

C06 – serverless Analytics con Amazon Web Services

Francesco Diaz

Amazon Web Services

 @francedit



The AWS analytics portfolio

DATA, VISUALIZATION, ENGAGEMENT, AND MACHINE LEARNING



ANALYTICS



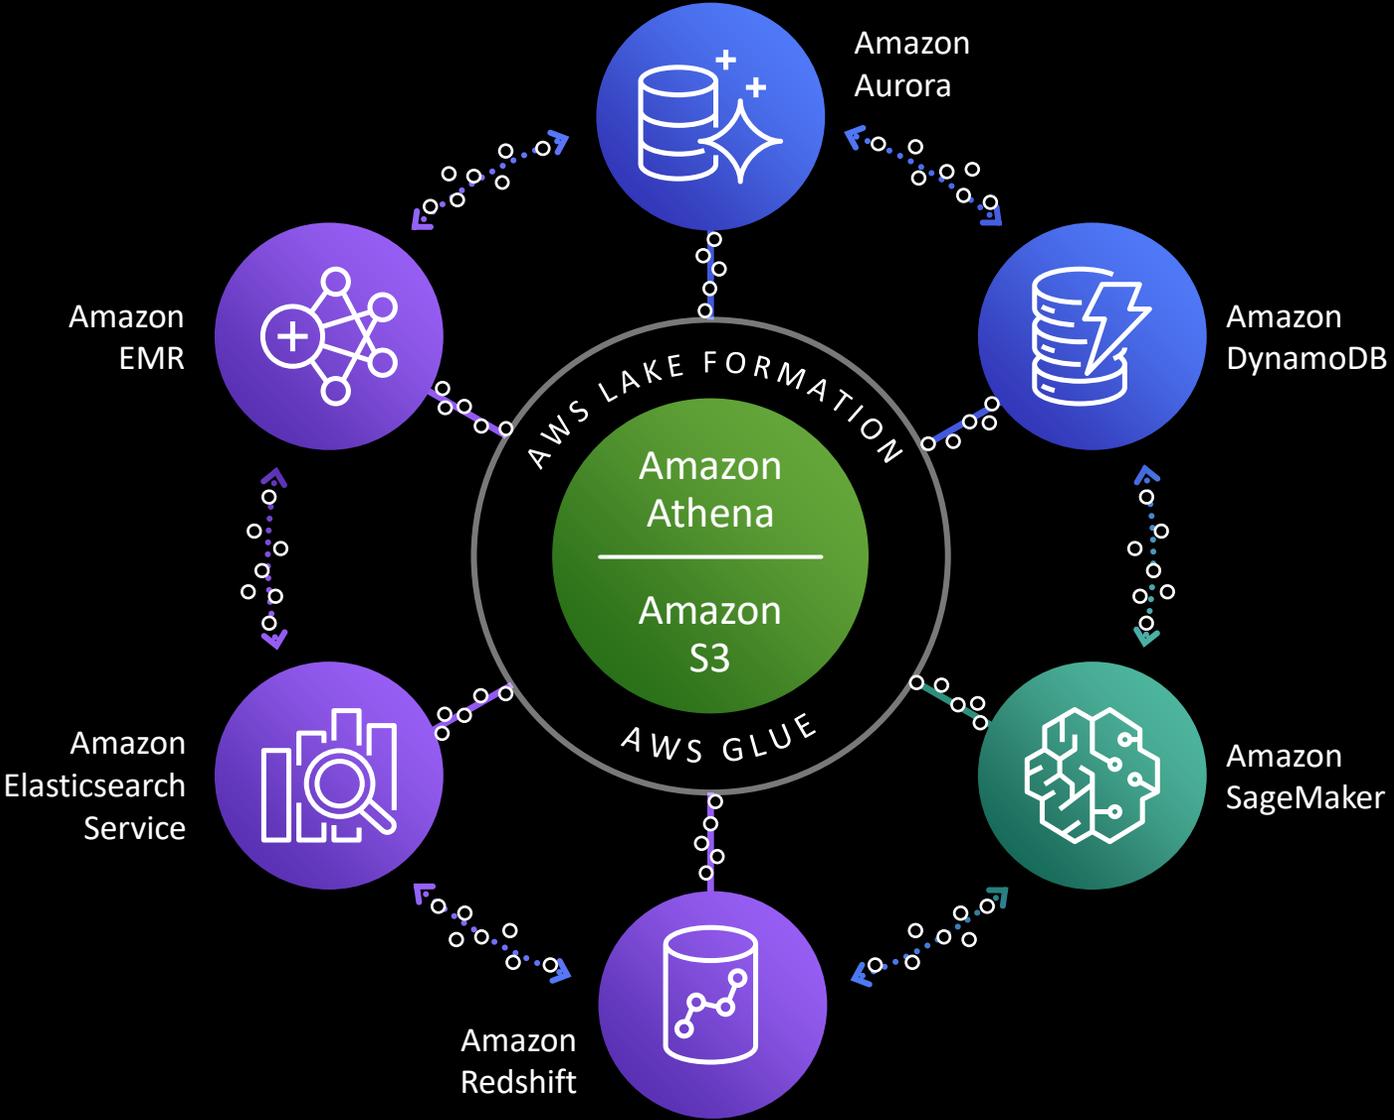
DATA LAKE INFRASTRUCTURE AND MANAGEMENT



DATA MOVEMENT

Database Migration Service | Snowball | Snowmobile | Kinesis Data Streams | Kinesis Data Firehose | Managed Streaming for Apache Kafka

Lake House architecture on AWS



SCALABLE DATA LAKES

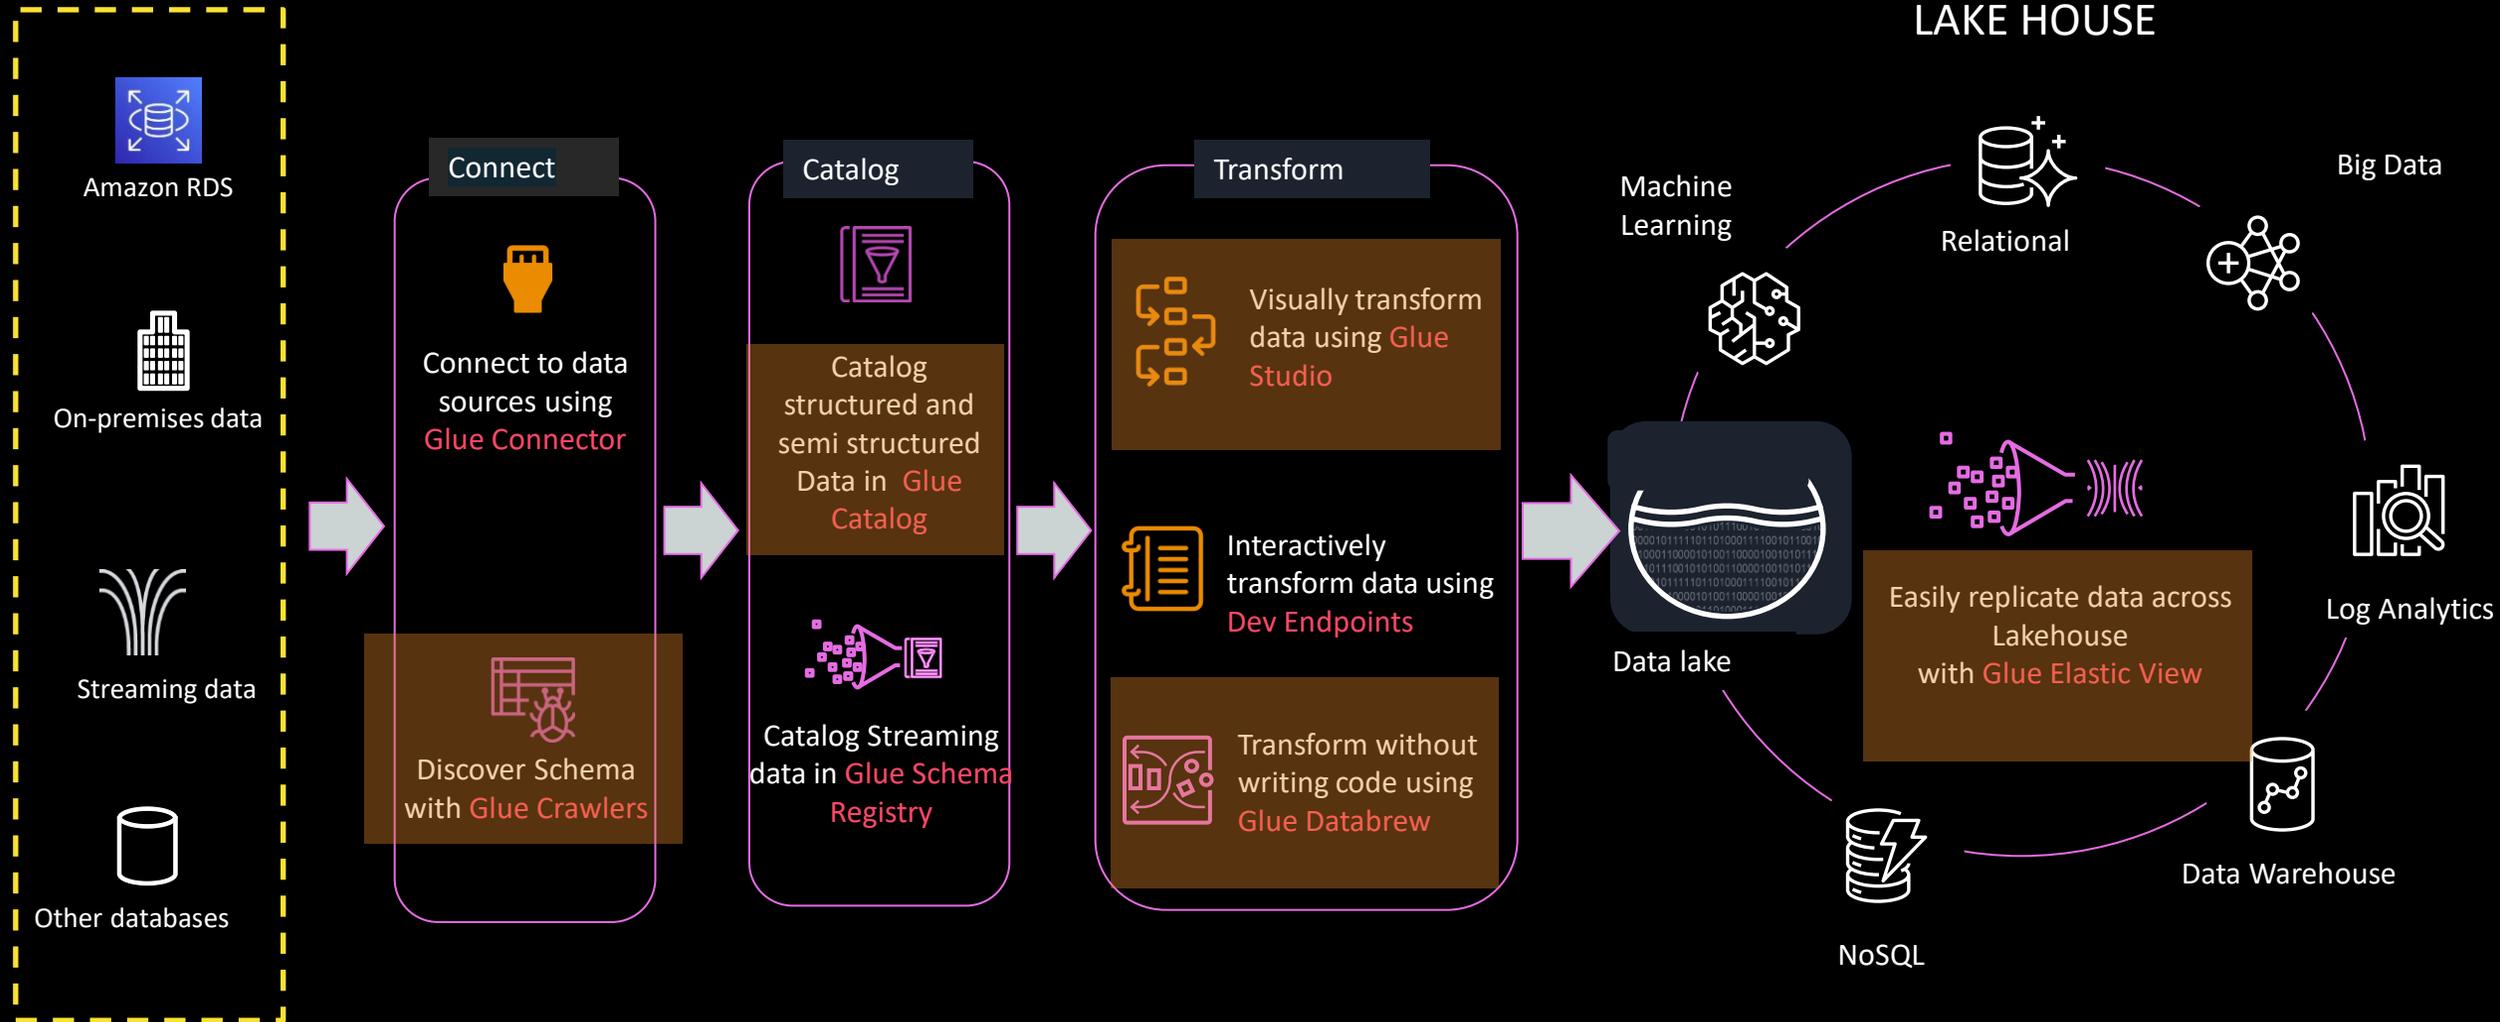
**PURPOSE-BUILT
DATA SERVICES**

**SEAMLESS
DATA MOVEMENT**

UNIFIED GOVERNANCE

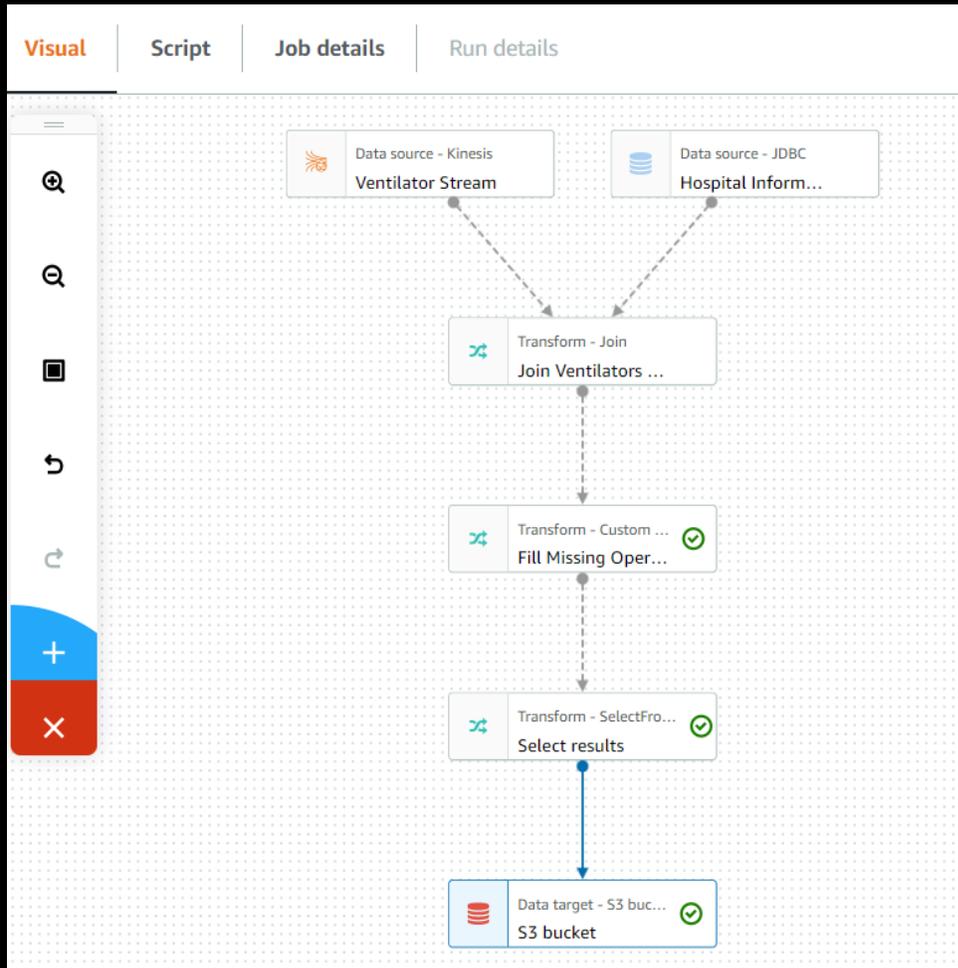
**PERFORMANT AND
COST-EFFECTIVE**

AWS Glue ecosystem



AWS Glue Studio

Easily author, run, and monitor AWS Glue ETL jobs



Author AWS Glue jobs **visually** without coding

Monitor **1,000s of jobs** through a **single pane of glass**

Distributed processing without the learning curve

Advanced transforms **though code snippets**

AWS Glue DataBrew – Visual Data Preparation

Built-in transformations

The screenshot shows the AWS Glue DataBrew interface for a project named 'Citibike-project'. The dataset is 'citibike' with a random sample of 2,500 rows. The main view displays a table with columns: 'start station name', '# start station latitude', '# start station longitude', '# end station id', 'end station name', and '# end station latitude'. Each column has a histogram and summary statistics (Unique, Total, Min, Median, Mean, Mode, Max). The table lists various station names and their corresponding coordinates and IDs.

start station name	# start station latitude	# start station longitude	# end station id	end station name	# end station latitude
Greene Ave Throop Ave	40.6894932	-73.942061	494	W 26 St 8 Ave	40.74734825
Stanton St Norfolk St	40.720747	-73.986274	3077	Stagg St Union Ave	40.70877084
9 Ave W 22 St	40.7454973	-74.00197139	3709	W 15 St 6 Ave	40.73804614
Carmine St 6 Ave	40.73038599	-74.00214988	483	E 12 St 3 Ave	40.73223272
Lexington Ave E 26 St	40.741459	-73.983293	435	W 21 St 6 Ave	40.74173969
5 Ave E 78 St	40.77632142	-73.96427393	3159	W 67 St Broadway	40.77492513
E 7 St Avenue A	40.72621788	-73.98379855	4110	Avenue D E 10 St	40.72506551
W 92 St Broadway	40.7921	-73.9739	3623	W 120 St Claremont Ave	40.8109494
W 74 St Columbus Ave	40.7785669	-73.97754961	167	E 39 St 3 Ave	40.7489006
W 95 St Broadway	40.7937704	-73.971888	3294	E 91 St Park Ave	40.7835016
W 20 St 11 Ave	40.746745	-74.007756	3812	University Pl E 14 St	40.734814
Maiden Ln Pearl St	40.70706456	-74.00731853	447	8 Ave W 52 St	40.76370739
Central Ave Weirfield St	40.69055	-73.91181	3042	Fulton St Utica Ave	40.6794268
Bayard St Baxter St	40.71602118	-73.99974372	401	Allen St Rivington St	40.72019576

Recipe steps

The screenshot shows the 'Recipe (5)' interface for a 'chess-project-clean-recipe' (Version 2.0). It lists five applied steps:

1. Filter values by `white_rating`
2. Filter values by `black_rating`
3. Group by `winner`, `victory_status` and create `winner_count` from `COUNT(winner)`
4. Filter values by `victory_status`
5. Replace text `mate` with `checkmate` in `victory_status`

Amazon Athena

Query data in S3 using SQL



Serverless

Quickly query S3 data without managing infrastructure, and pay only for the queries you run



Open and standard

Use ANSI SQL for querying with support for Parquet, CSV, JSON, Avro and other standard data formats



Fast interactive performance

Parallel execution to deliver most results within seconds, with no cluster management required



Cost effective

Pay only for queries run and save 30–90% by compressing, partitioning, and converting your data into columnar formats

Amazon Athena

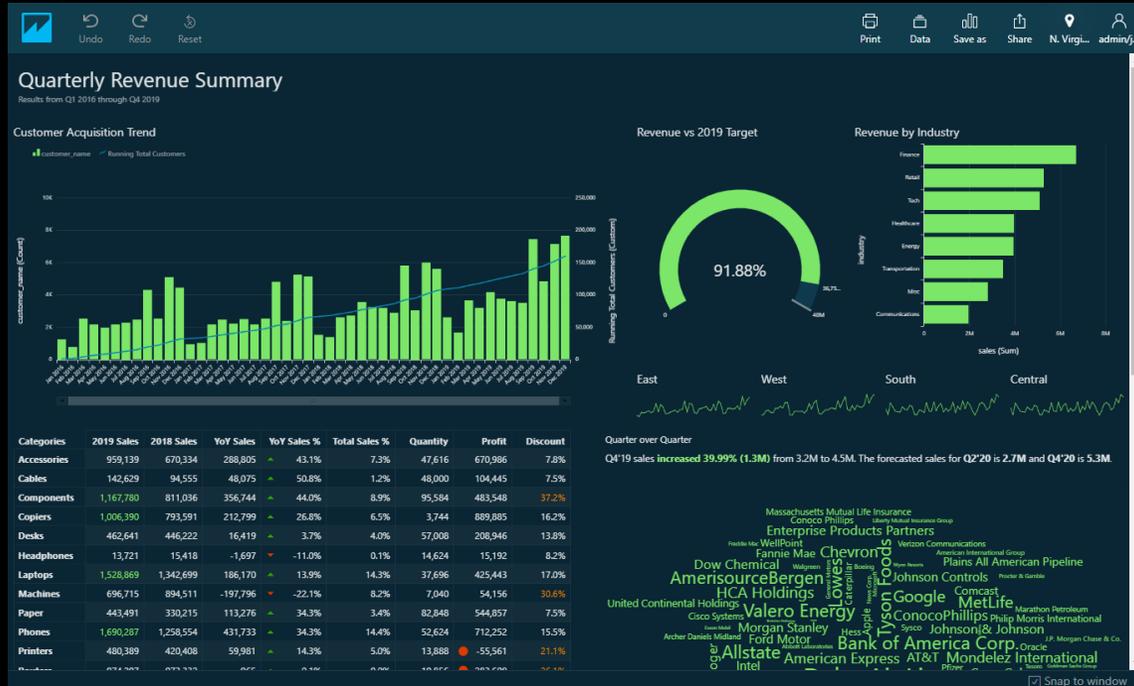
- User Defined Functions (UDF)

```
USING EXTERNAL FUNCTION UDF_name(variable1 data_type [, variable2 data_type] [, ...])  
RETURNS data_type  
LAMBDA 'lambda_function'  
SELECT [...] UDF_name(expression) [...]
```

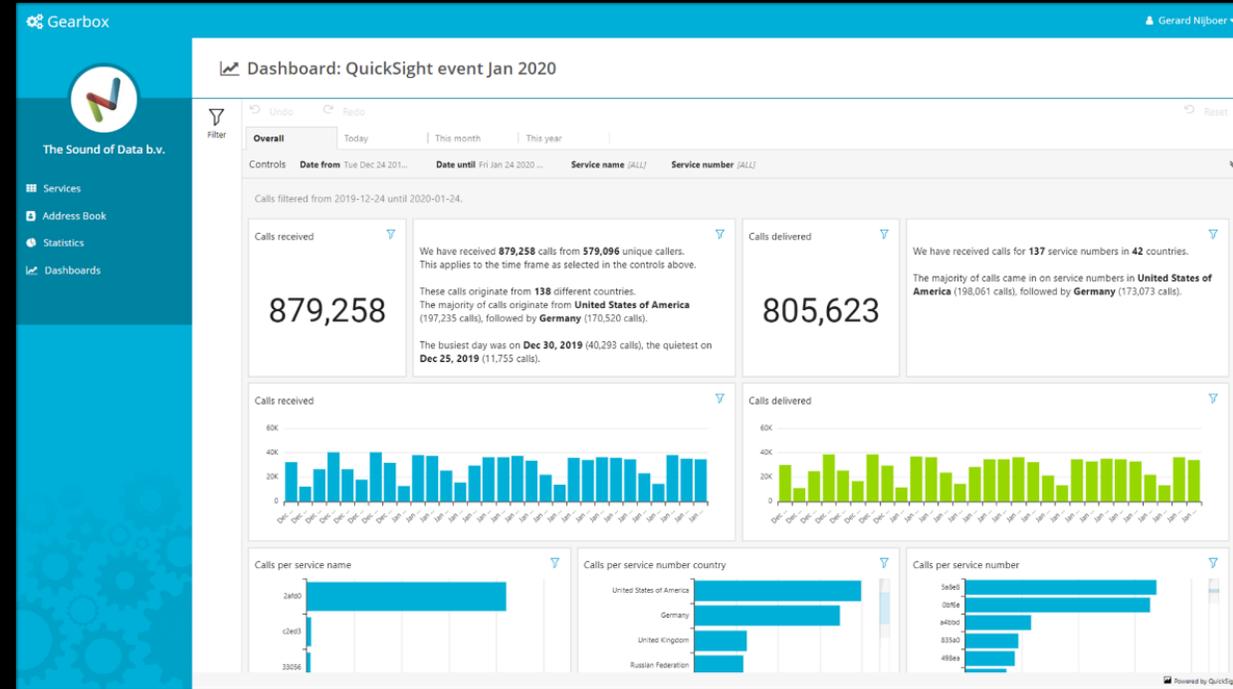
- Amazon Athena Federated Query - support for running SQL queries across relational, non-relational, object, and custom data sources
 - data source connectors on AWS Lambda
- Using Machine Learning with Amazon Athena

```
USING EXTERNAL FUNCTION ml_function_name (variable1 data_type [, variable2 data_type] [, ...])  
RETURNS data_type  
SAGEMAKER 'sagemaker_endpoint'  
SELECT ml_function_name()
```

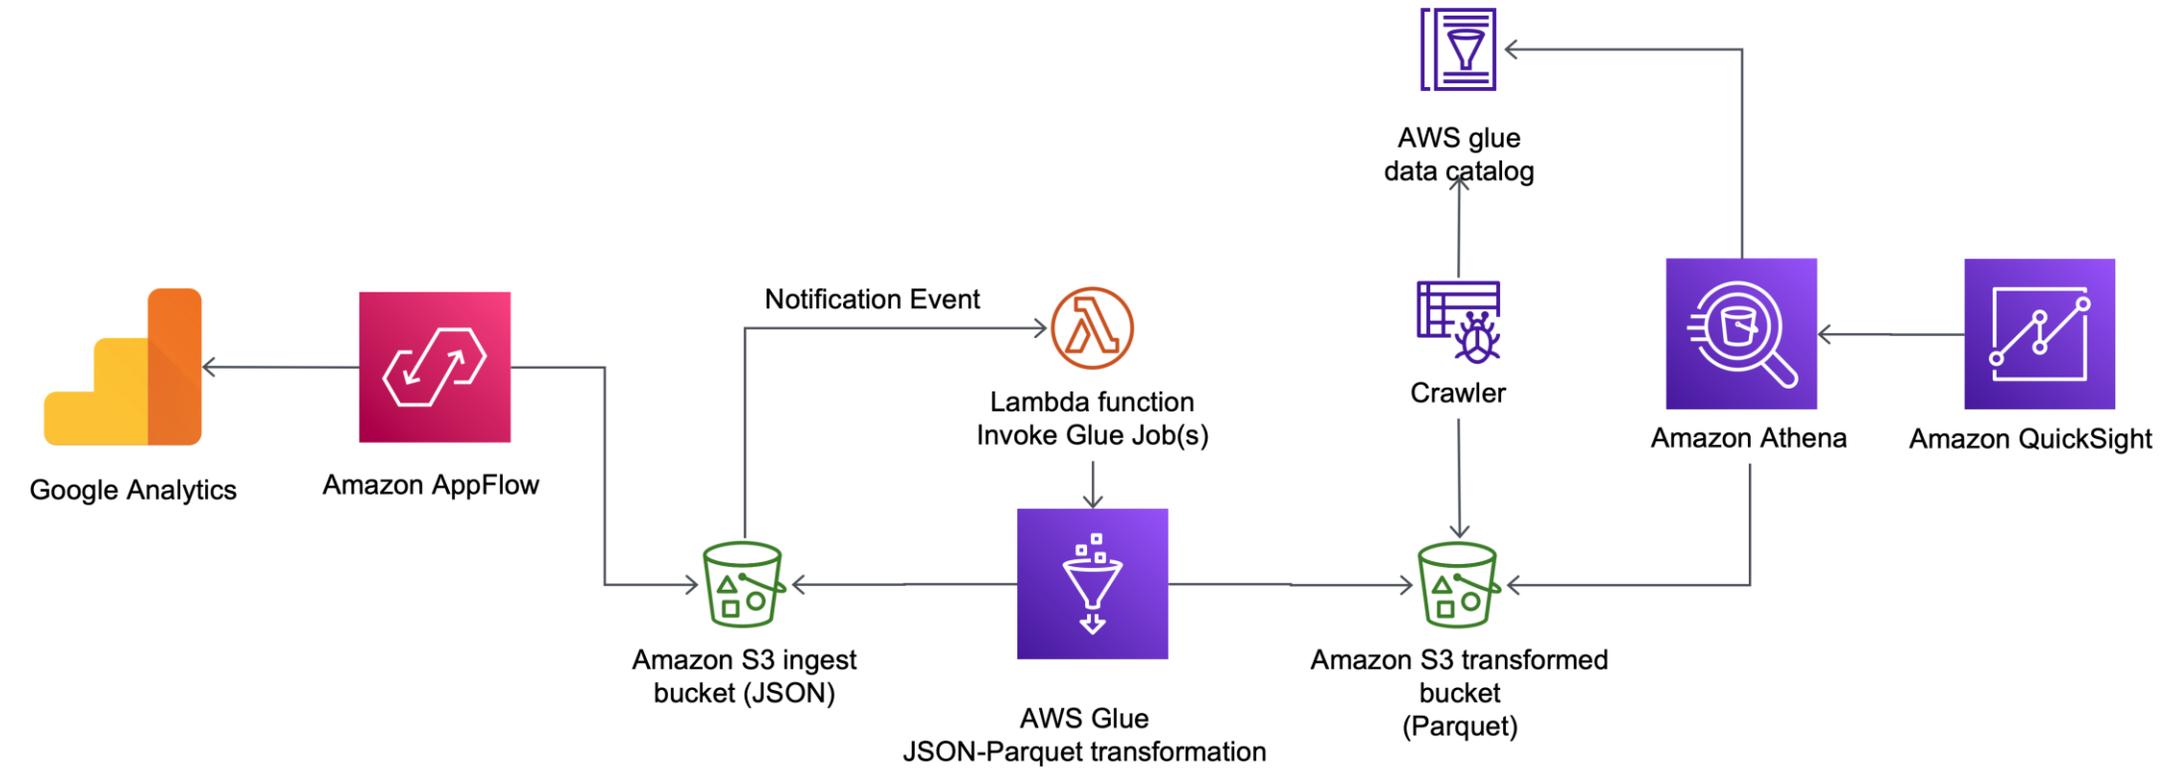
Amazon QuickSight



Corporate & self-service BI



External BI / Embedded BI



Demo

References

- Cloud Formation template and Workshop: <https://appflow-immersionday.workshop.aws/en/googleanalytics.html>
- Using Athena with JDBC: <https://docs.aws.amazon.com/athena/latest/ug/connect-with-jdbc.html>
- Create a Glue crawler: <https://docs.aws.amazon.com/glue/latest/ug/tutorial-add-crawler.html>
- Getting started with Glue Studio: <https://docs.aws.amazon.com/glue/latest/ug/tutorial-create-job.html>
- Amazon QuickSight workshop: <https://learnquicksight.workshop.aws/en/>

Grazie!

- Il materiale sarà online nei prossimi giorni su <http://www.communitydays.it>