PURESTORAGE®

Reliably deploy and operate databases on Amazon EKS

Leverage Portworx© by Pure Storage and AWS to manage your mission-critical data services

As modern applications running on Kubernetes increasingly move from stateless to stateful, developers need to maintain access to key databases like MySQL, PostgreSQL, and MongoDB. A traditional approach to managing these databases relies on scripting and numerous repetitive and manual activities to ensure databases remain online and in support of mission-critical applications. These highly manual activities are prone to human error, excessive downtime, and can't be easily automated with existing infrastructure automation tools.

While Kubernetes provides some support for such automation, teams need an enterprise-grade Kubernetes data services platform to ensure the reliability and performance of these databases in production at scale.

Enter Portworx by Pure Storage.

Automated database management

Portworx solves these problems by providing your enterprise with a single data management layer that automates much of the process, helping modern apps get into production faster. Your administrators can leverage this consistent management framework, to manage the entire lifecycle of a database, from delivering self-service database access to developers to automating day-2 operations like migration and scaling:

- **Developer self-service:** Portworx makes it easy to offer Database as a service (DBaaS) instances to developers so they can access the services they need without being expert in each service and its related storage requirements. The Portworx Kubernetes Data Platform allows admins to centrally manage storage and resource quotas, centrally apply policies, and optimize the underlying storage, making it simple to provide developer self-service access while maintaining control.

- **Cluster migration:** Using PX-Migrate, enterprises can easily migrate data, application configuration, and Kubernetes objects across Kubernetes clusters seamlessly and with minimal effort. This functionality supports a number of key use cases with databases including: blue-green deployments, moving from DevTest to production with no manual overhead, and moving workloads from Amazon Elastic Kubernetes Service (Amazon EKS) to AWS Outposts and vice versa while maintaining data consistency.

- **Scalability:** PX-Autopilot allows enterprises to set volume and cluster resize policies that automatically integrate with your monitoring tools and dynamically scale your Portworx storage pool on Amazon Web Services (AWS) when it's needed. This means you are charged only for the cloud storage you consume.

> "**By using Portworx,** we're able to run those stateful applications similar to how we would run a stateless environment but still provide the resiliency and redundancy that we need."
>
> **-Rob Cameron,** SRE, Roblox

Comprehensive high availability, data protection and security for databases

For many DevOps teams, their number one concern is the availability and security of their database. Without reliability and high availability across different failure zones (pod, node, AZ, etc.) mission-critical container workloads are at risk of data loss or a data breach. As more mission-critical workloads move to container environments, Kubernetes alone is insufficient to address these concerns. Companies need to invest in a purpose-built platform, like Portworx, designed from the ground up to manage and protect critical data in modern apps running on Kubernetes.

Portworx by Pure Storage offers a robust set of data availability and security features for databases:

- Portworx provides enterprises with a single, application-aware data services layer that automatically spans AWS Availability Zones (AZ). This allows administrators to efficiently run and replicate database volumes across Kubernetes clusters spanning AZs to ensure high availability.
- Portworx also extends built-in database security policies with volume and container granular encryption and role-based access control (RBAC) protocols. This helps to ensure only authorized users have access to the data.

Portworx enhances these performance and security features with disaster recovery, backup, and restore functionality that has been optimized for AWS:

- PX-Backup provides enterprises with container-granular point-and-click backup and recovery to or from any Amazon Simple Storage Service (Amazon S3) endpoint globally. This allows administrators to ensure compliance and regulatory requirements are met while providing fast recovery to any Amazon EKS cluster.
- PX-DR automates DR for any database across Regions and AWS Availability Zones, providing enterprises with up to Zero RPO DR within metro areas as well as flexible DR solution across countries or the globe.
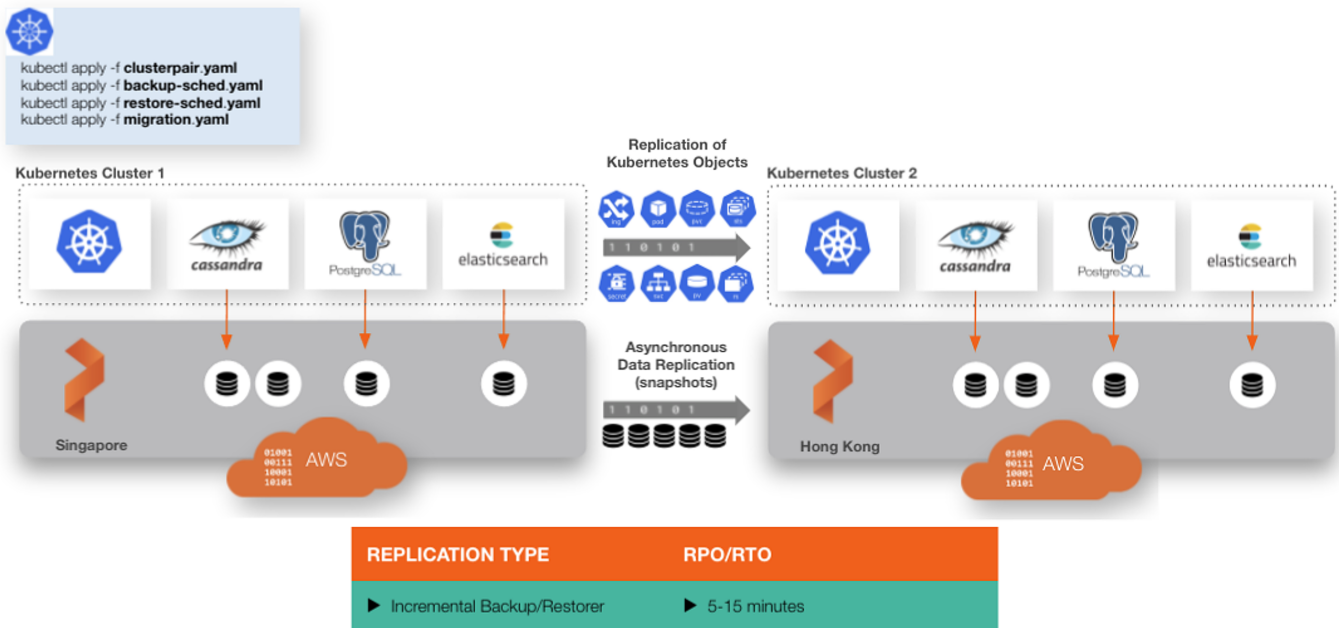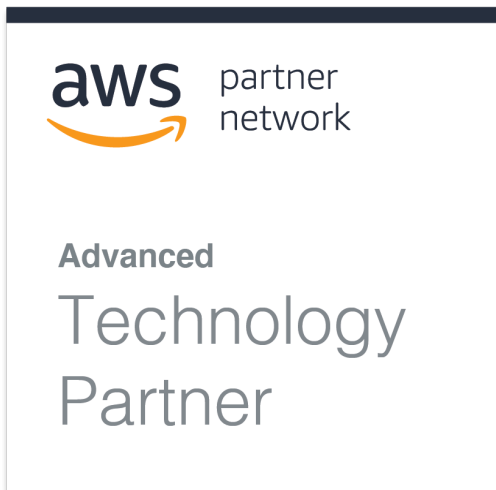
```
kubectl apply -f clusterpair.yaml
kubectl apply -f backup-sched.yaml
kubectl apply -f restore-sched.yaml
kubectl apply -f migration.yaml
```

Figure 1: Replication of Kubernetes objects

Learn more about Portworx

Additional resources

- **Running production databases on Amazon EKS**

- **How to run MySQL on Amazon EKS**

- **Run MongoDB, Cassandra, ELK Databases in Containers I Portworx**

Visit Portworx in AWS Marketplace

[insert publication number and date here]