

Assignment 3

FIT5225 2024 SM1

PixTag: An AWS-powered Serverless Image Storage System with Advanced Tagging Capabilities

1 Synopsis

This assignment aims to build a AWS-based online system that allows users to store and retrieve images based on auto-generated and manual tags. The focus of this project is to design a serverless application that enables clients to upload their images to public cloud storage. Upon image upload, the application automatically tags the image with the objects detected in it, such as person, car, etc. Later on, users can also modify and add new tags to the images. They can also query images based on the objects present in them. To achieve this, the application provides users with a list of image URLs (or thumbnails) that include the specific queried objects.

2 Group Assignment

The majority of software developers will work in a team at some point in their career. Therefore, to prepare students to experience software development as a team, this assignment is designed as a group project. Students will be placed into software teams of up to four members to work on implementing the system described above. In this assignment, students need to organize their team and their collective involvement. There is no designated team leader, but teams may decide to establish processes for agreeing on work distribution. Understanding the dependencies between individual efforts and their successful integration is crucial to the success of the work and for software engineering projects in general. If teams encounter any “issues”, they should inform the teaching team as soon as possible, and help will be provided to resolve them. The due date for this assignment is firm, and there is no ‘**special consideration**’ for granting an extension. If a team member fails to complete their work, other team members should handle their tasks and report this to the teaching team as soon as possible. We will evaluate individual involvement in the project, and if a team member fails to satisfy their tasks within the team, they may be penalized heavily, regardless of the project’s success.

3 Assignment Description

Teams should develop an AWS-based solution that leverages services such as S3, Lambda, API Gateway, and database services (e.g., DynamoDB) to build a system for automated object detection tagging and query handling. The teams should produce a solution that enables end-users to upload their images into an S3 bucket. Upon uploading an image to a designated S3 bucket, a lambda function is automatically triggered, which uses the Yolo object detection feature to identify the objects in the image and stores the list of detected objects along with the image’s S3 URL in a database. Furthermore, the end-user should be able to submit queries to an API endpoint using API Gateway to search for tagged images (more details to come). Table 1 provides a glossary of terms used in the assignment description.

Table 1: Glossary of terms

Term	Meaning
team	A group of 4 students who are doing the project
queries	You will implement multiple APIs to query different information including a request message by the end user to find list of images with specific tags, find images with similar tags of an image, update tags of an image, and delete record of an image.
tags	list of objects detected in an image, e.g., person, car, and cat.
federated authentication	Federated authentication is a Single sign-on (SSO) mechanism that allows you to use authentication credentials of external identity providers such as Google or Facebook to access your AWS services such as Lambda, S3 or DynamoDB.
Peer assessment	Each team member is expected to complete a confidential report and evaluation on his/her role in the project and the experiences in working with other team members on the demonstration day.

3.1 Authentication and Authorisation

Security is one of the most crucial aspects of developing cloud applications. When your application is publicly exposed, you must ensure that your endpoints and resources are safeguarded against unauthorized access and malicious requests. AWS, through its Cognito service, provides a straightforward, secure, and centralized approach to protect your web application and its various resources from unauthorized access.

To leverage the AWS Cognito service, first, you need to create a user pool that stores user credentials. Then, you need to create and configure a client app that provides access to your application and/or other AWS services to query and use the user pool. Finally, you have options to communicate with the AWS Cognito service and perform authentication and/or authorization. You can use the AWS JavaScript SDK to access the user pool and identity provider(s) that you have defined earlier. Please note that AWS Amplify JavaScript Library to initialize the authentication module of your application might be an option. However, if you use your AWS academy account, AWS Amplify is not available. For further info on this you can check forums discussions.

Your application should support the following workflow and features:

- Detect whether a user is authenticated or not. If the user has not signed in, access to all pages/endpoints except the sign-up service needs to be blocked, and the user should be redirected to the “sign-up.html” page to register a new account. For each new account, you need to record the user’s *email address*, *first name*, *last name*, and *password*. Cognito will take care of sending an email to new users, asking them to verify their email address and change their temporary password.
- Your application should include a login page that allows users to sign in. After successful authentication, users should be able to upload images, submit queries, view query results, and sign out of the application. All of these services must be protected against unauthorized access. You can implement login and sign-up web pages using either the Hosted UI feature of Cognito or your own version that calls Cognito APIs.
- Uploading files to an S3 bucket, invoking Lambda functions to execute the business logic of your

application, and accessing the database for data storage and retrieval all require fine-grained access control permissions that you need to set up via IAM roles and appropriate policies. It is important to note that IAM roles in AWS Academy have several limitations. Therefore, you should carefully consider how to perform authentication and authorization in your system while taking these limitations into account.

As an **optional** feature, you can add federated authentication using AWS Cognito Identity Pools to your application and earn bonus marks (up to 5 points shall be awarded if you add federated authentication to your project). For this purpose, you need to create a Facebook or Google app that serves as an external identity provider and authenticates users on behalf of your application, then forwards authentication tokens to your application. Note that having external authentication providers in your project is **not** mandatory. Since federated authentication might be challenging and maybe impossible with your AWS academy account, I **strongly** recommend that you finish the requirements of the assignment first, and then, if you have extra time, work on this feature.

3.2 Image Upload

Your solution should provide a mechanism to upload an image to an S3 bucket. Uploading an image to an S3 bucket can be done either through an API Gateway endpoint (using POST REST APIs) or it can be done directly using AWS SDKs (for instance, boto3 if you are using Python). Whenever an image is uploaded to the S3 bucket, your system must trigger an event and invoke a Lambda function.

Building Images Thumbnails: Upon uploading an image to S3, a Lambda function should make a thumbnail for that image. Creating a thumbnail for an image involves resizing the original image to a smaller dimension while maintaining its aspect ratio, followed by a compression to reduce file size. You can use OpenCV library for this purpose. You are expected to configure the triggers and grant the required Amazon resource permissions (execution roles) for the Lambda function.

Another Lambda function should be triggered to read the uploaded image, detect objects in the image, and save the list of detected objects (called tags) along with the S3 URLs for that image and its thumbnail in an AWS database for future queries. Please note that you should update your Yolo script that you were given in your **first assignment** to suit the AWS Lambda function. This includes removing any Flask-related code, incorporating the Lambda function and required libraries to read the image from S3 buckets, and storing the S3 URLs and tags in the database. You may create a separate S3 bucket to store the Yolo and other config files, which can be referenced in your Lambda function.

You may also ignore detected objects with an accuracy of detection below a specific threshold (e.g., 0.6). Please note that **Amazon Rekognition** is a specialised AWS service for identifying objects in images. However, in this assignment, you are **NOT** allowed to use the Amazon Rekognition service.

3.3 Queries

Your solution should provide APIs which allow following queries.

1. **Find images based on the tags:** The user can send a text-based query to request URLs of images that contain specific tags with a minimum repetition number for each tag (e.g., “person, 1”, “person, 2, car, 1”). You are expected to create an API Gateway with a RESTful API that allows users to submit their requests, such as GET or POST requests.

Your application might send a list of tags via specific GET parameters in the requested URL, for example:

```
https://jyufwbyv84.execute-api.us-east-1.amazonaws.com/dev/search?tag1=cat&tag2=car
```

or it can be a POST request with a JSON object including a list of tags and their counts. A sample JSON object for a query request is given below:

```
{
  "tags": [
    "person",
    "desk",
    ...
    "cat"
  ]
}
```

A response should include the list of URLs to all thumbnail images that contain **all** the requested tags in the query. This can be a JSON message similar to the following:

```
{
  "links": [
    "https://pixtag.s3.amazonaws.com/image1-thumb.png",
    "https://pixtag.s3.amazonaws.com/image59-thumb.png",
    "https://pixtag.s3.amazonaws.com/image180-thumb.png"
  ]
}
```

Your query may require to trigger one more Lambda function that receives a list of tags and finds s3-url of thumbnail images containing **all tags in the query** from the database, i.e., logical **AND** operation, not **OR** operation between tags.

Please note that duplicate tags should not affect the results.

For example: {"tags": ["person", "person"]} results in the same response as {"tags": ["person"]}.

In your UI instead of showing the s3-urls of thumbnails you can preview the thumbnails found as the results. The user can request the full size image by clicking on the thumbnails or sending another query to get the full size image.

2. **Find images based on the thumbnail's url:** Users can input the s3-url of a thumbnail into the system and receive the corresponding full-size image s3-url.
3. **Find images based on the tags of an image:** The user can send an image as part of an API call. The list of all objects (tags) in the sent image is discovered and then all the images in *PixTag* storage containing those set of tags are found. Finally, as a response, the list of s3-URLs of the matching images thumbnails (similar to to the previous section) are returned to the user. You should make sure that the image uploaded for the query purpose is not added to the database or stored in s3.
4. **Manual addition or removal of tags with bulk tagging:** Your solution should also provide an API that allow end-user to add or remove tags of images. You are expected to create a POST API which allow users to submit their requests.

A sample JSON message sent to add/remove tags is:

```
{
  "url": [
    "https://pixtag.s3.amazonaws.com/image1-thumb.png",

```

```

        "https://pixtag.s3.amazonaws.com/image60-thumb.png",
        "https://pixtag.s3.amazonaws.com/image23-thumb.png",
    ]
    "type": 1, /* 1 for add and 0 for remove */
    "tags": [
        "person",
        "alex"
    ]
}

```

“type” can be set to 1 or 0 for adding or removing a tag, respectively. The above request adds “person” and “alex” tags to the tag list of the images in the list of URLs. If “type” is set to 0, the tags are removed from the tag list of the image. If a tag is not included in the list of tags requested for deletion, you can simply ignore it in the request.

5. **Delete images:** The user can communicate the URL list of image thumbnails to an API and the system should remove the images and their thumbnails from s3, and all relevant entries from the database.

3.4 Tag-based Notifications

Users can receive email notifications or updates related to images based on specific tags. For example, users can receive notifications when new images with certain tags are added or updated to the system. You can use AWS SNS for this purpose.

3.5 User Interface

You can design a simple user interface (UI) (We suggest web-based GUI). But UI can be of any form that includes the following: upload images, submit queries, and view query results. A UI makes your system easier and more enjoyable to use. However, you have the option not to create a UI for application or have UI for some parts and not the others. If you opt to ignore UI or full-fledged UI, you can write scripts (e.g., Python) to handle communications with your application APIs. Please be aware that you might need to manually copy credentials provided by the identity pool to your scripts each time if you choose to work with the later option.

4 Final Reports

You should write an **individual** and **team** report on the developed application. You should use a highly-legible font type such as: *Arial*, *Helvetica*, and *Times New Roman* with font size 12 points.

4.1 Team Report - 750 words

Please prepare a team report based on the following guidelines:

- Include an architecture diagram in the team report. For the architecture diagram, you must use AWS Architecture Icons (more info can be found here: <https://aws.amazon.com/architecture/icons/>).
- Include a table to describe the role of each team member in your team report, You can provide a three-column table in your report that shows student name and id, percentage of contribution and elements of the project the member contributed (maximum of 100 words per member).

- Report includes a simple user guide for testing your application (you should keep your application up and running for two weeks after the submission deadline).
- Your report should include a link to the source code (GitHub or Bitbucket). All students must commit their code to a private code repository rather than delegate this to a single team member. This can also provide an evidence base if teams run into “issues” and how you contributed to the project. You should share your private repository with the whole teaching team. Please do not use a public repository to avoid any **plagiarism**.

Note that the tables, architecture diagram, screen shots of UI, and references are excluded from the word count limit.

4.2 Individual report - 750 words

Please prepare an individual report based on the following guidelines:

- In your individual report, you should describe your role in the project, and how you contributed to the delivery of the system (150 words).
- Include a personal reflection on the team work your experience in working with other individual team members in your group. Comment on your teammates’ work, challenges the entire team faced, and how well the team worked (150 words).
- You should answer the following questions (a paragraph with 100 word each):
 1. What sort of design changes you will make to reduce chance of failures in your application?
 2. What are the design changes you will make to increase the performance of your applications in terms of response time, query handling, and loading images?
 3. What are the updates you will perform to your application if you have users from all over the world?

Note that individual report should be done with little or no consultation with other team members. Please be careful about plagiarism and collusion.

5 Submission

You need to submit the following via Moodle:

- Your team report in **PDF format**. **Only one of the team members submit this report**. Please do not forget to include your team members’ name and student ids. One submission per team!
- Your individual report in **PDF format**. **Every team members should submit this report**. Please do not forget to include your name and student ids.

6 Demonstration Schedule, Venue and Peer Assessment

Teams are required to give a demonstration of the working application via Zoom. All team members are required to be present at the demonstration time slot to answer questions. The presentation should start with the architecture of your application as per your report along with the design and implementation choices made. Each team has up to 3 minutes to present their architecture. This is followed by the demo of application as per examiners’ guideline for the next 10-15 minutes and Q&A. We will announce the

demo date and time closer to the due date. We will ask you to select a time slot during the working hours to join a Zoom meeting.

As it is mentioned earlier. This is a group assignment. If one of the team members is experiencing difficulties or has special circumstances, the rest of the team should proceed with the project in the absence of their fellow member. This aligns with industry practices where the entire project cannot be delayed due to a team member's departure or absence.

At the end of the session, teams will participate in peer assessments regarding their fellow members' contributions and involvement in the project. This feedback may be gathered through a Google form. Each team member is required to submit a confidential report detailing their individual role in the project and their experiences collaborating with team members. These reports will be treated with utmost confidentiality, alongside your final report.

The purpose of this evaluation is not to assign blame, but rather to ensure that all team members have an opportunity to provide comments and to guarantee active participation from everyone. We encourage all team members to collaborate effectively towards achieving the project's objectives to optimize your team's performance. Should you observe a team member encountering difficulties with their tasks, please offer assistance. If a team member fails to fulfill their assigned tasks or deviates from the project's expectations, we urge you to promptly inform the teaching team. Alternatively, you can utilize this opportunity to bring the matter to our attention.