

Serverless Data Querying with Amazon Athena

12 Practical Uses of Amazon Athena Number 6 Will Shock You!

Paul Pilone Paul@Element84.com Element 84 / NASA







Introduction

- Who Am I?
- Why Were We Looking Into Athena?

Searching for Answers

- What is Athena?
- How Do We Get Started?
- How Do We Structure Our Data?
- How Much Does Athena Cost?
- Is Our Data Safe?
- But Wait, There's More?
- What Else Should We Know?
- Questions?





Who Am I?



Hi! I'm Paul

- Senior Software Engineer at Element 84
- Background in cloud architecture, integration, and development
- Full stack developer with experience in web and mobile
- *Not* a data scientist
- Representing the Cloud Giovanni team and some of the work we have been doing the past couple of months



Why Were We Looking Into Athena?



Giovanni queries, analyzes and performs algorithms against NASA's data from various satellite and surface observations



WHY WERE WE LOOKING INTO ATHENA?

EOSDIS

Tasked with finding a cloud-based solution that would not sacrifice performance

Services considered for data querying and processing:

- Athena
- Spark on Amazon EMR (Elastic MapReduce)
- Amazon Redshift



What Is Athena?



- Interactive query service launched by Amazon in November 2016
- Queries data in Amazon S3
- Little-to-no ETL (Extract, Transform, Load) requirements*
- Queries are written in standard SQL

SELECT time, SUM(weighted) / SUM(weights) AS average FROM ...

*Not required, but can be useful



- Backed by Presto
- Leverages existing AWS resources
- Serverless
- Supports data partitioning
- Table modification statements use standard Apache Hive DDL



Schemas are applied-on-read

- Table definitions are applied to data at query execution
- No data loading or transformation required
- Modify or delete table definitions without impacting data
- Single data set; multiple table definitions



How Do We Get Started?



All you need to do is:



• Store data in a supported format in a S3 bucket

• Create and define tables within Athena

START QUERYING!



3

How Do We Pick a File Format?



The right file format will vary for each project

Things the Giovanni team considered:

- Raw data
- File size
- Preprocessing time and cost (ETL)
- Readability
- Egress
- Ease/frequency of updates



Supports querying against the following formats:

- CSV
- Apache Parquet
- TSV
- Snappy
- JSON

- Zlib
- Text
- LZO
- Apache ORC
- GZIP



The Giovanni team focused on GeoJSON and Apache Parquet formatted datasets

Geojson

- Lightweight
- Text-based
- Human-readable

Apache Parquet

- Columnar
- Compact
- Open source

Data was originally available in NetCDF which was converted to GeoJSON and Parquet



Sample Query

EOSDIS

Algorithm calculates the weighted mean for cloud area

fraction between Jan 2003 and Apr 2003

```
SELECT time, SUM(weighted) / SUM(weights) AS average
 1
 2
 3
  FROM
 4
 5
   SELECT time,
 6
     variable * COS(RADIANS(lat)) AS weighted,
     COS(RADIANS(lat)) AS weights
 8
 9
   FROM merra time parquet
10
  WHERE time >= 1041379200
11
   AND time <= 1041465599
12
   AND lon \geq -113
   AND lon \leq -110
13
14
    AND lat >= 41
15
     AND lat \leq 44
16
  )
17
18 GROUP BY time ORDER BY time ASC
```

HOW DO WE PICK A FILE FORMAT?

- The same data has been uploaded to an S3 bucket in both JSON and Parquet formats
- Let's compare executing the algorithm against each format

MERRA-2 Cloud Area Fraction	File Size	Data Scanned	Query Time
JSON	4.58 MB	4.58 MB	1.89s
Parquet	528 KB	132 KB	.88s



How Much Does Athena Cost?



Price per Query

\$5 per TB of data scanned

Better cost and performance by compressing, partitioning, and converting your data into columnar formats



MERRA-2 Cloud Area Fraction	File Size	Data Scanned	Query Time	Cost/ Query	Cost/ Month
JSON	4.58 MB	4.58 MB	1.89s	\$0.0000238	\$1.43
Parquet	528 KB	132 KB	0.88s	\$0.000000615	\$0.037

Based on ~60k queries per month



Optimize data for scanning

Smartly structured and partitioned data can drastically decrease cost of queries.

AIRX Cloud Area Fraction	File Size (Parquet)	Data Scanned	Cost/ Query	Cost/ Month
Feb 2003	6.9 GB	1.01 GB	\$0.005	\$300
All Time	6.9 GB	6.9 GB	\$0.034	\$2,040

Based on ~60k queries per month



Comparing Cost

Assumption:

 Small EMR Spark cluster of 1 Master and 2 Slaves (m.xlarge EC2 instances)

Service	Hours in Service	Data Scanned/ Query	Cost Model	Cost/ Month
Spark (EMR)	2,160	N/A	\$ per Hour	\$432
Athena	N/A	1.01 GB	\$ per TB Scanned	\$300

Based on ~60k queries per month



Is Our Data Safe?



Two ways to keep your data secure with Athena:

1. Set access control policies using:

- AWS Identity and Access Management (IAM)
- S3 bucket policies
- Access control lists (ACLs)
- 2. Utilize S3 and Athena encryption
 - Query encrypted data
 - Encrypt query results



Configuring Encryption

- Easily encrypt query results
- Encryption type is set with table definitions

Settings		
Query result location	s3://aws-athena-encrypted/ Example: s3://query-results-bucket/folder/	0
Encrypt query results		
Encryption type	SSE-KMS 🔻 🚯	
Encryption key	Enter a KMS key ARN	▼ 1 Create KMS key
KMS key ARN	arn:aws:kms:us-east-2:////////////////////////////////////	1fc-1
		Cancel Save



But Wait, There's More?



Query Result Caching

- Cached queries and results are stored in S3 buckets
- Past queries can be pulled up instantly without having to rescan any data
 - Saves money as no data is scanned

History

Search for name, query, etc.

🔷 Query submitted time	🔷 Query	Encryption type ()	State	Run time(s)	Data scanned	Action
2017/07/17 16:54:45 UTC-4	SELECT time, SUM(weighted) / SUM(weights) AS average FROM (SELECT time, variable * COS(RADIANS(la	N/A	SUCCEEDED	0.84	621.76KB	Download results
2017/07/17 16:51:07 UTC-4	SELECT time, SUM(weighted) / SUM(weights) AS average FROM (SELECT time, variable * COS(RADIANS(la	N/A	SUCCEEDED	68.9	999.03MB	Download results



Athena and QuickSight

- QuickSight: Amazon's cloud-based analytics service for visualizing data
- Simply enable QuickSight access to S3 buckets where Athena data resides



But Wait, There's More?

What Else Should We Know?



Observations from the Giovanni team

- At the time of prototyping, the only access to Athena was by using the AWS Console or the JDBC driver
 - In May 2017 Amazon added support for running queries via the REST API using the AWS SDK available in Java, .NET, Node.js, PHP, Python, Ruby, Go and C++
- Migrating current data into Parquet proved to be no small task



Limitations

- Complex geo-spatial queries may prove difficult
- Algorithm must be expressed in SQL
- No API for reporting query status
- Service limits
 - Concurrency limits
 - Query timeout
 - Number of databases
 - Tables per databases
 - Number of partitions





Questions?

Paul Pilone paul@element84.com

Giovanni https://giovanni.gsfc.nasa.gov/giovanni/

Mahabaleshwa Hegde, Christine Smit, Hailiang Zhang, Maksym Petrenko, Andrey Zasorin, Keith Bryant

Athena https://aws.amazon.com/athena



Additional Information



Converting NetCDF to JSON and Parquet

- Lambda triggered to write NetCDF file to Parquet
- Updated converter to run on Giovanni servers
 - $_{\odot}\,$ Writes metadata to DynamoDB and data file to S3 $\,$
- Roughly 1 day conversion for 10 year daily data
- 36 year daily data took > 1 week
 - Better performance if writing to multiple Parquet files





Athena's Encryption Support

- Added the ability to read and write encrypted data in April 2017
- Both server-side and client-side encryption are supported

S3 Encryption Options	Read from Table Data	Read/Write on query result	Require table property flag "has_encrypted_data"=true
Server-Side Encryption with Amazon S3–Managed Encryption Keys (SSE-S3)	Yes	Yes	No
Server-Side Encryption with AWS KMS– Managed Keys (SSE-KMS)	Yes	Yes	No
Client-Side Encryption with AWS KMS– Managed Keys (CSE-KMS)	Yes	Yes	Yes
Client-Side Encryption with Master Key (CSE-C)	No	No	N/A
Server-Side Encryption with Customer- Provided Encryption Keys (SSE-C)	No	No	N/A

