


Industry trends for database deployment on Containers and Kubernetes



Organizations are adopting containers and Kubernetes to develop cloud-native applications and modernize existing applications. Data is a critical component of all applications, and deploying databases on containers and Kubernetes allows for faster access to data and scalability.

Gartner Peer Insights and Red Hat surveyed 200 tech leaders across the globe to understand how they're adopting databases on containers and Kubernetes, and what they are considering as they deploy crucial workloads.

Data collection: January 26 - February 27, 2022

Respondents: 200 tech decision makers

Organizations are adopting multiple approaches to deploying databases on containers and Kubernetes

Deploying databases on containers and Kubernetes is very popular and respondents have many ways of consuming them. Database-as-a-service offerings from cloud vendors (61%) and deploying cloud-native databases such as Crunchy Data, Couchbase, MongoDB etc. (52%) are the most popular ways of doing so. Additionally, respondents often use more than one of these methods to serve their purpose.




Adoption of databases on containers and Kubernetes will grow in 24 months

More organizations are shifting towards adopting databases on containers and Kubernetes. There is an increase in both the respondents that plan to deploy a cloud-native database and those that plan to consume a DBaaS offering from independent database vendors. Respondents that plan on connecting to databases outside containers and Kubernetes are also decreasing, pointing to more organizations wanting to adopt databases on containers and Kubernetes.




Databases on containers and Kubernetes are being utilized to build cloud-native apps and for AI/ML

While a majority of respondents have started adopting databases on containers and Kubernetes, few are far along in this process. Since implementation can be challenging, organizations need support from platform providers and database providers to modernize and build new cloud-native applications with databases on containers and Kubernetes.



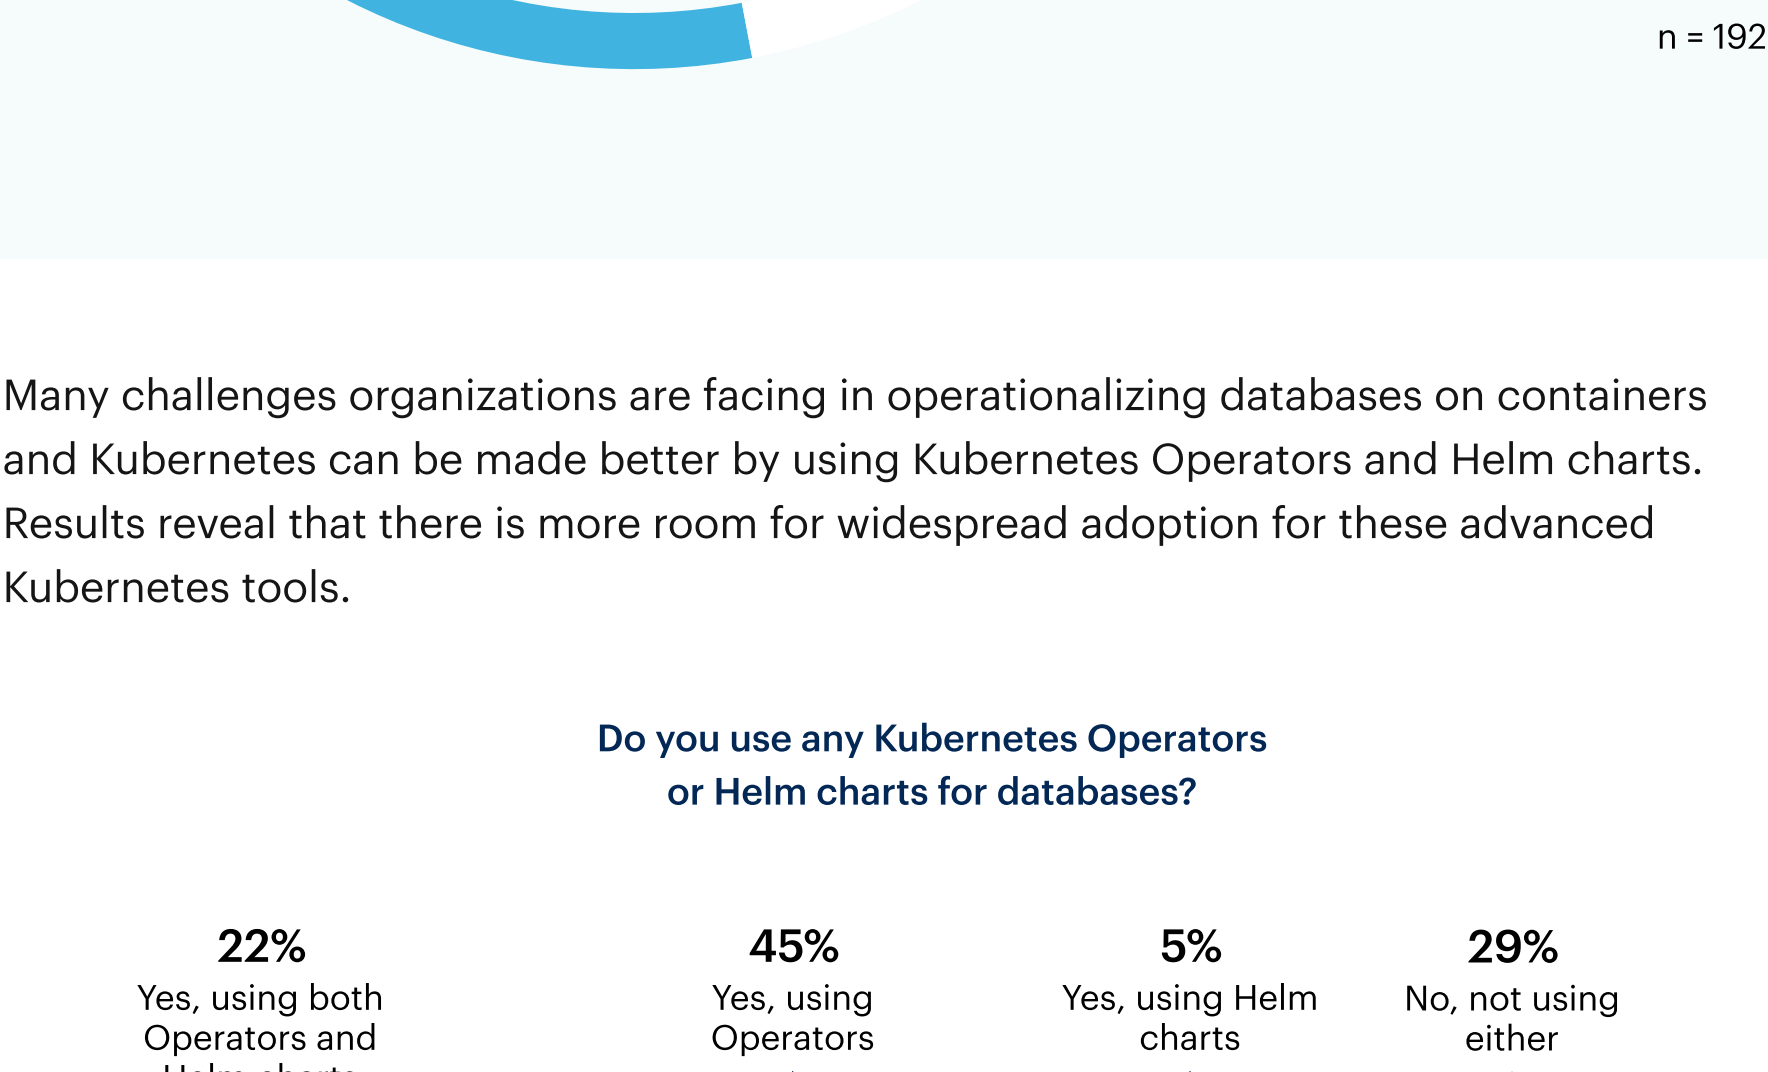
Databases are being adopted on containers and Kubernetes for both transactional and analytical use cases. A significant number of respondents have also begun to modernize their traditional databases.

Databases are being adopted on containers and Kubernetes for both transactional and analytical use cases. A significant number of respondents have also begun to modernize their traditional databases.



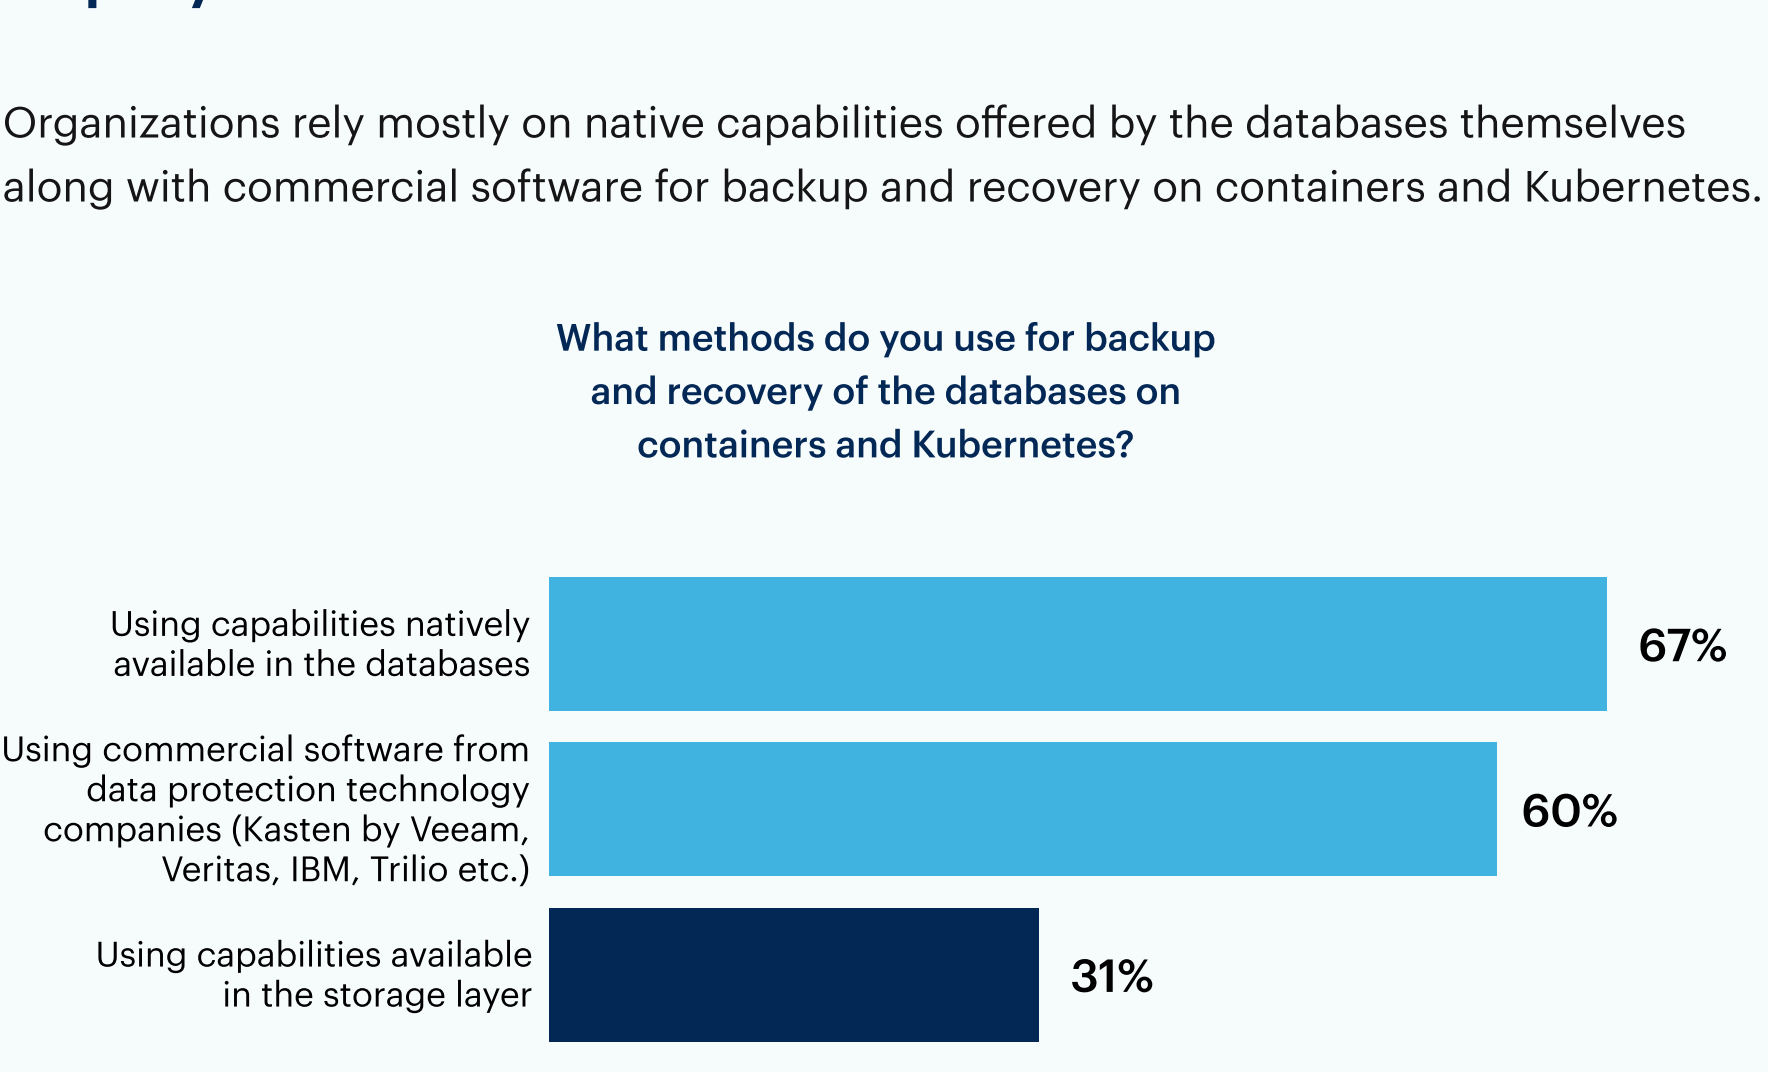
Modernizing non-cloud enabled databases is a priority for tech leaders

As organizations modernize their applications, over half of respondents (51%) plan on containerizing non-cloud enabled databases, which may lead to a greater adoption of cloud-native databases and DBaaS.



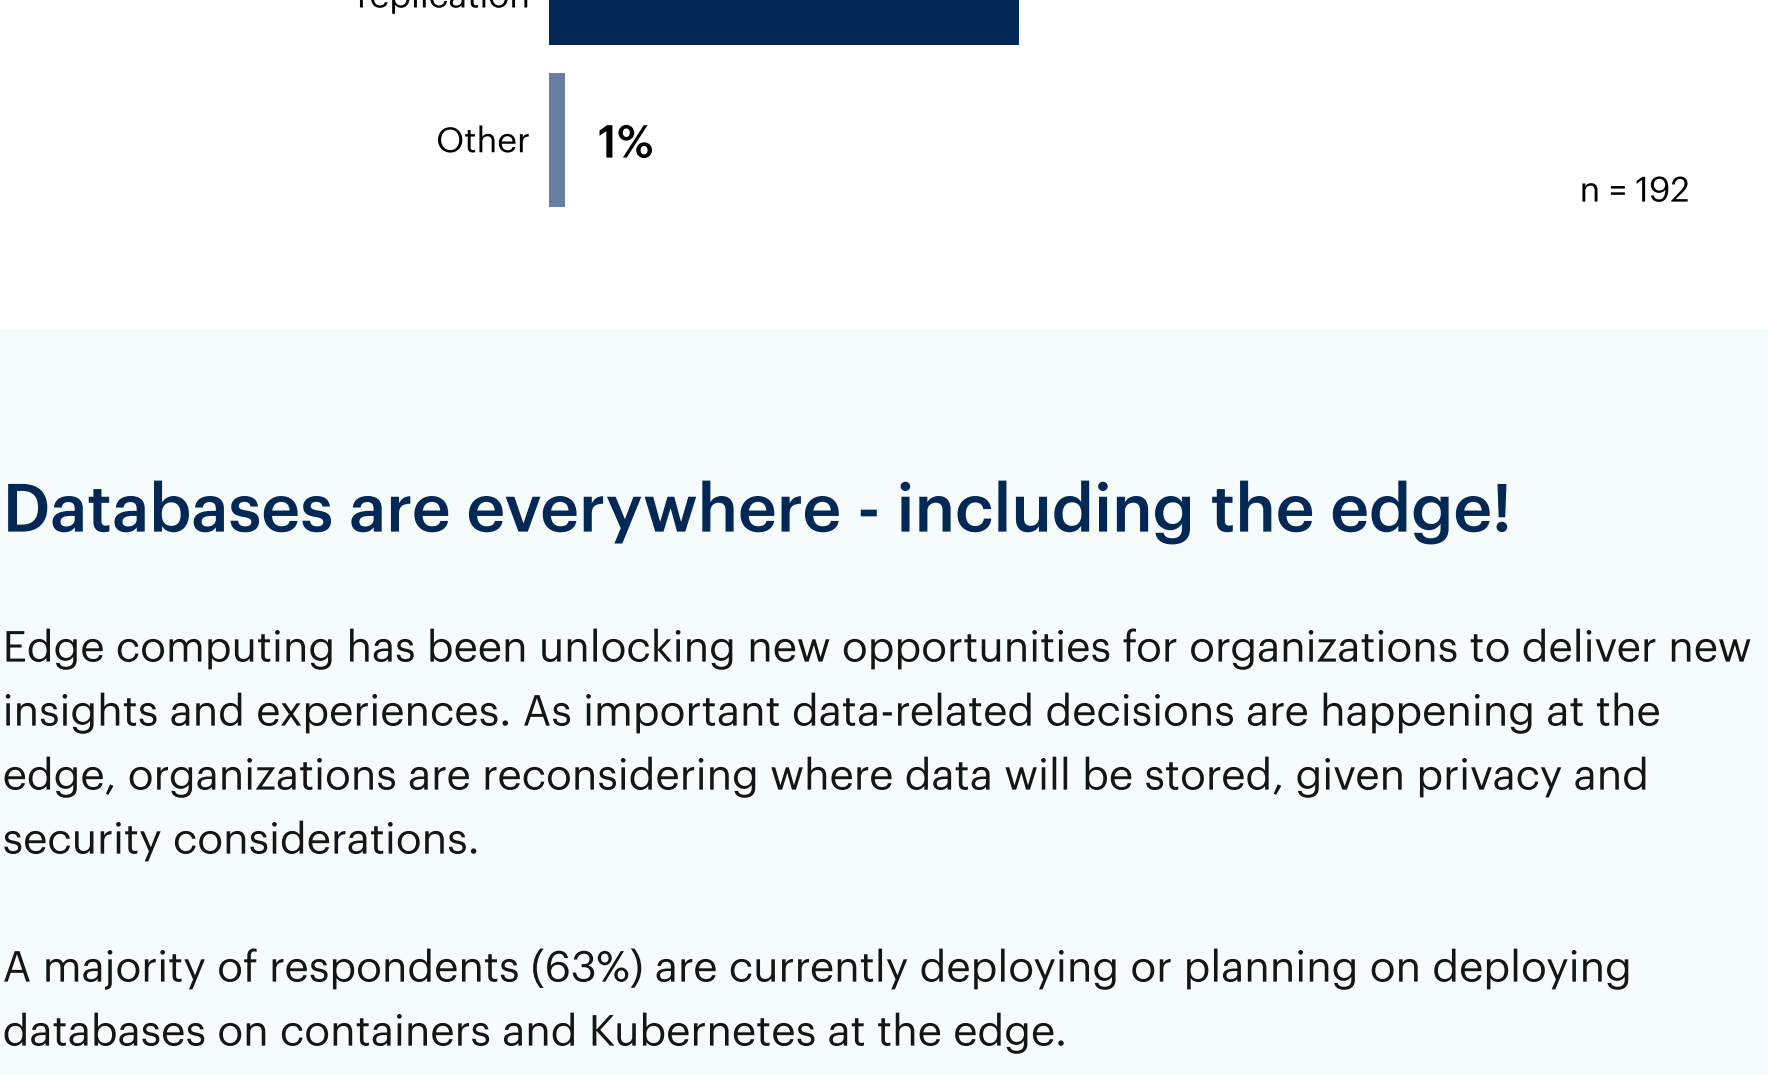
Organizations face challenges with operational tasks like installations and updates that Kubernetes Operators can help solve

Organizations face typical day-1-2 operational challenges as they ramp up databases on containers and Kubernetes.



Many challenges organizations are facing in operationalizing databases on containers and Kubernetes can be made better by using Kubernetes Operators and Helm charts. Results reveal that there is more room for widespread adoption for these advanced Kubernetes tools.

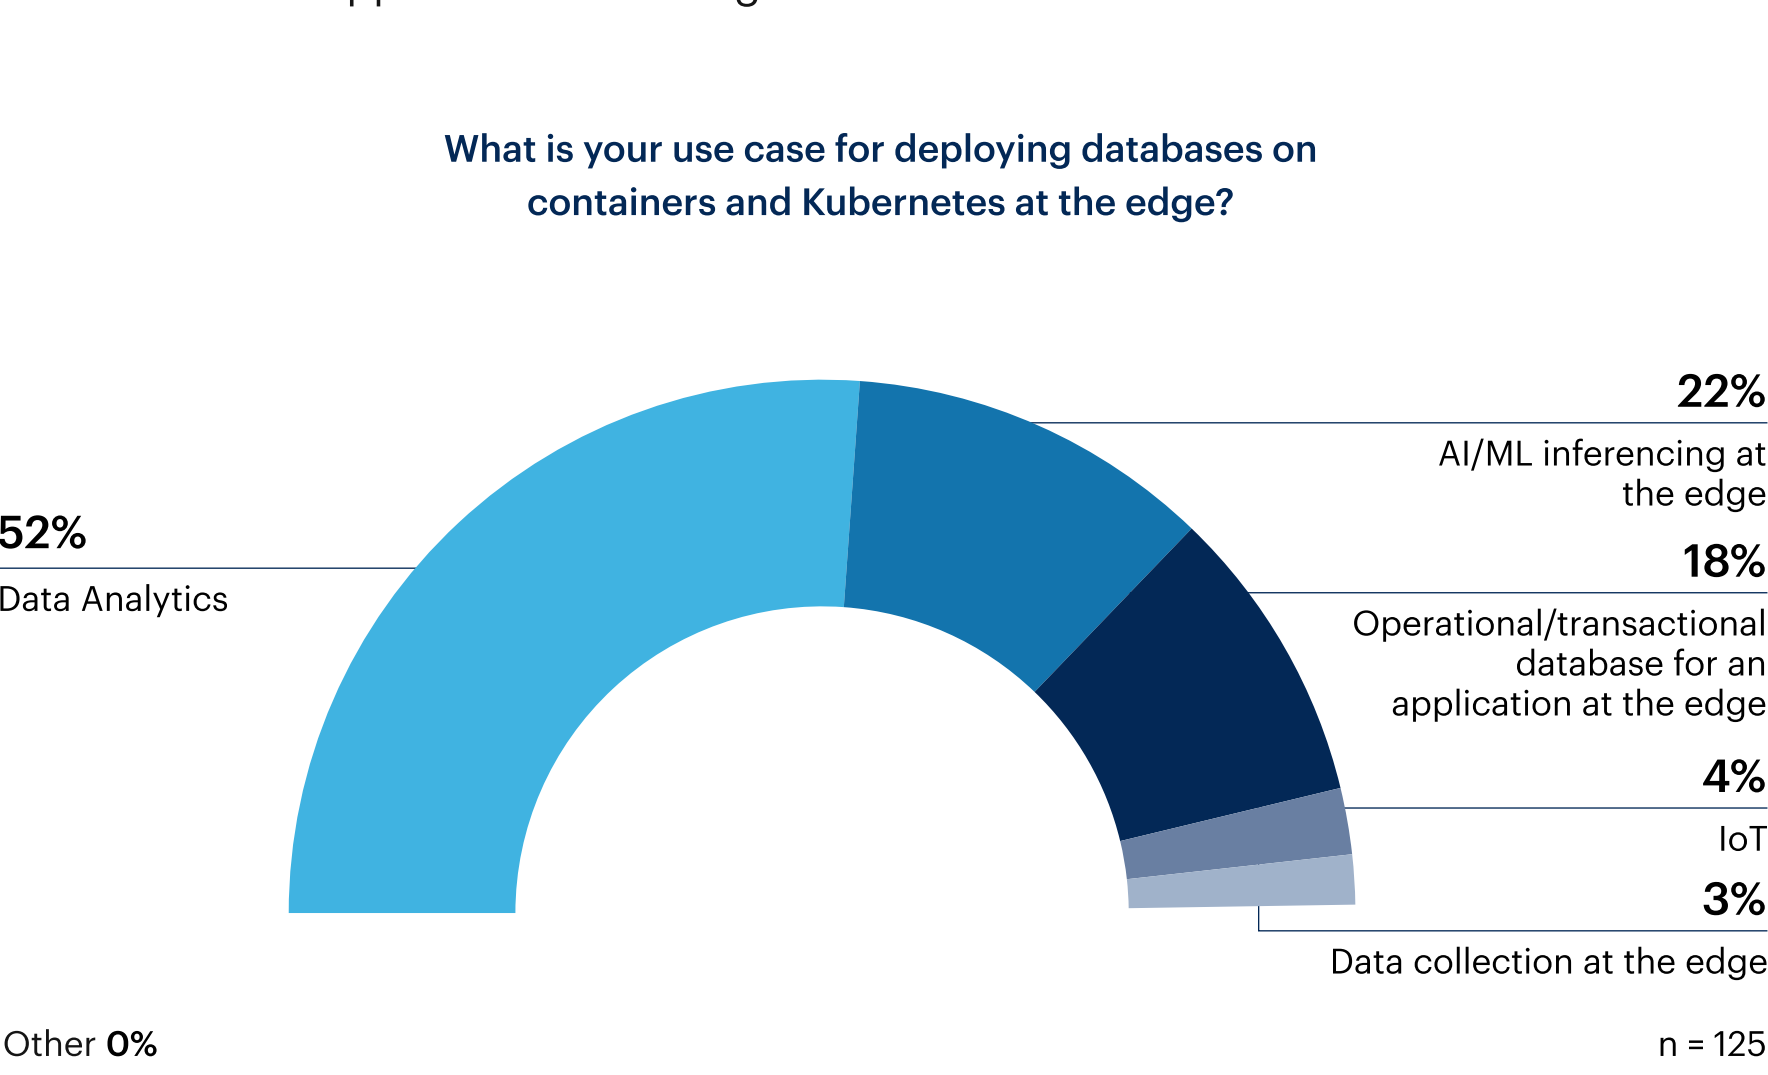
Many challenges organizations are facing in operationalizing databases on containers and Kubernetes can be made better by using Kubernetes Operators and Helm charts. Results reveal that there is more room for widespread adoption for these advanced Kubernetes tools.



Kubernetes Operators, and Helm charts help automate day-1 and day-2 operations, such as installation, configuration and updates and upgrades, as well as help with lifecycle management of applications. Most database providers have Kubernetes Operators.

Backup and disaster recovery are important considerations when selecting databases and deployment environments

Organizations rely mostly on native capabilities offered by the databases themselves along with commercial software for backup and recovery on containers and Kubernetes.




What methods do you use for disaster recovery?



Databases are everywhere - including the edge!

Edge computing has been unlocking new opportunities for organizations to deliver new insights and experiences. As important data-related decisions are happening at the edge, organizations are reconsidering where data will be stored, given privacy and security considerations.

A majority of respondents (63%) are currently deploying or planning on deploying databases on containers and Kubernetes at the edge.



Of those who plan to deploy databases on containers and Kubernetes at the edge, almost three-quarters (74%) plan to use it for either data analytics or AI/ML inferring at the edge. In contrast, only 18% need an operational/transactional database for an application at the edge.

As you can see, Kubernetes users are frequently mixing multiple methods of integrating data services into their clusters, with a combination of cloud databases, databases deployed directly through Kubernetes and also connecting to VMs outside of the cluster running data services. A platform that enables all three is key to developer agility.

Red Hat OpenShift supports a vast breadth of databases, from self-managed cloud-native databases, to Database-as-a-Service offerings from both database vendors and cloud providers. The integrations with database technology partners also allows native backup and recovery capabilities directly within OpenShift. Red Hat's [OpenShift Database Access](#) also simplifies access to managed cloud databases directly from OpenShift cloud services to further accelerate how quickly developers bring database-centric applications to production.

Respondent Breakdown

