

# END USER PERSPECTIVES ON RUNNING STATEFUL APPLICATIONS ON KUBERNETES

September 2021



Kubernetes is continuing to play a huge role in this era of digital transformation that accelerated when the pandemic began. Increasingly mission critical applications such as databases are managed using Kubernetes in order to increase business agility and application resilience. The business results of doing so are impressive. Apps are delivered faster, scale easier, enjoy greater security and, often, do so at reduced cost. But the transition to new application management patterns is not easy. Enterprises face challenges meeting business requirements like data protection, mobility and security, with common solutions such as vendor support proving costly when many different data services are involved. For greater adoption of Kubernetes to continue, enterprises will need to solve the challenges around data protection, mobility, capacity management and support that come with enterprise applications.

This survey was conducted among 500 IT professionals in the U.S. and the U.K. who are knowledgeable of their company's Kubernetes usage, to assess how adoption of running stateful applications on Kubernetes has shifted, the benefits and efficiencies it has on their companies' bottom lines, and what the future may hold.

## KEY FINDINGS:


- **An overwhelming majority (99%, n=500)** of respondents agree that running stateful applications on Kubernetes have impacts to their company's bottom line, for the following reasons:
  - **Over half (55%)** agree it allows them to scale their applications faster and allows their developers to be more efficient
  - **Over half (54%)** agree it allows them to develop applications more quickly
  - **Half (50%)** agree it allows them to increase security by leveraging automation
  - **36%** agree it reduces IT budget
- **More than half (55%, n=500)** of respondents run stateful applications on Kubernetes instead of using a managed database-as-a service offering from the public cloud because they **need more control over their data services**.
  - Additional reasoning (n=500) includes needing more customization options (47%), needing to run data services on-prem (44%), data service is not available as a managed service (42%) and the cost of a cloud database-as-a-service is too high (41%)
- When asked what stateful applications respondents currently run on Kubernetes (n=500), **41% indicated MySQL**, followed closely by GitLab and PostgreSQL (31%), RabbitMQ (28%), MongoDB (27%) and Jenkins, TensorFlow and Cassandra (26%), among others.
- **More than half (55%, n=500)** of respondents noted Backup & Restore is a key requirement for stateful applications running on Kubernetes, followed by Data Mobility and Capacity Management (49%), High Availability (48%), Multi-Cloud/Multi-Region Support (45%), Encryption and Disaster Recovery (43%) and Role-Based Access Controls (38%).
  - Among these requirements (n=499), **29% cited Data Mobility** as the requirement they most struggle to achieve, followed closely by High Availability and Capacity Management (28%), Backup & Restore and Multi-Cloud/Multi-Region Support (26%), Disaster Recovery (25%) and Role-Based Access Controls and Encryption (22%).

55%




of respondents noted Backup & Restore is a key requirement for stateful applications running on Kubernetes


- **Almost half (46%, n=500)** of respondents cite Data Protection as the biggest operational challenge when managing stateful services, followed by Managing Backups (41%), Service Discover & Ingress to Database Shards and Managing Storage (38%), Version Upgrades (36%), Operational Knowledge of Each Specific Data Service (35%) and Deployments (30%).
- When asked what methods of support respondents use for data services running on Kubernetes (n=500), **72% indicated Support from the Vendor Most Closely Aligned with the Data Service**, followed by Relying on our Internal Expertise (62%) and Crowd Sourcing from the Open-Source Community (52%).
  - Among those who indicated Support from the Vendor Most Closely Aligned with the Data Service, when asked what challenges they have getting support for data services running on Kubernetes from data services vendors:



agree that managing a support contract with different vendors for different data services is expensive




agree that there is added complexity to their team due to different vendor approaches to Day 2 Ops



agree that the cost of supporting many database instances is too high

- **Over half cite Increasing Agility (58%, n=500) and Increasing Resiliency (52%, n=500)** as the biggest drivers behind their team's decision to build and deploy stateful applications on Kubernetes, followed by their App Team's Desire to Run Stateful Apps on Kubernetes (48%), Reducing Maintenance Burden (45%) and Reducing Costs (43%).
- **On average, over half (53%, n=500)** of respondents' workloads running on Kubernetes are stateful.


- On average, the typical DevOps or IT Ops engineer saves **17 hours a week** (n=500) due to operational efficiencies from running stateful applications on Kubernetes.



**17hrs a week**

saved by the typical DevOps or IT Ops engineer due to operational efficiencies

- **A majority (87%)** of respondents expect the percentage of stateful workloads over the next 12 months to increase. Only 9% expect it to stay the same, followed by 4% who expect it to decrease.



**87%**

of respondents expect the percentage of stateful workloads over the next 12 months to increase.

## METHODOLOGY

The Pure Storage Survey was conducted by Wakefield Research ([www.wakefieldresearch.com](http://www.wakefieldresearch.com)) among 500 IT Professionals, employed full-time in the US (250) or UK (250) at companies of 500+ employees, who self-identify as possessing significant knowledge of their company's Kubernetes usage and currently running stateful applications on Kubernetes. The survey ran between September 9th and September 17th, 2021, using an email invitation and an online survey. Quota were set for 250 respondents in the US and in the UK. Results of any sample are subject to sampling variation. The magnitude of the variation is measurable and is affected by the number of interviews and the level of the percentages expressing the results. For the interviews conducted in this particular study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 4.4 percentage points from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample.