**Name of the customer (Internal or affiliate case studies will not be accepted)**

Cencosud Argentina

**AWS Account ID (Will be used to verify AWS service usage)**

257946566499

**Problem statement/definition**

Cencosud Argentina is a retail company that has several subdivisions according to different business lines. In this case we worked specifically with the MDH division. MDH data structures were supported by Teradata, and Cencosud wanted to migrate their data seamlessly to a licence free solution. The processes managed or supported by the Teradata infrastructure where mainly data warehouse storage and ETLs for Easy and Blaistein bussiness divisions, and this data was used for OLAP purposes.

Thus, the requirement was to migrate the current Teradata structure, both the storage and the ETL processes, to a license free environment capable of supporting their current processes and give further functionality, like real time processing.

**What you proposed**

Once we approached Cencosud and studied their cases and their current solution, we assessed the ETL processes and their general architecture, we designed a complete solution that would allow them to migrate their current use cases to the cloud. In this context, there were several considerations to take into account:

- The Data Warehouse, used to make fast OLAP queries, is going to be migrated to a Redshift Cluster.
- The ETL process is going to be managed through Spark scripts being done through an EMR cluster. Scripts will be managed through Airflow and Livy.
- For archival, long term storage, and data lake structure, we proposed the use of S3.
- For queries on S3 data we proposed the use of Athena.

There are several other services related to the more inner structure of the solutions that will be expanded in the next sections.

**How AWS services were used as part of the solution**

For data storage and data lake support: S3
The ingestion of data was done through: Data Migration Service (DMS) and AWS CLI
The quality of data and ETL processes where carried out with EMR con spark
To manage the metadata and the data catalog, we used: Glue y Crawlers
To orchestrate and check logs and data flows: eventos de CloudWatch, AWS Lambda and SNS

To query data, explore and several use patterns: AWS Athena
For data governance: IAM and policies
To support OLAP queries to data: Redshift

**Third party applications or solutions used**

Apache Airflow
Apache Spark
Apache Livy

**Start and end dates of project (Case studies must be for projects started within the past 24 months, and must be for projects that are in production)**

Start: Marzo 2020
End: Junio 2020

**Outcome(s)/results**

With the Cloud DW, Cencosud was able to continue their daily operations, their decision making process, supported by OLAP queries and dashboards created by their BI area, in a more reliable and stable infrastructure.

Cencosud achieved the same operational efficiency they had with Teradata significantly reducing their costs. They also gained the capability to handle real time processing and to easily upscale when needed. They have future projects to introduce real-time processing of data to carry out certain advanced analytical studies like fraud detection.

**Lessons Learned**

The development process was carried out under a really tight budget. This challenged us to keep the solution as cost optimized as possible. It was an interesting challenge, that made us manage an evolving architecture and infrastructure, monitoring as closely as possible the resources usage.

## 2.2Architecture Diagrams

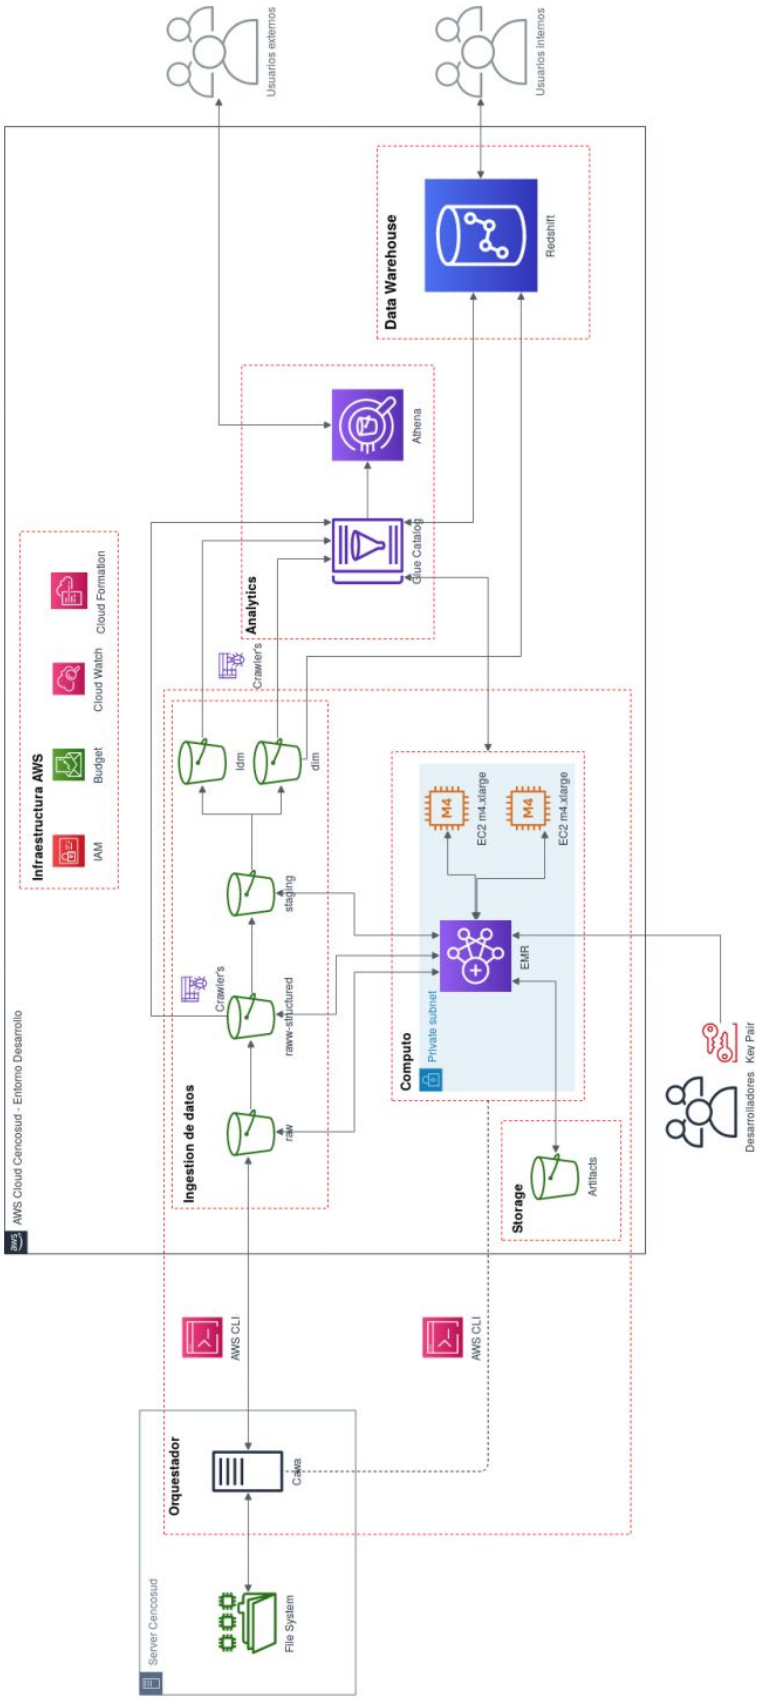

The selected region is North Virginia (us-east-1), with subnets ranging through different Availability Zones. The CIDRs used were:

- 172.17.0.0/24
- 10.0.0.0/16

El acceso a las distintas áreas y servicios de AWS desde la consola se hace a través de single sign in, específico para Cencosud. Además, el acceso externo a los servicios de forma programática sólo puede hacerse a través de una conexión VPN o de las conexiones internas propias de las subredes de la VPC.

**RED-001 - Data storage is optimized based on metrics and patterns**

Sorting keys were selected by first designing the most common queries to be made, thus allowing to better partition the data and optimize the query process. The queries to be made on this data are, mainly, related to datetimes columns and to country columns. The queries are mostly based on timestamp, so the datawarehouse was structured separating countries in different tables and sorting by timestamp.

Compression was managed by Redshift automatically in the following ways:

- Columns that are defined as sort keys are assigned RAW compression.
- Columns that are defined as BOOLEAN, REAL, or DOUBLE PRECISION data types are assigned RAW compression.
- Columns that are defined as SMALLINT, INTEGER, BIGINT, DECIMAL, DATE, TIMESTAMP, or TIMESTAMPTZ data types are assigned AZ64 compression.
- Columns that are defined as CHAR or VARCHAR data types are assigned LZO compression..

The maintenance of the Redshift cluster includes regular vacuuming, according to the frequency and quantity of data loaded. For this particular case, and according to the work hours and peak consumption we suggested a nightly schema.

**RED-002 - Redshift Database User Access Management and Security is following best practices**

The access to Redshift is managed through IAM. Permissions and passwords policies are currently audited and managed by the IT personnel of Cencosud, according to the company standards and, in general, in accordance to AWS best practices. There are no temporary credentials. During the development process, we set up the following roles:

- Role for EMR to read/write on Redshift Cluster and to access S3 to read and write files during the ETL processes.
- Different Roles for Glue to crawl and read write in order to carry out the processes related to certain buckets.
- Role for EC2 to access EMR cluster (Airflow connection to Livy in EMR)
- I am users with permission to access Athena and make queries.
- A role for services and applications that would query the data in Redshift , only giving permissions to list and read data.

Data in Redshift cluster was encrypted using KMS service.

**RED-003 - Workload Management is configured properly to meet application needs**

Workload Concurrency settings were set according to CencoSud requirements. At the moment of handling the production stage the concurrency is pretty low since the queries carried out are simple and fast. However we have set a queue order priority, handling

queries with low requirements in a low memory queue and heavier ones in a high memory queue. The total concurrency level is set at 12, with two queues with 5 concurrency and one with 2 concurrency. To check workload management we evaluated the following metrics:

- Query CPU time
- Blocks read
- Query execution time
- Query queue time
- CPU usage
- Return row count
- Segment execution time
- Query priority

**RED-004 - Solution Composition Requirements**

CencoSud use case involved the migration of a Teradata data warehouse to Redshift.

Te development process involved the coding of all ETL processes carried out through Teradata, using spark as a framework. Spark jobs were then carried out by an EMR cluster, starting with two slaves nodes and one Master, since the workload was not so intense. The ETL processes moved data from a Raw S3 zone to a trust S3 zone, and after that to the Redshift Cluster. This cluster was used for OLAP queries by end users.

The acceptance tests included evaluating the results from ETL processes so that the times and transformations were exactly the same, or better. The queries performance was also evaluated. Every aspect of the architected solution outperformed, at least in a 5%, performance to the previous Teradata solution.

## General

### ACCT-001 - The root user is secured

Root User has not been assigned access keys and has been only used to manage the first accounts, admins and power users. MFA is enabled on the root user.

Each user has been assigned a user, and services that interact with other services have been assigned roles. Policies and group permissions have been elaborated following the best practices of the well architected framework.

### ACCT-002 - Account contact information is set

Account contact information is correctly set up to a corporate email address.

### ACCT-003 - AWS CloudTrail is enabled

Cloudtrail has been enabled as a best practice, since it allows to manage and control the work history, contributing to the prevention and debugging of every process related to the use of AWS Services.

**Operational Excellence**
Requirements in this category relate to the ability of the APN Partner and the customer to run and monitor systems to delivery business value and to continually improve supporting processes and procedures.

**OPE-001 - Metrics are defined for understanding the health of the workload**

In order to better understand the health of the workload, both cloudwatch and cloudtrail logs were analyzed. We designed metrics that allowed to assess the number of lambda functions that were successful, the number that failed, most common fails, times consumed by each process, amount of data processed and stored by each process, and services up and down time, with the additional metrics of cpu and I/O utilization, throwing alerts on peaks.

**OPE-002 - Workload health metrics are collected and analyzed**

As stated before, the metrics are collected through cloudwatch and cloudtrail and stored in an s3 bucket for further analysis. This analysis is crawled by glue and queried through athena. We have recommended the use of AWS QuickSight in order to facilitate the reporting of these metrics. The AWS Personal Health Dashboard is also used to monitor certain workload health metrics, due to the easiness of use.

**OPE-003 - Operational enablement**

The handover process was done through documentation and a personal knowledge transference, including operation related to redshift management, the etl processes, logs generated and how to read and understand them. There were several training sessions carried out with the IT people responsible for the maintenance of these tasks.

**OPE-004 - Deployment testing and validation**

A Development / Production schema was used to test the correct implementation of every part of the process, including ingestion and etl jobs. Development was carried out and, later on, after testing occurs in the development environment, the development team merges work to the production environment and tests are done in the new environment again. This schema was carried out through out the development process, and certain automation was done through cloudformation

We handled over the process to Cencosud IT teams once we reached production stage for the whole process, training them in the use, monitoring and further development in case they wanted to add new functionalities.

**OPE-005 - Code assets are version controlled**

The code is being version controlled through the use of git and bitbucket as cloud repository.

**OPS-006 - Application and workload telemetry**

Redshift and the ETL processes carried out log to Cloudwatch and the logs are also stored in S3. The logs include execution details to facilitate the debugging process. Redshift metrics logged include cluster status, concurrency and cpu utilization.

**Security - Identity and Access Management**

Requirements in this category focus on best practices around AWS Identity and Access Management (IAM) and other identity and access management systems owned by the APN Partner.

**IAM-001 - Access requirements are defined**

The account provided by CencoSud has been created by the IT team and they manage de permissions. Roles and additional permissions are also managed by CencoSud team. We have recommended to give us specific permissions, according to the specific needs for every task carried out.

**IAM-002 - Grant least privileges**

Policies created were given fine grained access to the services and permissions needed. As an example, we submit an example policy to unload data from Redshift to S3:

```
{
  “Version”:”2012-10-17″,
  “Statement”:[
    {
      “Effect”:”Allow”,
      “Action”:[
        “s3:PutObject”,
        “s3:DeleteObject”,
        “s3:ListBucket”
      ],
      “Resource”:[
        “arn:aws:s3:::redshift-testing-cencosud*”
      ]
    }
  ]
}
```

**IAM-003 - Static AWS Access Keys are not used for programmatic access.**

There are several ETL processes that use API calls, and they are authenticated through different Roles specifically adjusted to the use case. The IAM roles are defined by CencoSud IT Team.

**IAM-004 - Unique non-root credentials are used for interactive access.**

Each developer access the account with a different IAM user and credential. Credentials were not shared between individuals.

**Security - Networking**
Requirements in this category focus on security best practices for Amazon VPC and other network security considerations.

**NETSEC-001 - Security groups are tightly scoped.**

Security groups are restricted to the minimum access policies. An example security group, for airflow, receives inbound traffic only from the IP that will connect to it vía VPN and has outbound to EMR cluster by IP.

**NETSEC-002 - Data that traverses the Internet is encrypted in transit.**

There are no endpoints traversing internet.

**NETSEC-003 - Data stores are in private subnets.**

This project works with RedShift and S3 as data stores, and they are in all in private subnets.

Security - IT Operations
Requirements in this category focus on IT security operations best practices including logging, monitoring, incident response, and data classification.

**SECOPS-001 - Cryptographic keys are managed securely.**

Data is encrypted through the managed AWS services. No user provided keys were used.

**AWSAPI-001 - Official AWS SDKs are used to call AWS API endpoints.**

Programming Language used was python, pyspark and scala. The access to aws services were done vía AWS CLI.

**Reliability**

Requirements in this section focus on the ability of the solution to prevent, and quickly recover from failures to meet business and customer demand.

**REL-001 - Deployment automation.**

The Deployment has initially been developed through cloud formation and is now in process of being migrated to terraform, due to client requirements.

**REL-002 - Availability requirements are defined for the solution.**

In relation to individual zone failure, RTO and RPO is automatically managed by AWS since it uses replication and continuous backups to enhance availability and improve data durability and can automatically recover from node and component failures.
In case of availability zone disruption, the RTO and RPO is managed through snapshots storage, with a periodicity according to the requirements set by the client. In this case, there is an automated snapshot done daily. In case of a availability zone disruption, the RTO is around 2 hs and RPO max is less than 24 hs, since is the snapshot available. The process taking into account the restoration of the snapshot saved in S3.

**REL-003 - The solution adapts to changes in demand.**

The provided solution adapts to changes in demand. ETL jobs are carried out mostly through Spark Jobs, orquestated by Airflow and run on an emr cluster. This jobs are limited by the size of the EMR cluster, and this cluster has been created with 2 nodes, according to the budget limitations of the client. It will be scaled by the client according to the needs and budget. In the same sense, S3 allows scalability in a transparent way. Redshift nodes where not set on auto scaling since the client wanted to manage them manually in order to keep the cost as tight as possible

**Cost Optimization**

Requirements in this category relate to the APN Partner's ability to help customers run systems that deliver business value at the lowest price point.

**COST-001 - Total cost of ownership (TCO) analysis or cost modeling was done.**

The TCO initially submitted to the client included the cost estimation over several services that were, later on, tailored to the reduced budget and specific requirements of the processes. The improved and adjusted TCO for the actual production usage includes 4 dc2 large nodes at a monthly cost of 730 U$S monthly. We suggested the use of upfront payments to reduce costs. The TCO also includes an EC2 40 hs per week at around U$S 20 monthly, and the use of a two nodes m4.xlarge EMR cluster, at a cost of around U$S 130 monthly. Costs are being monitored and estimated on a per request basis, according to the client needs.