

## SOLUTION BRIEF

# Distributed High Availability:

## Always-on Operations for the Agentic Enterprise



### Executive summary

The modern enterprise runs in real time. Transactions are processed in milliseconds, customers expect instant responses, and critical systems operate around the clock. In this environment, downtime is not a technical setback; it is a business failure. The infrastructure has one job: never to fail the business that depends on it.

The financial and reputational stakes have never been higher. A single outage can trigger lost revenue, customer attrition, regulatory scrutiny, and lasting brand damage. High availability is no longer about minimizing disruption; it is about eliminating it.

Yet most enterprises are running on architectures built for a different era. Maintenance windows, failover delays, and manual intervention are embedded in traditional high availability operations, resulting in constraints that are increasingly incompatible with the demands of always-on, globally distributed operations.

The good news is that there is a better way. EDB Postgres® AI delivers distributed high availability, providing the always-on infrastructure that businesses demand.

## Distributed high availability redefines what enterprises should expect from their data infrastructure

EDB Postgres AI (EDB PG AI) eliminates downtime as a constraint, replaces complexity with simplicity, and enables global operations without compromise. The demands placed on enterprise database infrastructure have fundamentally changed. Applications must be available around the clock. Data must be accessible globally while remaining compliant locally. Operations must scale without proportional increases in cost or complexity. And downtime, whether planned or unplanned, must be invisible to the business. Traditional high availability (HA) architectures were never designed to meet these requirements.

By delivering distributed resilience to Postgres, EDB PG AI addresses each of these demands. Continuous availability is built in, not bolted on. EDB's active-active clustering eliminates the reactive failover cycles that introduce risk and latency. Operational complexity is replaced by a unified control plane that gives teams visibility and automates provisioning to simplify management across the entire system. Global scale is achieved without sacrificing data sovereignty, enabling organizations to deliver low-latency experiences to users worldwide, while meeting the strictest compliance requirements. And distributed consistency mechanisms ensure that every transaction is guaranteed, thus protecting the integrity of the data that enterprises depend on.

With distributed HA, EDB PG AI helps enterprises move beyond management availability and toward guaranteeing it as a business outcome. Because in today's environment, anything less is no longer acceptable. The result is a platform that is not only highly available but inherently reliable, scalable, and adaptable to meet the demands of a real-time, always-on world.

### EDB PG AI Customers



## The 4 structural limitations of traditional high availability

- **Traditional PostgreSQL architectures are reactive.** They rely on primary-standby replication models in which a single node handles all write operations while secondary nodes remain idle until a failure occurs. When a failure does occur, the system must promote a standby node to primary, a process that introduces latency, operational risk, and potential data inconsistency.
- **Management is complex,** because these architectures also depend on a collection of loosely integrated tools for monitoring, failover orchestration, backup management, and upgrades. Each tool introduces its own configuration requirements, failure modes, and operational overhead. The result is a fragmented system that requires constant attention from highly skilled database administrators and site reliability engineers.
- **These systems struggle to meet modern requirements** in several key areas. Failover times are often measured in minutes rather than seconds, resulting in unacceptable service interruptions. Maintenance operations typically require planned downtime, forcing organizations to schedule disruptive outages. Global scalability is limited by centralized write architectures, which introduce latency for users in distant regions.
- **Cloud-managed database services cannot solve the local residence requirements.** Data is often required to reside within vendor-controlled environments, limiting the ability to meet strict sovereignty and compliance requirements. Additionally, these services may not provide the level of transparency or configurability required for mission-critical workloads.

Traditional high availability architectures expose organizations to unacceptable operational risk and cost, including failovers measured in minutes, fragmented tooling that demands constant expert intervention, and architectures that force trade-offs between uptime and compliance. These structural limitations translate directly into revenue loss, regulatory exposure, and an inability to scale without proportional increases in complexity and cost. Today's modern businesses require an always-on global operation with strict data sovereignty.

# Beyond uptime: Five ways EDB PG AI redefines high availability

EDB PG AI distributed high availability works differently from traditional Postgres architectures—and that difference matters.

In a traditional setup, one node handles all writes while others sit on standby, waiting for something to go wrong. When it does, the system scrambles to promote a standby, redirect traffic, and restore consistency. That's a process measured in minutes, not seconds, and one that carries real risk of data loss.

EDB PG AI takes a different approach. Every node is active. Every node is synchronized. When a failure occurs, there is nothing to promote and nothing to restore—traffic simply routes to another active node, and the system keeps running.

Data consistency is protected by quorum-based commit protocols, which ensure that every transaction is confirmed across multiple nodes before it is considered complete. This means that even if a node or network connection fails, no data is lost and no inconsistency is introduced.

The result is a system that doesn't recover from failure but one that never stops running in the first place.

## Pillar 1: Continuous operations without interruption

One of the most significant advantages of EDB PG AI is its ability to support continuous operations without downtime. Because all nodes are active and synchronized, there is no need to pause operations for maintenance or upgrades. Rolling upgrades can be performed incrementally across the cluster, allowing the system to remain fully operational throughout the process.

Failover is similarly seamless. In traditional architectures, failover involves detecting a failure, promoting a standby node, and redirecting traffic—a process that can take minutes and introduces the risk of data loss. In an active-active system, failover simply involves rerouting traffic to other active nodes, a process that can be completed in seconds with no interruption to service.

This capability is particularly important for mission-critical applications, for which even brief interruptions have significant consequences. Every minute of unplanned downtime carries measurable costs in lost transactions, damaged customer trust, and potential SLA penalties. By eliminating both planned and unplanned downtime, EDB PG AI gives organizations the confidence to run their most critical workloads with near-zero downtime. This makes continuous availability an operational reality, not an aspiration.



## Pillar 2: Global performance and data sovereignty guaranteed

As enterprises expand globally, they must balance performance with regulatory compliance. Users expect low-latency access to applications regardless of location, while regulators require strict control over where data is stored and processed.

EDB PG AI addresses both requirements through a write-local, read-global architecture combined with granular data placement controls. Data can be selectively replicated across regions, allowing organizations to ensure that sensitive information remains within designated jurisdictions while still enabling global access where appropriate.

From a technical standpoint, this architecture reduces latency by allowing users to interact with local nodes rather than routing all requests through a centralized primary database. At the same time, distributed consistency mechanisms ensure that all nodes maintain a synchronized view of the data.

This combination of performance and control transforms data sovereignty from a competitive liability into a strategic advantage by enabling organizations in regulated industries to confidently expand globally, meet compliance mandates, and deliver the low-latency experiences users expect, all without compromise.

## Pillar 3: One control plane to run it all

Another key innovation offered by EDB PG AI is the consolidation of all high availability functions into a single control plane. Traditional high availability environments require separate tools for monitoring, failover management, replication configuration, and upgrades. This fragmentation increases complexity and introduces additional points of failure.

EDB PG AI replaces these disparate tools with a unified management interface that provides visibility into the entire system. Graphical topology maps allow operators to visualize cluster configurations, while point-and-click controls simplify deployment and management tasks.

This unified approach reduces the need for specialized expertise and accelerates deployment timelines. According to internal analysis, organizations can deploy geo-distributed clusters up to 75% faster compared to traditional approaches.

From a business perspective, this translates into faster time to value and lower operational costs by reducing the specialized expertise required to manage complex database infrastructure, compressing deployment timelines from weeks to days, and freeing engineering teams to focus on innovation rather than operations.

## Pillar 4: Integrity that is architected, not assumed

For industries in which data integrity failures are business crises, EDB PG AI leaves nothing to chance.

Every transaction is confirmed across multiple nodes before it is considered complete, ensuring consistency is never dependent on a single point of truth. Quorum-based commit protocols eliminate split-brain scenarios, so every user always sees the same data. Commit-at-most-once mechanisms prevent duplicate transactions during failover, a nonnegotiable requirement for payment systems whose errors have direct financial consequences.

These are not incremental improvements over traditional PostgreSQL high availability. They are architectural improvements that provide the level of assurance organizations need without the cost, inflexibility, and vendor lock-in associated with proprietary systems, which have traditionally been the only way to meet these requirements. EDB PG AI delivers the transactional integrity enterprises demand, without sacrificing the openness, flexibility, and control of PostgreSQL.

## Pillar 5: Resilience that isn't a feature, it's your competitive edge

EDB PG AI transforms resilience from a reactive necessity into a proactive strategic capability. By eliminating downtime, simplifying operations, and ensuring data integrity, it enables organizations to focus on delivering value rather than managing infrastructure.

This shift has significant implications for business performance. Continuous availability supports uninterrupted customer experiences, while simplified operations reduce costs and improve efficiency. The ability to scale globally without compromising performance or compliance enables organizations to expand into new markets with confidence.

Data integrity failures aren't just technical incidents; they're business crises that can trigger regulatory scrutiny, financial loss, and irreparable reputational damage. EDB PG AI delivers the transactional guarantees organizations need to protect revenue, satisfy regulators, and maintain customer trust, without locking them into costly proprietary systems.

## Compare your options

Organizations evaluating high availability strategies must consider how different approaches align with their business and technical requirements.

	<b>EDB PG AI</b>	<b>Cloud-managed high availability (RDS, Aurora, Azure)</b>	<b>Distributed database alternatives</b>	<b>Traditional high availability (primary-standby + tools)</b>
<b>Strengths</b>	Active-active, zero downtime, unified control, sovereign deployment	Simplified operations, managed services	Native scalability and resilience	Familiar architecture
<b>Limitations</b>	Minimal learning curve for Postgres teams (90%+ Postgres compatible)	Limited sovereignty, vendor lock-in, cloud outage risks	Requires application rewrites, new operational models	Slow failover, downtime for upgrades, high complexity
<b>Business impact</b>	Maximum uptime, reduced complexity, fastest time to value	Moderate efficiency with availability trade-offs	High migration cost and disruption	High operational cost and downtime risk

\*Competitive comparisons are based on publicly available information and are subject to change as vendor offerings evolve and new information is made available. All product names, trademarks, and registered trademarks are the property of their respective owners.

## Why this matters strategically

The enterprises winning the next decade will be built on real-time, AI-powered platforms that never stop, and the data infrastructure underneath them cannot afford to either.

EDB PG AI eliminates downtime as a category of risk, delivering a distributed, always-on data platform that gives global teams the sovereignty, consistency, and scale to operate without limits, without compromise, and without interruption. The result isn't just a more reliable database. It's a platform that guarantees your business.

### EDB Postgres AI: The sovereign data and AI platform for the agentic enterprise

EDB PG AI brings together a unified data layer, governance, sovereign control and orchestration, and an agent runtime environment, giving enterprises a trusted foundation for AI on infrastructure they own and control. The platform unifies transactional, analytical, and AI workloads in a single Postgres-based architecture—eliminating ETL, data movement, and operational fragmentation. And you choose where and how to deploy: on-premises, cloud, managed, or certified appliance.

The outcome: production-ready sovereign AI in days or weeks, not months.



EDB Postgres® AI (EDB PG AI) is the sovereign data and AI platform for the agentic enterprise. Built on Postgres, the world's leading open source database, EDB PG AI unifies transactional, analytical, and AI workloads in a single governed architecture, on-premises and across clouds. To learn more, visit [www.enterprisedb.com](http://www.enterprisedb.com).

© EnterpriseDB Corporation 2026. All rights reserved.