

CASE STUDY

Delivering on the AWS Promise:

60 Times Faster at 10% of the Cost

Delivering on the AWS Promise: 60 Times Faster at 10% of the Cost

AT A GLANCE

The Customer

Telarix, the market leader in telecom interconnect business optimization

The Opportunity

Significantly reduce AWS cost to support market expansion

The Solution

dbSeer showcased that by re-architecting Telarix's event processing engine and leveraging AWS elastic services and open-source technologies, unlimited scalability can be achieved at a fraction of the cost

The Results

- 90% cost reduction
- 60X faster processing

The Customer

Telarix is the world's leading provider of business-to-business and OSS/BSS management solutions to improve efficiency and productivity. Through software and services, Telarix empowers over 4,000 communication companies worldwide with unprecedented levels of visibility and control over their entire voice, video, data, and SMS business.

The Opportunity

To provide actionable business intelligence, Telarix processes hundreds of billions of events (Call Detail Records), which contain data such as call duration and time of day. The results are needed as soon as possible to facilitate near real-time optimized trading, billing, and routing in accordance with the many agreements between telecom companies. Using their existing event-processing engine, Telarix requires large, dedicated clusters to generate the results in a timely manner. The near-infinite elastic resources available at AWS presented an opportunity to validate significantly cutting the event-processing cost, as well as to reduce the time required to produce the results. This would enable Telarix to better serve their customers and expand their business.

Distributed Architecture for Processing Call Detail Records

"Traditionally, Telarix uses commodity servers to do this extensive processing. If we needed to double speed, we had to double the hardware, which required doubling costs. While it was possible to expand capacity, it came at consistently high costs," said Aravind Venkateswaran, Chief Technology Officer, Telarix.

Maximum workload may occur only once or twice per month, but to meet that requirement, Telarix had to size the infrastructure accordingly. Unfortunately, this required paying for full capacity at all times, despite only using a small fraction in most cases.

Telarix migrated the event-processing engine to Amazon Web Services EC2 instances to reduce costs, but did so by copying and moving their existing on-prem infrastructure. While this “lift and shift” approach translated costs from capital to operating expenses, without rearchitecting, the company was unable to maximize savings and achieve additional benefits.

The Solution

Telarix chose dbSeer to help with the proof-of-concept work and rearchitect their AWS infrastructure. The new architecture changed what had been a SQL process running on Microsoft SQL Server into a Python-based process built using open source libraries, Linux based EC2 instances, and various other AWS services to perform the same tasks. By using AWS Simple Work Flow (SWF), AWS Simple Queue Service (SQS), and AWS Simple Storage Service (S3), dbSeer created a system that could deallocate compute resources based on data volume and availability of those resources.

Whenever unprocessed files are placed into storage in S3, messages are sent via SQS to a controller instance that makes a decision about whether additional compute resources (EC2 instances) are necessary; if so, how many; and then assigns the files to the individual instances according to user-defined rules. Once the files are downloaded, the Python code performs the same data transformations as the SQL process had, before writing back the results to an output folder also located in S3.

To recreate the SQL data processing procedures, a series of JSON files are used to define the inputs, outputs, processing logic, and reference data associated with each individual function. These JSON files, when imported into the processing program along with the data files, quickly query the data set and perform the necessary transformations and enrichments. The JSON files also allow for a high level of flexibility and customization within a structured format.


The Results

By these means, dbSeer achieved significant results in the realms of efficiency and elasticity. The automatically scaling nature of the solution created a balance between processing many loads on fewer servers, or accelerating the processing by distributing workloads across even more servers. With less server idle time, the new architecture provided increased speed while at the same time decreasing costs.

Cloud Architecture that Achieved True Elasticity

dbSeer’s solution was able to achieve true elasticity. The amount of compute resources used from hour to hour, or workflow to workflow, was able to scale up with increases in the volume of data, and then scale back down once a peak period had passed. This dynamic allocation of resources is a hallmark of cloud solutions, and one of the main benefits to switching to a cloud based architecture. It eliminates the need to over-invest in hardware to handle expected but infrequent periods of high traffic.


The elasticity of the system also makes it easier and cheaper for Telarix's data volume to grow in the future, without the need for additional investments in hardware to achieve the same throughput. As Mr. Venkateswaran said, "The new proposed architecture can truly scale without limitations. Currently, we have common hardware and common software in the middle. These shared components could create bottlenecks, preventing us from scaling reliably and securely, which enables us to move even more of our compute processing to AWS."



90% Cost Reduction, 60 Times Faster Processing

The new AWS processing engine was tested on a variety of server configurations and volumes.

As expected, the architecture demonstrated linear scalability, as the processing speed increased proportionally with the volume of servers added. Because of the low cost of the Linux based instances, and the volume focused pricing model that AWS offers, the processing cost per million CDR records could be reduced by up to 90%. Furthermore, the ability to distribute the workflow across hundreds of low cost, efficient machines meant Telarix can potentially see a processing speed that will be up to 60 times faster than their current architecture. Finally – and most importantly – the output generated by dbSeer's proposed new architecture was compared against output from the prior system, and the data was 100% validated, showing no changes..



Additional Benefits of the AWS Cloud

There were also some less obvious benefits with migrating to the new, more heavily cloud based architecture. The first gain will be in durability. With S3, Amazon guarantees 99.999999999% durability of your data. This is achieved by replicating the data across multiple data centers, providing significantly safer storage than relying on an individual server instance.

Another potential benefit of the new architecture is the ease of deployment and modification. Each EC2 instance runs as an exact copy of a single Amazon Machine Image (AMI). The deployment team can spend time configuring just one server, and then be sure that every worker machine running will be created and configured with the same specifications, instructions, and files. Similarly, when making changes to the process, the deployment team only has to make changes on one instance, and then create a new AMI, which will propagate any changes to all instances created with that new image. This significantly simplifies the process of ensuring that your worker processing machines remain up to date.

Finally, even further cost savings are possible if Spot Instances are used, instead of On-Demand Instances. Amazon offers a variety of pricing models for purchasing EC2 servers. On-Demand instances, which were used in dbSeer's initial testing, offer a set hourly price for their usage. However, Amazon also offers Spot Instance pricing, which charges the user a fluctuating price based on the availability and demand for Amazon compute services at the time of request. dbSeer found that Telarix could save an additional 70% by using Spot Instances for its workloads. Even though spot instances pose the risk of being terminated with spot price changes, the new decoupled architecture enables the usage of spot instances due to its graceful, fault tolerant nature.