

# Production-Real Testing Without Data Risk

## Executive summary

Digitail's testing cycles either missed real-world edge cases against synthetic data, or copied production databases and risked exposing sensitive pet-parent and patient records. Threading Clouds designed and automated a secure data masking pipeline on AWS, giving Digitail's engineers production-realistic test environments on demand — with zero PII exposure and zero manual effort.

## The challenge

### Edge Cases

Synthetic test data missed real-world edge cases that surfaced only in production. Bugs reached end customers and eroded clinic trust.

### GDPR Risk

The alternative of copying production databases exposed pet-parent PII to dev environments. Every manual copy was a GDPR liability.

### Slow Cycles

Preparing safe test databases was manual and time-consuming. Developers waited hours for usable environments, slowing release cadence.

## Production-Realistic Testing on AWS, Zero PII Risk

Digitail needed production-realistic test data without GDPR exposure. Threading Clouds built an automated masking pipeline on AWS using DMS, custom PII stripping, S3 lifecycle, Jenkins — turning hours of manual prep into a single click, with zero sensitive data in non-prod.

## DIGITAIL

### About the customer

Founded in 2018, Digitail offers an all-in-one, cloud-based practice management software designed for veterinary clinics. The platform streamlines clinic operations, enhances patient care, and facilitates seamless communication between veterinarians and pet parents through its integrated pet parent app and AI Assistant.

### Ionut Calara

CTO, Digitail

*Threading Clouds' deep expertise in AWS, Kubernetes, and Terraform accelerated our projects. Their documentation made complex processes accessible. The collaboration has had a lasting impact on our operations and our uptime.*



## Ionut Calara

CTO, Digitail

*“Their ability to quickly understand and address our challenges has been remarkable. Beyond their technical expertise, their dedication to thorough documentation has been invaluable, making complex processes accessible and easier to manage for our team.*”

## Why Threading Clouds?

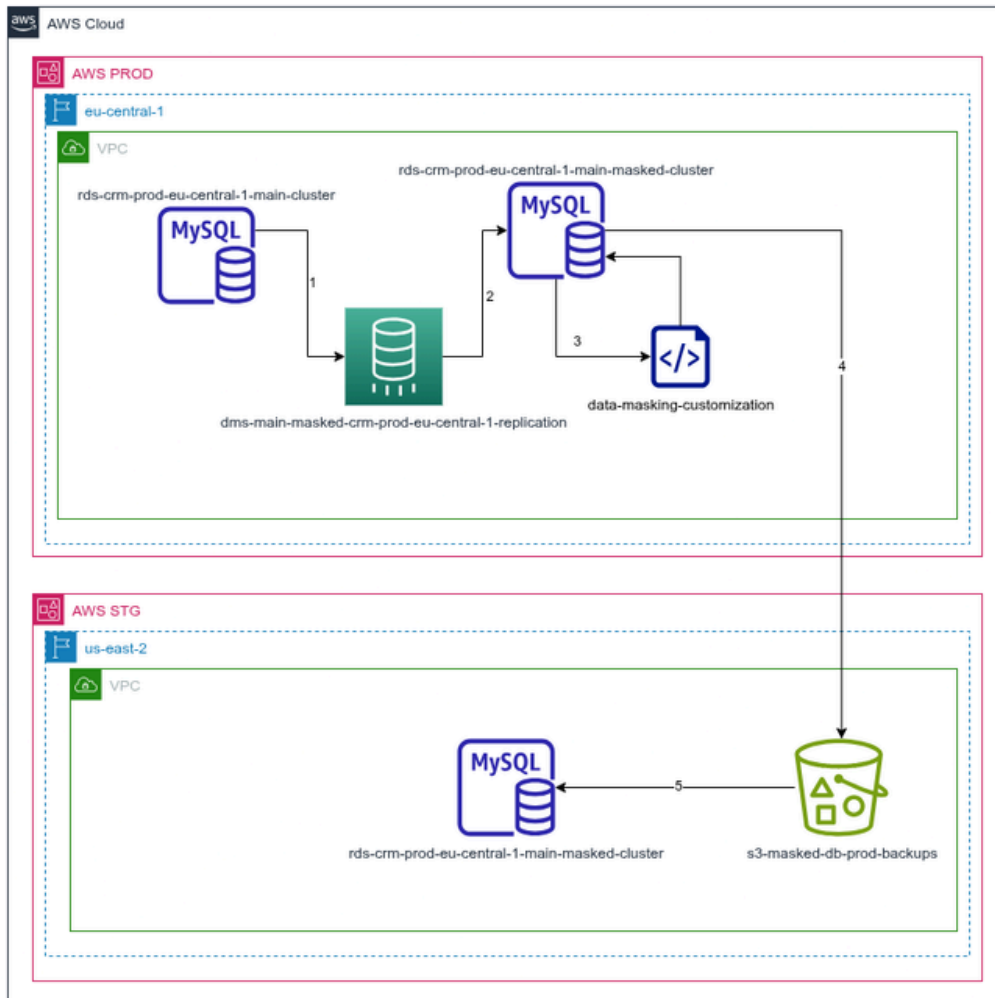
Digitail had already tried the obvious path — standard AWS tools for database replication and masking. They hit a wall. DMS couldn't mask with the granularity GDPR demanded, and no off-the-shelf product handled their schema complexity without breaking indexes and referential integrity. Threading Clouds was brought in specifically because we design custom solutions where packaged tools have failed. The combination of deep AWS database expertise (Aurora, DMS, S3 lifecycle), masking engineering, and Jenkins automation meant Digitail got a single partner who owned the problem end-to-end — from replication through PII stripping to on-demand restore. As an AWS Select Tier Services Partner, Threading Clouds operates this pipeline as part of an ongoing managed service, evolving the masking rules alongside Digitail's data model and the regulatory landscape.

## The solution

Threading Clouds built a fully automated pipeline that solves all three challenges in a single orchestrated flow on AWS. Production data is replicated via AWS DMS, then passed through custom masking scripts that strip sensitive pet-parent and patient PII while preserving the data relationships and edge cases that make test environments realistic. Database indexes and schema integrity are migrated separately so masked copies perform like production, not like degraded snapshots. The masked output is compressed, stored in S3 with lifecycle policies that control costs automatically, and available for on-demand restoration into a consolidated Aurora cluster in the staging account. Jenkins orchestrates every step end-to-end — replication trigger, masking execution, S3 staging, Aurora restore — with no manual intervention, no compliance gaps, and no developer waiting on the data team. The pipeline runs as a managed service: Threading Clouds owns operations, monitors the runs, and evolves the masking ruleset as Digitail's schema changes.

## Architecture Overview

The pipeline spans two AWS accounts for staging and production. In the production account, AWS DMS replicates the live cluster into a dedicated masked cluster, where custom masking logic strips pet-parent and patient PII while preserving the indexes, relationships, and edge cases that keep the data production-realistic. The masked cluster is then exported to an S3 backups bucket whose lifecycle policies age out old copies automatically, and from there the dataset is restored on demand into a consolidated Aurora cluster in an isolated staging account, giving engineers a clean, GDPR-safe environment with no path back to the original records. Jenkins orchestrates every step, so the entire flow runs end-to-end without manual intervention.





## Results

The bugs-in-production problem disappeared. With test environments that mirror real-world data patterns, Digitail's QA cycles now catch edge cases before they reach end customers, eliminating the complaint pattern that had been eroding clinic trust. GDPR compliance risk dropped to zero across all non-production environments — every database copy is automatically masked before it leaves the pipeline. Test database preparation collapsed from hours of manual work into a single Jenkins trigger, removing the bottleneck that had been holding back release cadence. Engineering time previously lost to database wrangling is now redirected toward features that matter to Digitail's veterinary customers.

## Benefits

Digitail now ships with confidence. Product quality improved because engineers test against data that behaves like production, catching issues synthetic data never could. The compliance burden shifted from a manual, error-prone responsibility to an automated guarantee baked into the pipeline — no human in the loop, no chance of an unmasked copy slipping into dev. Storage and compute costs are controlled automatically through S3 lifecycle policies and consolidated test clusters, so infrastructure spend scales predictably as data volumes grow. Engineering time that was previously lost to database preparation is redirected toward building features that matter to Digitail's veterinary clinic customers. Because the pipeline runs as a Threading Clouds managed service, Digitail's team carries zero operational burden — masking rules evolve alongside the data model without taking time away from product work.

### Ionut Calara

CTO, Digitail

*“Threading Clouds support has not only accelerated our projects but also enriched our team's skill set, demonstrating the power of collaboration and knowledge sharing. Their contributions have left a lasting impact, and we are grateful for the expertise and insights he has brought to our team. Collaborating with Threading Clouds through Giorgio has been a pleasure, and his work been quite impactful for our team day to day activity and our uptime.”*



## Next steps

With the automated masking pipeline operational, Digitail is positioned to extend production-realistic testing across new product lines and database schemas as the platform grows. Threading Clouds continues to evolve the masking rules alongside Digitail's data model and the evolving GDPR landscape, ensuring the solution stays ahead of both business requirements and compliance demands. Next on the roadmap: extending the pipeline to support per-clinic data scoping for tenant-isolated testing, integrating the masking flow with Digitail's CI pipeline so feature branches can spin up production-realistic environments automatically, and refining the masking ruleset as new sensitive data categories enter the schema.

## Key highlights

### ▶ Quality Customers Can Feel

Production-realistic test data surfaces the edge cases synthetic data misses. Digitail's QA cycles now catch real-world issues before release, protecting clinic trust and reducing support burden.

### ▶ Compliance

Every database copy is automatically masked before it leaves the pipeline. GDPR compliance shifted from a manual responsibility to a built-in guarantee, eliminating the risk of sensitive data in non-prod.

### ▶ Engineering Time Back on Features

Test database preparation went from hours of manual work to a single automated trigger. The time engineers spent wrangling data now goes toward building the features Digitail's customers need.

## About Threading Clouds

Threading Clouds is an AWS Partner specializing in cloud consulting, DevOps and platform engineering, reliability and operations, and innovation in advanced technologies including generative AI and blockchain infrastructure for SaaS, fintech, and regulated industries. We embed with engineering teams to innovate with confidence by designing, implementing, and managing secure, scalable, and cost-efficient cloud infrastructures.

