Zerto
a Hewlett Packard Enterprise company

Zerto for Kubernetes Architecture Guide

Backup, Disaster Recovery, and Mobility for Kubernetes

Version 1

Abstract

The purpose of this document is to provide architectural guidance to anyone who is designing or implementing Zerto for Kubernetes production environment that is optionally paired with a secondary Kubernetes cluster for disaster recovery, with either site targeting a long-term retention repository for backups.

It includes a full overview of all design principles and architecture considerations for anyone looking to implement or anyone who has already deployed Zerto for Kubernetes.

The architectures in this guide show only a glimpse of what Zerto for Kubernetes can support. Having a true software-only, scale-out architecture and a platform- and storage-agnostic approach allows you to protect your Kubernetes environments with unrivaled granularity and industry-leading RPOs and RTOs and with application-centric data protection as code.

Zerto for Kubernetes Overview

With Zerto for Kubernetes, you can integrate continuous backup, disaster recovery, and mobility into the application life cycle. Zerto for Kubernetes easily protects any containerized application and its associated components running on Kubernetes as well as its persistent data. The foundation is built upon Zerto's industry-leading continuous data protection (CDP) technology that has been protecting 9,000+ customers' virtual machines for over a decade. CDP provides unrivaled recovery speeds and greater reduction of data loss and downtime, achieving RPOs of seconds and RTOs of minutes. Centralizing key metrics and information into Zerto Analytics allows organizations to have complete visibility across multi-site and multi-cloud environments, giving organizations confidence that SLAs and compliance needs are met.

Zerto for Kubernetes is currently available for:

- Azure Kubernetes Service
- AWS Elastic Kubernetes Service
- Google Kubernetes Engine
- IBM Cloud Kubernetes Service
- Red Hat OpenShift
- VMware Tanzu

Benefits of Continuous Data Protection vs. Snapshot Technology

Zerto is built on a proven CDP technology foundation that enables your organization to keep containers running 24/7 no matter what disruption or threat arises. CDP helps you say goodbye to periodic, snapshot-based technology, to deliver always-on replication for disaster recovery and backup with industry-leading RTOs and RPOs.

- Thousands of restore and recovery points, seconds apart, to recover cloud native applications without using snapshot copies
- Application-centric recovery for accelerated RTOs with quick and consistent recovery of complex applications across multiple Kubernetes deployments
- Instant journal-based recovery for a simple recovery experience that doesn't impact performance

Disaster Recovery

Organizations of all sizes that are running containerized applications are having to reconsider how they achieve disaster recovery in a Kubernetes environment. Legacy tools are available but rely on outdated technologies that do not meet the rigorous 24/7 operational demands of Kubernetes and containerized platforms. Legacy tools typically rely on snapshot technology which will take periodic points in time with lengthy RPO metrics and complex restore workflows. Alternatively, tools that protect the underlying storage are also an option but limit your platform choices dramatically and promote vendor lock-in. These solutions add significant overhead and complexity by needing additional tools, skillsets, and resources to manage, which does not align with the benefits of Kubernetes orchestration.

Zerto for Kubernetes is purpose-built for Kubernetes to be the simplest, most powerful disaster recovery solution for modern applications. By including all the replication, recovery orchestration, and automation in one simple software platform, users can recover all their Kubernetes objects from source to anywhere for resilience and portability.

Through native integration into all supported platforms, Zerto for Kubernetes not only allows replication and recovery between any persistent storage, but it also protects across and between multiple on-premises and public cloud Kubernetes environments. Zerto for Kubernetes delivers CDP at scale with deep integration into the Kubernetes stack for native protection of applications, offering data protection as code.

Backup

Backup is an essential part of the IT strategy for containerized applications, but with the rapidly changing landscape and increasing threats to data and applications, can we still rely on legacy backup technologies? The costs of downtime and data loss continue to rise. To avoid the impacts of disruptions and productivity and revenue loss, organizations require the ability to resume operations immediately after a disruption and completely minimize data loss. To do this, continuous data protection, simpler recovery workflows, and the ability to restore whole applications consistently is needed.

When we look at the legacy backup technologies currently protecting one of our most valuable assets—data—not much has changed over the last decades. The process remains the same: periodically take a copy of the data that has changed in our production environments and store it in another, secondary location. This introduces large gaps of missing data and complex architectures to perform a copy, move it, and store onto a repository.

To achieve far greater granularity and zero impact on containers and their underlying storage, modern backup strategy requires users to evolve from periodic or scheduled backup technologies to continuous backup. By using Zerto CDP, you can deliver RPOs of seconds by continuously replicating (block level) every change that is being generated in near real time. All these replicated changes are then stored in a journal which allows you to not only recover to the latest point in time, but also offers almost near-zero data loss. The outcome is the ability to safely rewind to any point in time, ranging from seconds to years.

Application Mobility

A key part of cloud native applications is the ability for them to be portable and to avoid traditional infrastructure problems such as vendor lock-in. Zerto for Kubernetes lets you migrate your entire containerized application with its persistent data to any Kubernetes platform of your choice, on-premises or in the cloud, with zero risk and zero data loss. This ability gives organizations flexibility and choice when it comes to the target Kubernetes environment and the underlying persistent storage platform. Zerto for Kubernetes integrates directly into the Kubernetes stack, ensuring there is no vendor lock-in and data protection is not reliant on the underlying storage technology.

Key Zerto Differentiators

- Always continuous data protection – journaling, replication, and granular restores
- Continuous backup, disaster recovery, and application mobility in one scalable, software-only platform – quick deployment for multiple protection and mobility use cases
- Technology, platform, and storage agnostic – allowing freedom of choice and avoiding legacy infrastructure issues such as vendor lock-in
- Data protection as code – purpose-built on Kubernetes for Kubernetes – build data protection into your app using the tools you already know
- Application-centric – can protect your entire containerized application, including persistent data, as a single entity with application consistency
- Simplicity at scale – limitless scale for protecting your cloud native workloads

Analytics & Monitoring

As data centers become more complex, you need more visibility and control of your protected IT assets across your private, public, and hybrid cloud environments. Zerto Analytics, included in the Zerto platform, is a secure SaaS-based offering that requires no further configuration and provides a single, comprehensive overview of your entire protected multi-site, multi-cloud environment no matter what infrastructure you are running, VMs or containers. Zerto Analytics delivers real-time and historical analysis of the health and protection status of your applications and data by using metrics such as average RPO, network performance, and storage consumption. Built-in intelligent dashboards enable you to spot trends, identify anomalies, and troubleshoot issues. Whether your data resides on-premises or in the cloud, you can confidently monitor the real-time health and protection status of applications and data. Zerto Analytics informs you to make better decisions to achieve an efficient, resilient mode of operation.

Core Components of the Zerto Platform

Zerto Kubernetes Manager (ZKM)

A containerized application that manages everything required for the replication between the protected and recovery clusters, apart from the actual replication of data. ZKM interacts with Zerto Kubernetes Manager Proxy (ZKM-PX) as its proxy to a Kubernetes cluster and virtual replication appliances (VRAs) to orchestrate replication. ZKM manages the entire environment and therefore it is required to have only one instance of it.

Zerto Kubernetes Manager Proxy (ZKM-PX)

A containerized application which serves as a communication proxy between the Kubernetes cluster and the VRA. Zerto requires one ZKM-PX installed per cluster.

Virtual Replication Appliance (VRA)

A DaemonSet which is automatically installed on each cluster node. VRAs manage the replication of data from the protected cluster to the recovery cluster.

Networking

When replicating between clusters, Zerto requires ingress to manage all cross-cluster communication. If replication is done within the same cluster, there is no need for this component.

Keycloak

Keycloak is an open-source identity and access management tool which is used for user and component authentication. It is deployed automatically as part of the ZKM installation, and only one instance is required.

Virtual Protection Group (VPG)

A key failing of many traditional approaches to backup and disaster recovery is an inability to protect applications that exist across multiple virtual components. These traditional solutions protect a single VM or a single storage volume and have no way to ensure application consistency between protected components. With cloud native applications, many components make up an application and can span across multiple volumes and deployments. Zerto takes an application-centric approach to protecting your cloud native application, utilizing data protection as code for your whole cloud native application. Your cloud native application and its persistent data can be protected as one single entity. With Zerto, your whole application is recovered to a single point in time, ensuring all components and the data are in sync.

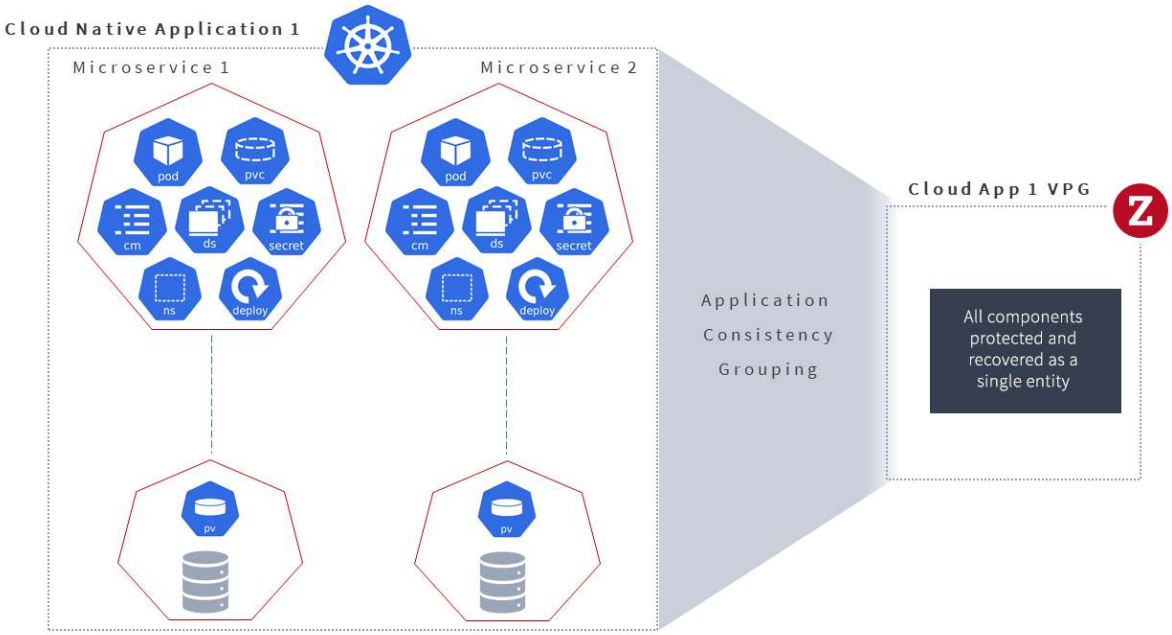

Zerto Journaling Technology

Zerto CDP stores all replicated data in a journal. The journal stores all changes for a user-defined period and allows you to recover to any point in time within the journal, ensuring your RPO is always as low as possible. Every write to a protected Kubernetes application is replicated by Zerto locally and/or remotely and written to a journal managed by a virtual replication appliance (VRA). In addition to the writes, every few seconds all journals within the VPG are updated with a checkpoint time stamp. Checkpoints are used to ensure write order fidelity and application crash consistency. Recovery can be performed to the last checkpoint or a user-selected checkpoint. This enables recovering individual applications or entire sites either to the previous crash-consistent point in time or, for example, when the container is attacked by a virus or ransomware, to a point in time before the attack.

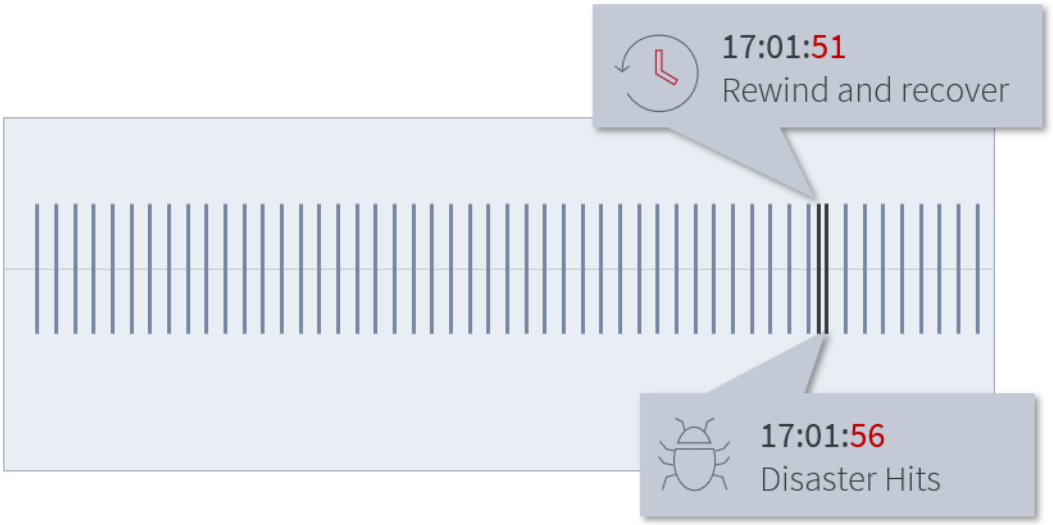

Long-Term Retention (LTR) Repositories

In addition to disaster recovery and local continuous backup scenarios, organizations that have compliance requirements need to retain data for long periods of time using cost-effective storage. Traditional methods of providing LTR protection have impacted performance in the production environment, often disrupting user experiences. Zerto LTR offloads point-in-time copies of data already protected in your existing CDP journal to secondary storage targets as often as you want without impacting production workloads. Zerto LTR allows you to store data from any point in time for days, weeks, months, or even years. LTR repositories can be stored on low-cost cloud storage, including Amazon S3 and Azure Blob Storage.

Reference Architectures

In this section, three example Zerto configurations highlight the different capabilities of Zerto. These are intended as guides only to help visualize the benefits that Zerto can provide your organization while also demonstrating the simplicity of the Zerto platform.

Use Cases

All three example configurations support the following use cases. Where there are unique differences, they will be highlighted under the relevant architectures.

Disaster Recovery

Any disruption on the production site, whether loss of power, network, or otherwise, can be quickly resolved within just a few minutes with recovery of all Kubernetes objects in a cluster to a point in time just seconds before the issue occurred. *Example: Recovery of your Kubernetes workloads and associated objects within minutes after a regional outage.*

Backup & Recovery

Rewind and recover your Kubernetes deployments to just before an issue occurred. This can be a single deployment or a group of deployments that are recovered back to the exact same point in time with application consistency. *Example: Recovery of your Kubernetes deployment after a corruption to your persistent volume storage.*

Ransomware Attacks

Recover after a ransomware attack, minimizing data loss to seconds and downtime to minutes. Kubernetes objects can be recovered into the same Kubernetes cluster using local replication or can be recovered into a remote cluster. *Example: Recovery of your containerized workloads and associated objects to a separate Kubernetes cluster after a ransomware attack, allowing you to recover quickly while also investigating the attack in the original environment if needed.*

Long-Term Retention

Allow application-consistent points in time of your deployments to be stored in long-term retention for backup use cases but also with unlimited retentions to allow for archiving use cases. *Example: Storing an entire Kubernetes cluster in Amazon S3 for seven years to comply with regulatory requirements.*

Application Mobility

Eliminate risk and minimize migration windows by utilizing Zerto for Kubernetes to migrate whole Kubernetes applications and their persistent data to any supported Kubernetes platform. Guaranteed zero data loss and migration windows measuring in just minutes, Zerto for Kubernetes gives unrivaled mobility to your containerized applications.

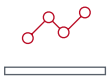

Analytics Across All Deployments

A single SaaS-based analytics platform providing complete data analysis across all your sites. This provides a single view that simplifies management and monitoring without added cost. *Example: Report the exact RPO of your virtual protection group and check how much storage each deployment is consuming.*

Architecture 1: Disaster Recovery & Mobility

The disaster recovery reference architecture depicted in Figure 1 below shows a set of proven practices for setting Zerto for Kubernetes where there are production Kubernetes clusters with replication set up to remote Kubernetes clusters.

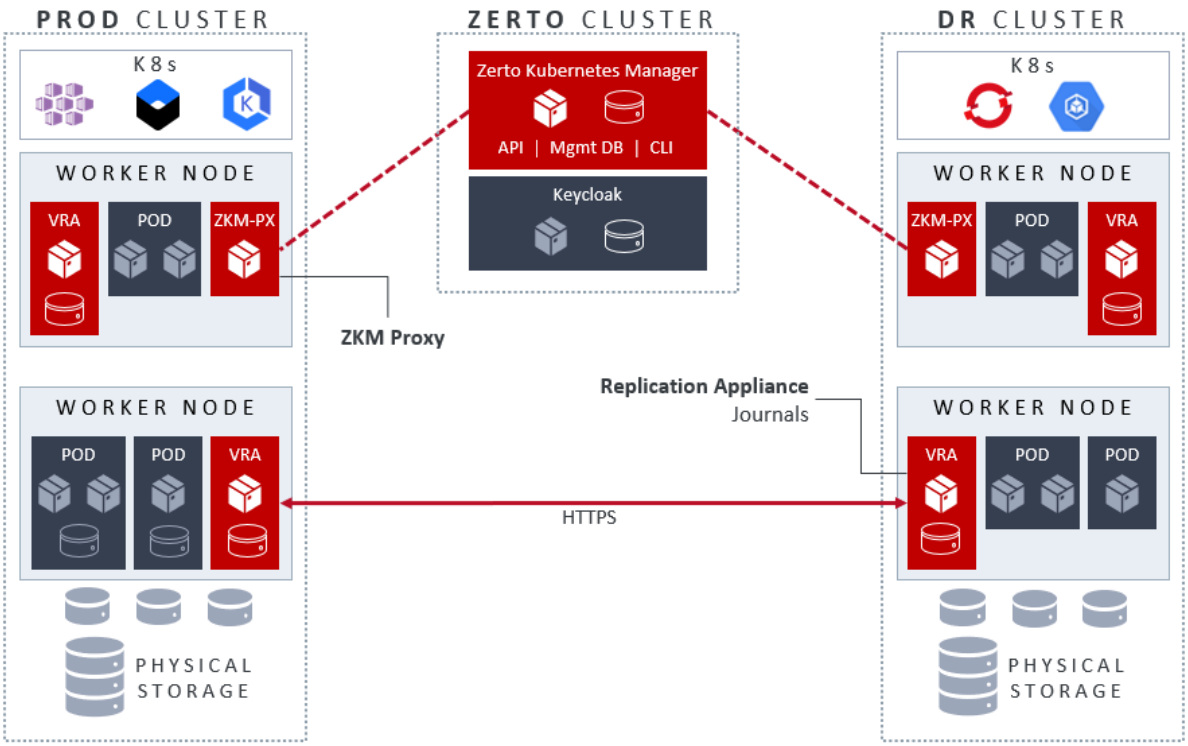

Figure 1

Description

Protected Kubernetes deployments are grouped in VPGs with application consistency across all the deployments. A remote journal is configured on the remote cluster and used for short-term recovery scenarios where recovery granularity of just seconds can be achieved. All the management components for Zerto for Kubernetes are homed in the target or isolated Kubernetes cluster so that any large-scale issue will not have an impact on Zerto for Kubernetes. The target clusters can be any supported Kubernetes environments so there is no need to stick with a single provider/vendor.

Architecture 2: Local Continuous Backup with Long-Term Retention

The reference architecture depicted in Figure 2 below shows a set of proven practices for setting up Zerto for Kubernetes with local replication and a long-term retention repository for backups and archiving.

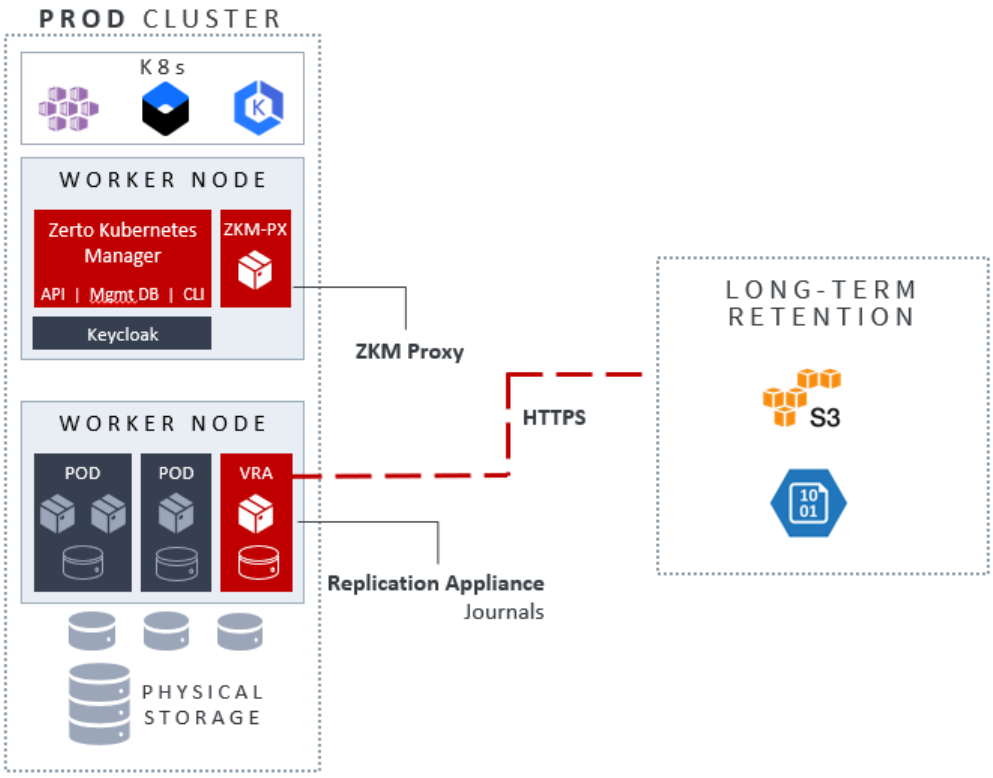

Figure 2

Description

In this architecture, the VPG is set up to replicate the Kubernetes deployments to the same Kubernetes cluster where they are currently running. This allows continuous backups to be stored in the journal without impact to any production components. This architecture provides the ability to recover quickly and easily with just seconds of data loss and granularity of recovery over a seven-day period. The long-term retention feature allows you to store Kubernetes deployment backups in a separate repository away for the Kubernetes ecosystem, allowing the application-centric backups to be stored for longer period of times with reduced costs for compliance purposes. Long-term retention is connected via HTTPS, ensuring encryption in flight. The Zerto for Kubernetes management components reside in the production Kubernetes cluster in the above architecture but could also reside in their own cluster if required.

Architecture 3: Local Continuous Backup, Disaster Recovery, and LTR

The reference architecture depicted in Figure 3 below shows a set of proven practices for setting up the Zerto platform just as the previous local continuous backup and disaster recovery reference architecture, but with the remote LTR target being the public cloud.

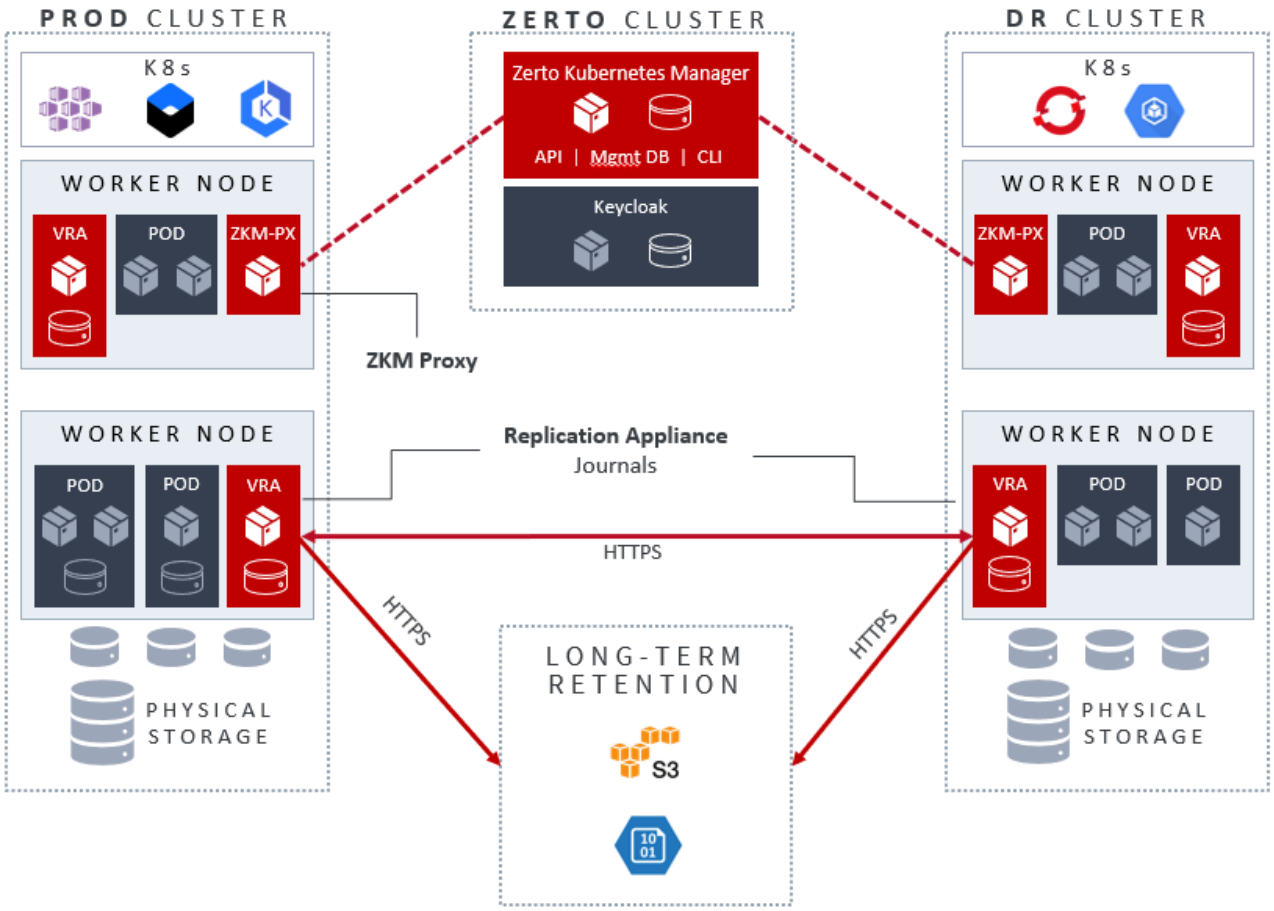

Figure 3

Description

In this configuration, the Zerto for Kubernetes management components are located on a separate Kubernetes cluster, neither housed in the production nor DR clusters. We have some VPGs set up for local continuous backup and others with remote DR replication enabled. This mesh architecture allows us to store local continuous backup VPGs and DR replica data VPGs in long-term retention, both transferred from the journals at their respective locations. This architecture allows users to backup Kubernetes deployments on a continuous basis and replicate different Kubernetes deployments to another cluster. With long-term retention enabled at both sites, we can store copies from both sites into cloud-based object storage away from Kubernetes infrastructure while maintaining application consistency between deployments.

Additional Resources

The resources below provide more detailed prerequisite, guideline, and sizing information.

Zerto – Quick Start: Zerto for Kubernetes Environments

Zerto – Zerto for Kubernetes Online Help

Please visit https://www.zerto.com/myzerto/technical-documentation for all Zerto's technical documentation.

About Zerto

Zerto, a Hewlett Packard Enterprise company, empowers customers to run an always-on business by simplifying the protection, recovery, and mobility of on-premises and cloud applications. Zerto's cloud data management and protection platform eliminates the risks and complexity of modernization and cloud adoption across private, public, and hybrid deployments. The simple, software-only platform uses continuous data protection at scale to converge disaster recovery, backup, and data mobility. Zerto is trusted by over 9,500 customers globally and is powering offerings for Microsoft Azure, IBM Cloud, AWS, Google Cloud, Oracle Cloud, and more than 350 managed service providers. **www.zerto.com**

IS-15910