

Easily Operate a Database-as-a-Service Platform

Leverage Portworx® by Pure Storage® to manage your production-grade, mission-critical data services.

Modern applications built with microservices require access to SQL and NoSQL databases like MySQL, PostgreSQL, Cassandra, and Couchbase, and to streaming, search, and AI/ML data services like Elasticsearch, and Kafka. A traditional approach to managing these data services relies on scripting and numerous repetitive and manual activities, resulting in significant effort to ensure databases remain online and support mission-critical applications. These highly manual activities are prone to human error and excessive downtime, and can't be easily automated with existing infrastructure automation tools.

While Kubernetes provides some support for such automation, running data services on Kubernetes remains complex. DevOps teams today struggle to manage the numerous data services deployed on Kubernetes across the enterprise, as each has its own set of complex requirements around performance, high availability, data protection, data security, and more. Instead of delivering new features and applications to market, DevOps teams instead spend much of their time firefighting deployments and managing the operations of these data services.

What DevOps teams need is an as-a-service experience that delivers reliable, performant data services for their apps running on Kubernetes, without the vendor lock-in that accompanies existing database-as-a-service (DBaaS) platforms hosted by cloud service providers. [Enter Portworx Data Services.](#)

The First Database-as-a-Service Platform for Kubernetes

Portworx Data Services solves these problems by giving DevOps engineers the industry's first DBaaS platform for Kubernetes. With Portworx Data Services, databases and other stateful services can be deployed and managed on any Kubernetes cluster with a single click, giving enterprises the advantages of a managed database offering without the



Fully-Automated Management

Fully automated Day-2 operations, so DevOps teams can focus on using data services, not managing them.



Broad Catalog of Data Services

Choose from a wide variety of data services including Cassandra, Couchbase, Elasticsearch, Kafka, MongoDB, Postgres, and more.



Single-Click Deployment

Deploy an enterprise-grade, production data service quickly and easily.



drawbacks of vendor lock-in. By focusing on delivering three key elements, Portworx Data Services provides a DBaaS platform that makes running data services on Kubernetes dramatically simpler:

Single-click deployment: Modern applications are composed of dozens or even hundreds of microservices supported by data services such as databases, streaming and message queues, search and AI/ML pipelines. Managing so many types of data services is complex, but this complexity is magnified by the number of database instances used to manage and scale across disparate test, dev, and production environments, availability zones, and clouds. Portworx Data Services enables enterprises to offer single-click deployment for developers and ops teams alike. With the click of a button, DevOps engineers can deploy a production-grade, managed data service on Kubernetes. Teams can easily configure data services with automated high availability, backups, encryption, and more. Developers benefit by having self-service databases and unified APIs, while DevOps teams can deploy a data service on-demand, accelerating the continuous integration and delivery (CI/CD) pipeline and application delivery.

Fully automated management: Managing multiple data services across hundreds of microservices is time-consuming. Each data service takes a different approach to configuring key business requirements—and scaling these data services to support multiple applications across a modern Kubernetes environment is nearly impossible.

Portworx Data Services brings simplicity into this complex world, removing the burden of management from DevOps teams by introducing comprehensive automation. Enterprises can gain fully automated Day-2 operations, including monitoring, backups, high availability, disaster recovery, migration, auto-scaling, and security. Portworx Data services removes the complexity of the modern application environment—enabling DevOps teams to focus on building new applications, not managing them.

Broad catalog of data services: The open nature of Kubernetes has fueled a dynamic and ever-changing ecosystem of technologies that extend and enhance cloud-native applications. Cloud Native Computing Foundation recognizes over 100 different data services including over 50 databases. With Portworx Data Services enterprises can leverage this vibrant ecosystem with the industry's broadest catalog of options including Cassandra, Couchbase, Elasticsearch, Kafka, Redis, RabbitMQ, MongoDB, and PostgreSQL.

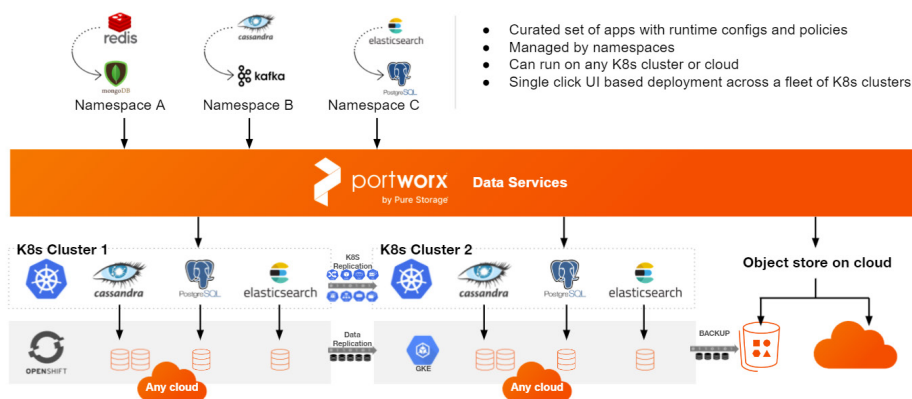


Figure 1. Portworx Data Services database management

Additional Resources

- Learn more about [Portworx](#).
- Explore what [Portworx Data Services](#) can do for your organization.

purestorage.com

800.379.PURE

