

A photograph of a man and a woman in an office environment. The man is seated at a desk, looking at a laptop screen. The woman is standing next to him, leaning over and pointing at the screen. They both appear to be smiling and engaged in their work. The office has a modern, industrial feel with concrete walls and a desk lamp.

HOW TO FUTURE-PROOF BUSINESS PERFORMANCE WITH POSTGRES HIGH AVAILABILITY IN THE CLOUD

Whether you're thinking about moving on from Oracle or migrating to open source Postgres databases in the cloud, it's critical to identify ways to future-proof business performance. Essential considerations include assuring that the databases driving all cloud applications have the correct availability characteristics to retain customers, ensure uninterrupted revenue and optimize business processes.

As a large-scale organization supporting global teams and customers, it can be a challenge to maintain 24/7 high availability and keep application performance

high, especially with teams and databases spread out across different continents. At the same time, your R&D platforms may not need the same high availability Postgres environments as those supporting customer-facing financial applications.

Achieving high availability is complex in Postgres database environments. A skilled, experienced Postgres vendor can help you evaluate your application and business requirements and define suitable Postgres availability configurations to meet them—without unnecessary cloud infrastructure provisioning that drains your budget.

What you need to know about running high availability Postgres environments in the cloud

The best time to start thinking about high availability is before it's needed. Waiting until a crisis hits is the worst time, especially with more than two-thirds of all outages [costing more than \\$100K](#). An excellent place to start with cloud high availability is by understanding how critical a given application is to your business. As a result, focus on three key areas when running Postgres in hybrid and multi-cloud:

- **Application needs.** How long can the company manage during an application outage, and what are the financial and corporate impacts?
- **Regional and infrastructure resilience.** Does the application need to be resilient to regional outages? Think about the IT infrastructure required to support cloud deployments. Start by asking how resilient the infrastructure is and whether redundancy is needed for database instances and replication.
- **Support.** Unlike on-premises environments, cloud deployments depend on cloud service providers (CSPs) and their platforms. It's, therefore, critical to understand what kind of support and SLAs can be expected from your cloud provider.

Even your best-planned and well-executed DR with SLAs is expensive

It's easy to understand that mission-critical workloads need high availability, but there may be other considerations. For example, issues with non-mission critical applications running too slowly or experiencing downtime can lead to lost revenue and other business-impacting consequences. Additionally, the cost of any downtime can vary by industry and be significant.

While today's organizations and their customers expect maximum resilience from the cloud infrastructures supporting applications operating in their Postgres databases, a CSP region does not provide SLAs higher than 99.99% for their computing environments.

It's important to note the differences between "disaster recovery" capabilities within cloud infrastructures supporting Postgres and the "resilience" needed to minimize the true impacts of unanticipated downtime on business operations. Even with well-designed DR plans, outages will occur, so focusing on how to reduce associated impacts should be prioritized.

In addition to real-world financial loss, an extended downtime incident can lead to negative fallout or diminished brand reputation. Regional and regulatory compliance concerns can also be drivers that factor into high availability considerations.

Since any cloud-based Postgres-as-a-Service deployment will involve a cloud service provider and cloud-based application, it's also essential to think about a few other aspects when it comes to availability, including:

- What [type of resilience is required](#) for your workload from both zonal and regional failures, and which CSP do you use?
- What kind of support and SLA does your CSP provide?
- How long can you withstand a database, network or storage failure in cloud environments supporting Postgres database operations?

Additionally, beyond cloud operations, you should evaluate whether your disaster recovery plans factor issues like a regional outage in your CSP environment. These issues also existed when running on a legacy database, and they remain considerations when moving workloads to the cloud.

How EDB positions you for building highly available apps

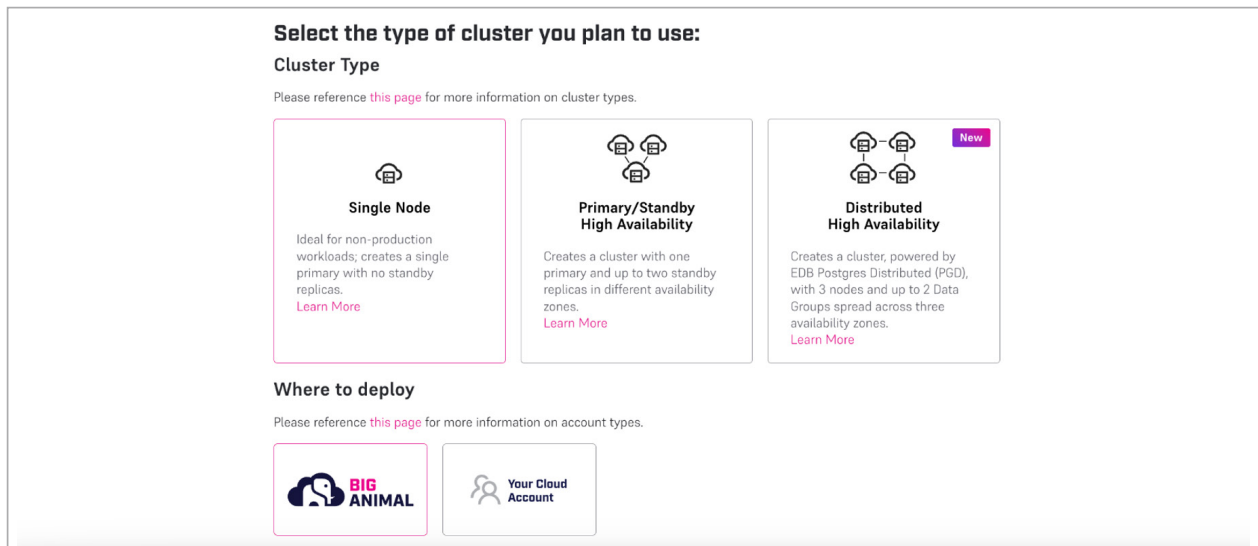


Figure 1: The BigAnimal console provides intuitive Postgres high availability and cloud service configuration options.

In helping you address these challenges, BigAnimal now provides distributed high availability for Postgres applications as a fully managed service in the cloud, powered by EDB Postgres Distributed. Running EDB Postgres Distributed on BigAnimal provides you with the ability to deploy Postgres in geo-distributed (active/active) architectures. With EDB, you can empower developers to build highly resilient Postgres applications on any cloud, giving you the power to build applications with active-active use cases and choose the best environment for your data.

EDB provides flexible Postgres high availability requirements in cloud environments, including the following:

NON-MISSION-CRITICAL APPLICATIONS

Non-mission-critical applications include pre-production, internal, and non-revenue applications that won't stop the business or significantly impact revenue generation if they are down for a given period. In general, these types of applications might have a 99.5% availability requirement.

- Non-mission-critical applications can be deployed on single node configurations, with an uptime SLA of 99.5%. No failover or redundancy is needed, and organizational Postgres deployments in the cloud can start small and grow.

MISSION-CRITICAL APPLICATIONS

Mission-critical applications usually include revenue-generating and customer-facing applications upon which the business depends. If these applications aren't up and running, a company loses money, customers, or both. These applications would typically have a 99.99% availability requirement.

- Mission-critical applications require significantly more technical planning than non-mission-critical apps. Their advance uptime requirement of 99.99% necessitates distinct capabilities, such as automatic failover, configurable replications, the ability to start with one replica and scale to two, and RTO for failovers of under 60 seconds.

APPLICATIONS THAT CAN'T GO OFFLINE

Applications that require near-zero downtime include emergency services and revenue-critical applications.

- Technical view. Applications requiring distributed high availability of 99.995% are explicitly designed for increased resilience during unplanned outages. They may include a geo-distributed architecture with multi-write capabilities. They are also resilient to zone-specific or region-specific failures because of their globally distributed architecture.

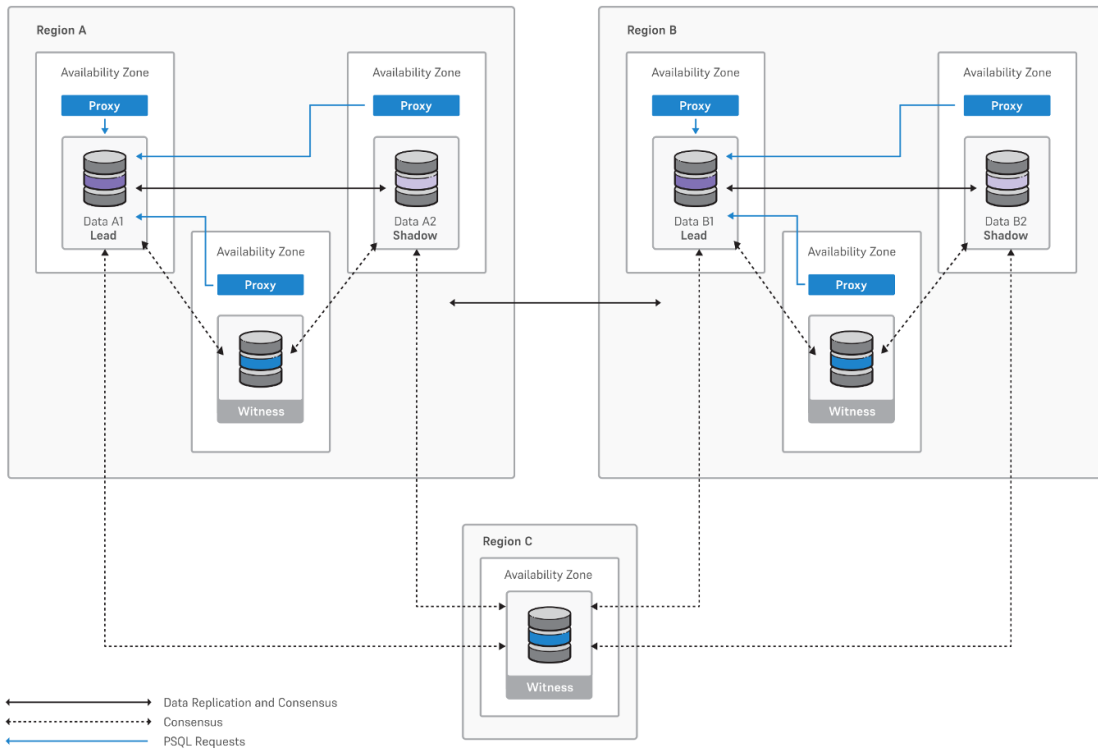


Figure 2: EDB Postgres Distributed on BigAnimal, showing a distributed high availability, multi-region configuration, with three Data Nodes (x2) and one Witness.

Right-sizing Postgres availability in the cloud

With EDB Postgres Distributed on EDB BigAnimal, you can:

- Meet the needs of globally distributed customers and employees
- Improve business continuity across multiple data centers
- Protect against unplanned outages that cause downtime, with up to 99.995% availability.

Historically, it was nearly impossible to achieve active-active architectures in the cloud, especially with Postgres. With EDB Postgres Distributed on BigAnimal, you have even more freedom to run distributed Postgres with the confidence of distributed high availability, everywhere you need it.

With EDB Postgres Distributed on BigAnimal, active-active architectures distribute the workload across multiple database nodes to achieve higher availability and unlock capabilities enterprises need to build modern applications with confidence and ease. In case of node failure, workloads are automatically rerouted to operational instances. This balancing across multiple resources benefits enterprises by ensuring uninterrupted service and improving uptime.

- ✓ Active-Active Setup
- ✓ Flexible Deployments
- ✓ Efficient Logical Replication
- ✓ Customer's Choice of Durability
- ✓ Robust Conflict Handling




Figure 3: Running EDB Postgres Distributed on BigAnimal improves application uptime, enhances geo-distributed data capabilities and simplifies maintenance in any cloud.

Improving disaster recovery and maintenance processes in the cloud

EDB Postgres Distributed on BigAnimal provides best-in-class *high availability* and disaster recovery without the “heavy lift” needed to do it in-house. The solution’s multi-region, always-on architectures ensure that disaster recovery is always operational. With EDB Postgres Distributed on BigAnimal, you benefit from built-in survivability and instant application switchover when needed, protecting brand reputation, avoiding disruptions and revenue loss, and keeping customers satisfied. You can leverage advanced durability options to control replication lag to meet your RPO, ensure transactions are only applied once or select the level of redundancy needed for your business requirements.

EDB’s geo-distributed (active/active) architecture delivers essential benefits, including improved application performance through localized database reads and writes. It also enables you to manage protected data and meet business or regional data compliance requirements by enabling replication at the schema, table, or transactional level. In addition, it has robust, configurable conflict handling. Use of active/active architectures also enable you to evolve Postgres DevOps, easing maintenance and release of blue-green, canary, and rolling deployments.



Get the Postgres high availability your organization requires

Running EDB Postgres Distributed on BigAnimal provides you with a fully managed Postgres database equipped to meet business and government high availability requirements, with the flexibility to operate on AWS, Google Cloud or Microsoft Azure—your cloud or ours. EDB BigAnimal is the only Postgres cloud solution, built by Postgres experts, that provides you with the ability to:

- Move off Oracle database with minimal changes
- Rapidly build cloud Postgres applications that are highly available
- Support enterprise-scale requirements
- Deploy on any cloud

EDB, with our EDB Postgres Distributed technology running on the BigAnimal managed service, can help you navigate these challenges.

Experience all the high availability
features and functionality of running EDB
Postgres Distributed on BigAnimal

START TODAY
with \$300 in credits.



ABOUT EDB

EDB provides enterprise-class software and services that enable businesses and governments to harness the full power of Postgres, the world's leading open source database. With offices worldwide, EDB serves more than 1,500 customers, including leading financial services, government, media and communications and information technology organizations. As one of the leading contributors to the vibrant and fast-growing Postgres community, EDB is committed to driving technology innovation. With deep database expertise, EDB ensures extreme high availability, reliability, security, 24x7 global support and advanced professional services, both on premises and in the cloud. This empowers enterprises to control risk, manage costs and scale efficiently. For more information, visit www.enterprisedb.com.