds this policy to the S3 bucket it automatically creates:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AWSCloudTrailAclCheck20150319-06449c8a-4c59-
4189-b67a-be8073c460b1",
            "Effect": "Allow",
            "Principal": {
```

```
                    "Service": "cloudtrail.amazonaws.com"
                },
                "Action": "s3:GetBucketAcl",
                "Resource": "arn:aws:s3:::aws-cloudtrail-logs-
737911634202-a1d07d70",
                "Condition": {
                    "StringEquals": {
                        "AWS:SourceArn": "arn:aws:cloudtrail:eu-west-
2:737911634202:trail/management-events"
                    }
                }
            },
            {
                "Sid": "AWSCloudTrailWrite20150319-d322cb63-aa73-417c-
8a15-596512fe57fb",
                "Effect": "Allow",
                "Principal": {
                    "Service": "cloudtrail.amazonaws.com"
                },
                "Action": "s3:PutObject",
                "Resource": "arn:aws:s3:::aws-cloudtrail-logs-
737911634202-a1d07d70/AWSLogs/737911634202/*",
                "Condition": {
                    "StringEquals": {
                        "AWS:SourceArn": "arn:aws:cloudtrail:eu-west-
2:737911634202:trail/management-events",
                        "s3:x-amz-acl": "bucket-owner-full-control"
                    }
                }
            }
        ]
    }
```

Let's take a moment to analyse this. We have two actions (`s3:GetBucketAcl` and `s3:PutObject`). Take a moment to spot those in the above.

Why CloudTrail needs PutObject is obvious. But GetBucketAcl is more puzzling.

Could I just copy and paste this policy? Suppose I wanted to do it myself - to create a bucket manually using the AWS CLI, for CloudTrail to dump logs into. Is there anything in the policy I would have to change? Take a moment to take a look.

I would need to change:

(i) The values for the `Resource` key

One value is this:

```
"arn:aws:s3:::aws-cloudtrail-logs-737911634202-
a1d07d70/AWSLogs/737911634202/*",
```

And that is the name of this specific S3 bucket. So, we'd need to change that.

Let's take a look at the condition key. Its value is:

```
{
                "StringEquals": {
                    "AWS:SourceArn": "arn:aws:cloudtrail:eu-west-
2:737911634202:trail/management-events",
                    "s3:x-amz-acl": "bucket-owner-full-control"
                      }
```

What this is doing is ensuring that the person meddling with the bucket has a specific ARN (SourceArn), or Amazon Resource Name.

Notice where is says `trail/management-events`. This means that a specific, named Trail must be the person meddling with the bucket.

So, we'd need to change this ARN, if we were to copy and paste this policy. We'd need to ensure the ARN is the ARN for *our* trail, and not *that* trail.

(ii) The value of the `Condition` key

Those are the two things we need to change.

We need to have a policy which has this part:

```
"Action": "s3:PutObject",

"Resource": "arn:aws:s3:::trailbucketdemo2015",
```

And also this part:

```
"Action": "s3:GetBucketAcl",

"Resource": "arn:aws:s3:::trailbucketdemo2015",
```

Here is the completed policy with the parts it ought to have. I am also going to add more appropriate Sid values:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "MyAclCheckStatement",
            "Effect": "Allow",
            "Principal": {
                "Service": "cloudtrail.amazonaws.com"
            },
            "Action": "s3:GetBucketAcl",
            "Resource": "arn:aws:s3:::trailbucketdemo2015",
            "Condition": {
                "StringEquals": {
                    "AWS:SourceArn": "arn:aws:cloudtrail:eu-west-
2:737911634202:trail/mytrail20254679"
                }
            }
        },
        {
            "Sid": "MyPutObjectStatement",
            "Effect": "Allow",
            "Principal": {
                "Service": "cloudtrail.amazonaws.com"
            },
            "Action": "s3:PutObject",
            "Resource": "arn:aws:s3:::traildemobucket2015",
            "Condition": {
                "StringEquals": {
                    "AWS:SourceArn": "arn:aws:cloudtrail:eu-west-
2:737911634202:trail/mytrail20254679",
                    "s3:x-amz-acl": "bucket-owner-full-control"
                }
            }
        }
    ]
}
```

Adding Policies to Buckets using the AWS CLI

How do we add an IAM policy to an S3 bucket using the AWS CLI?

I take a look at the AWS CLI reference, online, which lists all the commands. We can use `aws s3api put-bucket-policy`.

Note - you do not just paste the policy into the `--policy` parameter. You must place the path to a file in this place.

I try this command:

```
aws s3api put-bucket-policy --bucket trailbucketdemo2015 --policy
"C:\Users\samja\OneDrive\Documents\teacher_training\AWS_tutoring\aws-
tube\cloudtrail_demo\cloudtrail_demo_policy.txt"
```

It does not work. I am told:

```
An error occurred (MalformedPolicy) when calling the PutBucketPolicy operation:
Policies must be valid JSON and the first byte must be '{'
```

What could the problem be? I've looked at the file containing the policy and it *does* begin with an opening curly bracket.

Perhaps I need to add file://

However, this certainly looks strange, with the double-forward-slash before the upper-case C.

Let's try this:

```
aws s3api put-bucket-policy --bucket trailbucketdemo2015 --
policy
file://C:\Users\samja\OneDrive\Documents\teacher_training\AWS_
tutoring\aws-tube\cloudtrail_demo\cloudtrail_demo_policy.txt
```

We get an error:

```
An error occurred (MalformedPolicy) when calling the
PutBucketPolicy operation: Policy has invalid resource
```

I put the policy into ChatGPT and it seems that I have made a mistake. I have mis-quoted the name of the bucket at one point. The name is `trailbucketdemo2015`.

ChatGPT has also given me some other feedback. The `s3:PutObject` action applies to objects, not the bucket itself. Therefore, the correct ARN should include the `/*` suffix. We should **not** have written:

```
"Resource": "arn:aws:s3:::trailbucketdemo2015"
```

But rather this:

```
"Resource": "arn:aws:s3:::trailbucketdemo2015/*"
```

We have added a forward slash and an asterisk. Before we were specifying the bucket as the resource. But the `PutObject` action applies to *objects*, so we need to specify *objects* using the `"Resource"` key, not buckets. By using the asterisk we specify objects - all of them.

Another way we can improve our policy is by adding `s3:ListBucket` permissions.

Now, why will this help? Granting this permission helps CloudTrail provide clearer error messages and quicker resolution of issues. With `s3:ListBucket`, CloudTrail can validate the path and avoid silent delivery failures.

I correct some comma and indentation issues. Now I try the command once more, and this time the command below works:

```
aws s3api put-bucket-policy

--bucket trailbucketdemo2015

--policy
file://C:\Users\samja\OneDrive\Documents\teacher_training\AWS_tutoring\aws-
tube\cloudtrail_demo\cloudtrail_demo_policy.txt
```

Now that we have added a policy to our S3 bucket, which allows CloudTrail to meddle with it, we can have another go at creating a trail.

Creating a Trail

Let's try this command:

```
aws cloudtrail create-trail
      --name mytrail20254679
      --s3-bucket-name trailbucketdemo2015
```

It works. We get back this response:

```
{
    "Name": "mytrail20254679",
    "S3BucketName": "trailbucketdemo2015",
    "IncludeGlobalServiceEvents": true,
    "IsMultiRegionTrail": false,
    "TrailARN": "arn:aws:cloudtrail:eu-west-2:737911634202:trail/mytrail20254679",
    "LogFileValidationEnabled": false,
    "IsOrganizationTrail": false
}
```

Now let's enable logging.

I'm going to run this command:

```
aws cloudtrail start-logging --name mytrail20254679
```

You get nothing for the response.

Now we want to perform an action that triggers logging. Consider this command:

```
aws s3 mb s3://trailbucketdemo2015
```

The fact that this command will fail (because we already have this named bucket) does not really matter. We will see:

```
make_bucket failed: s3://trailbucketdemo2015 An error occurred
(BucketAlreadyOwnedByYou) when calling the CreateBucket operation:

Your previous request to create the named bucket succeeded and you already own it.
```

We have made an API call. So, there should be a record of this in the trail.

Investigating

Run this command:

```
aws cloudtrail get-trail-status --name
mytrail20254679
```

You'll see a response which looks something like this:


```
{
    "IsLogging": true,
    "LatestDeliveryTime": "2025-03-11T22:19:30.178000+00:00",
    "StartLoggingTime": "2025-03-11T22:14:26.331000+00:00",
    "LatestDeliveryAttemptTime": "2025-03-11T22:19:30Z",
    "LatestNotificationAttemptTime": "",
    "LatestNotificationAttemptSucceeded": "",
    "LatestDeliveryAttemptSucceeded": "2025-03-11T22:19:30Z",
    "TimeLoggingStarted": "2025-03-11T22:14:26Z",
    "TimeLoggingStopped": ""
}
```


Now let's list the logs in the S3 bucket:

```
aws s3 ls s3://trailbucketdemo2015/AWSLogs/737911634202/CloudTrail/
```

Notice how we specify the precise folder(s) within the bucket.

When I try this command, I get an unsatisfactory response:

```
PRE eu-west-2/
2025-03-11 21:40:17          0
```

What does this mean? What does **PRE** mean?