LARGE MANUFACTURER SECURES DATA MIGRATED TO CLOUD WORKLOADS

One of the largest makers of energy technologies in the world, this manufacturer provides gas and electric equipment globally. The organization's application IT department is a key business partner that creates and maintains core systems for managing complex customer contracts.

Aging Applications and Database Scalability Limits

Everytime an equipment is bought by one of their customers, it is usually accompanied by a 15-30 year services contract with managed services and other varying commercial terms. The services are critical to ensuring turbines and power equipment are running optimally.

The application managing these customer contractual services had been in operation for more than 15 years. It was built using Java and Spring framework and was only designed to handle a couple of hundred contracts at the time.

Over its use, the application had to evolve to support increasingly complicated contracts with more variables. For example, customers with multiple fleets wanted to be able to optimize parts management between sites. The on-premises Oracle database reached over 3 TB in size. The reporting database grew to another 1 TB. It became almost impossible for the application to scale to support modern requirements.

Millions of dollars of potential revenue were being left on the table, and the business felt significantly hindered by the application. Additionally, parts could not be optimized as needed, and the resulting outages were a direct cost for the organization, as per SLAs with their customers.

Data Lakes and Cloud Transformation

Senior architects in the department decided that they would “lift and shift” to Amazon Web Services (AWS). They started building a data lake which needed to host production data plus development, testing, and staging environments. It also needed to support very large reporting instances.

However, a wrinkle appeared in their modernization quest. The legal team objected to data being stored in public clouds. Customers also did not like the thought of their data being exposed in the cloud. The organization in the past has been subject to many cyberattacks, similar to most organizations of their size. There was also concern that cloud administrators would have access to any clear data, and any breach of cloud infrastructure would expose sensitive data.

After exploring multiple options, the technical and legal teams came to the agreement that if data was properly encrypted in cloud databases, they could move forward with the project. Customers also approved of their data being stored in the cloud if it was encrypted. A key challenge now was finding an encryption solution that could scale to support over 8,000 tables and many terabytes of data without a significant performance impact.

The data team began planning migration of their Oracle data to Postgres in AWS RDS. The data had to be encrypted before it ever touched the cloud. Hence, they evaluated multiple data protection solutions. Across the organization others had used tools such as pgcrypto, however, those suffered from significant performance and scalability issues.

Highly Performant, No-Code Data Encryption

After a rigorous evaluation the team chose Baffle Data Protection for Databases as their data encryption solution. Compared to other approaches, this provided them with several key advantages:

- **Security** - offered role-based data protection, protecting their crown jewels from even cloud administrators
- **Scalability** - Baffle was able to handle their current and future data needs with ease
- **No- code** - applications did not have be changed or specially coded, all other solutions required an API that applications have to program for and maintain
- **Control over keys** - unlike other solutions Baffle did not require that encryption keys or tokens be stored by them
- **Vendor access** - similarly, Baffle never has access to their applications or data
- **Confidential computing** - encrypted data is processed seamlessly by commercial off-the-shelf analytics and reporting tools without having to decrypt the data
- **Customer service** - the team at Baffle was highly responsive and collaborative, helping accelerate deployment and resolving all issues expeditiously

With Baffle’s unique data proxy approach, their on-premises data is encrypted on-the-fly as it moves through AWS DMS. By being able to select only the sensitive data fields/columns to encrypt, they were able to improve performance and efficiency by a factor of four.

Reporting is now performed in Tableau, having moved from a Java based system which had gotten expensive to manage and maintain. A DMS replicator copies production data to the reporting instance which is also in Postgres. Since the data is still encrypted, they then use Baffle’s Advanced Encryption to consume the data in Tableau without having to decrypt it. This provides the highest level of security with data encryption while at rest, in transit, and now also in use.

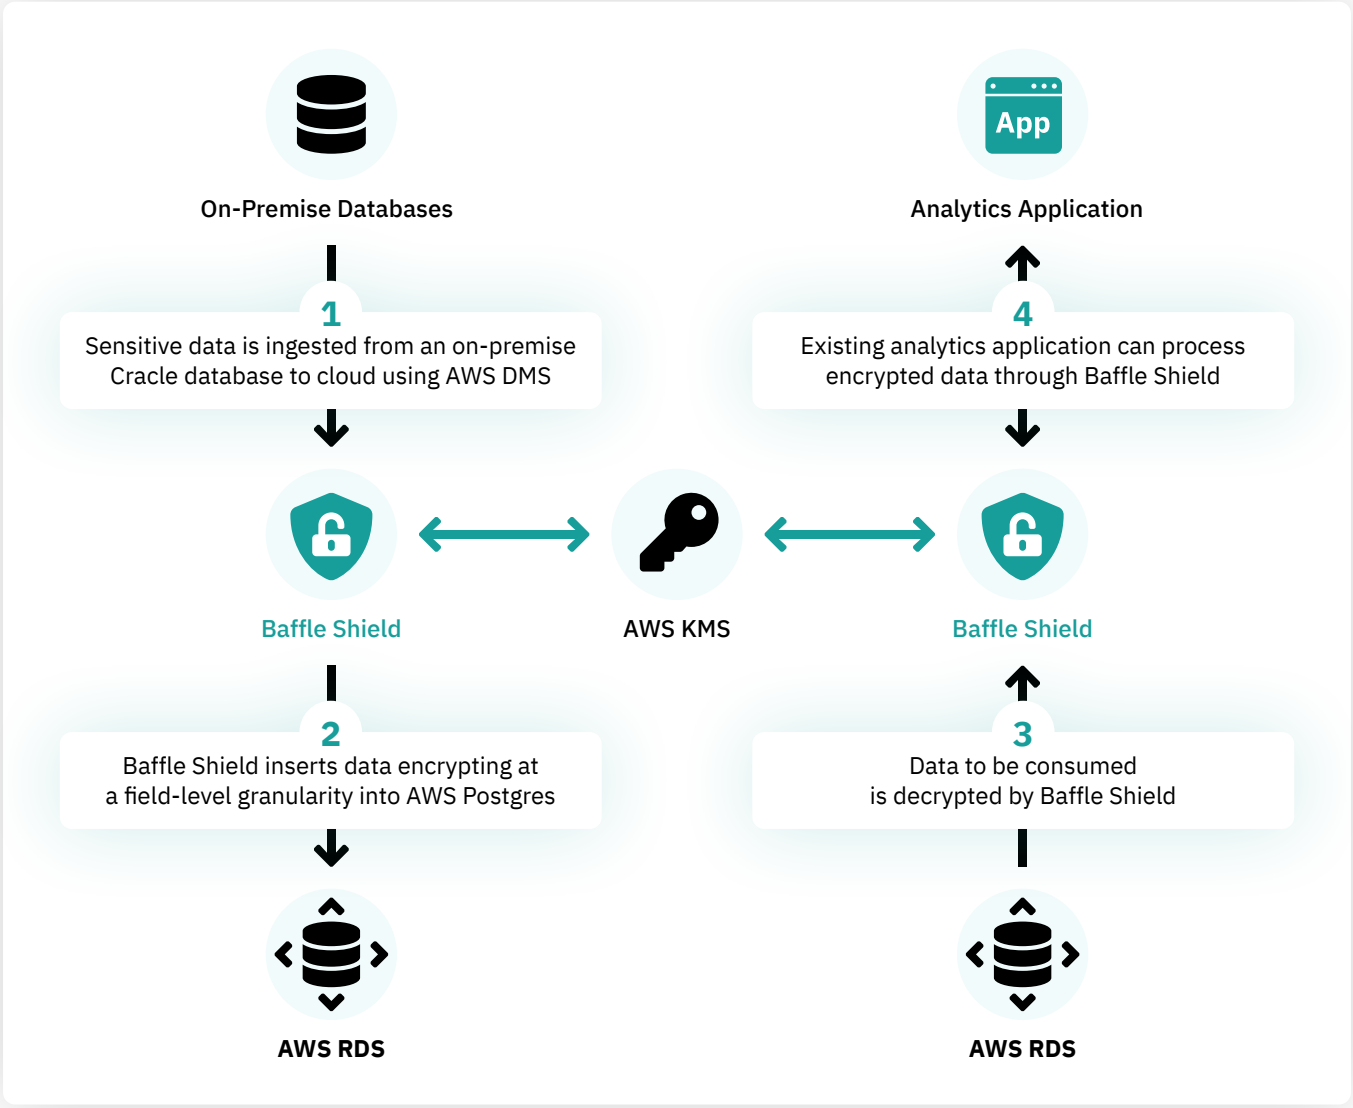

The new architecture has been successfully deployed and has been in production for several months. The IT, security, legal, and business teams have all gleaned tremendous value and are very supportive of it. In fact, they are now planning to apply the same approach to numerous other applications across the organization in the near future.

About Baffle and AWS

Baffle and AWS have collaborated closely to create the leading data protection solution for cloud data stores. Baffle is the leading provider of no-code data protection solutions, is a proud member of the Amazon Partner Network ISV Accelerate program, and is available on the AWS Marketplace.

aws Qualified Software

aws PARTNER

aws PARTNER Security Software Competency

info@baffle.io
https://baffle.io
3979 Freedom Circle, Suite 970
Santa Clara, CA 95054