



CUSTOMER SUCCESS STORIES

Italy's Largest Bank Modernizes Sovereign Banking Infrastructure with EDB and CloudNativePG



CUSTOMER: INTESA SANPAOLO

EDB customer since 2021

Stefano Morgano

Database Architect, Intesa Sanpaolo Group Services

CHALLENGE: Legacy Oracle infrastructure and centralized database clusters limited scalability, automation, and resilience for Kubernetes-based microservices.

EDB SOLUTION: CloudNativePG operator deployment, EnterpriseDB (EDB) automation and enterprise support

RESULTS:

- 500+ Postgres® instances on IaaS with Ansible
- Phased adoption of CloudNativePG
- Enabled self-healing failover and recovery
- Implemented separation-of-duty security model
- Established sovereign, Kubernetes-native database operations

OVERVIEW

Modernizing banking infrastructure for the cloud-native era

Across Europe, banks are modernizing infrastructure to support cloud-native applications while maintaining strict governance over critical systems. Kubernetes has become central to that shift, enabling automation and scalability without relinquishing operational control over data environments. For financial institutions operating under strict regulatory oversight, modernization must also preserve sovereignty to ensure that infrastructure remains secure, auditable, and institutionally governed.

For Intesa Sanpaolo, that transformation operates at national scale. Serving approximately 13.9 million customers through more than 2,600 branches, the bank runs one of the most extensive technology estates in European banking, including one of Italy's largest Oracle Exadata environments supporting mission-critical services.

As digital banking services expanded, provisioning timelines lengthened, replication processes strained infrastructure resources, and failover required coordinated manual intervention. These constraints limited the bank's ability to efficiently scale Kubernetes platforms, microservices, and sovereign operational models.

Breaking free from centralized database constraints

The bank's legacy estate was deeply tied to centralized database architecture. As the largest Oracle Exadata user in Italy, it relied on complex legacy packages that could not be easily reengineered or migrated, limiting flexibility for new development initiatives.

Operational processes reflected that rigidity. Failover events required high-pressure manual coordination to prevent service disruption. Multisite replication posed additional friction, with 2–3 TB data transfers taking a week or more using VM-based systems. Centralized clusters also constrained scalability, locking infrastructure resources to monolithic hardware environments while driving high operational cost.

These structural limitations created growing misalignment between legacy database operations and modern application delivery models.



Designing databases for microservices and Kubernetes scale

As development teams adopted Kubernetes and microservices to support new banking services, the database infrastructure evolved. The bank shifted toward an independent instance model, moving away from shared clusters to dedicated Postgres environments provisioned per application.

This approach aligned database deployment with containerized application lifecycles, enabling services to scale independently while improving workload isolation. Distributing database environments across different server architectures also improved resource efficiency and reduced reliance on centralized hardware.

“We needed databases that could operate at the same speed as our application platforms. That required an architecture built on automation, independence, and Kubernetes-native operations.”

Stefano Morgano
Database Architect, Intesa Sanpaolo Group Services

CloudNativePG and EDB enable a sovereign, self-healing platform

To operationalize this model, Intesa Sanpaolo began adopting CloudNativePG, a Kubernetes-native operator designed to manage database lifecycle operations directly within container environments. CloudNativePG integrates storage orchestration, replication, failover, and recovery into Kubernetes control planes, enabling declarative database management aligned to cloud-native operating models.

The bank expanded from 10 shared databases to more than 500 independent Postgres instances aligned to application workloads. Parallelized archiving capabilities eliminated serial processing bottlenecks, improving performance for high-volume data operations.

To support enterprise deployment, Intesa Sanpaolo partnered with EDB to implement an Ansible-based automation framework and engage dedicated professional services teams. This collaboration accelerated rollout, standardized configurations, and ensured alignment with governance and compliance requirements.

Security architecture evolved alongside operations. A strict separation-of-duty model ensures that database administrators manage declarative YAML configurations while operating system and cluster-level access remain segregated, reinforcing sovereign control over critical data systems.

Automation reshaped resilience as well. Failover processes that once required manual coordination now occur automatically, often overnight, with teams reviewing alerts rather than managing live incidents.

“EDB’s support allowed us to automate resilience at a level that wasn’t possible in our previous environment. We’ve built a database platform designed to scale, recover, and operate under our own governance model.”

Stefano Morgano

Database Architect, Intesa Sanpaolo Group Services



EDB PostgreSQL AI is the first open, enterprise-grade sovereign data and AI platform, with a secure, compliant, and fully scalable environment, on premises and across clouds. Supported by a global partner network, EDB PostgreSQL AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs where, when, and how they need them. For more information, visit enterprisedb.com.

© EnterpriseDB Corporation 2026. All rights reserved.