

PostgreSQL users seek high availability for distributed environments



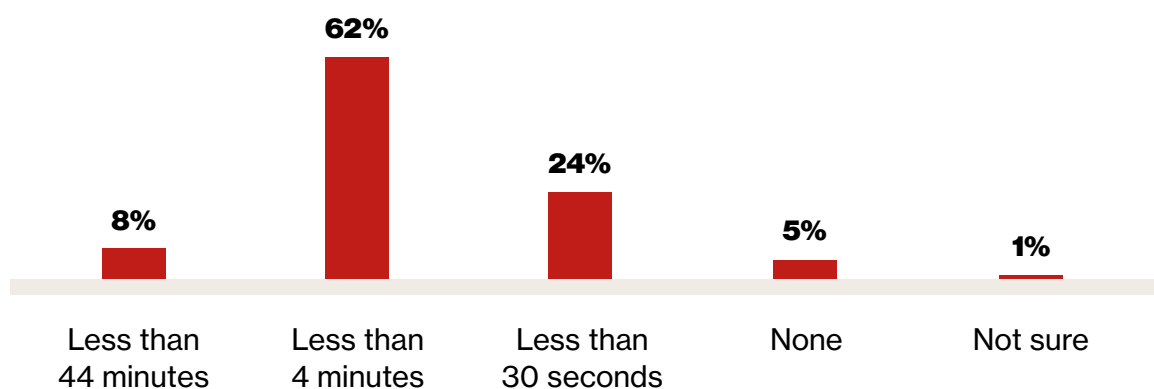
Technology decision-makers have moved many critical business processes to the cloud. Yet cloud service outages are far from rare. There are dozens of significant outages annually and, inevitably, a handful are extremely serious, lasting many hours and disrupting businesses large and small. Outages bring risks of data loss, including vital transaction records, not to mention missed business opportunities, poor customer service, and reputational damage.

A new survey from Foundry¹ shows that IT leaders are acutely aware of the risks of cloud region failure: 82% cite it as a concern – and in the last year, 21% experienced it directly. However, high availability isn't free, and the high cost deters many organizations from taking the necessary steps to protect business

continuity. What's needed is a high availability strategy that is affordable, yet effective in overcoming cloud or system outages.

To strike the right balance between low cost, reliability and high availability, many enterprises rely on PostgreSQL as a strategic database platform. A new survey from Foundry finds that 37% of PostgreSQL deployments are for mission-critical workloads. Because many PostgreSQL deployments are part of a hybrid database environment and are customer-facing, near-continuous uptime is mandatory. In fact, the vast majority of organizations using PostgreSQL (91%) require 99.99% uptime, equivalent to no more than four minutes of downtime per month. And, significantly, 24% of enterprises need downtime below 30 seconds.

How much downtime is acceptable for your mission-critical applications per month?



Multi-master provides effective failover

The fact that 21% have experienced a cloud failure in the past year is worrisome, since 35% of implementations are customer-facing and 37% of respondents rely on PostgreSQL for mission-critical applications, while 30% are standardizing on PostgreSQL for most workloads. Simply put, these enterprises need a better strategy to ensure availability.

To that end, multi-master replication (MMR) is generating strong interest (47%) among organizations that deploy applications across multiple cloud regions. As the name suggests, a multi-master implementation establishes several master databases spread across different geographic regions. Each database can handle read and write traffic simultaneously, ensuring eventual data consistency through bi-directional replication (BDR) and conflict resolution. By contrast, in a traditional monolithic architecture, only one database can perform read-write operations, while distributed nodes are read-only.

It's important to note that in a multi-master configuration, any of the several master databases can perform failover if one of the others goes offline. Consequently, all nodes collectively enable higher uptime than is possible in a monolithic database implementation.

Investing in high availability

Survey respondents recognize the need for high availability, and many are purchasing high-availability products for PostgreSQL environments.



These include:

- Cloud-native tools: **41%**
- Commercial products: **33%**
- Patroni (open source): **29%**

As seen above, the highest percentage of respondents are opting for cloud-native tools. While this might seem sensible, doing so fails to take into account the likelihood of a cloud region failure, even though 21% have experienced such a failure in the past year.

Still, the variety of strategies reflected in the survey shows that IT decision-makers are open to different approaches to providing high availability. Respondents anticipate a number of benefits from a high-availability or distributed PostgreSQL solution, in particular, led by increased uptime or zero downtime operations (53%) and faster disaster recovery or resilience (43%). However, cost savings compared to proprietary databases (36%) are also sought by respondents, as is the ability to deploy databases across multiple or hybrid cloud environments (40%).

How does your organization prioritize the following benefits of a high-availability or distributed PostgreSQL solution?



pgEdge: standards-based high-availability

Faced with the need for high availability in a global environment, and aware of potential cloud region failure, PostgreSQL users are exploring multiple failover strategies. Thanks to its native multi-master architecture, pgEdge delivers high availability across multi-cloud and multi-region environments.

As distributed environments increase in complexity, the ease of deployment of pgEdge and its ability to simplify complex environments are both highly important. And because pgEdge is fully standard PostgreSQL and 100% source available, it delivers cost savings while preserving choice among vendors, helping users avoid vendor lock-in.

Most importantly, the multi-master architecture of pgEdge gives businesses the confidence to deploy PostgreSQL in mission-critical environments that deliver the highest level of customer experience.

Learn how to confidently
deploy **Distributed**
PostgreSQL with pgEdge

1

Foundry Research MarketPulse Survey, May 2025. All respondents use PostgreSQL, many in multiple use cases. Top five industries are Services, Manufacturing, Retail, Banking, Consumer packaged goods. Company size: 100 – 10,000 + employees. Respondent job titles: 13% C-level; 43% Director; 75% IT; 12% management; 9% R&D.

© 2025 IDG Communications, Inc.

CIO

SPONSORED BY

