Tutorial 5 – Introduction to Lambda II:
Working with Files in S3 and CloudWatch Events

Disclaimer: Subject to updates as corrections are found
Version 0.12
Scoring: 20 pts maximum

The purpose of this tutorial is to introduce the use of the Amazon Simple Storage Service (S3) from AWS Lambda to support receiving, processing, and/or creating files. Additionally, this tutorial introduces combining multiple Lambdas into a single Java project. The tutorial also describes how to configure a CloudWatch event rule to trigger a Lambda function in response to an S3 Write Data event that is tracked by setting up a CloudTrail log "trail".

Goals of this tutorial include:
1. Create a new CreateCSV Lambda function to write a file to S3.
2. Create a new ProcessCSV Lambda function to read a file from S3.
3. Package these two Lambda functions into a single Java project to produce a single, composite jar file. The concept of a composite JAR provides the basis for setting up the "Switchboard" architecture by simply adding additional flow-control code.
4. Create a CloudWatch event rule to trigger the ProcessCSV Lambda function as a "target". The event rule is triggered when a file is uploaded to an S3 bucket by the CreateCSV Lambda function. This event is provided by setting up a CloudTrail log trail to track S3 Write Data events. CloudWatch event triggers provide a means to implement **asynchronous application flow control** as in:

For example, Lambda functions can be triggered in response to data being available in S3.

1. Create a new SAAF Lambda function template application

On your laptop, create a new directory for the project files and clone the git SAAF project to start:

```
git clone https://github.com/wlloyduw/SAAF.git
```

Refer to Tutorial #4 for information on the "SAAF".
Also refer to Tutorial #4 regarding selection of Java version (8 or 11) for AWS Lambda.

2. CreateCSV Lambda Function

In the SAAF project, rename the HelloPOJO.java class to CreateCSV.java.

If using the Netbeans IDE, in the project, right click on the classname, and select "Refactor | Rename". Rename the "HelloPOJO" class to "CreateCSV".

IMPORTANT TIP USING NETBEANS: **Java Import statements can be automatically added to your code by pressing CONTROL – SHIFT - i for imports... This saves from having to manually determine and add the import statements at the top of the class file.**

If using the command line rename the file, and use a text editor to modify the class name inside CreateCSV.java to match "CreateCSV".

```
cd {base directory where project was cloned}/SAAF/java_template/src/main/java/lambda
mv HelloPOJO.java CreateCSV.java
```

In the Request.java class in the same directory, add 4 input parameters for this Lambda function. Define getter and setters methods accordingly:

Property Name	Property Type
bucketname	String
filename	String
row	int
col	int

See existing properties with "get" and "set" methods in the Request.java class as an example. These properties will allow a client to request the creation of a new CSV file. The file is stored in the S3 Bucket described by "Bucketname". The filename is described by "Filename". The CSV file will consist of comma-separated random numbers (range 1 to 1000). Row and Col specify the number of total rows and columns in the CSV file.

In the CreateCSV class handleRequest() method, consume the request variables into local variables:

```
    int row = request.getRow();
    int col = request.getCol();
    String bucketname = request.getBucketname();
    String filename = request.getFilename();
```

Then, generate the random matrix of values, storing each row in a separate String:

2

```
        int val = 0;
        StringWriter sw = new StringWriter();
        Random rand = new Random();

        for (int i=0;i<row;i++)
            for (int j=0;j<col;j++)
            {
                val = rand.nextInt(1000);
                sw.append(Integer.toString(val));
                if ((j+1)!=col)
                    sw.append(",");
                else
                    sw.append("\n");
            }
```
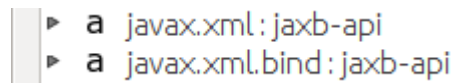
The next step is to write Java code to generate the S3 bucket file.

Working with the S3 Java API requires adding the Java support library.
In maven, the library can be added in the pom.xml file.
The pom.xml file is under the java_template directory.

The dependencies can *alternatively* be added through the Netbeans IDE, by RIGHT-clicking on dependencies, and select "Add Dependency", and then in the Query box type "**aws-java-sdk-s3**".
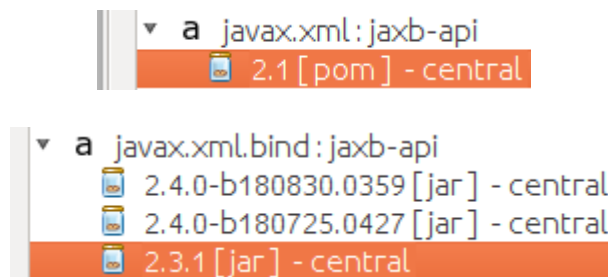
Once found, select the version, such as ~ "1.11.306"…
This automatically includes all Java jar libraries required to work with Amazon S3.

Two more dependencies are required. Select the newest non-beta version. Beta versions have a "b" in the name. Search for "**jaxb-api**". Add "javax.xml: jaxb-api" and "javax.xml.bind".

newest non-beta versions:

Alternatively, if not using Netbeans, the dependency can be added in the pom.xml file in between the tags "<dependencies> </dependencies>" directly as follows:

```
        <dependency>
            <groupId>com.amazonaws</groupId>
            <artifactId>aws-java-sdk-s3</artifactId>
            <version>1.11.306</version>
        </dependency>
        <dependency>
            <groupId>javax.xml</groupId>
            <artifactId>jaxb-api</artifactId>
            <version>2.1</version>
            <type>pom</type>
        </dependency>
        <dependency>
            <groupId>javax.xml.bind</groupId>
            <artifactId>jaxb-api</artifactId>
            <version>2.3.1</version>
        </dependency>
```

Note that these dependencies are for Java 11. If deploying with Java 8 they may be somewhat different.

Once S3 dependencies have been included in the project, acquire the StringBuilder output as a Byte Array, and use this to create a new input stream. Then create metadata for describing the file to be written to S3 and create the new file on Amazon S3:

```
byte[] bytes = sw.toString().getBytes(StandardCharsets.UTF_8);
InputStream is = new ByteArrayInputStream(bytes);
ObjectMetadata meta = new ObjectMetadata();
meta.setContentLength(bytes.length);
meta.setContentType("text/plain");

// Create new file on S3
AmazonS3 s3Client = AmazonS3ClientBuilder.standard().build();
s3Client.putObject(bucketname, filename, is, meta);
```

The response object of the CreateCSV service will populate the "value" attribute. The value will be a string that describes the CSV file which is created. Modify the response.setValue() to:

```
response.setValue("Bucket:" + bucketname + " filename:" + filename + " size:" +
bytes.length);
```

Once these changes are completed, compile the project.

2. Deploy CreateCSV Lambda Function

Next, using the AWS Management Console, create a new CreateCSV Lambda service. Refer to Tutorial #4 to review this procedure.

Be sure to specify the function handler:

```
lambda.CreateCSV::handleRequest
```

For Tutorial #5, choose only **_one_** method for invoking your Lambda functions: either via curl and the API Gateway (REST endpoints), or using the "aws lambda" CLI. ***It is not necessary to configure both methods.***

If choosing curl for function invocation, be sure to configure API Gateway URLs (POST) for your Lambdas. Refer to tutorial #4.

Next, modify the callservice.sh to call your new CreateCSV service.

3. Prepare to call CreateCSV Lambda Function: Create S3 Bucket

Begin by copying and adapting the callservice.sh script from tutorial 4. This is in SAAF/java_template/test/callservice.sh. Copy this script to the tutorial 5 project and redefine the input JSON object as follows:
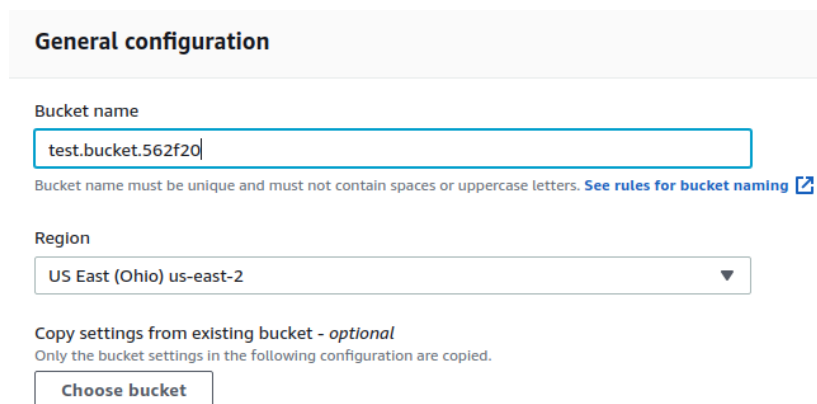
```
json={"\"row\"":50,"\"col\"":10,"\"bucketname\"":\"test.bucket.562f20.aaa\"","\"fi
lename\"":\"test.csv\""}
```

Next, create an S3 bucket named "test.bucket.562f20.aaa" on the AWS management console replacing "aaa" with your initials.

Navigate to the "S3" Simple Storage Service from the dropdown list of services, and then inside S3 select the button:

Create the bucket as follows:

S3 Bucketnames must be unique.
For the bucket name use: **test.bucket.562f20.aaa**

Replace "aaa" with your initials.
If the name is still not unique, modify as needed until it is unique.
Revise the bucket name above in the JSON accordingly.
Press the [NEXT] button.

For all remaining options in the Create Bucket wizard, accept the default values.

Next, it is necessary to configure permissions for Lambda to access your S3 bucket.
In the AWS Management Console, navigate to Lambda, and inspect your
"Basic Settings" for the CreateCSV Lambda function. Scroll down to **Execution role**:

Execution role
Choose a role that defines the permissions of your function. To create a custom role, go to the **IAM console**.

🔘 Use an existing role

⚪ Create a new role from AWS policy templates

Existing role
Choose an existing role that you've created to be used with this Lambda function. The role must have permission to upload logs to Amazon
CloudWatch Logs.

| service-role/createCsv_562-role-ml6mtzud | ▼ | C |

View the createCsv_562-role-ml6mtzud role on the IAM console.

Below the existing role, Lambda provides a BLUE link to open your security role called:
"**View the createCsv_562-role-ml6mtzud role** on the IAM console". Note this is a
dynamically generated security role name. Click on this link to navigate to your role to
grant Lambda permission to access your S3 bucket.

Without the helpful link, this role can also be by going to the upper right-hand corner,
selecting your name, and in the drop-down choosing "My Security Credentials". Then on
the left-hand side go to "Roles", and find and click on the security role name.

Once you're looking at the Role Summary, select the button:

Attach policies

Search for the policy: "**AmazonS3FullAccess**" and select the button below:

Attach policy

The policy should then appear as added to the Role.
This grants any Lambda function with the Role permission to work with S3.
More fine grained security policies can be specified as needed.

4. Test your CreateCSV Lambda Function

To reduce/change verbosity (the number of metrics returned from SAAF) feel free to replace the call to **inspector.inspectAll();** with another option, or comment out **inspector.inspectAll(); and inspectAllDeltas();** for development purposes right now. See tutorial #4 – page 21 for a list of inspect methods.

Now, using the `callservice.sh` script, invoke your Lambda function.
Try creating different sizes of CSV files by increasing or decreasing the values for "row" and "col". Please note, for creating large CSV files, it may be necessary to increase timeout values in the API-Gateway and/or Lambda as creating large CSVs is slow.

```
$ ./callservice.sh
{"row":50,"col":10,"bucketname":"test.bucket.562f20.aaa","filename":"test.csv"}
Invoking Lambda CreateCSV function using API Gateway

real 0m10.662s
user 0m0.092s
sys 0m0.008s

CURL RESULT:
{"value":"Bucket: test.bucket.562f20.aaa filename:test.csv size:1938","uuid":
"eaa34121-d5fd-43b4-a19c-ebc381db5c56","error":"","vmuptime":1540528993,
"newcontainer":1}
```

Next, verify that the CSV file has been created in your S3 bucket.

First try this using the AWS CLI. Try out the following commands. Adjust bucketnames as needed:

```
$ aws s3 ls test.bucket.562f20.aaa
2018-10-25 21:44:42       1938 test.csv

$ aws s3 ls s3://test.bucket.562f20.aaa
2018-10-25 21:44:42       1938 test.csv

$ aws s3 cp s3://test.bucket.562f20.aaa/test.csv .
download: s3://test.bucket.562f20.aaa/test.csv to ./test.csv

$ cat test.csv
38,869,146,8,578,793,8,713,581,259
49,994,324,882,412,287,402,428,401,922
971,584,184,972,717,611,14,660,978,867
...
```

Note that "s3://test.bucket.562f20.aaa/test.csv" is considered a URI or a Uniform Resource Identifier which is analogous (similar to) a URL.

The "aws s3 ls" command doesn't require "s3://", while "aws s3 cp" does.
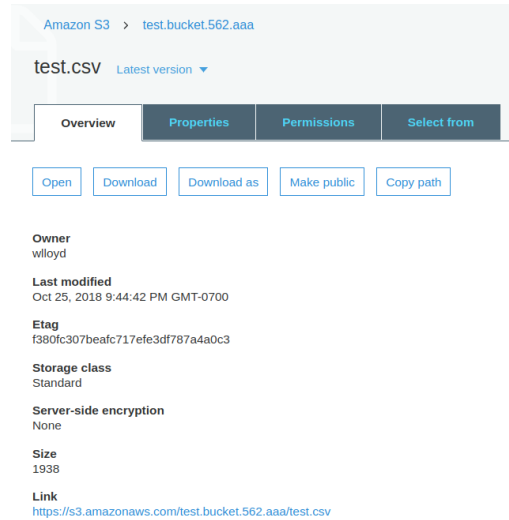
Next, using the S3 GUI, inspect your bucket and verify that the test.csv file exists.

Click your bucket name in the GUI.

Then click on the filename:

It is possible to download the file here, but without allowing public access to your bucket (not recommended) the **web link** at the bottom does not work:

5. Create ProcessCSV Lambda Function

Next, make a copy of the "CreateCSV" class called "ProcessCSV".
If you're using Netbeans, right click on "CreateCSV" and select "Refactor | Copy".
Or alternatively press "ALT-C".

A refactor popup appears:
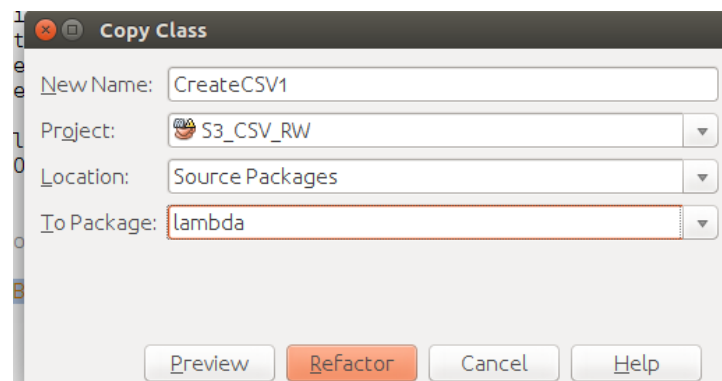
Rename the class to "ProcessCSV".

If not using Netbeans, copy "CreateCSV" and rename to create a new class called "ProcessCSV". Adapt the Lambda function to read a CSV file called "filename" from the S3 called "bucketname".

Adapt the example code provided from the URL and in the box below to read a file from S3 line-by-line. URL: https://blog.webnersolutions.com/use-aws-lambda-function

Here is the most relevant sample code for this activity:

```
AmazonS3 s3Client = AmazonS3ClientBuilder.standard().build();

//get object file using source bucket and srcKey name
S3Object s3Object = s3Client.getObject(new GetObjectRequest(srcBucket, srcKey));

//get content of the file
InputStream objectData = s3Object.getObjectContent();

//scanning data line by line
String textToUpload = "";
Scanner scanner = new Scanner(objectData);
while (scanner.hasNext()) {
     textToUpload += scanner.nextLine();
}
scanner.close();
```

For the s3Client.getObject() request, you'll need to provide the bucketname and filename to the API call. Add properties to the Request class for bucketname and filename. Be sure to add getters/setters for bucketname (getBucketname(), putBucketname) and filename (getFilename(), putFilename()) in the Request class. Then use these get methods to obtain the bucketname and filename and provides these to s3Client.getObject().

The ProcessCSV Lambda function should add all numbers from the CSV file to calculate the average value and total value of all elements.

Read each line of the CSV file to parse each individual comma-separated value.
Add the total of all values using a Java "long" primitive variable:

```
long total;
```

Track the total number of elements processed.

At the end, use a Java "double" primitive, to calculate the average value for all elements in the entire CSV file:

```
double avg;
```

Add a logging statement to print the average value to the AWS Lambda log:

```
LambdaLogger logger = context.getLogger();
logger.log("ProcessCSV bucketname:" + bucketname + " filename:" + filename + "
avg-element:" + avg + " total:" + total);
```

Finally, adjust the "value" property of the response object:

```
r.setValue("Bucket: " + bucketname + " filename:" + filename + " processed.");
```

Note:
AWS Lambda functions automatically receive permission to write out the CloudWatch logs.

This permission is included in the **AWSLambdaBasicExecutionRole** policy that is automatically created when you create the Lambda function.

Policies for the role assigned to your function can be viewed by using the link at the bottom of "**Basic Settings**". For example, it may be called: "**View the createCsv_562-role-ml6mtzud role** on the IAM console." at the bottom of Basic Settings. Click this link to inspect security policies assigned to your function. By clicking on a policy and clicking the "**{ } JSON**" button, a JSON description of the policy appears. An "**Edit policy**" button allows the JSON to be changed. Permissions are highlighted in the JSON:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "logs:CreateLogGroup",
            "Resource": "arn:aws:logs:us-east-2:465394327572:*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "logs:CreateLogStream",
                "logs:PutLogEvents"
            ],
            "Resource": [
                "arn:aws:logs:us-east-2:465394327572:log-group:/aws/lambda/createCsv_562:*"
            ]
        }
    ]
}
```

Alternatively, you could add the **CloudWatchEventsFullAccess** Policy to your Role. Refer to how we previously added the **AmazonS3FullAccess** policy. Adding this policy adds more permissions than needed and is less secure and is not recommended.

Now, compile the project.
This will create a JAR file with both Lambda functions.
From this JAR file, a switchboard **could be** implemented by having the Handler method call different code based on the Request.java inputs within the same package.

Deploy the new Lambda function using the jar file, and specify the new handler for ProcessCSV:

Handler   Info

lambda.ProcessCSV::handleRequest

6. Automatically Trigger ProcessCSV when CreateCSV creates a file in S3

Next, we'll create a CloudWatch Event rule to fire the ProcessCSV Lambda function to run whenever a specific file is placed into the S3 bucket.

In the AWS Management Console, search for the "CloudTrail" service.
CloudTrail is considered a Management & Governance Tool:

Management & Governance
AWS Organizations
CloudWatch
AWS Auto Scaling
CloudFormation
CloudTrail
Config
OpsWorks
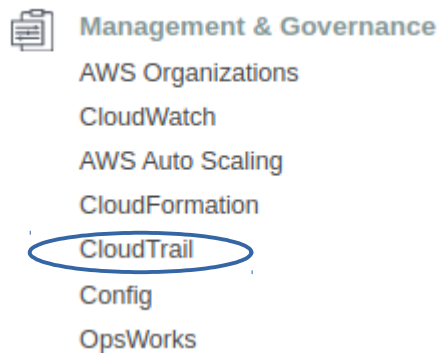
On the left-hand slide, select "**Trails**":

Dashboard

Event history

Insights

Trails

And then click the button:

Create trail

Create a trail as follows:

Create Trail
Trail name: **s3_1**

Storage Location
                ** *Create a new independent bucket for logging* **
Select "Create new S3 bucket" (*select radio button*)

Trail log bucket and folder:

Can just use *the automatically generated unique name that is provided*, *or provide your own name.*

Log file SSE-KMS encryption: *uncheck this – this should be disabled*

Then click the **NEXT** button.

On the "Choose log events" screen, first uncheck **Management events**.

Then select "Data events".

In the Data events dialog box, configure an S3 Data Event by filling the "Data event: S3" form.

Start by Unchecking: "All current and future S3 buckets" !
We don't want to create a trigger for ALL buckets !
Just the bucket we just created where CSV files are published.

Now fill in the "Individual bucket selection" to associate a CloudTrail with an S3 data event (e.g. *adding files to your bucket*).

For the bucket name, **Browse** to find your bucket, or type in your bucket's unique name: test.bucket.562f20.aaa

Uncheck Read. We only want to call a Lambda function when files are written to the bucket.

Then click the **NEXT** button.

Review the settings, and the create the trail:

Next, navigate to the "**CloudWatch**" service, also a Management tool:

On the left-hand side, near Events, select "Rules":
Click the "Create Rule" button:

Create rule

Then configure a rule as follows:

(X) Event Pattern (*select this*)
Service Name: Simple Storage Service (S3)
Event Type: Object Level Operations

The message notifying about the requirement to configure CloudTrail should appear:

AWS API Call Events sent by CloudTrail will only match your rules if
you have trail(s) (optionally with event selectors) configured to
received those events. See CloudTrail for further details.

(X) Specific Operation(s) (*select this*)
In the drop down list, select the **PutObject** operation (*search to find "PutObject"*)
Specific bucket(s) by name:
test.bucket.562f20.aaa (*replace with your custom bucket name for Lambda)*
        ** *this is the Lambda bucket, not the logging bucket from immediately above* **

Next, on the right-hand side, click the "Add Target" button:

⊕ Add target*

Select: Lambda Function
Function*: ProcessCSV
> Configure Input:
(X) Constant (JSON text)   (*select this*)
Provide JSON:

```
{"bucketname":"test.bucket.562.aaa","filename":"test.csv"}
```

Once everything is configured and ready, scroll down and in the lower right click:

Configure details

For Step 2, provide a rule name:

Name*: processCSV_rule

Then complete creation of the CloudWatch rule by pressing:

Create rule

13

Now, run your callservice.sh script again to call createCSV.

When createCSV finishes, the creation of a file in your S3 bucket automatically triggers the processCSV Lambda function to process the file !! The processCSV results will be written to the CloudWatch log.

7. Submitting the tutorial

To submit the tutorial, submit the results of the **ProcessCSV log file,** when **calling CreateCSV** to generate a 50x10 CSV file (500 total elements, 50 rows by 10 columns) with average elements selected randomly by the code provided above ranging from (1..1000). In the AWS Management Console, navigate to AWS Lambda.

Go to your "ProcessCSV" function.

Click the "Monitoring" tab:
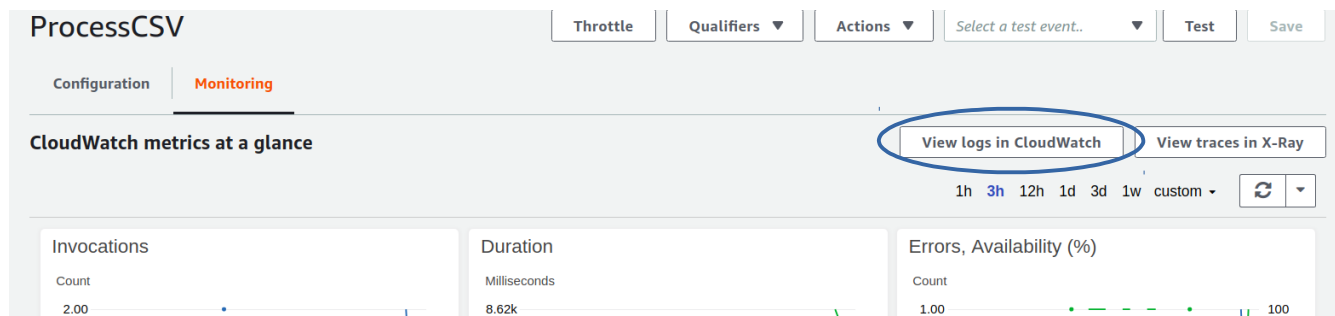
Then click the [View logs in CloudWatch] button.

Look through the log entries.
The top row should be your log entry where CreateCSV triggered the execution of ProcessCSV.
Click on this entry.
Now, CAPTURE THE SCREEN:

Expand all ● Row ○ Text

Time (UTC +00:00)	Message
2018-10-26	
	No older events found at the moment. Retry.
▸ 08:15:09	START RequestId: 4083c7e4-d8f7-11e8-ab1a-a55a35987a02 Version: $LATEST
▸ 08:15:16	ProcessCSV
▸ 08:15:16	END RequestId: 4083c7e4-d8f7-11e8-ab1a-a55a35987a02
▸ 08:15:16	REPORT RequestId: 4083c7e4-d8f7-11e8-ab1a-a55a35987a02 Duration: 7070.67 ms Billed Duration: 7100 ms Memory Size: 640 MB Max Memory U
	No newer events found at the moment. Retry.

Using the CTRL-PrintScreen button to capture the entire screen,
or use CTRL-SHIFT-PrintScreen to draw a box around the relevant section of the screen to copy to the Clipboard an image of your log with the avg-element and total values for ProcessCSV.

In OpenOffice, Microsoft Word, or Google Docs, paste this image into a document.

Create a PDF file of the document, and submit this PDF file to Canvas.

Optional Tutorial Activity – Retrieve Bucket Name and Filename Dynamically

One shortcoming of the CloudWatch Event target here was that we used a hard-coded value for the bucket name and the file name to pass to our ProcessCSV Lambda function. Ideally, we would like ProcessCSV to know the name of any new files added to S3 so our ProcessCSV will dynamically process the new file, and not a statically named one.

When the CloudWatch rule invokes the Lambda function, it is possible to pass the event object instead of a hard coded object. The event object contains the name of the S3 filename and bucket that was created to fire the trigger.
First modify the CloudWatch Rule. In "CloudWatch", go to "Rules" on the left-hand side. Select your rule. In the upper right-hand corner, select "Actions" and "Edit".

For the Target on the right, select "Configure input".
Change the selection to "**Matched Event**".
Click on the BLUE Configure Details button in the bottom-right of the screen.
Click on the BLUE Update Rules button to update the rule.

Selecting "Matched Event" will pass a JSON object describing the CloudWatch event that invoked your Lambda function to your Lambda Function.

Now, you'll need to modify your Java code to consume the bucket name and file name of the trigger.

Here is an example of the JSON object describing the CloudWatch rule that triggered the Lambda function as it is passed to the Lambda Function. The bucket parameters are highlighted:

{

15

```
    "detail-type": "AWS API Call via CloudTrail",
    "resources": [],
    "id": "8b933460-3063-7934-6f1c-d3835a33c987",
    "source": "aws.s3",
    "time": "2020-11-05T07:40:46Z",
    "detail": {
        "eventVersion": "1.07",
        "userIdentity": {
            "type": "AssumedRole",
            "principalId": "AROAWYW5J4AKBJ5W4ZKCI:createCsv_562",
            "arn": "arn:aws:sts::465394327572:assumed-role/createCsv_562-role-ml6mtzud/createCsv_562",
            "accountId": "465394327572",
            "accessKeyId": "ASIAWYW5J4AKOSEC7K6L",
            "sessionContext": {
                "sessionIssuer": {
                    "type": "Role",
                    "principalId": "AROAWYW5J4AKBJ5W4ZKCI",
                    "arn": "arn:aws:iam::465394327572:role/service-role/createCsv_562-role-ml6mtzud",
                    "accountId": "465394327572",
                    "userName": "createCsv_562-role-ml6mtzud"
                },
                "attributes": {
                    "creationDate": "2020-11-05T07:40:35Z",
                    "mfaAuthenticated": "false"
                }
            }
        },
        "eventTime": "2020-11-05T07:40:46Z",
        "eventSource": "s3.amazonaws.com",
        "eventName": "PutObject",
        "awsRegion": "us-east-2",
        "sourceIPAddress": "52.14.240.139",
        "userAgent": "[aws-sdk-java/1.11.306 Linux/4.14.193-110.317.amzn2.x86_64 OpenJDK_64-Bit_Server_VM/11.0.8+10-
LTS java/11.0.8 exec-env/AWS_Lambda_java11]",
        "requestParameters": {
            "bucketName": "test.bucket.562f20",
            "Host": "test.bucket.562f20.s3.us-east-2.amazonaws.com",
            "key": "test2.csv"
        },
        "additionalEventData": {
            "SignatureVersion": "SigV4",
            "CipherSuite": "ECDHE-RSA-AES128-GCM-SHA256",
            "bytesTransferredIn": 1943,
            "AuthenticationMethod": "AuthHeader",
            "x-amz-id-2": "JFVb8a4CFLjSITBcaFyIkqIZzW4X05dGL+JSzqFrFnpwICPBsSTLZy1VoTzJDOy/Y2ISGY9bGJw=",
            "bytesTransferredOut": 0
        },
        "requestID": "3B829E039B69C38B",
        "eventID": "e1e08ee7-f13d-49fd-a4a1-651208c41f99",
        "readOnly": false,
        "resources": [
            {
                "type": "AWS::S3::Object",
                "ARN": "arn:aws:s3:::test.bucket.562f20/test2.csv"
            },
            {
                "accountId": "465394327572",
                "type": "AWS::S3::Bucket",
                "ARN": "arn:aws:s3:::test.bucket.562f20"
            }
        ],
        "eventType": "AwsApiCall",
        "managementEvent": false,
        "recipientAccountId": "465394327572",
        "eventCategory": "Data"
    },
    "region": "us-east-2",
    "version": "0",
    "account": "465394327572"
}
```

The properties of interest are:

detail.requestParameters.bucketName
detail.requestParameters.key

The following source code demonstrates acquiring the bucketName and filename and printing them to the log file. This code can be modified to suit your purposes:

```
public HashMap<String, Object> handleRequest(HashMap<String, Object> request, Context context) {


//Collect initial data.
Inspector inspector = new Inspector();
inspector.inspectAll();

//*****************START FUNCTION IMPLEMENTATION*************************


Gson gson = new GsonBuilder().setPrettyPrinting().create();
LambdaLogger logger = context.getLogger();
logger.log ("bucketName=" + ((HashMap)((HashMap)request.get("detail")).get("requestParameters")).get("bucketName"));
logger.log ("filename=" + ((HashMap)((HashMap)request.get("detail")).get("requestParameters")).get("key"));
```

The event is provided as a nested HashMap object in Java. Each layer of JSON is a HashMap. Accessing the bucket parameters involves using the **detail** HashMap to access the **requestParameters** HashMap to read the **bucketName** and **key** fields.

Customize the sample code as needed to create a customized event handler that does not use a hard coded bucketname and filename.