PERCONA

DATABASE MANAGEMENT INSIGHTS:

Open Source Outpaces Proprietary Solutions

2023 Survey Report

Contents

Introduction

Databases are the lifeblood of business today. They keep commerce in motion, enable customer relationship management, support product and service management lifecycles, and help companies engage new prospects and close new deals. They are the most ubiquitous of business tools. So universal, in fact, that the only time most workers give any mindshare to the database is when the one they are relying upon isn't working for them.

But the story is very different for the DBAs, SREs, CIOs, CTOs, and other professionals who are responsible for any aspect of database operations, from procurement to provisioning to performance to security and more. They understand that their business goes as their databases go. So, we set out to hear from them about the state of database management today.

We surveyed nearly 300 database and other IT professionals and asked them to self-evaluate how their organizations are doing with their database operations. We asked them to relay the leading factors that steered them toward the database solutions, deployment models, and support and management processes they chose.

Their answers surprised us.

Executive summary and key findings

The results of this survey reveal a number of interesting trends in the database management space. Firstly, we've found that the survey confirms what we anecdotally knew to be true: Open source solutions — like MySQL and PostgreSQL — are the most popular and widely-adopted database management systems in the modern enterprise. While proprietary solutions like Oracle and Microsoft SQL Server still enjoy widespread use (especially within legacy enterprises), both have been edged out by open source alternatives—and this trajectory is likely to continue for the foreseeable future.

We also discovered that the database-as-a-service (DBaaS) deployment model now represents roughly a third of all enterprise database deployments — making it just as prevalent as on-premises and traditional cloud deployments. The emergence of DBaaS is another sign that organizations are looking for innovative approaches to their business operations.

While these findings represent novel trends in the database space, we also found that most organizations anticipate making few changes to their database management strategies in the near term.

Here's a closer look at our key findings

All else being equal, database professionals prefer community-supported open source solutions.

Community-supported open source databases continue to grow in popularity. Not long ago, proprietary solutions made up the lion's share of database instances. Now, the balance of power appears to be tipping in the opposite direction. Of course, pragmatism still reigns supreme in the enterprise as organizations most often choose a particular database based on the needs of the workload. However, when given a choice between proprietary and open source solutions, organizations overwhelmingly opt for open source. According to our respondents, community-supported open source remains the most popular licensing model, with roughly 42% of all database instances today being licensed in this way.

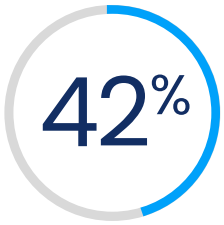

of all database instances today are being licensed via community-supported open source model

As of today, what percentage of your database instances are licensed in each of these ways?

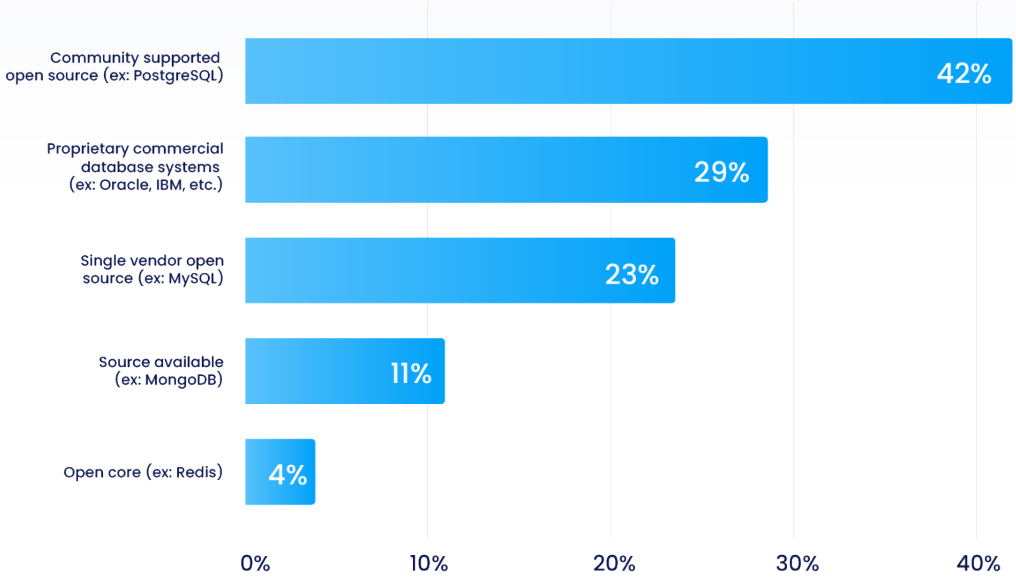

PERCONA

MySQL and PostgreSQL are the most widely-used relational database products today, beating out Oracle (and all other proprietary solutions).

When asked to identify the relational database products their organizations are currently using, respondents reaffirmed their preference for open source solutions. With 57% reporting using MySQL and 48% reporting using PostgreSQL (aka Postgres), these two open source solutions beat out both Oracle and Microsoft SQL Server to be the most widely-used relational database products on the market. In fact, open source database products represented three of the top five most widely used relational databases in our survey.

With 57% reporting using MySQL and 48% reporting using PostgreSQL (aka Postgres), these two open source solutions beat out both Oracle and Microsoft SQL Server to be the most widely-used relational database products on the market.

Which of the following relational database products does your organization use today? Mark all that apply.

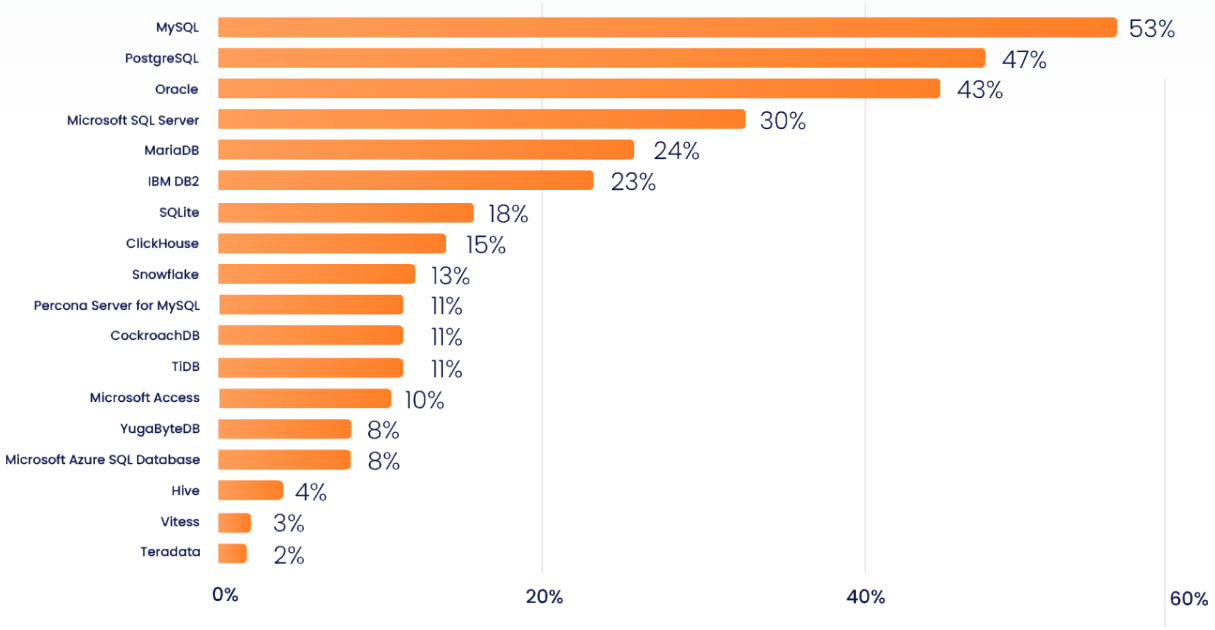

PERCONA

Virtually all high-performing organizations utilize third-party support and virtually all are happy with it.

As part of our analysis, we divided respondents into three groups — high-performers, middle-performers, and low-performers — based on self-reported assessments of their organizations' overall database capabilities. (You can read more about our methodology in the next section.)

We then compared the highest performers to the lowest and found significant variation between the two groups. When it comes to the use of third-party database support services, nearly all (99%) high-performing organizations reported receiving third-party support, and all (100%) reported being either somewhat or very satisfied with the service.

Rate the support you receive from each of these for your database management systems.

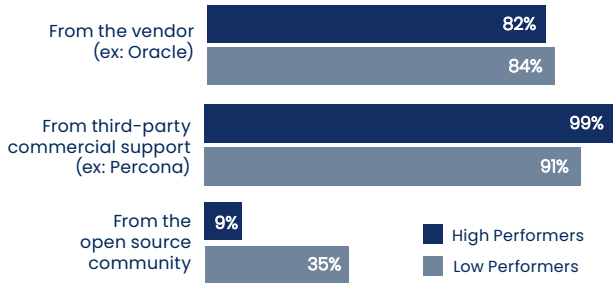

- Cloud-based Database-as-a-Service (DBaaS)
- Cloud, (using Infrastructure-as-a-Service (IaaS)
- On-premises or in a private cloud
- On-premise or private cloud Kubernetes
- Public cloud Kubernetes
- Other

The database-as-a-service (DBaaS) deployment model has arrived.

When determining their overall database strategy, one of the most critical factors organizations must decide is the deployment model. Our findings revealed that public cloud and on-premises/private cloud deployments are preferred roughly equally, which is consistent with long-standing trends favoring hybrid database architectures. However, what came as a surprise in our study was that DBaaS deployments were nearly as common as both IaaS and on-premises deployments.

This means, despite being a relatively new deployment model, DBaaS is already neck and neck with the other leading deployment options (i.e. on-prem and traditional cloud). The adoption of DBaaS solutions has been growing steadily since the first offerings became available in the marketplace. Now, it appears DBaaS has arrived.

PERCONA

Most organizations' near-term database strategy is to stand pat.

We also found that most organizations today aren't planning on making significant changes to their existing database strategy within the next three years. In terms of the types of deployments (i.e. on-prem vs. IaaS vs. DBaaS), the types of tools (i.e. proprietary vs. open source), and the specific products being used (e.g. MySQL vs. Postgres vs. Oracle), most respondents anticipate very little change in their organizations within the next three years. While the question of why this is the case remains a matter of speculation, it is possible that today's climate of economic uncertainty and post-pandemic disruption are weighing heavily on the minds of today's database decision makers.

NOW

As of today, what percentage of your database instances are licensed in each of these ways?

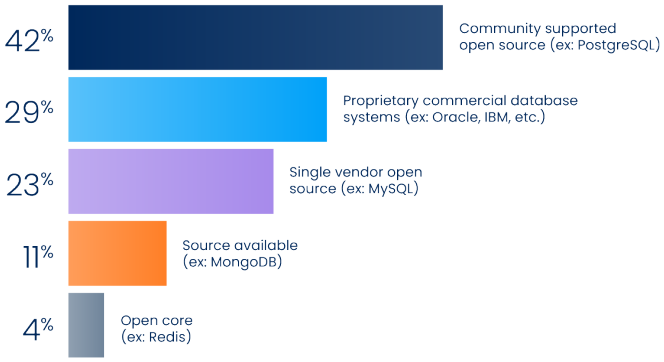

IN THREE YEARS

In three years, what percentage of your database instances will be licensed in each of these ways?

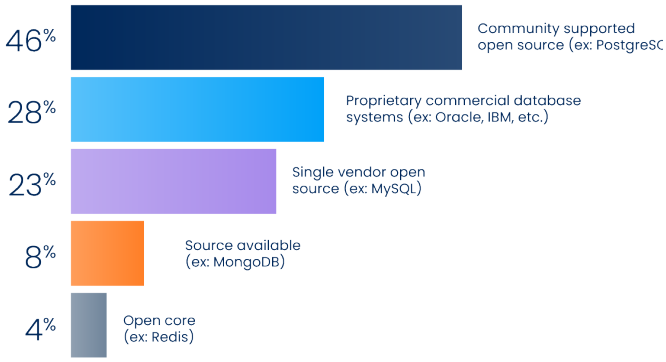

PERCONA

Self-reported company performance and survey methodology

It's helpful at this point to review the process we followed in our analysis of the data collected during this survey. We asked our 284 respondents to consider their organization's performance across 14 attributes of database operations:

1. **CapEx costs**
2. **OpEx costs**
3. **Security**
4. **Privacy**
5. **Scalability**
6. **Reliability**
7. **Compliance**
8. **Backup reliability and integrity**
9. **Performance**
10. **Developer agility**
11. **Synchronous replication of databases**
12. **Asynchronous replication of databases**
13. **Ability to identify issues before users**
14. **Time to fix issues**

On the surface, it appeared that most respondents felt their organizations were handling most aspects of database management somewhat or extremely well.

Please rate how well you are doing with each of the following database management metrics:

(Somewhat/Extremely Well)

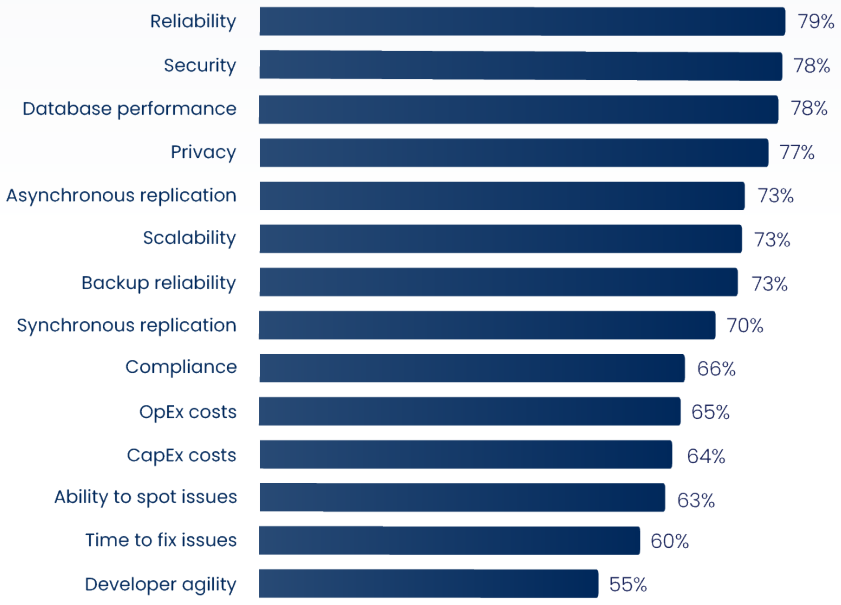

PERCONA

However, when we applied a scoring system to their responses, wherein each response to each criterion received a relative numeric value, we were able to triage the scores into high-, middle-, and low-performing groups. We then compared the responses of those participants who felt their organizations were performing well (henceforth referred to as "high performers") to the responses of those participants who felt their organizations were performing the least well ("low performers") to understand where each group was excelling and where each was struggling. These results painted a much more nuanced picture.

For example, in comparing the high performers with the low performers across these 14 attributes, the high performers were:

- **More than twice** as likely to say they are doing well with CapEx, OpEx, and Synchronous Replication;
- **Nearly twice** as likely to say they are doing well with Privacy, Scalability, Reliability, and Compliance.

Please rate how well you are doing with each of the following database management metrics:

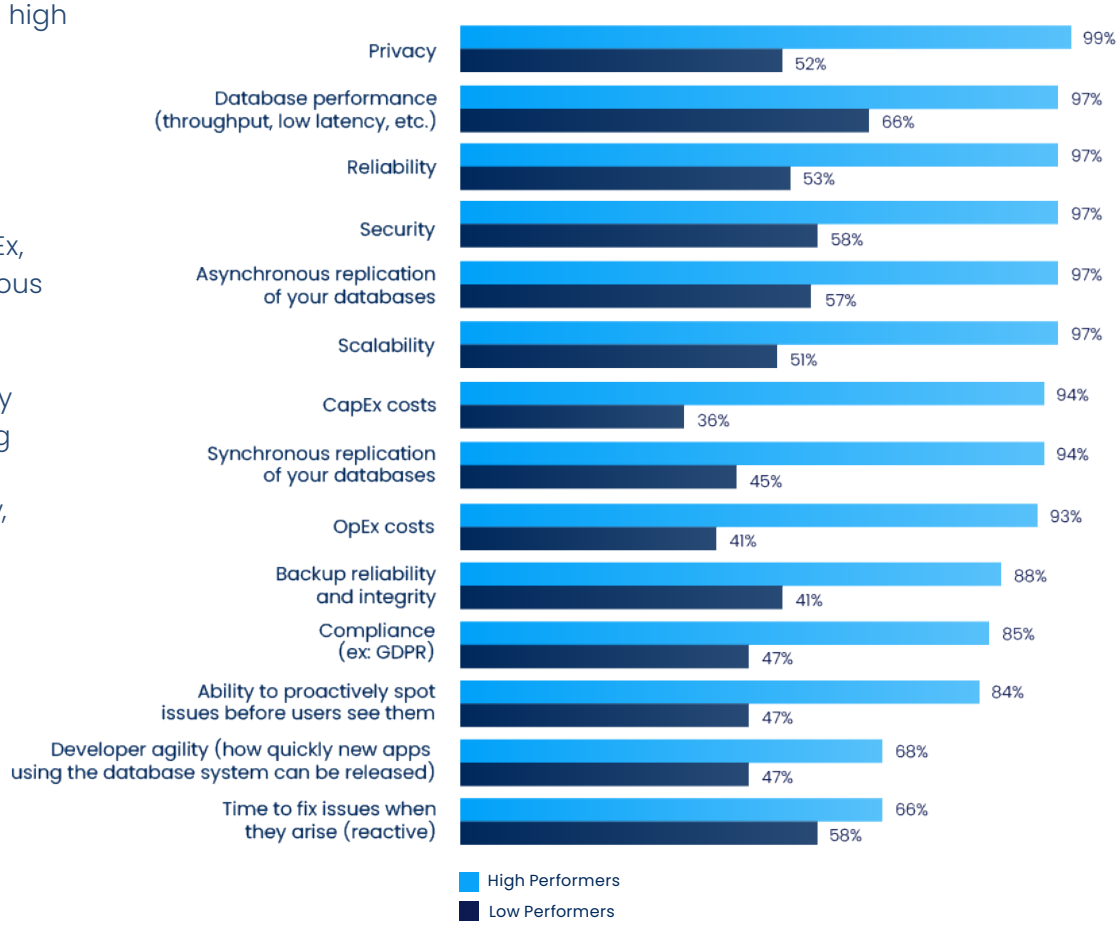

PERCONA

We also drew several additional comparisons: We compared the high-performers and low-performers when it comes to their in-house proficiency with database management. High-performing organizations were twice as likely as the low-performers to describe their organization's database management proficiency as "somewhat skilled" and nearly nine times as likely to characterize their organization's database management proficiency as "extremely skilled."

Rate your organization's skill level as it pertains to database management.

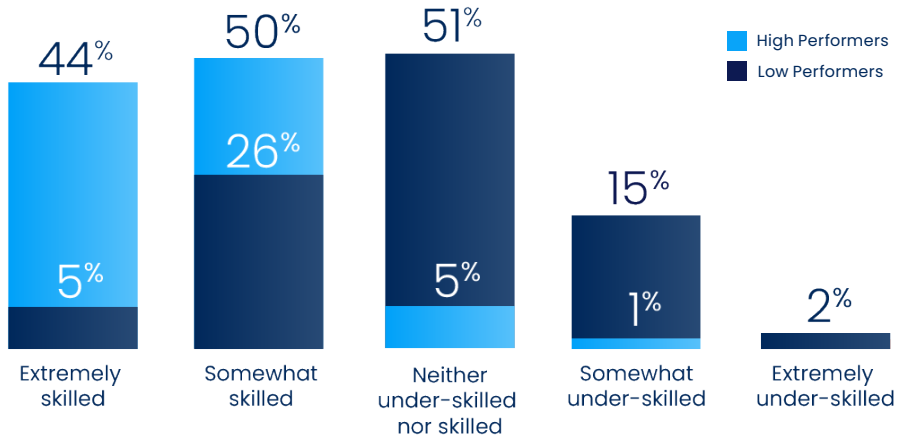

When it comes to staffing levels relative to the need for resources to address all aspects of database lifecycle management, the split between high- and low-performing organizations was fairly comparable, with respondents from high-performing organizations being three times as likely to say their organization was "somewhat overstaffed" and respondents from low-performing organizations being twice as likely to say they were somewhat understaffed.

Rate your organization's staffing level as it pertains to database management.

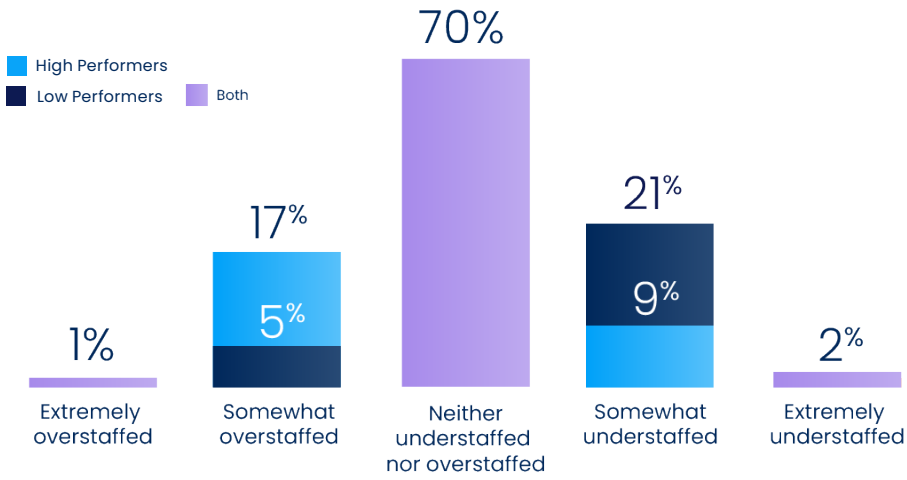

PERCONA

Additional findings: the state of database management

Here we take a more granular look into the complexities organizations face in managing database operations. In some cases, we've looked at the responses we received in totality. In cases where comparing the high-performing group with the low-performing group proved revealing, we've included that detail. In all cases, we've tried to present the most helpful look at the data so you can informally benchmark your organization's database operations against those represented in our study.

High performers opt for open source.

As reported earlier, open source is the preferred licensing model across all respondents. When we compare our stratified groups, however, we find a stark contrast:

2 : 1

High-performing organizations are twice as likely as their low-performing counterparts to be utilizing community-supported open source relational databases like PostgreSQL.

2 : 1

Among high-performing organizations, the preference for community-supported open source relational databases outstripped closed, proprietary database solutions like Oracle or IBM by two to one.

As of today, what percentage of your database instances are licensed in each of these ways?

- Proprietary commercial database systems (ex: Oracle, IBM, etc.)
- Single vendor open source (ex: MySQL)
- Community supported open source (ex: PostgreSQL)
- Source Available (ex: MongoDB)
- Open Core (ex: Redis)

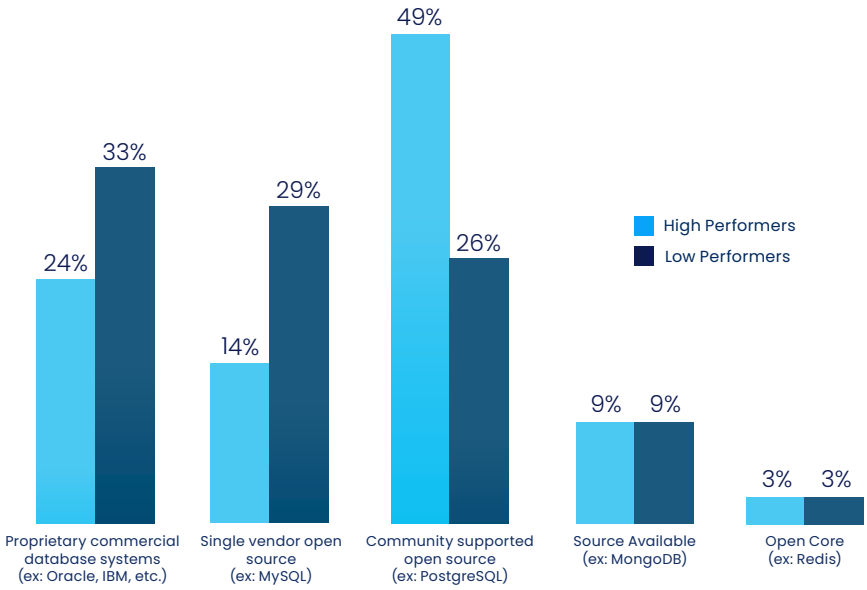

PERCONA

Security, reliability & consistency top the list of relational DB priorities.

Security, Reliability, Data Consistency, and Performance were the top features all groups prioritized in their relational databases.

How significant are the following factors in terms of driving you to implement and use relational database systems?

(Somewhat/Extremely Significant)

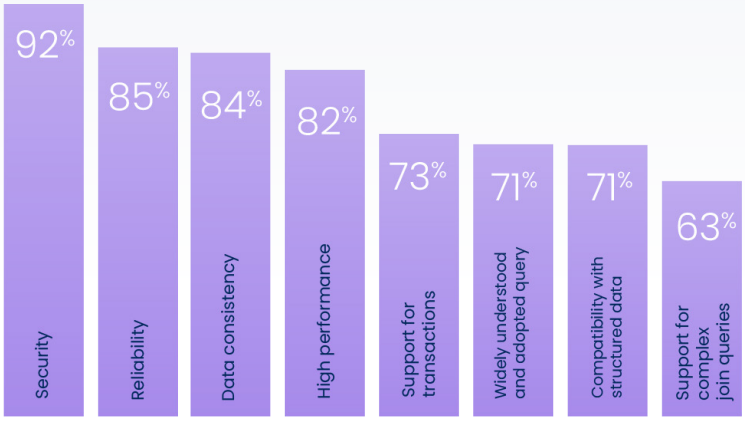

Motivations differ for proprietary and open source DB deployment.

When asked to consider the factors driving respondents to deploy either proprietary or open source databases, differences emerged:

A: Stability, Security, Regulatory Compliance, and Functionality received the highest priority for organizations deploying proprietary solution

How significant are the following factors in terms of driving you to use a proprietary database offering (ex: Oracle, IBM, etc.)?

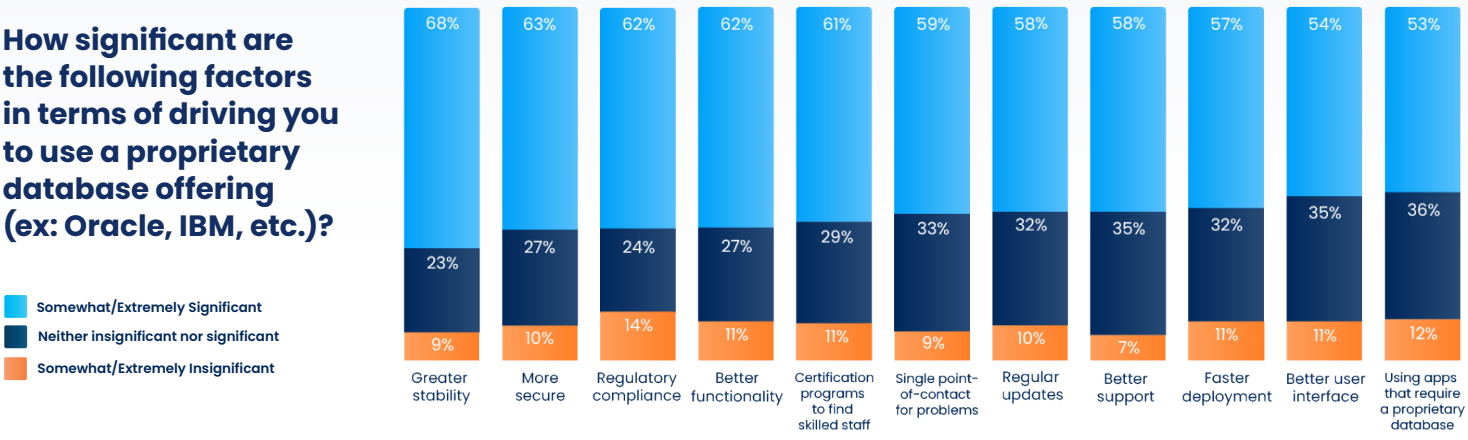

PERCONA

Organizations that give priority to Developer Preference clearly chose open source database solutions over proprietary solutions.

B: For those opting for open source databases, the top drivers were very different: Lower Costs, Developer Preference, and Faster Deployment were cited as the highest priorities.

How significant are the following factors in terms of driving you to use an open source database offering?

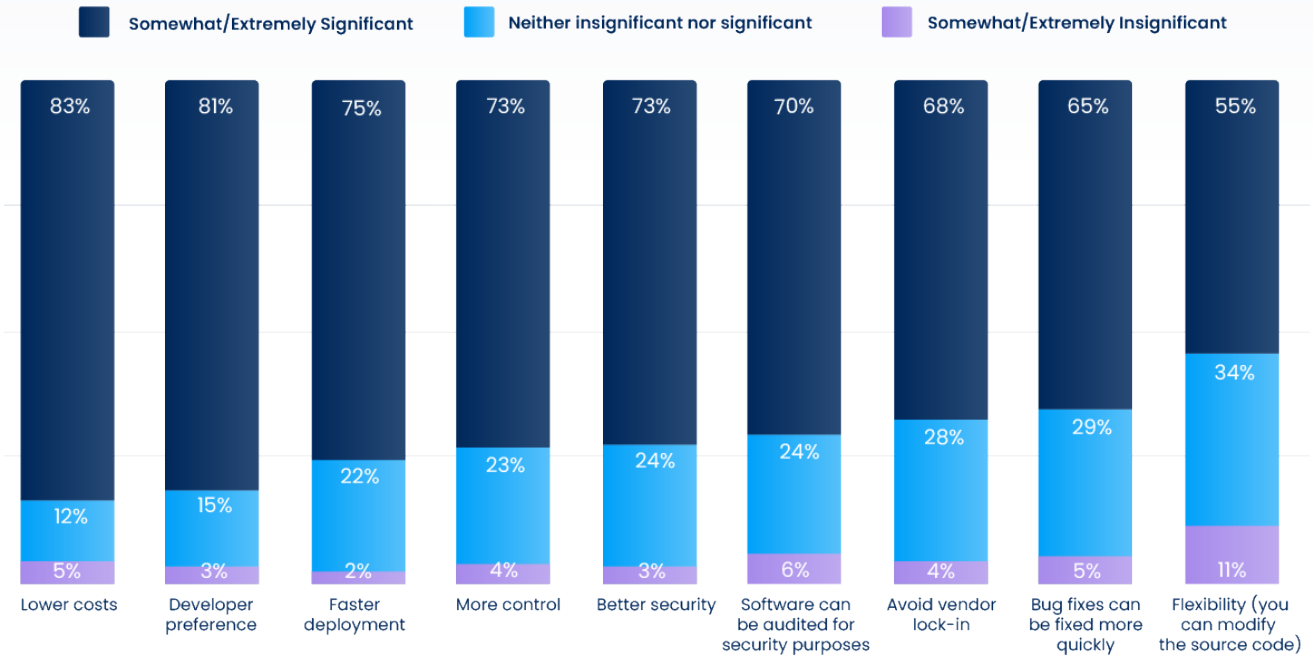

"Developers prefer open source software where they have choice and autonomy over how they can architect and build their applications. The flexibility and performance benefits aside, open source also helps companies save on costs compared to running proprietary databases, which can't be ignored in today's economy. However, companies can still improve how they run and manage those installations and that is where Percona can really help."

Ann Schlemmer | CEO at Percona

PERCONA

High-performers prefer IaaS.

As previously mentioned, when it comes to deploying database instances on-premises/in a private cloud, in the public cloud, or via DBaaS, there's an even split across all the respondents we surveyed:

A: Among all respondents with IaaS deployments, Faster Deployments, Superior Disaster Recovery, and Scalability were the top three drivers behind opting for IaaS.

B: Within just the high-performing group, Speed to Deploy and Develop remained among the highest priorities and grew in importance. Disaster Recovery and Lower Staffing Burdens rose to the top, as well.

How significant are the following factors in terms of driving you to deploy your database instances in the cloud using an IaaS (Infrastructure-as-a-Service) model?

(Somewhat/Extremely Significant)

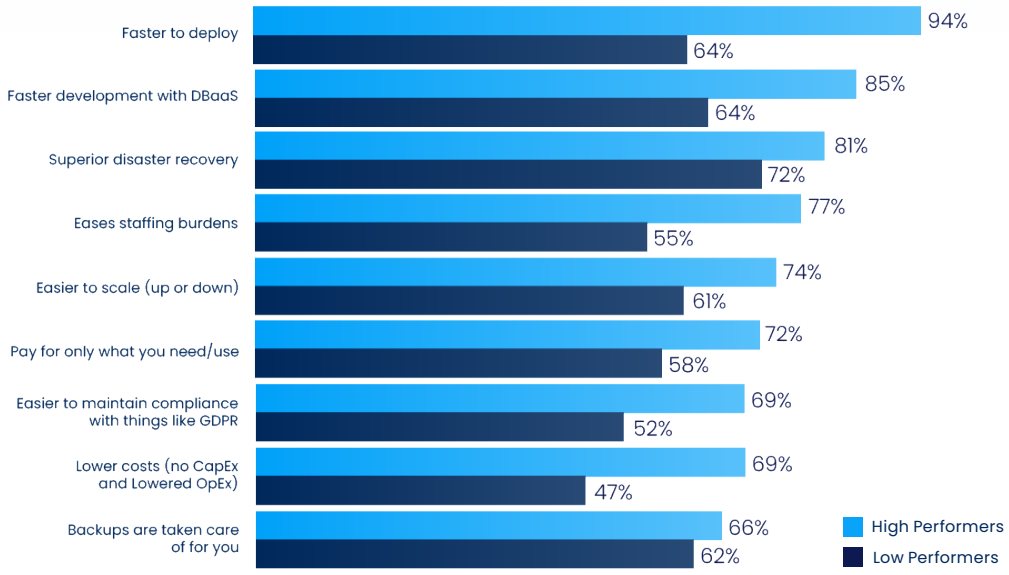

Speedy deployment and scalability top the list of DBaaS considerations.

Among all organizations opting for DBaaS, Faster Development, Superior Disaster Recovery, and Scalability were the most frequently cited criteria. The high-performing group prioritized the same four criteria, but they were on average 18% more likely to assign high priority to these criteria than the low-performing group.

PERCONA

How significant are the following factors in terms of driving you to deploy your database instances in the cloud using an DBaaS (Database-as-a-Service) model?

(Somewhat/Extremely Significant)

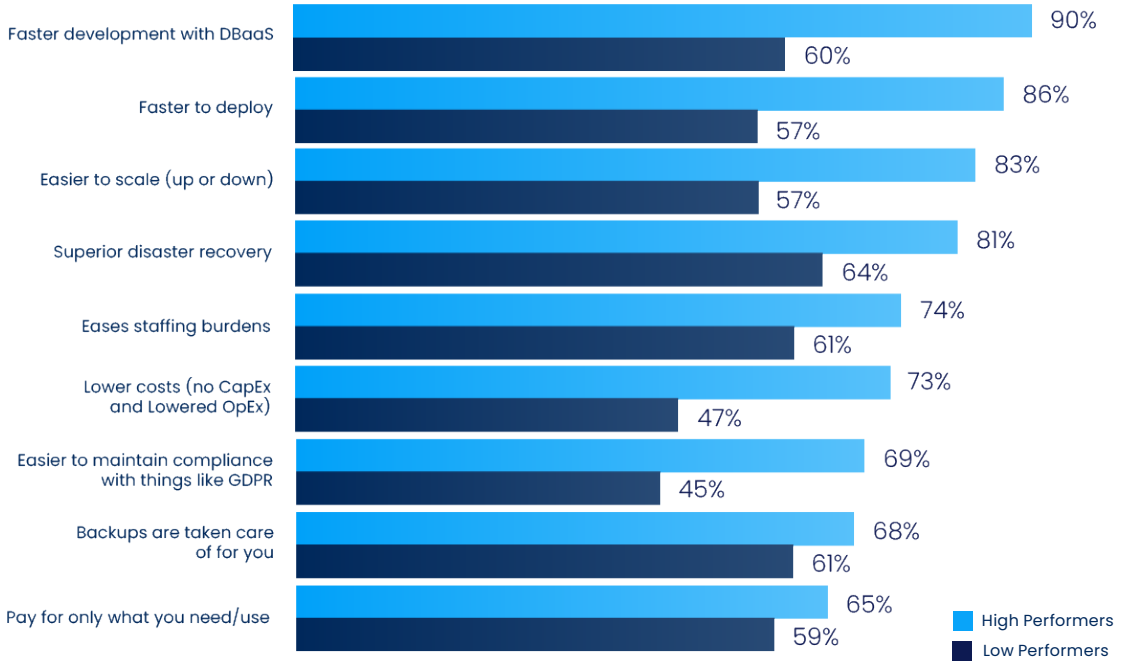

Cost savings, security, and sovereignty top the list of on-prem considerations.

Among respondents with on-premises or private cloud database deployments, Data Sovereignty, Reliability, Security & Privacy, and Lower Costs were most frequently cited as high-priority criteria. When comparing high-performing organizations with low-performing organizations, cost becomes a much more significant driver among high performers choosing on-premises deployments.

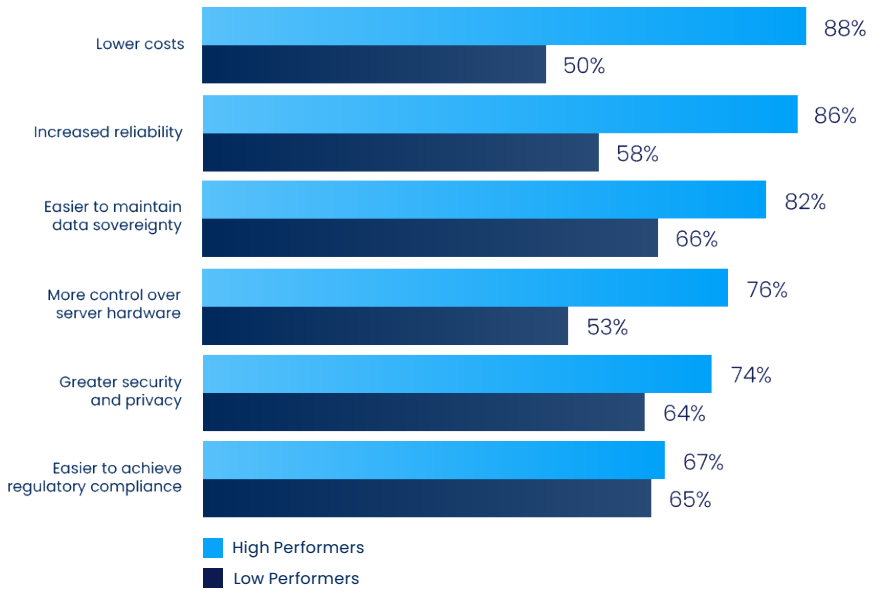

PERCONA

Third-parties are the preferred source of support.

When it comes to working with external partners for database management support, nearly all respondents engage Third-Party Commercial Support providers – more frequently than relying on Vendor support.

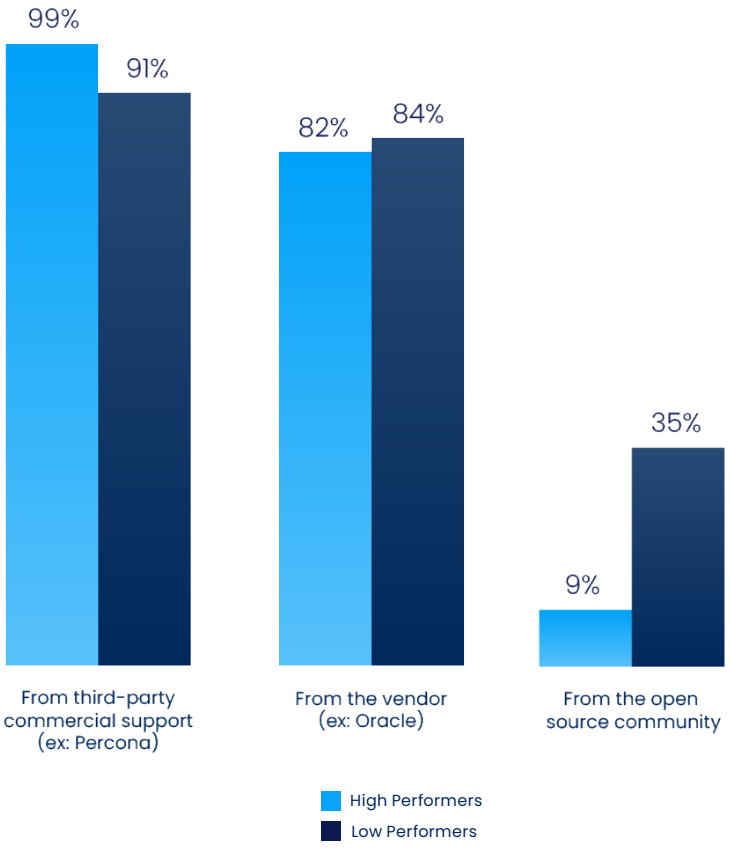

External Support Is Deemed Highly Favorable among High Performers.

Among the high-performing group, a full 100% of respondents rated the support they received from both Third-Party Commercial Support providers and database vendors as Somewhat or Extremely Good. Further, the biggest difference in support quality between the high-performing and low-performing groups lay in their experience with Third-Party Commercial Support, where the high-performing group was twice as likely to rate this form of support favorably as the low-performing group.

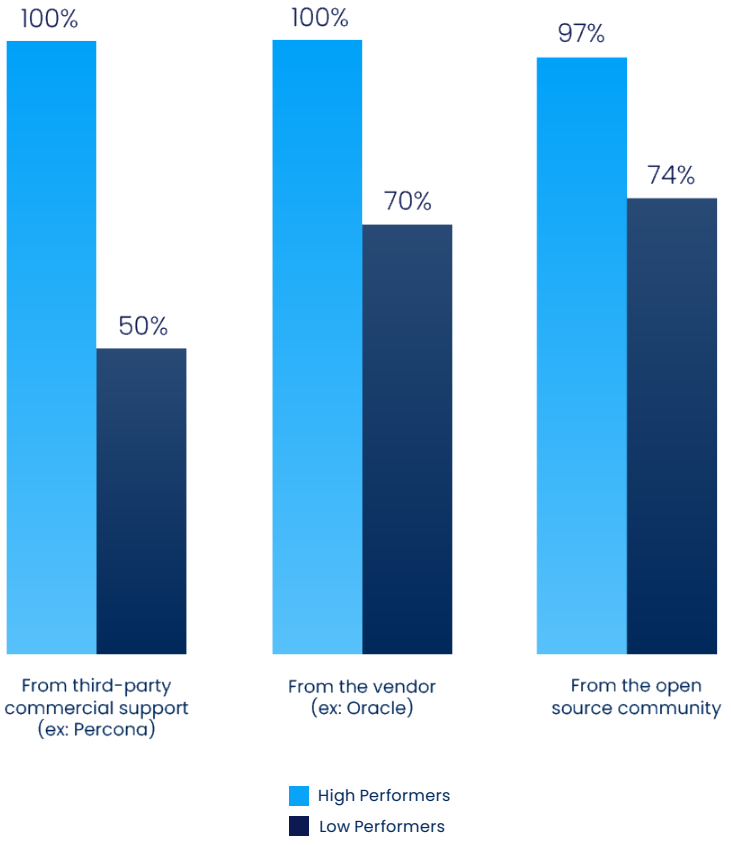

PERCONA

Database strategies in stasis - today vs. three years from now

We also asked respondents for projections of their database operations in three years. By most measures, the database market is currently in a period of stasis around database deployments and changes.

NOW

As of today, what percentage of your database instances are deploying in each of these ways?

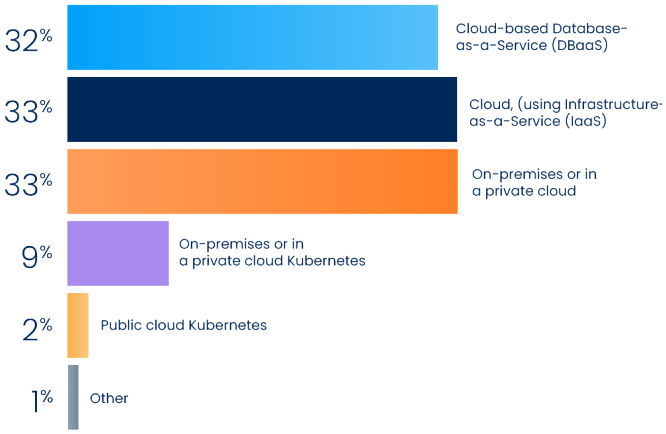

IN THREE YEARS

In three years, what percentage of your database instances will be deployed in each of these ways?

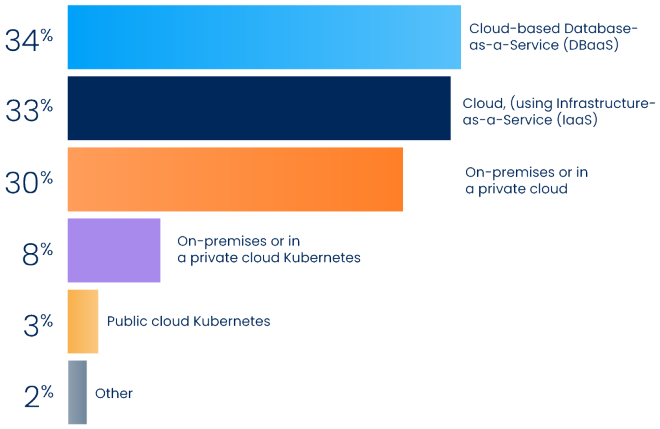

Firstly, we see that the types of deployments (e.g. DBaaS, Cloud, On-premises/ Private Cloud) being used across organizations (as a percentage of total deployments) are expected to remain largely unchanged over the next three years. Both currently and in three years' time, there is a rather even distribution between DBaaS (32% to 33%), IaaS (33% to 34%), and on-premises or private cloud deployments (33% to 30%) – with each model representing roughly a third of the organizations' total database deployments.

PERCONA

NOW

As of today, what percentage of your database instances are licensed in each of these ways?

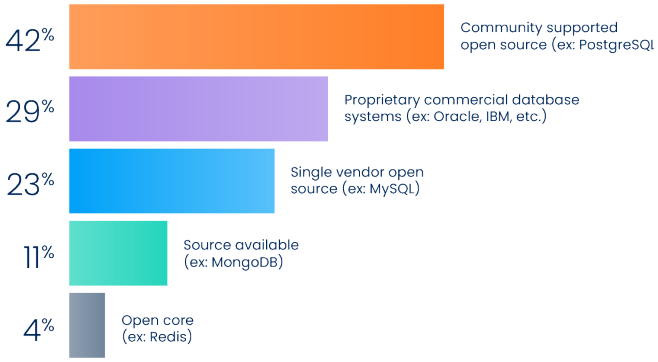

IN THREE YEARS

In three years, what percentage of your database instances will be licensed in each of these ways?

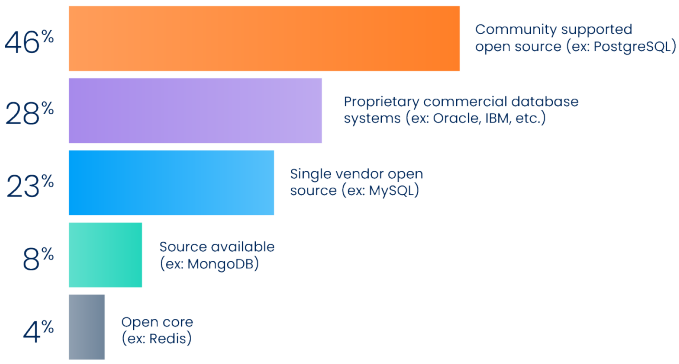

We see the same trend of general stasis when looking at the percentage of database instances deployed under various licensing models (e.g. open source, proprietary). Both today and in three years, we see that community-supported open source databases, such as PostgreSQL, represent the largest proportion of database instances in use.

Currently, open source databases represent 42% of respondents' deployments, and they are expected to account for 46% of total deployments in three years. Proprietary commercial databases accounted for the next largest percentage of database instances, albeit a significantly smaller share – 29% currently and 28% in three years.

NOW

Which of the following relational database products does your organization use today? Mark all that apply.

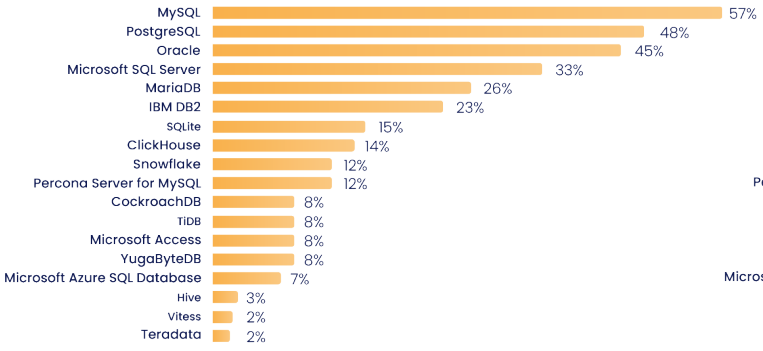

IN THREE YEARS

Which of the following relational database products will your organization use in three years? Mark all that apply.

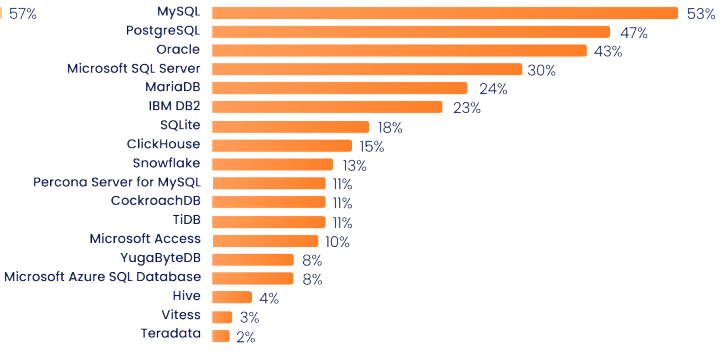

PERCONA

When looking at which specific relational databases were in use, we found that MySQL was the most widely used currently and likely to remain so, with 57% of respondents using the open source database now and only slightly fewer anticipating using it in three years. PostgreSQL was the second most popular relational database, with 48% of respondents using it currently and nearly as many anticipating using it in the near future.

The two most popular proprietary databases — Oracle and Microsoft SQL Server — followed in third and fourth place with 45% and 33% of organizations currently employing each, respectively. Like their open source counterparts, neither is likely to see significant changes in adoption over the next three years, with each expected to see a 2% decline.

"Our report indicates that database deployment strategies are expected to remain static for the next few years, with levels of usage remaining consistent across all categories of relational databases. We're seeing caution around making significant changes to deployments following the large-scale digital transformation undertaken during the pandemic and the uncertain economic climate the world is facing in 2023. Teams want to improve how they manage, monitor, and run those deployments rather than implementing more changes."

Joe Brockmeier, Head of Community at Percona

Although we don't know for certain, it's likely that the current holding pattern we're seeing in the database market can be attributed to economic uncertainty and post-COVID-19 recovery. With a recession looming and other macroeconomic challenges at play, it appears that most organizations are not planning to make significant changes to their database strategies over the coming three years.

Large enterprises enjoy increased access to tools and automation; startups do without oracle

When we compared the database strategies of large enterprises with the strategies of small and medium-sized businesses, much of what we found was rather intuitive. For example, larger organizations were more likely to deploy a greater number and variety of database instances than their smaller counterparts. Similarly, we saw that large enterprises are much more likely to use database monitoring and management solutions, with 93% of large organizations having such tools in place and only 57% of small businesses.

At the same time, small companies are much more likely to use on-premises and cloud services to host their databases (49%) compared to 26% of large enterprises. Large enterprises are also much more likely to use DBaaS offerings (32% of deployments). In all cases, one can imagine these differences can be chalked up primarily to larger organizations having larger budgets at their disposal.

However, some unexpected insights did emerge when comparing these cohorts. First, we saw that, while the use of Oracle was practically ubiquitous in this study, the only exception to this rule was among smaller organizations (<100 employees), of which just 11% are using Oracle. We also found that the proprietary licensing model was the second most popular among all cohorts except the smallest (<100 employees).

The price tag for proprietary database solutions is no doubt a factor for smaller organizations whose pockets are less deep, but we also know that smaller company size skews toward digital native startups, which are more likely to seek out more modern solutions like the open source alternatives to Oracle and other proprietary solutions. This is consistent with the commonly-held belief that Oracle and other proprietary databases are benefiting primarily from support from legacy enterprises.

Demographics

We surveyed 284 professionals with responsibilities related to database procurement, deployment, management, performance, and/or security.

The majority of respondents are based either in the U.S. or Canada (70%), and the remainder hail from more than 30 other countries around the world.

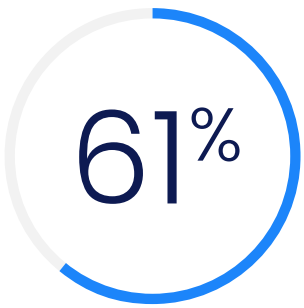

United States

Canada

Other Countries

Our survey drew in participants from organizations ranging in size from fewer than 100 to more than 5,000 employees.

46% fewer than 500 employees

34% 500 to 2,500

21% 2,500 or more

- Less than 100
- 100–499
- 500–999
- 1,000–2,499
- 2,500–4,999
- More than 5,000

PERCONA

Most respondents report working in the technology sector, but more than 20 other industries are represented in this report.

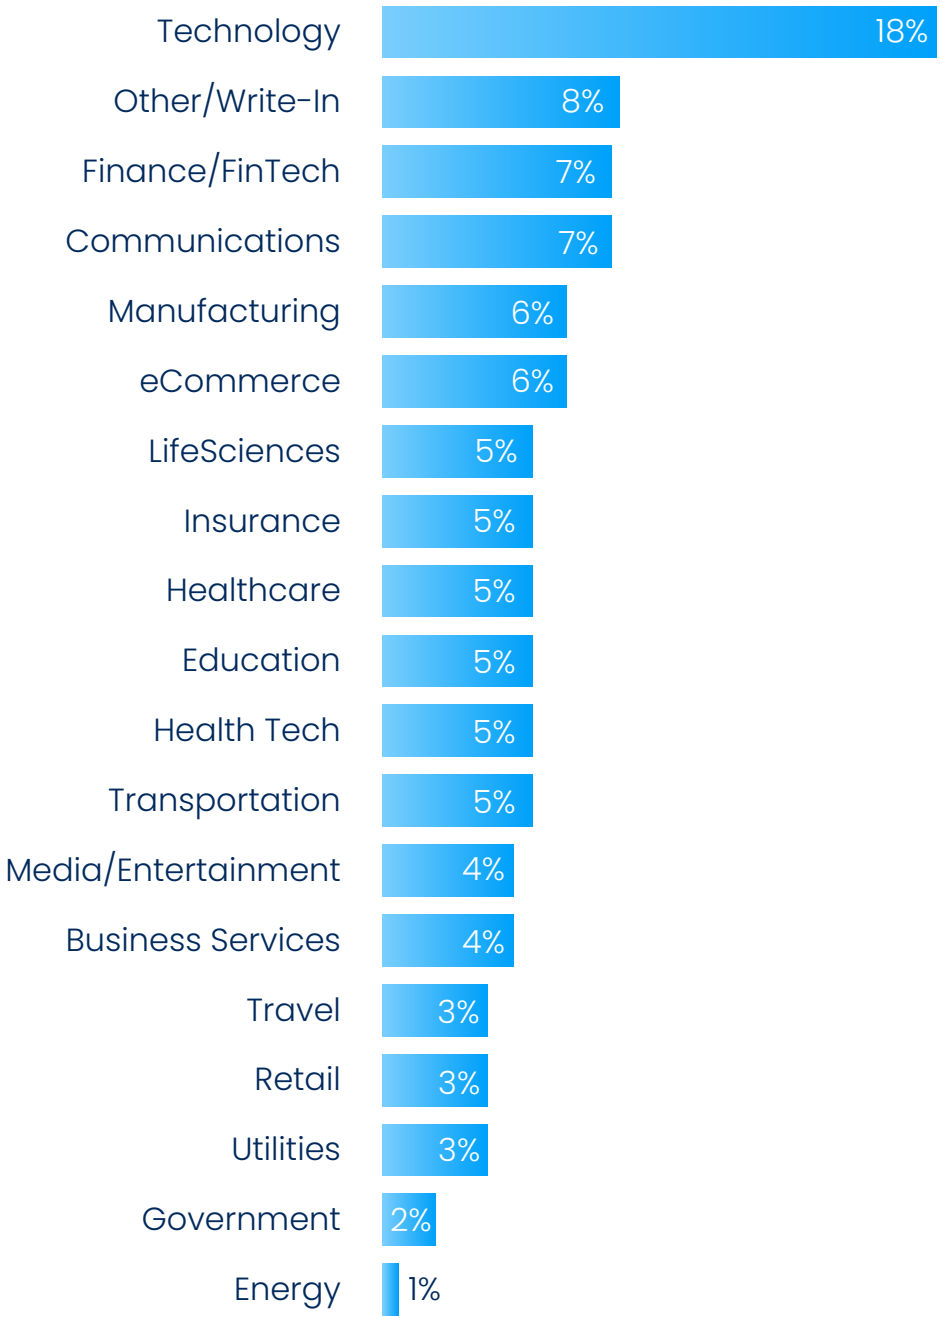

Overall, this comprehensive cross-section of respondents gives this report a well-balanced view of what's happening out in the field where critical database operations are central to keeping businesses moving forward.

PERCONA

About Percona

Percona is widely recognized as a world-class open source database software, support, and services company. The organization is dedicated to helping businesses make databases and applications run better through a unique combination of expertise and open-source software. Percona works with numerous global brands across many industries, creating a unified experience to monitor, manage, secure, and optimize database environments on any infrastructure.

Percona equips organizations with the freedom to choose, the freedom to create, and the freedom to make a difference – helping them scale and innovate with speed as they grow.

Databases run better with Percona.

For more information, visit **www.percona.com**

Learn more

2023 Survey Report