

The strategic value of migrating databases to EDB Postgres AI on Red Hat OpenShift



At a glance: why modernize the database layer now



What's driving change: proprietary databases are getting more expensive to run and harder to manage, with licensing, audits, and lock-in costs and risk increasingly difficult to control.



Why EDB Postgres AI: Enterprise-grade Postgres with the features, tooling, and support needed for large, business-critical environments – including compatibility that simplifies moving away from traditional databases.



Why containers now: Kubernetes has matured. It's no longer just for stateless apps – It can run databases reliably, with far more automation and consistency than traditional VM-based approaches.



What changes in practice: A cloud-neutral database platform that scales across environments and supports future needs, including analytics and AI.



Why Red Hat OpenShift: Many organizations already use Red Hat OpenShift as their standard platform. Extending it to databases avoids running separate stacks with different tools, teams, and operating models.



Where to start: Identify where existing database platforms are creating cost, risk, or friction, and assess which workloads are suitable for a phased move to EDB Postgres AI on Red Hat OpenShift.


The data layer is the modernization bottleneck

Across data-intensive, regulated industries such as financial services, organizations are under increasing pressure to modernize their technology estates while controlling cost, risk, and regulatory exposure. While many have modernised applications and general infrastructure, database platforms and operating models often stay anchored to legacy systems that are expensive, inflexible, and difficult to govern at scale.

Proprietary databases continue to drive escalating operating costs, audit exposure, and vendor lock-in. Over time, these pressures compound, limiting innovation and diverting budget away from strategic initiatives.

EDB Postgres AI on Red Hat OpenShift offers a proven alternative. Together, Red Hat and EDB provide an enterprise-grade, hybrid and multi-cloud capable data platform that delivers predictable costs, stronger data sovereignty, and a data foundation built around ease of use, enterprise grade security, and AI-driven innovation. In practice, organizations that move away from proprietary databases achieve significant reductions in database-specific operating costs.

For Red Hat OpenShift customers still running databases off-platform, EDB provides a supported, enterprise-grade path to bring the data layer onto Red Hat OpenShift without increasing operational risk.



The rising cost of staying on proprietary databases

For many organizations, proprietary databases represent significant and highly rigid costs within their IT infrastructure.

Annual licensing and maintenance costs typically rise by 5–10% year over year, creating an unsustainable long-term drain. Ongoing maintenance often becomes a larger burden than the original license investment itself. This, coupled with aggressive commercial enforcement, audits, and restrictive Unlimited License Agreements (ULAs) means even more financial uncertainty.

Operationally, these platforms are often tied to siloed teams, manual processes, and legacy VM-based infrastructure. As estates scale, complexity outgrows productivity, increasing people costs and slowing delivery.

For many organizations, the next renewal or audit means the question shifts from “Should we modernize?” to “How exposed will we be if we don’t?”

Why migrate to EDB Postgres AI on Red Hat OpenShift

Financial efficiency

There's a perception that migration is risky and complex, particularly for business-critical databases. In our experience, technical hurdles are rarely the blocker — it's the change management that determines success. With 50,000+ databases migrated, we treat migration as a business transformation, not just an infrastructure efficiency exercise. Our migrations start with a structured assessment of your existing environment. This covers everything from technical constraints to operational requirements and business priorities. This enables us to provide clear recommendations on what can move, how and when it should move, and where change may be needed before migration begins, in concert with a full ROI assessment as required.

EDB Postgres AI on Red Hat OpenShift replaces unpredictable licensing models with a simple subscription-based approach. It can reduce audit exposure, help avoid scaling penalties, and significantly reduce long-term operating costs.

[Customers who migrate routinely see 30–40% reductions in database-specific OPEX, mainly driven by removing maintenance fees and licensing constraints.](#)

For Red Hat OpenShift customers still running databases off-platform, EDB provides a supported, enterprise-grade path to bring the data layer onto OpenShift without increasing operational risk.

Infrastructure consolidation can also increase savings, by moving away from increasingly expensive proprietary stacks.

| The result: Lower, more predictable total cost of ownership, freeing up budget and the space to do more with less.


Are databases really moving to containers?

Yes, and increasingly at enterprise scale.

Historically, databases stayed outside Kubernetes because of storage immaturity, networking, and operational tooling. The last few years have seen ongoing engineering investment across the Kubernetes ecosystem that has fundamentally changed this reality.

Today, large enterprises are running mission-critical databases on Kubernetes, particularly where Red Hat OpenShift is already standardized as the enterprise application platform. While early adopters were digital-native organizations, regulated industries, particularly financial services, are now leading adoption as they safely modernize legacy estates.

Where organizations have already committed to Kubernetes as their future infrastructure platform, bringing the database layer onto Red Hat OpenShift isn't just experimental, but a logical next step.



Operational efficiency and automation at scale

Kubernetes-native automation fundamentally changes how databases are operated.

Rather than managing individual instances manually, teams define desired state and allow the platform to enforce it consistently across fleets. Patching, upgrades, high availability, and failover are all automated.

At scale, customers can report “click-of-a-button” updates across fleets, replacing 24/7 manual work. Entire patching teams may even be re-deployed as automation reduces toil and improves consistency.

The result: Database teams spend less time on repetitive maintenance and more time on higher value work.



Data sovereignty and regulatory alignment

Running Postgres on Red Hat OpenShift means more consistent data placement and governance policies across on-premises, private cloud, and public cloud environments. This means strong alignment with regulatory obligations around data residency, auditability, and security.

Crucially, organizations retain vendor neutrality and negotiating leverage, avoiding hyperscaler lock-in while preserving the flexibility to place workloads where they’re needed most.

The result: Stronger governance and regulatory alignment, without sacrificing flexibility or control.




AI-ready from day one

Modern data platforms must support more than transactional workloads.

EDB Postgres AI on Red Hat OpenShift provides a foundation for analytics and AI alongside core operational databases. Capabilities such as vector search, lakehouse integration, and retrieval-augmented generation (RAG) allow organizations to introduce AI without deploying unnecessary single-purpose databases.

This creates a clear evolution path: from application platform to unified data platform, to AI-ready architecture.


The result: A futureproof data foundation that supports AI initiatives without another round of platforming.



TCO and ROI: High-level value drivers

The business case for migrating to EDB Postgres AI on Red Hat OpenShift is driven by several factors:

- Software and maintenance savings through open-source Postgres and transparent subscriptions
- Infrastructure optimization, including consolidation on Red Hat OpenShift and traditional virtualization exits
- People efficiency, enabled by automation and fleet-level operations
- Platform consolidation, avoiding multiple databases to support AI and analytics



Why organizations should act now

- Contract renewals and ULAs that expose organizations to price increases and audits
- Data-center exits and cloud consolidation programs that naturally bring the database layer into scope
- Application modernization initiatives that require tighter alignment between applications and data
- AI strategies that depend on unified, governed data foundations
- Talent pressures that favor open platforms and modern tooling

Many organizations use renewal windows as the trigger to exit platforms or renegotiate. Delaying action means increased cost and risk.

Why this matters for regulated, data-intensive organizations

Heavily regulated industries such as financial services, government agencies and healthcare organizations face a unique combination of scale, regulatory pressure, and legacy complexity.

Large estates often comprise thousands or tens of thousands of databases. This amplifies both cost exposure and operational risk.

Red Hat OpenShift is already widely deployed as a trusted, consistent platform. Extending that platform to include the database layer allows financial institutions to apply consistent controls, automation, and security posture across their entire application and data ecosystem.

Red Hat and EDB: A unified platform for the next decade

Together, Red Hat and EDB provide a proven, enterprise-grade approach to modernizing the data layer.

EDB delivers enterprise-grade Postgres designed for mission-critical workloads, including Oracle-native compatibility that can reduce application rewrites by up to 95% when migrating from Oracle. Red Hat OpenShift provides a consistent hybrid cloud foundation with integrated platform services, running on the infrastructure of your choice.

For customers already running applications on Red Hat OpenShift, databases that remain off-platform present a clear opportunity. Standardizing Postgres with EDB allows organizations to extend platform value and support future innovation, including AI use cases, without ongoing modernization cycles.



www.enterprisedb.com

About EDB Postgres AI

EDB Postgres 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 Postgres AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs where, when, and how they need them.

© EnterpriseDB Corporation 2026. All rights reserved

