

Introducing Portworx Data Services



Today, we're excited to introduce Portworx Data Services (PDS), the industry's first database-as-a-service platform. Pure customers are reinventing their businesses and dominating their industries by building modern applications. These apps run in containers and on Kubernetes. They're backed by a growing number of SQL, NoSQL, search, streaming, and analytics data services.

PDS provides a production-grade, managed database-as-a-service experience for running these services so developers can focus on using data services, not managing them.

This blog post delves into why we built PDS and how you can use it to accelerate your digital transformation. Want to get started?

Modern Apps Need a New Approach

A movie streaming service is very different from an analytics platform used to manufacture vaccines. However, they are built using similar architectures. Microservices, containers, and Kubernetes form the basis of that architecture.

One key element of microservices is to use many different data services instead of a single monolithic database. Modern apps use many different, specialized data services, including SQL options like Postgres and [MySQL](#). They also use NoSQL options like MongoDB, Cassandra, Elasticsearch, Kafka, and many others.

But each of these data services has its own way of being deployed, its own way of monitoring, and its own way of managing backup and recovery. That puts significant strain on DevOps teams responsible for managing these services. They not only have many database instances to manage but also a large number of database types. As a result, DevOps teams are focused on firefighting deployments and ops rather than delivering innovation.

A self-service, on-demand way to provision these critical services make them easily scalable, available, and secure 24×7. That's PDS.

PDS: The Industry's First Database-as-a-Service Platform for Kubernetes

With PDS (and the click of a button), DevOps engineers can deploy a production-grade, managed data service on Kubernetes from the industry's broadest catalog of options.

Each data service instance can be easily configured with automated HA, backup and DR, encryption, migrations, and more. As a result, DevOps teams can focus on using data services, not managing them.

Why We Built Portworx Data Services

A [recent survey](#) that we conducted spoke to the challenges customers face running data services in production. Almost half of the 500 users surveyed (46%)

cited data protection as their biggest operational challenge when managing stateful services. This was followed by managing backups (41%), service discovery and ingress to database shards and managing storage (38%), version upgrades (36%), operational knowledge of each specific data service (35%), and deployments (30%).

It's no surprise that running data services is difficult. So we asked: How can Portworx® make it easier?

We started by looking at what customers were doing without Portworx. We realized that customers are faced with two unsatisfactory options for managing the data services backing their modern apps.

The first option was using a managed database service from a cloud provider. This is a common approach, but it limits customers to running specific versions of a database in a single cloud. It's also expensive.

This problem is so pronounced that in the survey mentioned above, over half (55%) of respondents choose to run stateful applications on Kubernetes instead of using a managed database-as-a-service offering from the public cloud. Why? Respondents cited the need for more control over their data services. Additional reasons included needing more customization options (47%), needing to run data services on-prem (44%), data service isn't available as a managed service (42%), and the cost of a cloud database as a service is too high (41%).

The second option is to hire specialized DBAs to handle operations. That can be challenging as well. Finding people with the appropriate expertise is hard, and they tend to be expensive to hire. Purchasing vendor support is also an option but becomes costly as you add many different data services.

In the survey, respondents indicated whether they tended to rely on 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:

- **69%** agreed that managing a support contract with different vendors for

different data services is expensive.

- **61%** agreed that there is added complexity to their team due to different vendor approaches to Day 2 Ops.
- **55%** agreed that the cost of supporting many database instances is too high.

Why Portworx Data Services Is Different

We built Portworx Data Services to provide a managed experience like you get in the cloud but without the drawbacks of other approaches.

PDS gives developers that self-service, on-demand experience to create the data services they need to build modern apps. And it dramatically simplifies operations by providing a managed, as-a-service experience.

In the end, teams have a consistent operational experience for all of their data services regardless of environment. They can easily solve HA, backup and DR, encryption, and migrations with the click of a button. If that sounds like what your team needs, learn more about [Portworx Data Services](#) and sign up for a trial.

Written By:

[Michael Ferranti](#)