Key takeaways from the webinar Reduce AWS costs of high volume ETL pipeline by up to 65%



Drivers & challenges for cloud migration

What is your main motive for migrating to a public cloud?

Scalability	47%
Long term cost structure	40%
Short lead time	13%
Peer pressure	0%

(Source: Sigmoid Poll, Dec 2020)

Which of the following challenges your organization is facing?



Sigmoid's cloud optimization framework

DEFINE

Requirement/scope creation

Security and cost optimization

Budget estimation

	· · · · · · · · · · · · · · · · · · ·		()	
	Tagging coverage	Services' spending	Process adherence	
-				
🖹 ANALYZE	Process gaps identification	Architecture technology options	Apps/data pipeline costs	
ြို့ ОРТІМІΖЕ	Standard housekeeping practices	Operational changes implementation	Architectural developments	
n GOVERN	Cloud cost accountability	Monitoring and anomaly detection	Cost to business value mapping	

Focus areas to maximize ROI from your AWS infrastructure



1) Matching supply with demand:

- Time the frequency of expensive ETL jobs to the analytics activities they support.
- Migrate infrequently accessed data in the warehouse from local storage to S3



3) Cost-effective resources:

Choose the right compute solutions

- Ensure the right storage class is chosen.
- Consider Spark instances for Transient workloads



2) Expenditure awareness:

Consider tools like

- AWS billing console, AWS cost Explorer & AWS Budgets
- Monthly AWS invoice, 3rd party AWS billing file analysis or DIY dashboards



4) Optimizing over time:

Optimization is an ongoing evolving exercise.

- Establish a workload review process.
- Review and implement services based on need

How we optimized our cloud cost & improved system performance - a case study

Sigview - Sigmoid's plug and play real time data analytics platform uses open source technologies with Cloud native services such as Amazon's Kinesis and EKS. We wanted to reduce our cloud costs and at the same time improve the query performance and overcome scalability issues.





- User filters migrated from Spark to Elasticsearch
- User searches returned within 4 sec on 1.5 BN records
- The infrastructure cost reduced by 25%

Migrated stateless UI microservices to Kubernetes with spot machine

Migrated from Kinesis to Kafka (Single-AZ) setup with Kinesis as backup

Began with Spark clusters with 30% on-spot, but later shifted to 100% spot machines for interactive queries

Reduced cost by 30%

Infrastructure cost reduction 40%
Management overhead reduced by 20%
Increased throughput to 750 mb/sec

Cost reduction by 50%



"We have seen that 1 in 5 companies are overspending on their cloud costs compared to their forecast"

Mayur Rustagi CTO and Co-Founder, Sigmoid

"We were able to enhance our customer experience significantly by improving the query performance by 40%"

Jagannath Nikam Enterprise DevOps Architect, Sigmoid





"Cost allocation is a shared tenancy model & requires a combination of logging and cost modelling to determine the correct price per transaction or workload"

Kousik Rajendran Solutions Architect, AWS

About Sigmoid

Sigmoid builds data solutions leveraging Open-Source and Cloud to power real-time intelligent decisions at scale. Our data engineers build and streamline data processing pipelines to continuously ensure fast, powerful, and reliable processing of large volumes of data. With Engineering and DevOps teams, we bring in the last mile connectivity to embed the data solutions in the existing production environments.

Our cloud capability









aws

For a free consultation to assess the cost saving potential of your AWS bill send us an email: marketing@sigmoid.com