DataStealth

Case Study: How a Global Insurer Protects Sensitive Data in Non-Production Environments

Without Breaking Downstream Functionality

datastealth.io | info@datastealth.io

Executive Summary

A leading Global Insurer, with multiple lines of business, required a consistent way to protect production data before it entered test, QA, UAT, and training environments.

An existing solution had already been deployed, but it broke downstream systems, distorted geographic and relational attributes, disrupted relationships between tables, and caused data to stop behaving like production data.

The customer selected DataStealth to deliver a Test Data Management (TDM) platform that could remove identifiable information without compromising the integrity, structure, or usability of the underlying data.

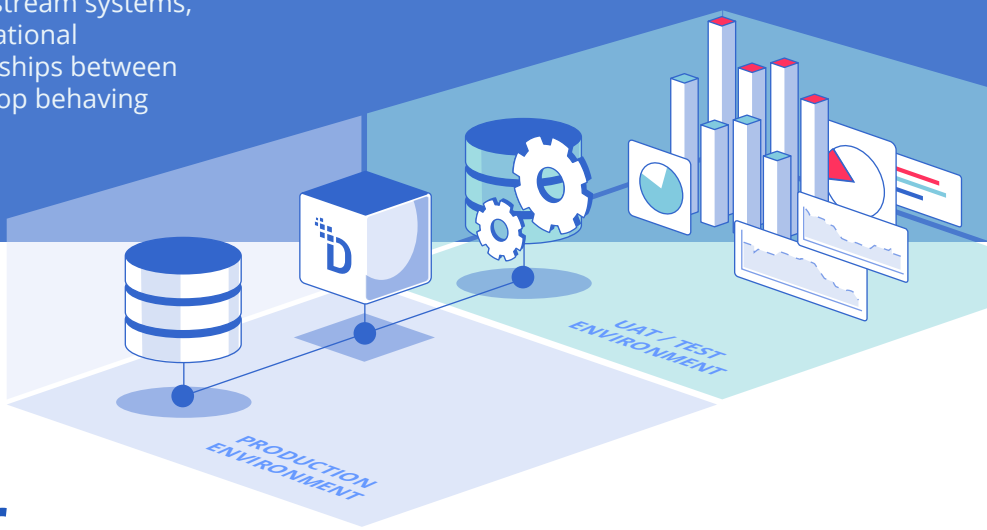

Customer Profile

The customer is one of the world's largest insurance company operating across 11 countries and serving 36 million policyholders across insurance, benefits, wealth, banking, and other areas.

This geographic and business diversity means this organization handles a significant volume of highly sensitive data, including Personal Health Information (PHI) and complex financial data, subject to global and local data privacy regulations (i.e., GDPR, HIPAA, and regional laws). Functional non-production environments are essential for training, onboarding, and change delivery at this scale.

Following a successful pilot in one line of business, our customer standardized on DataStealth as the enterprise-wide TDM platform.

Customer at a Glance:

- **Industry:** Insurance and Financial Services
- **Global Footprint:** 11 countries
- **Employees:** ~37,000
- **Assets Under Management:** $1.4 trillion
- **Data Ecosystem:** Hybrid (cloud, on-prem, mainframe, legacy, SQL)
- **Requirements:** Realistic UAT, QA, and training environments for thousands of staff.
- **Regulatory Landscape:** Privacy, PCI, data residency, cross-border data controls.

The Implementation:

DataStealth was deployed between production and non-production environments, protecting sensitive data by offering a combination of tokenization, encryption, and masking, as the sensitive data moved from production systems into lower-level non-production environments. This included both structured databases and a wide range of system files, including mainframe, copybook, CSV, and database export files, transferred over SFTP and FTPS. As a requirement of the deployment, DataStealth generated repeatable output across every database, file, and application, so all protected data maintained referential integrity, geographic accuracy, format consistency, and downstream functional behaviour.

The Outcome:

DataStealth replaced a legacy solution that routinely broke downstream systems with a robust platform that can preserve geographic, relational, and functional integrity across all databases and documents, eliminating the failures caused by inconsistent or outdated protection capabilities and logic, protecting data across legacy, mainframe, cloud, SQL, and SFTP/FTPS pipelines. Teams receive realistic, production-like test data without the operational risk or maintenance burden associated with manual scripts or older TDM processes.

The Bottom Line:

By eliminating the need for teams to build and maintain scripts, DataStealth reduced resource strain, prevented costly rework, and removed gaps created by frequent production schema changes. It eased the load on teams already facing budget constraints and competing priorities, delivering consistent protection without pulling strategic staff away from higher-value work. DataStealth now serves as the scalable foundation for TDM across the enterprise.

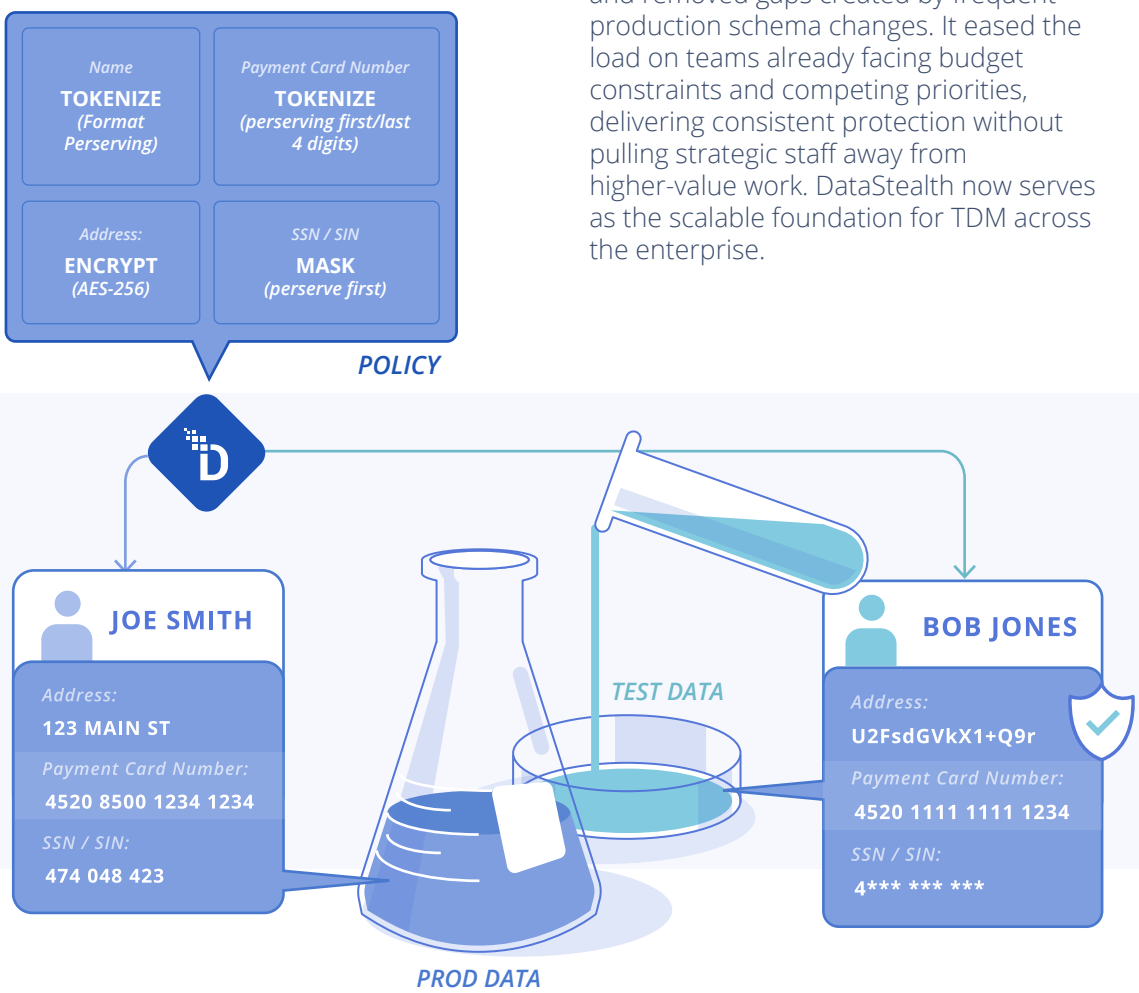

The Solution

Our customer implemented DataStealth as their TDM platform, positioned between production systems and downstream non-production environments.

Instead of relying on teams to clone production datasets and apply scripts after the fact, DataStealth protects data in transit, while being moved from the production to non-production environments, ensuring that unprotected information never reaches test, QA, UAT, or training systems.

The process is simple:

1. DataStealth reads the data directly from production sources, including large SQL databases, mainframe extracts, and file-based feeds.

2. Protection is applied in transit, using rules defined during the discovery and analysis phase.

3. Only protected, non-identifiable data is delivered into the target environment.

There is no duplication of production datasets, and no manual scripting required.

Referential Integrity

Our customer required test data that behaved exactly like production data across its highly interconnected systems, from core policy and billing databases to platforms, mainframe files, and downstream actuarial processes.

Maintaining referential integrity ensures that all related records across multiple tables remain properly linked after test data is protected. In practice, this means customer IDs, account numbers, and transaction references continue to match and connect correctly, allowing applications, reports, and downstream processes to function exactly as they do in production, without exposing any real sensitive data.

DataStealth also maintains domain-specific relationships critical for insurance. For example, customer addresses remain in the correct geographic region (i.e., the same FSA or area code), demographic attributes like dates of birth remain realistic (for example, same month and year, but slightly shifted day), and dependent/beneficiary linkages are preserved.

Training, pricing, adjudication, reporting, and customer service scenarios continue to reflect real-world behaviour, but without exposing real customer data.

Databases

DataStealth supports all of our customer's production SQL databases, including their central policy system, billing engine, benefits platform, and claims databases, which collectively contain billions of records. When test data is created or needs to be refreshed, our customer sends the production data to the non-production environment, and DataStealth applies the data protection in-flight, during the file transfer. There is no requirement for additional database environments for staging nor any additional database licenses in the staging environments.

Files

The customer's operations rely heavily on multiple types of files, including mainframe data, copybook, batch exports, and other file types across multiple legacy and distributed systems.

When test data is originally created, or needs to be refreshed, our customer sends the production data file to the non-production environment via their usual SFTP file transfer process, and DataStealth applies the data protection in-flight, in real-time, during the file transfer. There is no requirement for additional environments for storage or staging.

This ensures that:

- Batch processes relying on cross-system joins still run correctly.
- Mainframe outputs match the masked identities in distributed applications.
- Downstream actuarial and benefits receive coherent, consistent data.
- Multi-system pipelines (for example, a copybook file feeding a claims system that also references SQL tables) remain fully functional.

The result is a seamless test ecosystem: every system, application, and file receives protected data that behaves like production data, but without exposing any sensitive information.

Benefits Delivered

Production-Accurate Test Environments

The customer's UAT and QA environments now behave exactly like their production counterpart because the data in these lower-level environments looks, acts, and behaves the same as the real data, with one exception. The data in the non-production environments is not real data.

Testing, training, and development teams can run full workflows without any risk of data exposure because there is no real production data in non-production environments.

Referential Integrity Across All Systems and File Flows

Elimination of Script-Based Masking Overhead

Business units no longer maintain custom scripts that are fragile, difficult to update, and prone to missing new fields when production schema changes.

DataStealth centralizes all protection rules in a single platform, removing a major operational burden from already resource-constrained teams.

Faster Expansion Across Business Units

New teams are now onboarded through configuration rather than custom development. Because DataStealth already maintains global masking rules, additional systems such as cloud platforms, legacy apps, mainframe, and distributed databases can be added without rewriting code or re-engineering pipelines.

Data Sovereignty and Residency Enforcement

DataStealth applies protection within our customer's home region before the data is ever accessed by international teams. This approach allows our customer to leverage offshore development and operational resources without violating data sovereignty laws. Because the data residing in the non-production development environments is protected before it's accessed or moved, the sensitive clear-text data never leaves the home region, while offshore users can continue to build, test, and answer questions using functional, compliant data.

Reduced Operational Risk and Improved Reliability

By moving data protection into the flow of data, instead of relying on post-processing scripts in additional environments with additional storage and licensing needs, our customer eliminated the failure modes that previously caused broken UAT environments, rework cycles, and delays. The result is a predictable, repeatable TDM process that scales with the business.

Why DataStealth?

DataStealth's TDM solution protects sensitive data as it moves from production into non-production environments.

Instead of cloning databases and relying upon fragile, manually maintained scripts, DataStealth applies consistent, policy-based transformations in-flight, whenever data is moved from production to non-production, ensuring only protected data ever reaches UAT, QA, or training environments.

For a customer with billions of records, dozens of interconnected platforms, and both legacy and modern systems, this customer used DataStealth to create a single, reliable way to protect sensitive data across SQL, mainframe, and copybook files in the non-production environments, significantly reducing any risk of data exfiltration from these systems and environments.

By moving data protection into the data flow in a programmatic and consistent way, DataStealth gave our customer a predictable, compliant, and highly scalable TDM foundation that supports every business unit, every system, and every integration, today and in the future, as our customer continues to grow.

See How DataStealth TDM Works in Your Environment

DataStealth can help you eliminate data risks in non-production environments, replace fragile masking scripts, and enforce consistent data protection across databases, mainframe files, and cloud platforms.

See how data protection works in a real deployment.

SCHEDULE A DEMO →

Book a demo and walk through your exact use cases.

Our team will show you:

- How data protection in-flight eliminates manual scripts
- How integrity can be preserved across all your systems
- How fast new business units can onboard once DataStealth is in place