



UNIVERSITYof**ARKANSAS**

1871®

**University of Arkansas – CSCE Department
Capstone I – Final Proposal – Fall 2022**

Ponga Software Usage Tracking

Matthew Clemence, Dylan Vaughn, Ryan Drake, Julio Bonilla, Logan Reed

Abstract (Matthew Clemence)

In the world of software, the necessity of user data tracking is evident. Ponga wishes to nudge its users towards certain value propositions in order to increase their lifetime value. These value propositions comprise the number of uploaded photos, the number of shares, and the number of comments one user perpetrates. With a healthy mix of these three value propositions, a user's lifetime value would be considered satisfactory. The amount of time a user has existed also comes into play, but the numbers of their value propositions wouldn't reach a significant high without an extensive relationship with Ponga.

Our objective is to design a series of AWS Lambda functions and a database structure to handle the tracking of users. The data gathered on each user will be drawn on to assess their risk level as a customer and Ponga will act accordingly. To determine a user's risk level, we will perform calculations with a user's value proposition numbers. If the user's numbers are unsatisfactory in one area, then Ponga can push the user in a different direction. The significance of this problem lies in our unique solution to it. Our software is named Insight and will be unique in tracking user data by comparing value propositions over time via calculations involving n-dimensional geometry.

1.0 Problem (Matthew Clemence)

Ponga is a web application for photo identification, preserving family stories, and keeping track of people, places, and things. Ponga took off when they implemented a membership model on top of an already richly featured platform. Features include the ability to: highlight people and/or details in a photo with a zoom and click, attach a link to a photo, and play music alongside a photo. Ponga is for private interactions among families and not advertisers. Ponga is delving into the process of advertising themselves and that's where Insight comes in.

Every product that exists today embraces some sort of advertisement to bring in more customers. Ponga wants to not only bring in more customers but reinforce the ones they have. This can't be done without gathering appropriate data on their customers which portrays their use of the application over time. Ponga's problem can be demonstrated through a simple product such as Reese's. Reese's have two value propositions also interpreted as groups from which their customers come from. There is the group who enjoys chocolate, and the other group who enjoys peanut butter. For Reese's to be successful they must be able to increase the number of customers

who belong to both groups. Reese's problem is the same as Ponga's in that they need to know which customers lean towards one or more value propositions over others.

This is a problem every product on the market runs into, be it an application or candy. Not defining value propositions and grouping the customers by them can lead to stagnation in a product. Risky customers are those who do not fall into all groups of a product and as such are more likely to leave said product for another. In contrast, a healthy customer is one who falls into all groups and their respective scores in those groups are determined by the product. Without a deeper examination of their customers, Ponga can't influence risky customers into becoming healthy ones.

2.0 Objective (Julio Bonilla)

The objective of Insight is to design a solution that records customer activity, quantizes their activities, and generates meaningful data that can be used to analyze a customer's level of risk within an arbitrary window of time. Insight aims to track any change in customer behavior over time and advise Ponga about any potential customer risk, trends, or anomalies. Our design will involve a series of AWS Lambda functions coupled with DynamoDB as a database. One will retrieve data from a database table, and another will send data to a separate database table. One will convert Kinesis data streams to events. Lastly, one will calculate user attribute affinity based on the data sent and fill the table to be retrieved from. Our team aims to resolve the problem outlined above by instantiating APIs that are triggered by a user making any action on their account and organize that data so that Ponga can better manage that user.

The initial goal is not to weight value propositions, but to observe and collect user data; thereby allowing a thorough analysis of user patterns, preferences, and engagement levels. Using Amazon Kinesis Data streams, all data moving through our model can be timestamped and stored in each user record. The AWS platform will allow us to integrate services necessary to maintain a service level agreement prescribed by Ponga. The team will research the AWS platform and familiarize themselves with its functionalities. The team will also synthesize a solution to compare and quantify data points within an n th dimensional vector space. The team will remain in communication with the POC on a weekly basis to ensure we are operating within desired metrics.

3.0 Background

3.1 Key Concepts (Matthew Clemence, Dylan Vaughn, Logan Reed)

The meat and bones of Insight will be created using AWS (Amazon Web Services). AWS is a cloud provider that provides servers and services that can be used on demand and scales easily. AWS powers some of the biggest websites in the world, one being Netflix.

[6] AWS Lambda is one service we will be using. AWS Lambda introduces the idea of a serverless experience by employing virtual functions. Rather than being limited by RAM and CPU, these functions are limited by time. These functions are run on-demand, meaning there is no charge if the functions are not being used. When these functions are being used, we are billed by the 100ms. The most significant feature of AWS Lambda is automatic scaling. This means

that as many instances of a lambda function that we need can be run all at the same time. This service cuts out the need to manage the infrastructure of an application and allows us to focus solely on what the application does.

[2] DynamoDB is another amazon web service we will utilize. DynamoDB is a managed NoSQL database optimized for performance at scale. NoSQL stands for Not Only SQL. DynamoDB acts as a key/value lookup store, meaning values that are stored in the database are found through an API given a key. This is slightly different from the traditional SQL database with only rows, columns, and tables. A managed database is one with the storage, data, and compute services maintained by a third-party provider. DynamoDB can scale out horizontally by adding more nodes onto the cluster and separating data out onto those nodes. This mean that even if the size of the data grows, the performance will not see a major change. DynamoDB is a very safe service to use due to its high availability and durability. For availability, DynamoDB has a 99.99% guaranteed up time which translates to around 26 seconds of monthly down time. For durability, DynamoDB stores multiple copies of data across multiple nodes to avoid loss of data. One caveat of DynamoDB is that one should know how they want to access their data before using it, otherwise you may not be able to access data in the way you want when all is said and done.

[1] Amazon API Gateway is another service needed for Insight. With API Gateway, we can create, publish, maintain, monitor, and secure REST, HTTP, and WebSocket APIs at any scale. We will use Amazons API Gateway because it will easily integrate with other web services such as Lambda and DynamoDB. Our application will need to receive API calls from the Ponga application, and this is done through Amazon API Gateway.

[4] Amazon Kinesis allows for incoming data to be processed in real time without interrupting the database in any way that would prohibit other actions from being completed simultaneously. Using this, Insight would be able to analyze incoming data and trigger Lambda to begin working on manipulating data and sending it to a separate table. When the data arrives in the second table, Kinesis could trigger another action.

[3] Amazon EventBridge is a serverless event bus service that you can use to connect your applications with data from a variety of sources. EventBridge delivers a stream of real-time data from your applications, software as a service (SaaS) applications, and AWS services to targets such as AWS Lambda functions, HTTP invocation endpoints using API destinations, or event buses in other AWS accounts. We will use EventBridge to send a Kinesis data stream to multiple AWS Lambda functions.

[5] Amazon Simple Queue Service solves the problem of requests from the client to the application being lost due to the workload. Data sent and received between the client and API Gateway will be processed through SQS. This allows an application to take as many requests or data sent from the client as needed and process them when the API Gateway is available. This ensures that none of the client data is lost, or the API Gateway doesn't incur an error from an overload. The SQS allows for a manageable workload at any scale.

3.2 Related Work

Insight falls under the field of user data tracking applications. The most common ways to track user activity today include tools like Google Analytics and Search Console, click tracking, scroll tracking, and viewing session recordings of users. These methods all track user activity on a website or app in a rudimentary way that requires excess effort from the client.

[7] Google Analytics is a platform that collects data from websites and apps, then provides reports to the client. Setting up Google Analytics for yourself requires making an account and adding JavaScript measurement code to each page on your site. This tracks how many users visited a particular page or bought a certain item on a store page. It collects user browser information such as: language, type, device, and operating system. Once this data is collected by the JavaScript code, it is packaged and sent to Google Analytics to be processed into reports. These reports contain the aggregated data organized by user device and browser. The client can also filter the data sent to Google Analytics if they wish to protect private company data. The processed data is stored in a database where it can't be modified. This could be used to gather the data for our problem, but the client still needs to analyze it themselves.

[10] Click tracking is exactly what it sounds like. All user's clicks on a website are tracked and compiled into a report for the client. Clicks tracked include those on links and buttons. This is useful for website design and marketing. The reports contain events with click being its own event. The number of clicks per user is displayed, and each button/link has its own value. [9] Scroll tracking tracks how far a user scrolled down a web page. A trigger called the scroll depth trigger is used to track this. There are vertical and horizontal scroll depths. The client can enter percentages representing the page height or width to set trigger points on the page. The variables saved for future reference are scroll depth threshold (depth of trigger in percent), scroll depth units (pixels or percent), and scroll direction (vertical or horizontal). While our problem deals with marketing, we are tracking the use of core elements of an application which includes more than buttons and links.

[8] Session recordings are renderings of real actions taken by users as they browse a website. Recordings capture mouse movement, clicks, taps, and scrolling data for desktop and mobile devices. This kind of tracking can help improve the user experience of websites and applications. By watching session replays, the client can find patterns that explain the prevalence of risky users. Recordings also help catch any bugs causing users to get stuck, confused, and frustrated. The problem here is that the client must watch these recordings and spend time making sense of their customers actions. For a large-scale application, watching all the session recordings and putting together a trend would be more work than it's worth.

Overall, the problem with Google Analytics and session tracking is the amount of interaction required by the client and the nature of the processed data. Google Analytics just sorts the data for easier readability, but any trends need to be inferred by the client after receiving. Our application focuses only on the value propositions of an application/website and the user's trend with them. The client will provide as many value propositions as they want and connect the quantifiable actions taken to our API endpoints. Our processed data will determine a trend towards one or more value propositions. Once the client gets a read on their average customer, they can then add weights to the value propositions with triggers to notify the client when users

Ponga Data Tracking

enter a risky state. The session recordings would work for Ponga, but it is a rough solution to long-term problem. The difference between our data tracking and existing methods is the enhanced focus on value propositions and the use of N dimensional geometry for calculations on behalf of the client.

4.0 Approach/Design

4.1 Use Cases

Use Case: Submit Updated Metric

Author: Jerrel Fielder

Primary Actor: CLIENT application via *POST* method to API URL

Goal in Context: To provide an update to one or more user engagement metrics for tracking.

Preconditions: API should be available, and the call should be formatted correctly.

Trigger: The user takes a quantifiable action in the CLIENT application and the CLIENT invokes the API in order to log the action.

Scenario:

1. User takes quantifiable action in CLIENT application, such as uploads photographs
2. CLIENT counts the number of subject items in the user action
3. CLIENT invokes API URL providing:
 - a. User: <email>
 - b. JSON payload with an array of at least one key-value pairs:¹
 - i. Attribute ID: count
4. if the CLIENT passes a key for a customer that doesn't exist in the system, the system must create a record for that primary key and the data from CLIENT call becomes the T₀ data point.
5. SYSTEM uses current system time as the timestamp to be persisted with the event in the database

Exceptions:

1. Successful call returns HTTP status code 200
2. Incorrectly formatted API call from CLIENT returns HTTP status code 403
3. Unauthorized access attempt returns HTTP status code 404
4. Server error / down returns HTTP status code 500

Priority: High

Channel to actor: CLIENT application

Frequency of use: Up to many thousands of times per day.

Secondary actors:

Channels to secondary actors:

Open Issues:

¹ API should handle both positive and negative numbers for “count”, but the total on the user record should never be < 0

Use Case: Retrieve Historical Data For A User

Author: Jerrel Fielder

Primary Actor: CLIENT application via *GET* method to API URL

Goal in Context: To retrieve the historical attribution of a user.

Preconditions: API should be available, and the call should be formatted correctly.

Trigger: CLIENT application requires the historical attribute trajectory for internal processing and/or visualization. Results should always be presented in reverse chronological order (newest first, oldest last) and limited to the number of records specified in the OPTIONAL parameter "limit: n". If no "limit" parameter is included in the call, return all records in batches of 50 with a pagination ID in the return set.²

Scenario:

1. User is logged in and active in the CLIENT application
2. CLIENT needs raw historical data (visualization, etc.)
3. CLIENT invokes API URL providing:
 - a. User: <email>
 - b. OPTIONAL JSON payload specifying maximum array elements to return:
 - i. limit: n
4. Response Payload:
 - a. User: <email>
 - b. Records: <n records>

```
[
  {
    i. Timestamp: <created date>
    ii. Data: <Magnitude, All Angles>
  },
  ...
  N
]
```

Exceptions:

1. Successful call returns status code 200
2. Incorrectly formatted API call from CLIENT returns status code 403
3. Unauthorized access attempt returns status code 404
4. Server error / down returns status code 500

Priority: High

Channel to actor: CLIENT application

Frequency of use: Up to many times a day.

² <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Query.Pagination.html>

Use Case: Calculate Attribute Affinity And Persist

Author: Jerrel Fielder

Primary Actor: System

Goal in Context: To calculate the current attribution affinity for a given user.

Preconditions: None

Trigger: Upon the update of a user (via the "Submit updated metric" API), calculate the current attribute affinity and update the user record

Scenario:

1. Upon the update of a user (via the "Submit updated metric" API)³
2. Calculate the new resultant vector by adding T_n to T_{n-1}
3. Calculate the new state of the user affinity for T_n
4. If the T_{n-1} is different from the T_n state:
 - a. Generate webhook as described in the "Generate Webhook on Change Of State For A User" Use Case
 - b. Store the new state in a persistent store as the current state for the user

Exceptions:

1. Unsuccessful operation should be fully logged with:
 - a. User: <email>
 - b. 2 x n matrix containing all attribute values for T_{n-1} and T_n to
 - c. Timestamp of the failure
 - d. Any system or OS related error message
2. Log the webhook call and the entire payload

Priority: High

Channel to actor: System / Change Data Capture

Frequency of use: Up to many thousands of times per day.

Secondary actors:

Channels to secondary actors:

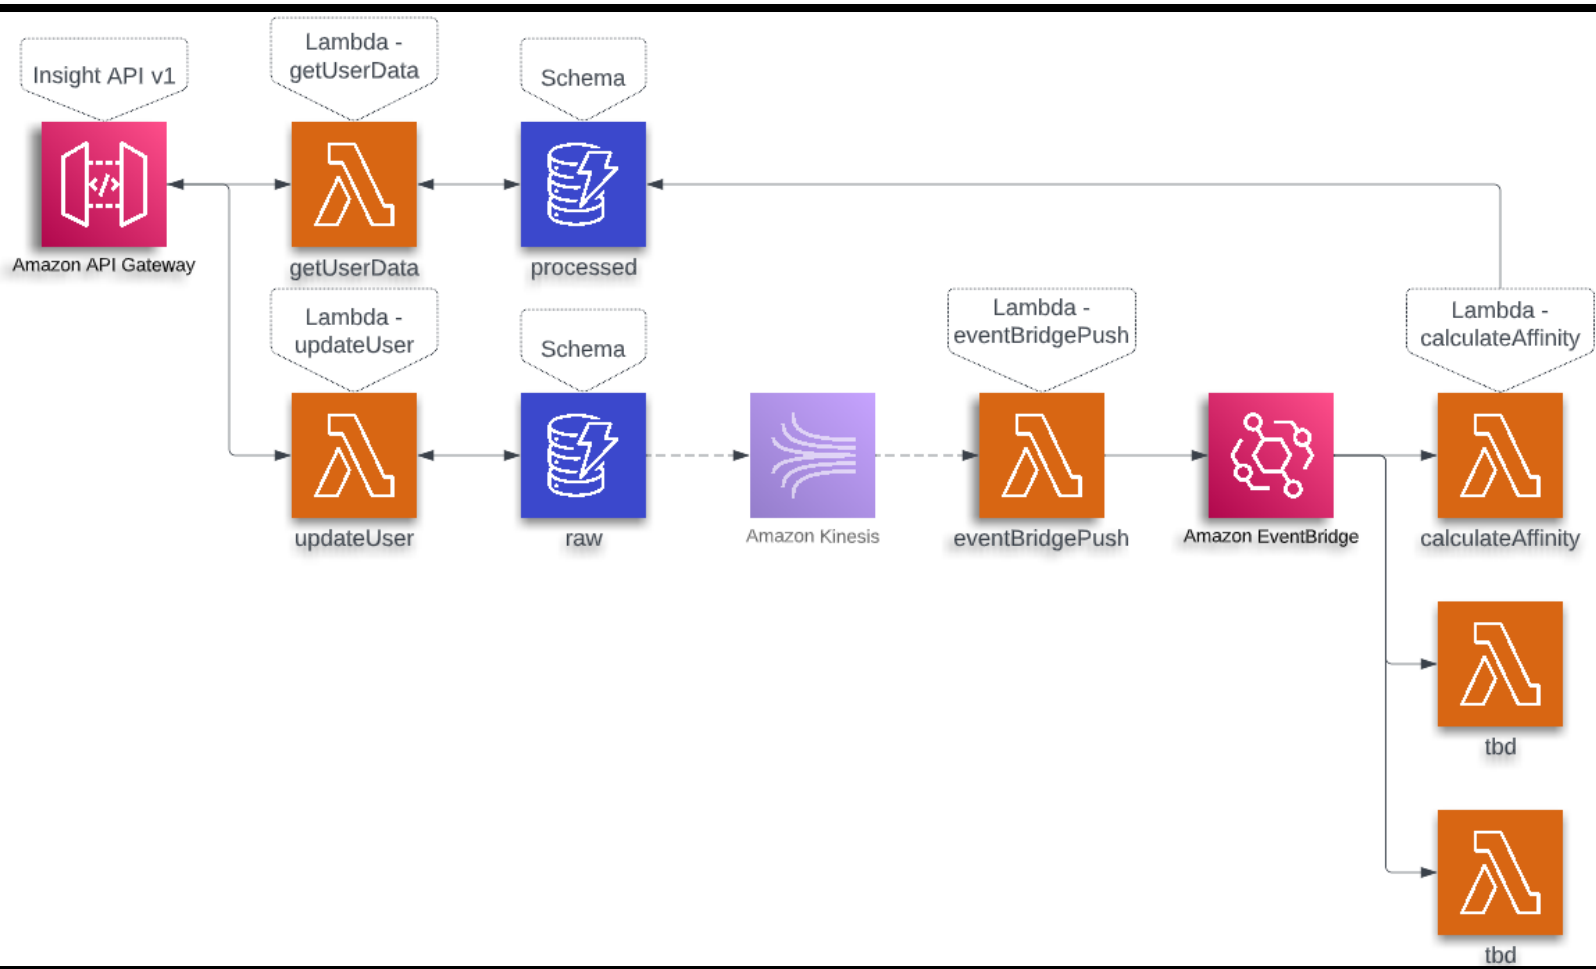
Open Issues:

³ Could be managed via a Change Data Capture trigger from the database. Implementation specifics to be discussed later.

4.2 Detailed Architecture (Julio Bonilla, Ryan Drake, Matthew Clemence)

This design is going to work on a platform of integrated web services hosted by Amazon. The client will call to the API Gateway depending on the use case. The API GET endpoint will trigger the `getUserData` lambda function, while the POST endpoint will trigger the `updateUser` lambda function to resolve the call. Our stretch goals for version 2 will be to integrate a queueing system for the calls that will be coupled with API Gateway and pushing the changed data event to additional lambda functions for more analysis. Our AWS Lambda functions will execute commands to access records from the database, calculate important data, and write both changes and current data to the database. The nature of the user-product model and access patterns is intricate; a data streaming service, Kinesis, will be used in conjunction with our design to monitor every event in our model.

Our design is structured around AWS services. Our APIs will handle calls from the client and have a high level of scalability. These services are as follows: API Gateway, Lambda, DynamoDB, Kinesis, EventBridge and SQS. Here is a diagram of how they will bridge together.



Flow Diagram for the Insight Version 1 Design

Ponga Data Tracking

These services flow directly from one to another; DynamoDB acting as the heart of our design while Lambda brings the services together. The first step in the data flow is a call to API Gateway from the client. A Ponga user will perform some sort of action triggering the client to send a properly formatted call to the respective endpoint. Insight API endpoints include a GET and POST. The GET endpoint receives query string parameters while the POST endpoint receives a JSON payload. Each endpoint launches an associated AWS Lambda function. These Lambdas process an event containing the JSON payload represented as a map. Insight has two maps, User and Record. The GET endpoint will take in a “User” parameter requiring the desired users email and a “Limit” parameter specifying the number of recent records to retrieve. Both are passed in as query string parameters of the GET request. If a request is made without a “Limit” parameter, you will receive all records in batches of 50 with a pagination ID in the return set. This array of user records is the User map. You cannot make a request without the “User” parameter. The POST endpoint will require a JSON payload named “Data”. Data is represented as a child map of the Record map. The Record map is what is provided by the client in the POST method.

User Map:

Field Value	Type	Description	Use
User	string	The unique email of the requested User. “Email”: “example@example.com”	Use this to retrieve a specific user programmatically.
Records	map	Holds data from records retrieved and their timestamp. “Records”: [{ “Timestamp”: “2022-07-03T23:12:08.000Z” “Data”: { “Magnitude”: 40, “Photos/Shares Angle”: 35, “Photos/Comments Angle”: 50, “Shares/Comments Angle”: 70 }, N Records }]	This field can be used to understand a User’s relevant data and analyze their change in data over time through multiple records.

Record Map:

Field Value	Type	Description	Use
User	string	The unique email of the new/existing User. "User": "example@example.com"	Use this to create a new user or update an existing one.
Data	map	Data that represents the value proposition quantity of a user. "Data": { "Photos": 4, "Shares": 60, "Comments": 20 }	This field holds a JSON payload of data to be submitted.

The second step is API Gateway triggering a Lambda function depending on the call it receives. This Lambda function will access a DynamoDB database table to either store or retrieve information. If storing information, a user record will be either created or updated in the Raw table. The change in this table is captured by Kinesis as a data stream and converted to an event by another Lambda function. This event is run through an EventBridge bus to another Lambda function that will do calculations on this new data and store the results in the Processed table in DynamoDB. If retrieving information, n user records will be returned to the client as JSON documents in batches of 50 with pagination ID. The Processed table stores all historical data on a user that has been processed, while the Raw table holds just the current raw data of each user. This is because the client only retrieves from the Processed table and not all records are needed to do the calculations of a user's affinity. The scalability of EventBridge paired with Lambda is why we prefer this design. In the future, anyone could come in with no knowledge of the connected design and work on a Lambda function with the purpose of utilizing the changed data event to add further functionality after the initial launch. This event can be pumped into as many Lambdas as needed. This kind of design suits a growing application perfectly and AWS eliminates a lot of work that would've been required.

Ponga Data Tracking

We will design the following Lambda functions: `getUserData`, `updateUser`, `eventBrigePush`, and `calculateAffinity`.

getUserData: This function receives an event payload from the API and returns data which gets queried from the database. The event payload is expected to have the following fields:

Field Value	Type	Description	How it can be used
id	String	The unique identifier for the user. Ex: "id": "12340987" "id": "realemail@email.com"	Use this to retrieve a specific user.
limit	Number	The number of records to be retrieved. Retrieves all records by default. Ex: "limit": 25	Use this to specify the number of records to retrieve.

First, the event payload will be parsed to extract the ID and limit fields. This function will then query the `Processed` table from the database using the ID as the primary key and retrieve a set number of records. By default, all data will be returned in batches of 50, but this can be changed using the limit field. After receiving the data, a response JSON payload will be created to be sent back to the client. The returned payload will include a status code, User map, and headers.

updateUser: This function receives an event payload from the API and updates information in the database. The event payload is expected to have the following fields:

Field Value	Type	Description	How it can be used
id	String	The unique identifier for the user. Ex: "id": "12340987" "id": "realemail@email.com"	Use this to retrieve a specific user.
data	Object	The data values to be added to the database. Ex: "data": { "x": 12 "y": 100 } "data": { "views": 3 "comments": 15 }	Use this to pass in new information to get added to the database.

First, the event payload will be parsed to extract the ID and data fields. This function will then write the data as an update to an entry in the Raw table or create a new entry if the user ID does not exist. A response JSON payload will be created to be sent back to the client. The response payload will include a status code and headers.

eventBridgePush: This function receives an event payload from a Kinesis data stream and invokes EventBridge. The event payload contains the data on a specific user from the Raw table before and after a change. Only the data after being changed will be formatted into an event for EventBridge. The event payload is expected to have the following field:

Field Value	Type	Description	Use
Event Pattern	String	<pre> The event pattern generated by AWS Kinesis: { "Records": [{ "kinesis": { "partitionKey": "partitionKey-03", "kinesisSchemaVersion": "1.0", "data": "SGVsbG8sIHRoaXMgaXMgYSB0ZXN0IDEy My4=", "sequenceNumber": "495451152434909850182800677149731445 82180062593244200961", "approximateArrivalTimestamp": 1428537600 }, "eventSource": "aws:kinesis", "eventID": "shardId-000000000000:49545115243490985 0182800677149731445821800625932442009 61", "invokedIdentityArn": "arn:aws:iam::EXAMPLE", "eventVersion": "1.0", "eventName": "aws:kinesis:record", "eventSourceARN": "arn:aws:kinesis:EXAMPLE", "awsRegion": "us-east-1" }] } </pre>	Use this to retrieve a specific event programmatically. This field holds a JSON payload of data through Kinesis to be captured by the 'Rule Filter' of EventBridge.

calculateAffinity: This function receives an event payload from EventBridge and queries the database to create an item in the Processed table for a specific user. The event payload is expected to have the following fields:

Field Value	Type	Description	Use
Event Pattern	String	<pre>{ "source": ["aws.kinesis"], "detail-type": ["AWS API Call via CloudTrail"], "detail": { "eventSource": ["kinesis.amazonaws.com"] } }</pre>	Use this to retrieve a specific event programmatically. This field holds a JSON payload of data through Kinesis to be captured by the 'Rule Filter' of EventBridge.
User	string	The unique email of the new/existing User. <pre>"User": "example@example.com"</pre>	Use this to create a new user or update an existing one.
Data	map	Data that represents the value proposition quantity of a user. <pre>"Data": { "Photos": 4, "Shares": 60, "Comments": 20 }</pre>	This field holds a JSON payload of data to be submitted.

DynamoDB provides several options for a database such as: a sort key, capacity mode, indexes, encryption, point-in-time-recovery, time to live, and stream details. The sort key allows for multiple records on the same user. We will have a sort key as one user will have multiple data entries. There are two capacity modes known as on demand and provisioned. On-demand simply bills for the actual reads and writes an application performs. Provisioned allows you to optimize your costs by allocating read/write capacity in advance. We will be using on-demand as we do not have subsequent knowledge for provisioned. Tables in DynamoDB can utilize local and global indexes. A local index can only be created at the time of table creation, and it allows for queries on attributes that are not part of the table's primary key. The global index can be created

after table creation but there is a limit on how many. It allows for queries on any attribute to find a specific value. We will determine which to use later in the design when both database tables are ironed out. Amazon provides encryption for their databases, and we will be utilizing it. Point-in-time recovery is a feature that provides continuous backups of DynamoDB data. Time to live is a feature that can put a lifetime on specific attributes or data in your tables. We will use point-in-time recovery and time to live to safeguard the database records and keep them fresh. Stream details refers to a data stream containing changes in data. DynamoDB provides a stream but can also connect to a Kinesis data stream. We will use the stream details to trigger events.

The Processed table holds all historical data on a user that is relevant to the client. As DynamoDB is a NoSQL database, the only columns of the table that are defined are the partition and sort key. Insight will only produce one attribute “Data” with multiple instances per user.

Partition Key: User ID - user’s email

Sort Key: Timestamp - ISO 8601 date and time convention

Example table with multiple records for bob:

User ID	Timestamp	Data
bob@example.com	2022-07-03T23:12:08.000Z	“Magnitude”: { “N” : “40” }, “Photos/Shares Angle”: { “N” : “35” }, “Photos/Comments Angle”: { “N” : “50” }, “Shares/Comments Angle”: { “N” : “70” }
bob@example.com	2022-09-03T21:10:05.000Z	“Magnitude”: { “N” : “42” }, “Photos/Shares Angle”: { “N” : “36” }, “Photos/Comments Angle”: { “N” : “48” }, “Shares/Comments Angle”: { “N” : “72” }
...

Ponga Data Tracking

The Raw table holds raw data on a user that is irrelevant to the client. As DynamoDB is a NoSQL database, the only columns of the table that are defined are the partition and sort key. Insight will only produce one attribute “Data” with one record per user.

Partition Key: User ID - user’s email

Sort Key: Timestamp - ISO 8601 date and time convention

Example table with one record for each user bob and bill:

User ID	Timestamp	Data
bob@example.com	2022-07-03T23:12:08.000Z	“Photos”: { “N” : “40” }, “Shares”: { “N” : “35” }, “Comments”: { “N” : “50” },
bill@example.com	2022-09-03T21:10:05.000Z	“Photos”: { “N” : “60” }, “Shares”: { “N” : “85” }, “Comments”: { “N” : “500” },
...

To utilize Insight with a large-scale application like Ponga, a queue for incoming API requests will be required. This is a stretch goal planned for version 2 of the design. Also, any additional Lambda functions to be added will be a part of version 2. A queue is needed because of the risk that API Gateway goes down, causing any pending requests to come back empty, and there is no way to handle lost requests. To fix this, we would connect a queue to store all API requests until API Gateway is ready to accept a new request. This can be done through Amazon Simple Queue Service (SQS). All requests will be queued to ensure that they eventually make it to API Gateway. SQS would act as a middleman for API Gateway and the Ponga client. Taking advantage of AWS to implement our API will allow for near infinite scalability and easy modulation in the future.

4.3 Risks

Risk	Risk Reduction
Setting up database schema	DynamoDB is a Managed NoSQL Database that is optimized for performance at scale and can store more complicated objects such as JSON files. We will have a complete design of the database schema before creating it.
Some queued requests could become lost before triggering a Lambda function	Have input data stored in an amazon SQS queue. Then the API listener can draw from it.
Rigid Design - We might want to add new functionality other than just calculating user affinity in the future.	The use of EventBridge allows us to add many more Lambda functions that take the same event as calculateAffinity.

4.4 Tasks

1. Meet with sponsor to discuss project
2. Research 3D vector spaces and AWS Services.
3. Understand AWS SQS, API Gateway, Lambda, DynamoDB, Kinesis – investigate API choices (http, rest, etc.), practice/learn programming lambda functions, investigate NoSQL and DynamoDB.
4. Understand the use cases given by sponsor. – Formulate questions for the sponsor as a team.
 - i. Submit Updated Metric
 - ii. Retrieve Historical Data
 - iii. Calculate Attribute Affinity and Persist – Address mathematics of the function to be created, so we can divide up work on it in the future.
 - iv. Generate Webhook – Is this needed?
5. High Level Design of project
 - i. API Gateway – Decide on API type and number of API being created

Ponga Data Tracking

- a. GET method – Do we use query string parameters or take in a JSON payload from the client?
 - b. POST method – What should the response payload look like?
 - ii. Lambda – Decide number of functions needed and any additions for future versions. Plan to stub out each lambda function for testing.
 - a. retrieve – Define event payloads. (Received and sent)
 - b. update – Define event payloads. (Received and sent)
 - c. calculateAffinity – Define event payloads. (Received and sent)
 - iii. DynamoDB – Decide number of tables and name them.
 - a. raw table – Define rough schema.
 - b. processed table – Define rough schema
 - iv. Kinesis - Study its connection to the raw table and how it will send data to the calculateAffinity Lambda.
 - v. SQS - Decide whether this will be part of version 1 of the project.
6. Fill out preliminary proposal and create presentation – Decide roles for the presentation and have everyone complete a section of the proposal.
7. Present High-Level Design to sponsor.
8. Decide on a Sprint planning method for next semester.
9. Prepare for long meeting on 11/11 and decide on Sprint planning method.
10. Understand Ponga’s use of AWS services and apply it to our own project.
11. Create contracts between the AWS services being used that serve as documentation on our project. Define what each block of the project requires and how the blocks communicate.
 - i. API
 - a. GET method – Document JSON map User. Display a sample request and response from the client.

- b. POST method – Document JSON map Record. Display a sample request from the client.
 - ii. Lambda – Define input event payloads and functionality.
 - a. retrieve
 - b. update
 - c. eventBridgePush – explain use of EventBridge
 - d. calculateAffinity
 - iii. DynamoDB – Define schemas
 - a. Raw
 - b. Processed
- 12. Complete preliminary report and create presentation – Decide roles for the presentation and have everyone complete a section of the report.
- 13. Setup AWS workspace for next semester using Cloud9 and GitHub.
- 14. Sprint 1 - Create DynamoDB database schemas and test them.
 - i. Raw
 - ii. Processed
- 15. Sprint 2 – Stub out and test Lambda Functions, then implement functionality.
 - i. retrieve
 - ii. update
 - iii. eventBridgePush
 - iv. calculateAffinity
- 16. Sprint 3 - Connect Lambdas to DynamoDB, Kinesis, and EventBridge.
 - i. Test and Debug updateUser → Raw → Kinesis → eventBridgePush → EventBridge → calculateAffinity → Processed

ii. Test and Debug Processed → getUserData

17. Sprint 4 - Create API and wire up endpoints to their respective lambdas

i. POST method

ii. GET method

18. Sprint 5 – Prepare all blocks of Insight for delivery and testing.

i. Update documentation with final design decisions and more specifics on code.

ii. Create account for testing Insight.

iii. Fill tables with data for testing use cases.

iv. Create step by step guide to AWS once logged in with sample API calls.

4.5 Schedule

Tasks	Dates	Details
Meet with sponsor	9/29	We will meet in JBHT 532 and use the whiteboard for a detailed description of the sponsors desired design.
Research	9/30-10/6	For this week, the team will each research AWS concepts and how they connect to our project.
Understand AWS services connection	10/7	Zoom call to discuss the use of each AWS service after having researched them.
Understand the use cases given by sponsor	10/8-10/14	Discuss use cases as a team and formulate questions for the sponsor. Zoom call at the end of the week to address questions.
High Level Design	10/15-10/23	Meet in discord several days to complete proposal and presentation.
i. DynamoDB	10/15	Decide on the number of tables. Discuss schema ideas for our DynamoDB tables.
ii. Kinesis and SQS	10/18	Understand how kinesis will work for our project. Decide on SQS being part of version 1.
iii. API Gateway and Lambda	10/21	Decide on API type and number. Make design decisions for the GET and POST endpoints.

Ponga Data Tracking

Proposal and Presentation	10/15-10/23	Complete required sections of the proposal while defining High Level Design. Create and practice presentation at the end of the week.
Present High-Level Design to sponsor	10/28	Zoom call with sponsor. Present current design to sponsor and get feedback. Discuss sprint planning.
Sprint planning	10/29-11/4	Research Sprint planning methods and applications for organization. Decide on one for next semester. 20-minute zoom call on 11/4.
Review and prepare	11/5-11/10	Prepare for meeting on 11/11 and decide on Sprint planning method. Look over proposal comments and fix errors.
Understand Ponga's AWS system	11/11	3 hours zoom call diving into Ponga's use of AWS and how to apply it to our project.
Documentation	11/14-11/26	Document all aspects of the Project and form Sprints around the documentation.
i. API Gateway	11/14-11/15	Document the requests and response of each endpoint. Formatting the payloads that the client send and receive.
ii. getUserData and updateUser lambda	11/18-11/19	Document the lambda functions to specify input events and formatting of output event.
iii. eventBridgePush and calculateAffinity lambda	11/25-11/26	Document conversion of data stream to event, and input/output event of calculation lambda.
iv. DynamoDB tables	11/17	Document database table schemas and connection to design.
Report and Presentation	11/26-11/28	Complete required sections of the report while defining detailed design and next semester sprint plans. Create and practice presentation at the end of the week.
Setup AWS workspace	11/29-12/8	Finalize AWS accounts and setup Cloud9 and GitHub.
Sprint 1 – DynamoDB Implementation	1/17-1/22	Create DynamoDB database schemas and test them.
i. Raw		
ii. Processed		
Sprint 2 – Lambda Implementation	1/23-1/29	Stub out and test Lambda Functions
i. Retrieve	1/30-2/12	Implement getUserData

ii. Update	1/30-2/12	Implement updateUser
iii. Event Bridge Push	2/13-2/26	Implement eventBridgePush
iv. Calculate Affinity	2/13-2/26	Implement calculateAffinity
Sprint 3 – Connect Code and Database	2/27-3/12	Hook up EventBridge and Kinesis, then test
Sprint 4 – API Implementation i. GET Method ii. POST Method	3/13-3/19	Create API and wire up endpoints to their respective lambdas
Sprint 5 – Prepare deliverables i. Update documentation ii. Setup testing account iii. Pre-populate tables iv. Create guide to test	3/20-4/23 3/20-4/2 4/3-4/9 4/3-4/9 4/10-4/23	Create an AWS testing account, tables, and guide to allow testing on individual functions and flows. Update Final Report Set up a testing account with all functions deployed Fill tables with sufficient data to test use cases. Steps guiding tester on the functionality of Insight on AWS

4.6 Deliverables

- Design Documentation: Google docs linked to a Lucid chart. These contain documentation on each major block of Insight.
 - i. For API this would be sample requests and responses.
 - ii. For Lambdas, we will provide input event payloads and functionality.
 - iii. For database tables, we provide the schemas.
- Login credentials to AWS account where all API, Lambdas, and DynamoDB tables have been deployed.
- Step by Step guide to testing Insight utilizing CloudWatch logs once logged into AWS with full access.
 - i. Pre-populated data in the Raw and Processed tables upon login.
 - ii. Tour of AWS blocks used for Insight.
 - iii. Series of API calls to test all use cases.
- API and Database Response Times: Timing data related to each APIs response times in connection with the database.
- Final Report: Contains background information, use cases, detailed project and database design, tasks, schedule, deliverables, key personnel, and references.

5.0 Key Personnel

Dylan Vaughn – Vaughn is a senior Computer Science major in the Computer Science and Computer Engineering Department at the University of Arkansas. He has completed Intro to Engineering. No relevant experience or internships. Was responsible for research and understanding AWS.

Matthew Clemence – Clemence is a senior Computer Engineering major in the Computer Science and Computer Engineering Department at the University of Arkansas. He has completed Software Engineering and Intro to Engineering. Was responsible for hosting weekly meetings and research.

Ryan Drake – Drake is a senior Computer Science major in the Computer Science and Computer Engineering Department at the University of Arkansas. He has completed Database Management Systems, Software Engineering, and Algorithms courses. Was responsible for researching and understanding AWS and the services we will be using.

Julio Bonilla – Bonilla is a senior Computer Science major in the Computer Science and Computer Engineering Department at the University of Arkansas. He has completed Database Management Systems, Algorithms, Software Engineering, Computer Networks, and Intro to Engineering. No relevant experience or internships. Was responsible for outlining the objective, research, unit testing, and coordinating team meetings.

Logan Reed – Reed is a senior Computer Science major in the Computer Science and Computer Engineering Department at the University of Arkansas. He has completed Database Management Systems, Big Data Management and Analyzation, and Software Engineering. Has experience in database and API relations in personal project. No relevant internships. Was responsible for outlining, researching understanding DB relations.

Jerrel Fielder – Industry Champion. Senior technology and business executive with 30 years in Enterprise and Consumer software in multiple roles: technology, technology leadership, architecture, consulting, sales, sales leadership, and founder.

6.0 Facilities and Equipment (Ryan Drake)

The utilities we will make use of are all the Amazon Web Services. The platform we are building will employ API Gateway, Lambda, DynamoDB, Kinesis, and EventBridge. These services will seamlessly flow together to accomplish the implementation of our solution.

API Gateway accepts requests to access the API and provides the promised services. API must be formatted correctly. Amazon Lambda is a dynamic service which can accomplish many tasks using programmable functions, and it allows for each service to be easily connected. DynamoDB is the NoSQL database that we will be using to store and access data. Kinesis will be used in conjunction with DynamoDB to detect changes in the tables so that certain functions can be called in response. EventBridge will handle the bus of events passed from Kinesis change data capture to Lambdas.

7.0 References

- [1] Amazon. “Welcome - Amazon API Gateway.” *Amazon Web Services*, 14 Nov. 2022, docs.aws.amazon.com/apigateway/latest/api/Welcome.html.
- [2] Amazon. “What Is Amazon DynamoDB? - Amazon DynamoDB.” *Amazon Web Services*, 15 June 2022, docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html.
- [3] Amazon. “What Is Amazon EventBridge? - Amazon EventBridge.” *Amazon Web Services*, 14 Nov. 2022, docs.aws.amazon.com/eventbridge/latest/userguide/eb-what-is.html.
- [4] Amazon. “What Is Amazon Kinesis Data Streams? - Amazon Kinesis Data Streams.” *Amazon Web Services*, 29 Nov. 2021, docs.aws.amazon.com/streams/latest/dev/introduction.html.
- [5] Amazon. “What Is Amazon Simple Queue Service? - Amazon Simple Queue Service.” *Amazon Web Services*, 20 Oct. 2022, docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/welcome.html.
- [6] Amazon. “What Is AWS Lambda? - AWS Lambda.” *Amazon Web Services*, 1 July 2022, docs.aws.amazon.com/lambda/latest/dg/welcome.html.
- [7] Google. “How Google Analytics Works - Analytics Help.” *Google Help*, 2022, <https://support.google.com/analytics/answer/12159447?hl=en>.
- [8] “Recordings: The Complete Guide.” *What Are Session Recordings (Session Replays) + How to Use Them*, Hotjar, 6 June 2022, <https://www.hotjar.com/session-recordings/>.
- [9] “Scroll Depth Trigger - Tag Manager Help.” *Google Help*, Google, 2022, <https://support.google.com/tagmanager/answer/7679218?hl=en>.

[10] Vanhee, David. "Google Analytics 4 Click Tracking – The Beginners' Guide." *Data Driven*,

Digital Mantis, Inc, 19 Aug. 2022,

<https://www.datadrivenu.com/how-to-track-clicks-in-googleanalytics-4/>.

Accessed 25 Nov. 2022.