



Data Warehouse Modernization on AWS

www.sigmoid.com





About Sigmoid

Sigmoid is an emerging leader in data engineering and AI solutions.



750+

Employees



Work with **30+**
Fortune 500 firms



>97%

CSAT score



200+

ML models operationalized



5000+

Data pipelines built

Backed by

SEQUOIA 

Awards and Recognition

500™

Technology **Fast 500**
2023 NORTH AMERICA
Deloitte.

***ISG**

Provider Lens 2023 Quadrant

Analytics Services
Rising Star, U.S. ★

Report releasing Jan 2024



America's
Fastest-Growing
Private Companies



Open Source data
solution provider
of the year



FORRESTER
Now Tech:
AI Consultancies,
Q1, 2021 Report

Major Contender in

EVEREST GROUP

Analytics and AI Services Specialists
PEAK Matrix (2022)

Offices



New York



San Francisco



Dallas



Lima



Bangalore



Amsterdam

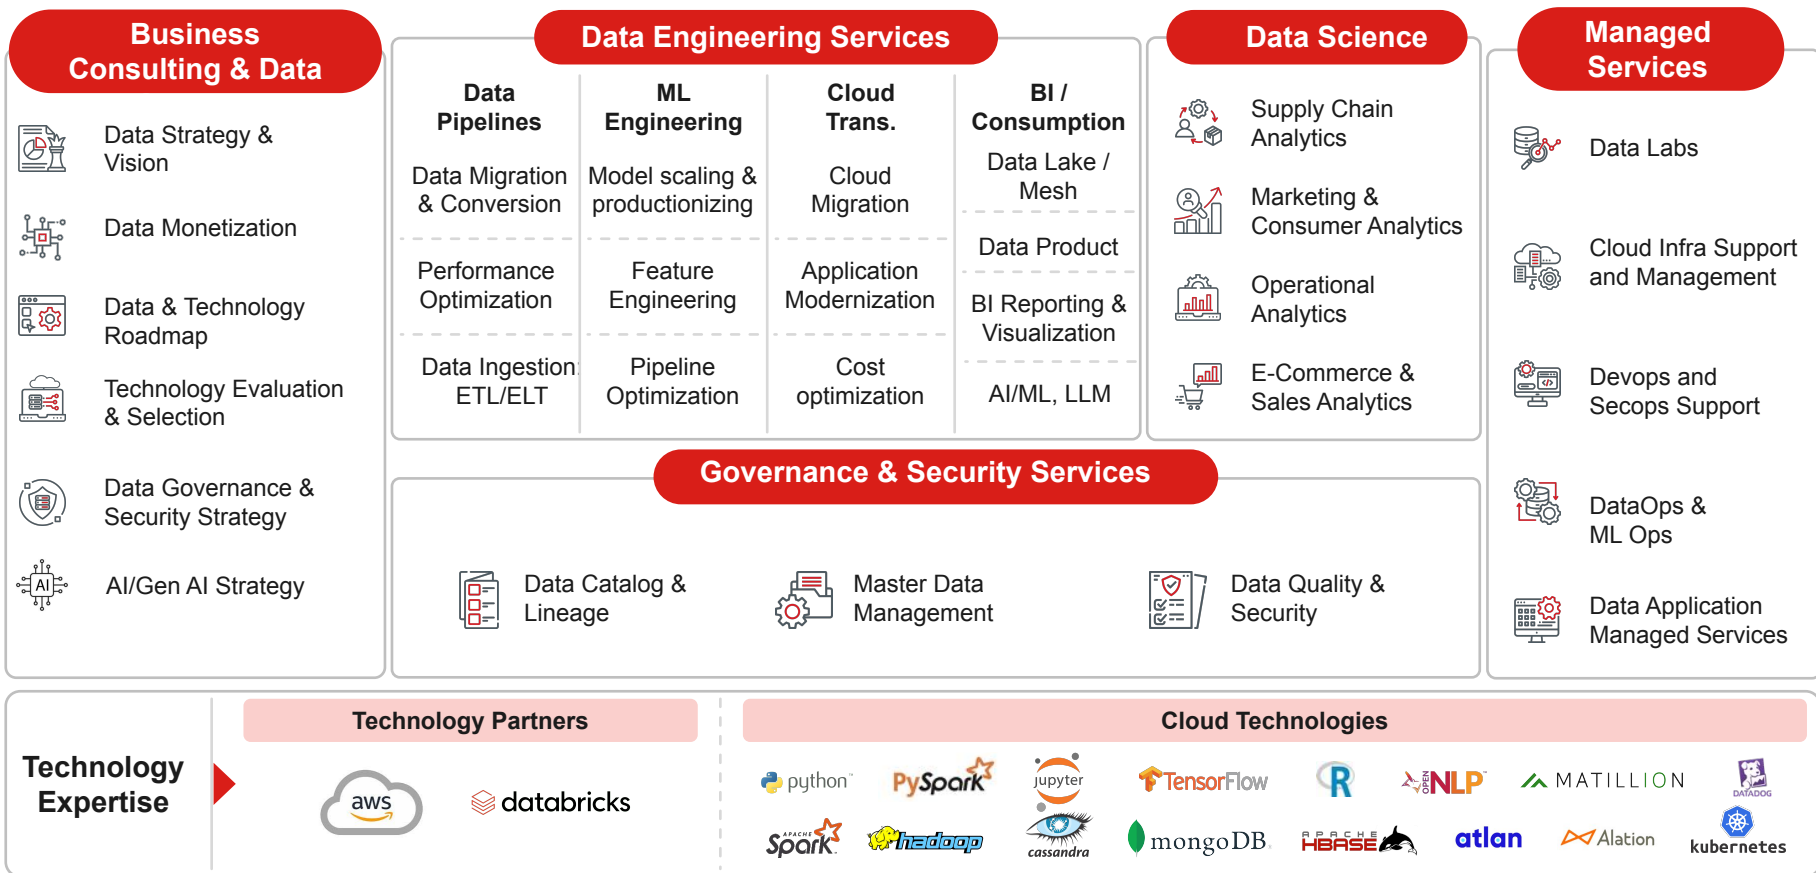


London



Sao Paulo

Enabling Business Transformation with Full-Service Capability Suite



Sigmoid Capabilities - Experience in implementing data solutions in AWS

Sigmoid has worked with more than Five large customers to design, build and deploy solutions in AWS

Data Processing & Transformation:

- **Amazon EMR:** Collaborative Apache Spark-based analytics platform used for big data processing and machine learning.
- **AWS Glue:** Fully managed ETL (Extract, Transform, Load) service for preparing and transforming data.

Data Analytics & Visualization:

- **Amazon Redshift:** Data warehouse service used for analyzing large datasets with either serverless or provisioned resources.
- **Amazon QuickSight:** Business intelligence tool for creating interactive visualizations and reports.

Machine Learning & AI:

- **Amazon SageMaker:** End-to-end platform for building, training, and deploying machine learning models.
- **AWS AI Services:** Pre-built AI services for vision, speech, language, and decision-making (e.g., Amazon Rekognition, Polly, Comprehend, Textract).

Security & Compliance:


- **AWS IAM (Identity & Access Management):** Identity and access management service.
- **AWS Organizations & AWS Config:** Governance and compliance tools for managing AWS environments at scale.

Data Storage & Management:

- **Amazon S3:** Scalable and secure data lake for storing large amounts of structured and unstructured data.
- **Amazon RDS:** Managed relational database service for structured data storage.

Data Ingestion & Integration:

- **AWS Data Pipeline:** Creating data workflows that move and process data across AWS services.
- **Amazon Kinesis:** Real-time data ingestion from applications, devices, or any streaming data sources.



Sigmoid's implementation of solutions in AWS involves leveraging a combination of services and tools tailored to specific business needs. Sigmoid would facilitate collaboration between data engineers, data scientists, business analysts, and other stakeholders to align the implementation with business goals and ensure success.

Problems We Solve for Our Clients in the Data Landscape



Data Pipelines

Business not able to respond in real time

Automating data ingestion & scalability of data pipelines

Data quality issues and trust in data

Maintaining and monitoring data pipelines



MLOps

High failure rate of deploying ML models in production

Creating and managing ML pipelines

Model Drift in Machine Learning

Longer development and deployment lifecycle



Cloud Data Warehouse


Technology selection for a use case

Data migration from on-premise to CDW

High TCO of a traditional data warehouse

Rigid and inflexible infrastructure


Services Spanning the Data Engineering and Analytics Value Chain



Data Warehouse Modernization



Steps for Data Warehouse Modernization



1. Assessing Current State

- Assessment of the organization's existing data warehouse infrastructure, data models, ETL/ELT processes and analytics capabilities
- Analyzing the query performance, data modelling, indexing strategies and resource utilization

2. Laying down Objectives & Requirements

- Collaborating with key stakeholders for defining the business objectives, priorities, and requirements for warehouse modernization

3. Defining the Target Architecture

- Defining and designing the target data architecture tailored to the customer's specific requirements and workload patterns, following AWS Well Architected Framework
- Suggesting best practises for data modeling, partitioning, indexing, and workload management

4. Cost Analysis & Optimization

- Performing a detailed cost analysis of the customer's AWS Data Warehouse compute, storage, and data transfer costs using Amazon Redshift, Amazon S3, and AWS Cost Explorer.

5. Continuous Monitoring and Optimization

- Defining the best practises for setting up monitoring and alerting mechanisms for tracking performance metrics and resource utilization
- Providing recommendations for fine-tuning based on changing workload requirements and data patterns

6. Summarize, Train & Knowledge Transfer

- Summarizing the findings of the assessment and presenting actionable recommendations
- Offering recommendations on the best practices for AWS Data Warehouse performance and cost optimization.

KPIs to measure the success of Data Warehouse

- **Load time** - measures how long it takes to load data from the source systems into the data warehouse
- **Query time** - measures how long it takes to execute queries and retrieve data from the data warehouse
- **Availability** - evaluates how often the data warehouse is up and running without any failures or downtime
- **Utilization** - assesses how much of the data warehouse resources (such as storage, memory, CPU, etc.) are being used and how efficiently



Data Warehouse Performance

- **Accessibility** refers to how easy it is for users and stakeholders to connect to and access the data warehouse from their preferred tools and platforms
- **Understandability** measures how easily users and stakeholders can comprehend the structure, content, and meaning of the data warehouse
- **Usefulness** looks at how well the data warehouse supports analytical and reporting needs
- **Satisfaction** evaluates how pleased users and stakeholders are with the functionality, performance, quality, and value of the data warehouse



Data Warehouse Usability



Data Warehouse Quality

- **Accuracy** - how well does the data match the source systems and real-world events
- **Consistency** - does the data conform to established standards and definitions
- **Completeness** - how much of the data needed to support analysis and reporting is available
- **Relevance** - does the data align with current and future business needs

- **ROI** - how much return on investment have you gained from your data warehouse in terms of increased revenue, reduced costs, improved efficiency
- **Alignment** - how well does your data warehouse support the strategic objectives and priorities of your organization
- **Innovation** - how much has your data warehouse enabled or facilitated new insights, opportunities, or solutions for your organization
- **Competitive Advantage** - how much has your data warehouse given you an edge over competitors in terms of market share, customer loyalty, or brand reputation



Data Warehouse Business Impact

Success Story - Top Fortune 100 retailer achieves near real-time visibility to customer insights with robust data warehousing infrastructure

The client is amongst the world's largest retailers with over 11,000 stores and 245 million customers visiting the stores on a weekly basis.

Case Background

- The existing process of reports generation was extremely manual and intensive, with their business users having to wait for several weeks to finally access them.
- They required an effective data warehousing solution to enable granular analysis on large volumes of data at faster response times, thereby expediting the process of customer insight generation.

Sigmoid's Solution Approach

- Sigmoid's solution enabled ad-hoc analysis on over 250 TB of customer and POS transaction data.
- Enabled slice-and-dice on data to deliver granular customer information such as weekly trends, penetration in terms of customer reach, etc.
- Highly effective data warehousing and management empowered 500+ business users to have full interactive access with a single source of truth (SSOT) to run ad-hoc queries at sub second response times, thereby saving 1000s of end-user hours.
- Completely automated process, removed existing manual labour, inter-team dependencies and delays in accessing data.

< 5 secs

Average Query
Response Time

60%

Faster reporting saves
1000s of end-user hrs

250 TB

Data volume processed

Data Sources:

Customer (Experian), POS transactions, Item data, SKU, Store, Calendar information, Weather channel, and others.

Engagement Models



Sigmoid's Engagement Models

Project Based



- Starts with consulting/scoping (2-3 weeks)
- Delivery Program Management
- Interim review
- Success criteria met and IP handover
- Option to continue with product support
- Fixed bid contract
- 3-5 months duration given complexity of problem

Benefits

- Cost effective
- KPI/SLA/Outcome driven
- Suitable for Fixed scope of work
- Less overheads

Staff Augmentation



- Understanding of skill requirements
- Profile match and rate card
- Onboarding and monthly billing
- Focused training based on client tech stack
- Project Management support
- 10% backup resources unbilled and trained

Benefits

- Scalability
- Flexibility in resourcing
- Ability to change/redefine scope

Hybrid-Flexi Model/Data Labs/CoE



- Mix of project and staff augmentation engagements
- Requirement gathering
- Requirement classification - as project or staff augmentation
- Joint delivery plan
- Secure resources internally from Sigmoid and bill monthly
- Dedicated PM, Engineering Managers
- Dedicated Management Consultant(s)
- Dedicated Team Leads and Product Owners

Benefits

- Cost effectiveness by focus on output
- Ability to change/redefine scope/Change requests
- Risk/Reward linked to KPI/SLA


Thank you



Email: abhishek.vora@sigmoidanalytics.com



Website: www.sigmoid.com




'India Future Unicorn Award'
in Data Science category by Hurun India


Global presence:



USA (NY, SF, Dallas,
Chicago)



EU (Amsterdam,
London)



India
(Bengaluru)



LATAM (Lima)