TECH
mahindra

# Kubernetes Application Migration with Persistent Volumes Across Hyperscalers

WHITEPAPER

# Abstract

The rapid growth of cloud computing has resulted in a diverse landscape of hyperscalers offering a multitude of services and platforms. Kubernetes, being the de facto standard for container orchestration, has enabled organizations to deploy and manage their applications efficiently. However, as businesses expand and diversify their cloud infrastructure, the need to migrate Kubernetes applications with their persistent volumes across different hyperscalers arises. This whitepaper presents a comprehensive guide on how to achieve seamless application migration with persistent data.

This whitepaper begins by exploring the challenges and complexities associated with Kubernetes application migration across hyperscalers, particularly when dealing with persistent volumes.

This whitepaper is intended for Solution Architects who want to migrate their Kubernetes (k8s) applications across hyperscalers. The goal of this whitepaper is also to assist Application Architects in developing their own k8s application migration strategy to any Kubernetes cluster.

## Key Focus Areas of Kubernetes Applications migration

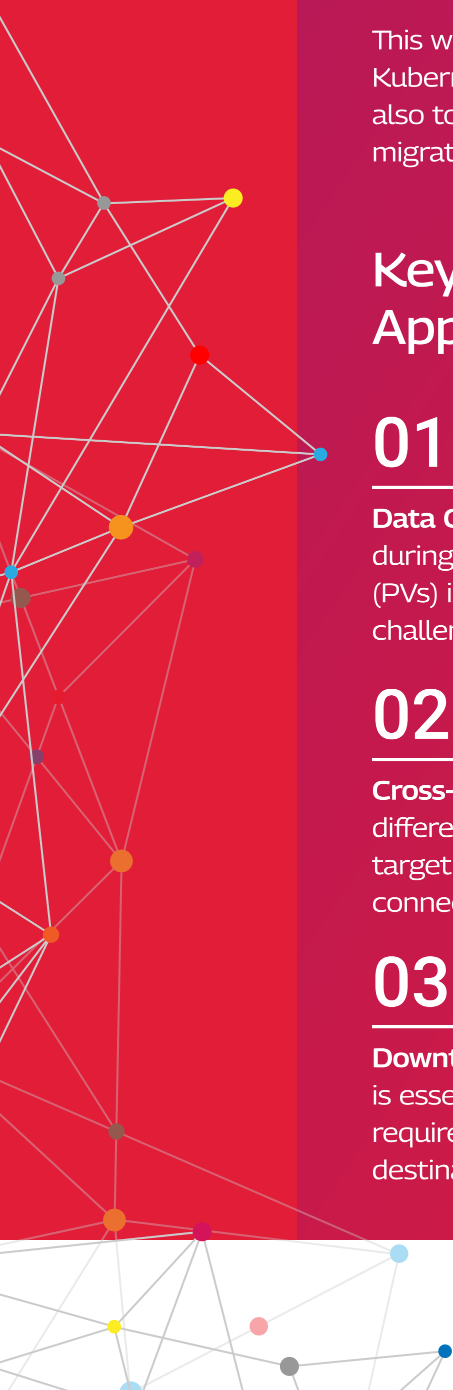

### 01

**Data Consistency and Integrity:** Maintaining data consistency and integrity during application migration is crucial. Ensuring that data in Persistent Volumes (PVs) is synchronized and intact across source and destination environments is a challenge that must be addressed.

### 02

**Cross-Hyperscaler Compatibility:** Migrating applications and data between different hyperscaler platforms requires compatibility between the source and target environments. Overcoming differences in storage configurations, network connectivity, and access controls can be complex and time-consuming.

### 03

**Downtime Minimization:** Minimizing application downtime during migration is essential for business continuity. Achieving zero or near-zero downtime requires careful planning and coordination between the source and destination environments.

## Introduction

As organizations embrace the scalability and flexibility of Kubernetes and leverage hyperscaler platforms for their infrastructure needs, the need for seamless application migration with data becomes paramount. This whitepaper explores the challenges and considerations involved in migrating Kubernetes applications with persistent volumes across hyperscaler environments using Portworx backup and highlights the industry landscape, discusses the key challenges faced during migration, and presents Portworx backup(px-backup) as a reliable solution for ensuring data integrity throughout the migration process.

## Industry Landscape

The adoption of Kubernetes for application deployment and management has witnessed rapid growth across industries. Kubernetes offers numerous benefits, including container orchestration, automatic scaling, and simplified application lifecycle management. Simultaneously, hyperscaler platforms such as Microsoft Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP) have emerged as leading infrastructure providers, offering scalable and on-demand resources for Kubernetes deployments.

However, migrating Kubernetes applications across hyperscalers presents unique challenges. The complexity arises when applications rely on persistent volumes (PVs) to store critical data. PVs provide a mechanism for Kubernetes pods to access and consume external storage resources. While Kubernetes supports migrating applications across clusters, moving the associated PVs seamlessly with the application presents significant hurdles. Organizations need a comprehensive migration strategy to ensure the integrity and availability of data during the migration process.

## Key Challenges

As containerization technology, led by platforms like Docker and Kubernetes, continues to revolutionize the way applications are deployed and managed, traditional backup methods face significant challenges in adapting to this new paradigm. Containerized apps are built to be flexible, dynamic, and scalable. They also are built of more than just data. They have application configurations and objects that need to be protected to be able to fully backup and restore quickly.

Traditional backups, which are machine-focused, simply aren't built to handle Kubernetes infrastructure. Traditional backup solutions simply capture the entire VM, which doesn't work with containerized applications.

To fully protect your Kubernetes data, you need a solution that backs up at the container level, and fully understands Kubernetes infrastructure. Traditional backup tools don't have Kubernetes namespace awareness and can't integrate with the Kubernetes API, leaving no alternative to taking a manual backup of every machine. This costs your team more time and can leave backups at risk for data loss and  human error. Beyond just costing more manual effort, traditional backup tools cannot fully backup an application. As stated earlier, you need to backup data, application configurations and objects to be able to restore fully and quickly.

Another big trend that we're seeing is that data protection is moving from IT to a shared responsibility among multiple teams, including application owners. Where traditional backup was centralized, container-aware backup provides application owners with self-service capabilities, including the ability to set their own backup policies and rules. This is huge, as it prevents IT dependencies and siloed responsibility.

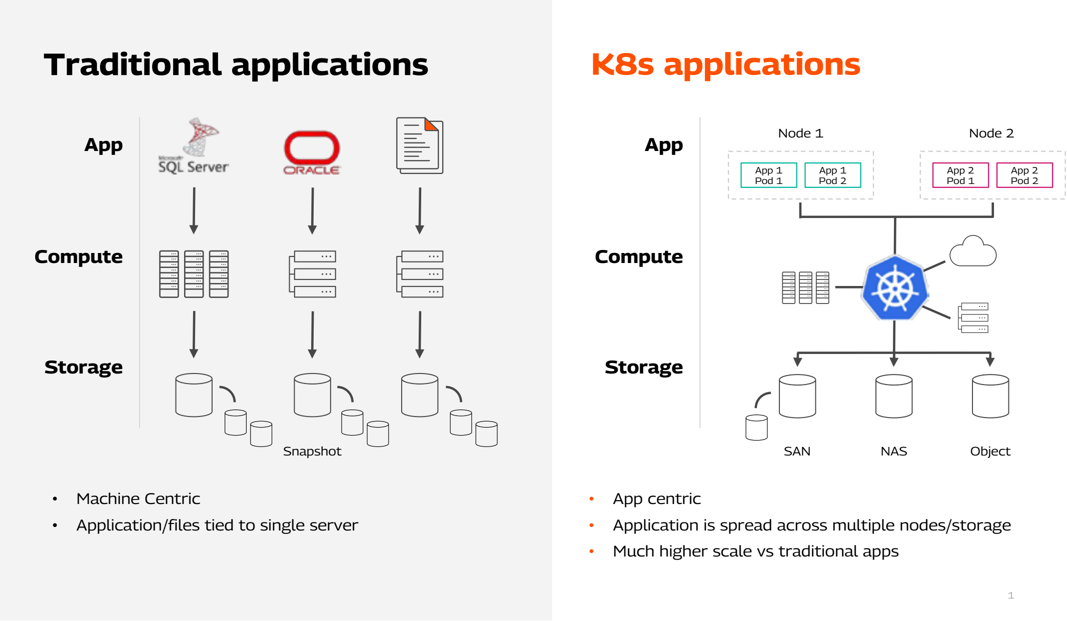

- Machine Centric
- Application/files tied to single server

- App centric
- Application is spread across multiple nodes/storage
- Much higher scale vs traditional apps

*Figure 1 Traditional Application and Kubernetes (K8s) applications storage*

Here, we see a further breakdown of traditional vs Kubernetes storage in Figure 1. Traditional applications, which run directly on physical or virtual machines, have simpler storage needs than distributed, containerized applications. Virtual machines (VMs) are more straightforward to backup, where an app is tied to a single server, and backing up the server is sufficient to fully protect the application. Containerized applications are different. Containerized applications are purpose-built to be highly dynamic, and traditional methods cannot support these modern architectures.

They're distributed by nature, backing up any given VM is likely to capture partial data from multiple applications, while failing to backup complete data from any single application makes targeted backup and restore very difficult and time-consuming with traditional backup. On top of that, backing up at the machine level, it could lump together applications that require different data protection policies.

To support fast recovery time objectives, backups must encompass the entire Kubernetes application. They also need to work at the container level rather than the VM level, and able to back up entire applications across VMs, including their configurations.

## Addressing the Challenges

To overcome the limitations of traditional backup methods for containers, new solutions have emerged that specifically cater to the containerized environment. These solutions are designed to address the unique characteristics and requirements of containers and container orchestration platforms. These solutions leverage container runtime APIs, Kubernetes-native tools, and container-aware backup agents to provide seamless and efficient backup and restore capabilities.

In this section, we will discuss about Portworx Backup (px-backup) and how this tool helps enterprises to migrate their Kubernetes workloads seamlessly.

Portworx Backup solutions adopt a container-centric approach, allowing granular backup and recovery of individual containers or groups of containers. They integrate with Kubernetes APIs to understand and leverage K8s concepts such as namespaces, labels, and configurations. By doing so, they can accurately capture the state of containerized applications and ensure consistent backups within the Kubernetes environment.

Furthermore, Portworx Backup solutions provide a user-centric backup management approach. They empower end users to define and manage their own backup policies, ensuring that one backup operation does not disrupt or compromise the backups of other applications. This decentralized approach enables greater flexibility and control over the backup process, aligning with the distributed and dynamic nature of containerized environments.

## Backup and Restore Kubernetes Apps Anywhere using Portworx Backup

Whether you use multi-cloud architecture or store your data across regions, PX-Backup enables easy migration for applications, including stateful apps, between any cloud, hybrid, or on-prem environment. To put it simply, you can back up your app in one environment and restore it in another. And all your Kubernetes application data, including any underlying infrastructure, is preserved.

Portworx Backup protects Kubernetes data anywhere. It helps enterprises to backup and restore their Kubernetes apps anywhere with a fast, easy, and secure platform.

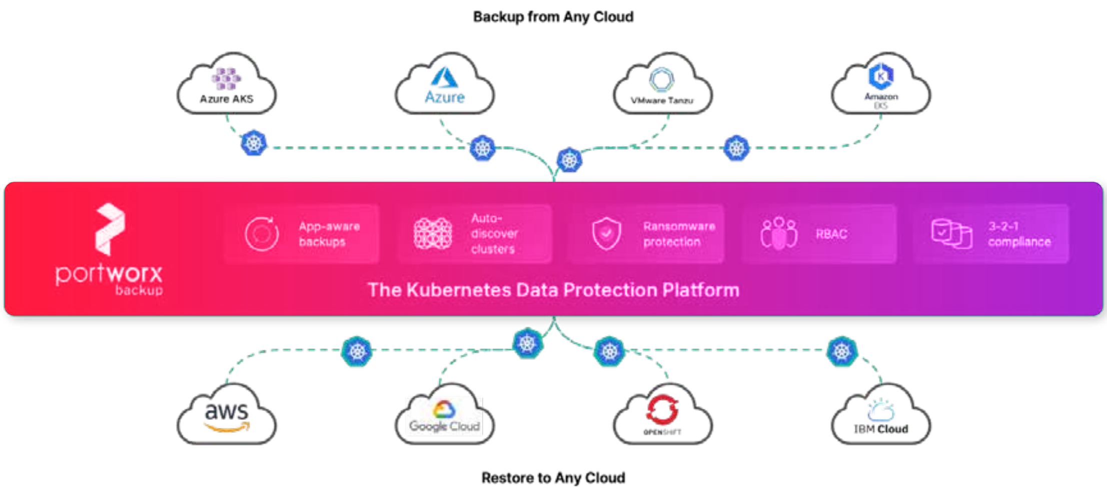

*Figure 2 Backup and Restore Kubernetes Apps Anywhere with a Fast, Easy, and Secure platform*

## Why Portworx Backup?

Portworx Backup (or px-backup) protects Kubernetes apps with the EASY button for Data Protection using container granularity.

## Self-service backup for app users in a few clicks

- Get started faster than ever on a fully managed data protection service with no install and easy cluster onboarding with auto-discovery
- Run apps in production with confidence with self-service management empowers app owners to centrally manage all backup and restore on their own
- Makes it easier to backup and protect data just the way app owners need
- Giving backup and restore policy setting to those who know best

## Recover and migrate applications anywhere

- Quickly backup and restore in a single click, and fully protect all associated application data, like configurations and objects with an app-aware, container-granular solution that was built to protect Kubernetes applications
- Backing up at the app level instead of the machine level allows for greater efficiencies–don't have to backup all these components separately in multiple machines, but also speeds up time to restore
- Easily migrate applications between clusters, clouds, and regions in minutes
- Any public cloud or on prem environment–provides a lot of flexibility in how and where to backup and restore

## Safe and secure data

- Guarantee protection against ransomware with S3 object lock and immutability
- Gain peace of mind with sophisticated RBAC controls and 3-2-1 compliance

## Ensure business continuity

- Portworx is able to protect your cloud-native workloads against failures or outages with near zero RTO and low RPO disaster recovery for mission critical data

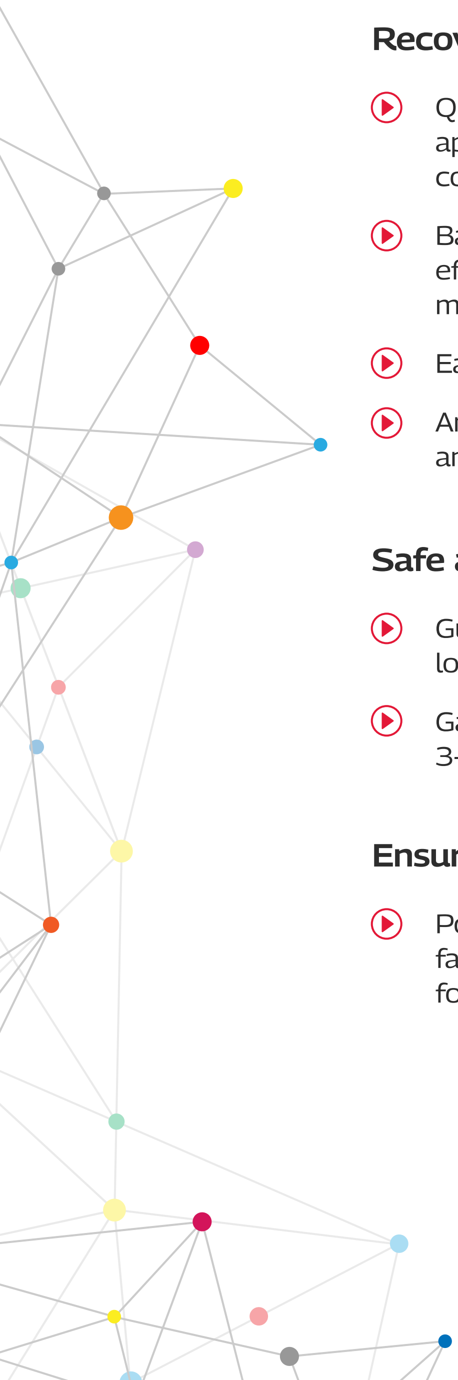

# App-granular Backup and Restore to Any Cloud

While traditional backup solutions lack application-aware protection for Kubernetes, Portworx Backup addresses this limitation effectively. With Portworx Backup, you gain access to single-click backup and restore capabilities that fully comprehend Kubernetes configurations, objects, and data. This ensures that Kubernetes backup and recovery become fast, comprehensive, and incredibly simple with just a single click.

Portworx Backup offers application-aware and container granular backup, eliminating the need to invest excessive time and effort in retrofitting machine-based solutions that are unsuitable for handling the scale and infrastructure of Kubernetes data. You can easily back up any application at the pod, tag, or namespace level, expediting the restoration process with complete app backups.

The following diagram Figure 3 demonstrates the cross-cloud mobility facilitated by Portworx Backup. It allows you to perform backups in an on-premises environment with distributed apps and select specific apps to back up in various cloud locations. For instance, you can restore one app in Azure while another in AWS. This data mobility empowers enterprises to enhance agility, meet compliance requirements, deliver superior service, and maintain better cost control.

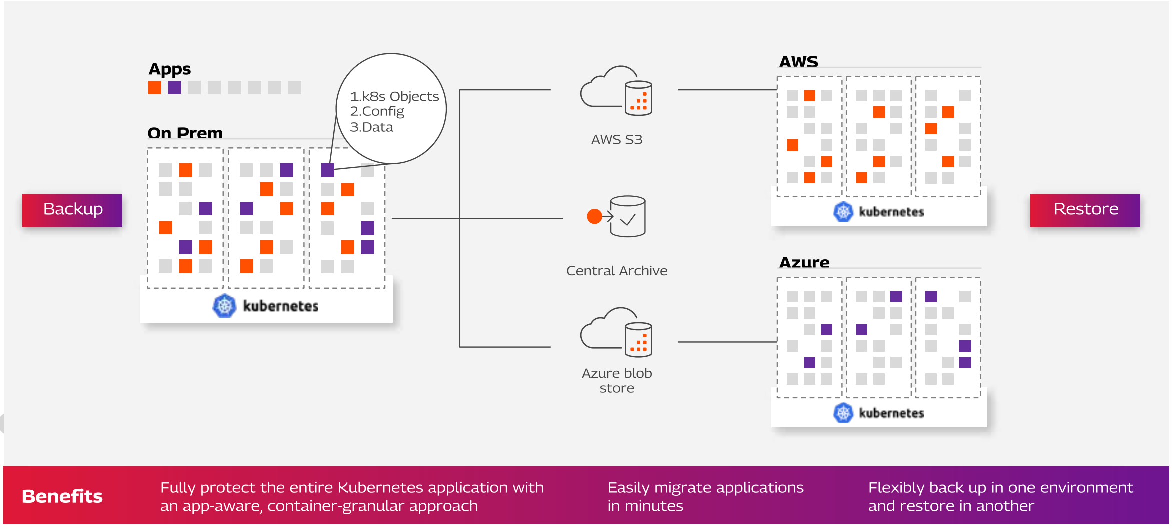

**Benefits**

Fully protect the entire Kubernetes application with an app-aware, container-granular approach

Easily migrate applications in minutes

Flexibly back up in one environment and restore in another

*Figure 3 Cross-cloud mobility with px-backup*

# Multi-tenancy and Enhanced RBAC with Portworx Backup (or px-backup)

Portworx backup provides a self-service, cloud-hosted control plane that enables organizations to manage their backup instances with ease.

It provides a multi-tenancy feature at all organizational levels, offering a comprehensive view of all backup instances.

| | |
|---|---|
| **Org Admin** | Setup Org<br>Add/manage users<br>Create/manage services |
| **Org User** | Create/manage services |
| **Service Admin** | Add/manage RBAC users<br>Create/manage cloud accounts<br>Create/manage backup locations |
| **App Admin** | Set schedules/rules<br>Use cloud accounts |
| **App User** | Backup/Restore their Apps |

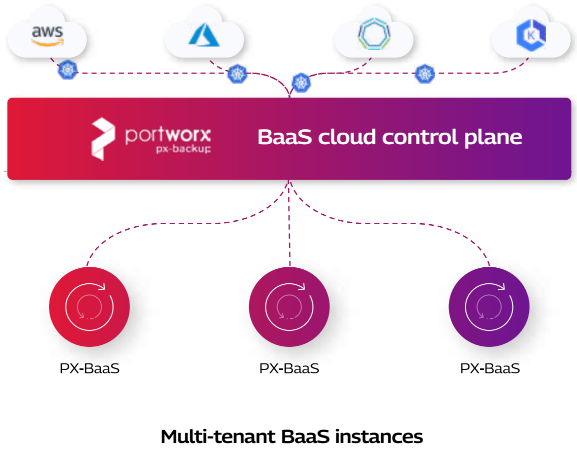

*Figure 4 Multi-tenancy at the org level offers single view for all backup instances*

# Conclusion

There is a rise in the adoption of Kubernetes and hyperscaler platforms for their application infrastructure. Hence the need for robust data protection and seamless application migration has becomes critical.  Portworx backup addresses the challenges in migrating Kubernetes applications with consistent volumes across hyperscalers, providing cross-platform compatibility, data protection, and zero-downtime migration capabilities. By leveraging Portworx backup, enterprises can migrate their critical applications while ensuring the availability and integrity of their data.

Traditional backup methods that were primarily focused on individual machines struggle to mitigate the unique challenges posed by containerized applications. Containers run across multiple machines that rely heavily on Kubernetes for orchestration and need a decentralized approach to backup. Specialized container backup solutions have emerged to effectively back up the containerized environments. These solutions are designed to be Kubernetes-aware, container-centric, and user-driven, enabling seamless backup and restore operations. As containers continue to gain prominence in modern application development, adapting backup strategies to accommodate their specific needs becomes paramount.

# Author

## Arunava Basu

*Solution Architect - Microsoft Cloud Business Unit*

Arunava Basu serves as the Practice Lead at Tech Mahindra's Microsoft Cloud Business Unit, overseeing the Application Infrastructure Modernization Tower, which encompasses Application & Infrastructure Modernization, DevOps & SRE, Automation, and Open Source technologies. With over 14 years of expertise, Arunava specializes in architecting cloud-native applications, migration and modernization roadmaps on Azure, and has taken on various roles such as Cloud and DevOps Architect and Infrastructure Architect. He excels in leading teams, mentoring members, and establishing strong client relationships, while also demonstrating a passion for automating processes. In his free time, he's a dedicated football enthusiast and a staunch Gooner (#COYG).

TECH
mahindra

Disclaimer. Brand names, logos and trademarks used herein remain the property of their respective owners.