TECHNICAL WHITE PAPER

Veeam Kasten *for Kubernetes*

Dispelling the Myths of Kubernetes Data Protection

By Alex Arcilla, Senior Validation Analyst
Enterprise Strategy Group

May 2024

Contents

Introduction

This Technical White Paper from TechTarget's Enterprise Strategy Group documents our evaluation of the Veeam Kasten Data Protection platform for backing up and protecting Kubernetes clusters. We reviewed how this platform is designed to better address the issues encountered when backing up and restoring containerized applications deployed and managed by Kubernetes as compared with traditional backup methods and procedures.

Background

Kubernetes appears to be maturing, as two-thirds (66%) of respondents to a survey conducted by Enterprise Strategy Group stated that they already use Kubernetes in their production environments.[1] This orchestration and management platform is also becoming an integral part of overall IT infrastructure, especially as organizations continue to deploy cloud-native applications. According to the same survey, nearly one in five respondents (19%) have 51 to 100 containerized applications managed by Kubernetes today, while over one in three (35%) have 21 to 50 applications.

Despite the growing preference for running business applications on containers managed with Kubernetes, less than half (48%) of respondents considered data protection during the architecture design process (see Figure 1). Instead, organizations are resorting to using traditional solutions for backing up and recovering containerized applications in Kubernetes environments or just considering what data protection technologies and processes to use after the fact. Unfortunately, these organizations do not consider the unique challenges they can face when backing up and protecting data in Kubernetes environments.

Figure 1. Considering Data Protection for Containerized Applications Was an Afterthought

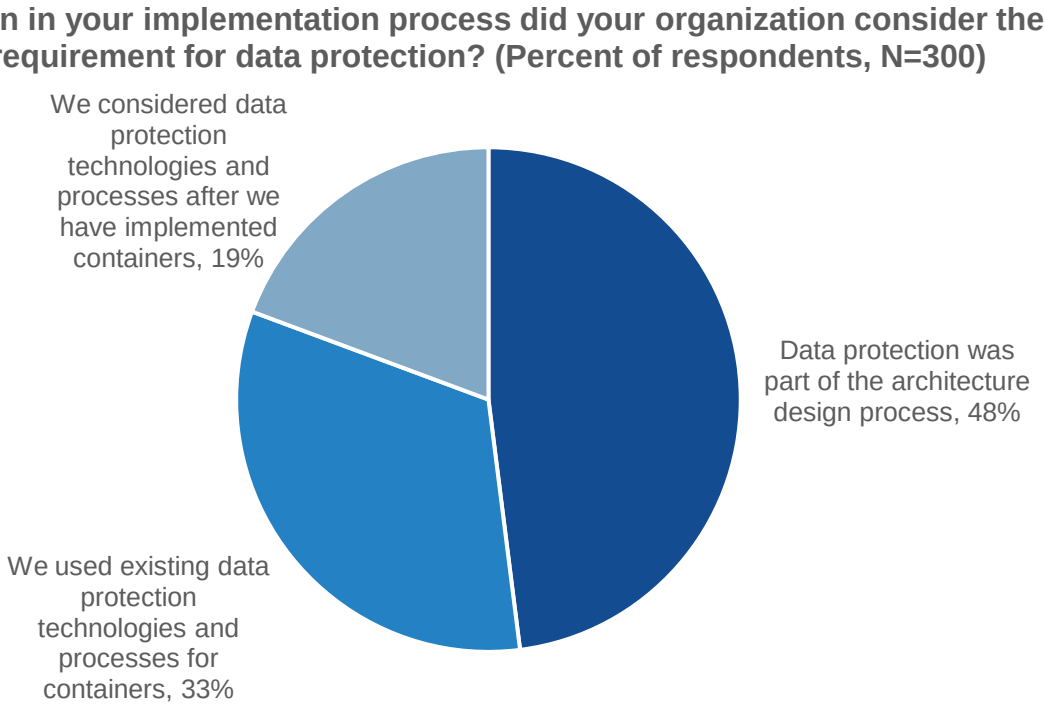

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

[1] Source: Enterprise Strategy Group eBook commissioned by Veeam, *Measuring the Current State and Momentum in the Enterprise Market for Kubernetes Protection*, April 2023.

Organizations that attempt to use existing backup and recovery (or "traditional") solutions and procedures within Kubernetes environments do not leverage the architecture of containerized applications or the automation and application orchestration that Kubernetes brings to the table. Using traditional solutions can inadvertently lead to excessive risk and long recovery times, which directly affect how fast businesses can achieve their short- and long-term objectives related to increasing revenue and maintaining reputation and brand equity.

When opting for an open source tool to develop backup and recovery solutions in-house, organizations face numerous drawbacks.

- The in-house Kubernetes team may lack the resources to effectively manage and update an open source data protection and management tool that aligns with new features and functionality released with services and solutions supporting Kubernetes applications. .
- Using open source tools requires writing numerous scripts tailored to the organization's IT infrastructure and relies on repetitive, yet error-prone, tasks.
- Open source code lacks the maturity required to operate at the scale required by enterprises and is inherently vulnerable to threats and attacks.

Veeam Kasten Data Protection Platform

Veeam Kasten is an easy-to-use, scalable, and secure system for backup and restore, disaster recovery, application mobility, and ransomware protection of Kubernetes applications. Veeam Kasten is purpose-built for Kubernetes, with an application-centric approach and deep integrations with relational and NoSQL databases and Kubernetes distributions, whether on premises or in the public cloud. This application leverages native Kubernetes APIs and includes features such as full-spectrum consistency, database integrations, automatic application discovery, multi-cloud mobility, and a powerful, intuitive, web-based user interface. As shown in Figure 2, Veeam Kasten provides strong support for data services, Kubernetes distributions, and a wide array of storage infrastructure. Veeam Kasten also offers several integrations with partners to help organizations implement active protection of Kubernetes environments.

Figure 2. Veeam Kasten Data Protection Platform

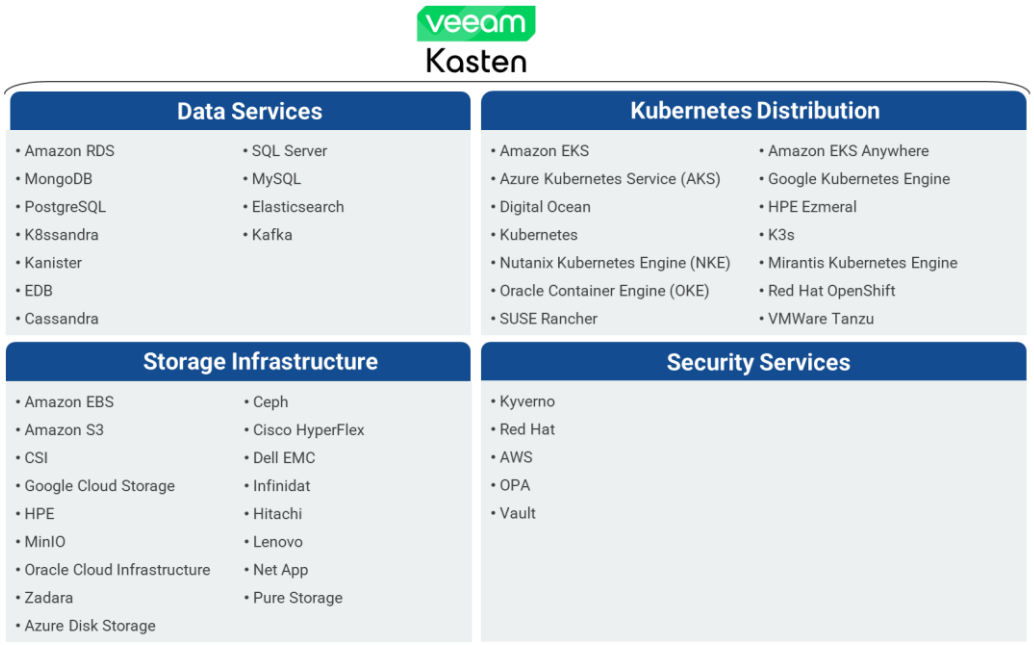

Source: Veeam and Enterprise Strategy Group, a division of TechTarget, Inc.

Veeam Kasten is constructed using cloud-native architectural principles. With this platform, organizations can automatically discover all applications running on Kubernetes clusters to facilitate proper backup and recovery. Veeam Kasten is designed to operate at enterprise scale, as this solution can execute multi-cluster management via a single, web-based user interface. Additionally, the management interface integrates monitoring, management, and alerts for all Veeam Kasten data management actions.

To decrease security risk, Veeam Kasten has been designed with numerous cyber-resilience capabilities, including encryption, role-based access control (RBAC), identity and access management roles (IAM), and ransomware protection via data immutability.

Enterprise Strategy Group Technical Evaluation

Enterprise Strategy Group evaluated how Veeam Kasten enables organizations to back up and recover the data associated with stateful and stateless containerized applications deployed and managed with Kubernetes. To illustrate this, we dispel common misunderstandings about the suitability of traditional methods and solutions for Kubernetes environments.

Myth: 'Using my existing backup and recovery solution is sufficient for a Kubernetes environment.'

Traditional solutions for business applications are organized and configured according to a three-tier architecture, in which application modules (e.g., client-side user interface, business logic, databases) are typically hosted within virtual machines (VMs) or public cloud instances, each with its own OS. Unless moved, the VM's location is static. Each server or instance is mapped to configured data stores, residing in either a local disk or external storage.

On the other hand, the code or microservices that comprise Kubernetes applications run on one or more containers and can be deployed in a virtual server environment, which is independent from the underlying physical infrastructure. For stateful applications, data is stored in persistent volumes (PVs), which are abstracted from the external storage infrastructure. The application and the data are essentially decoupled.

Figure 3. Differences Between VM-based and Containerized Applications

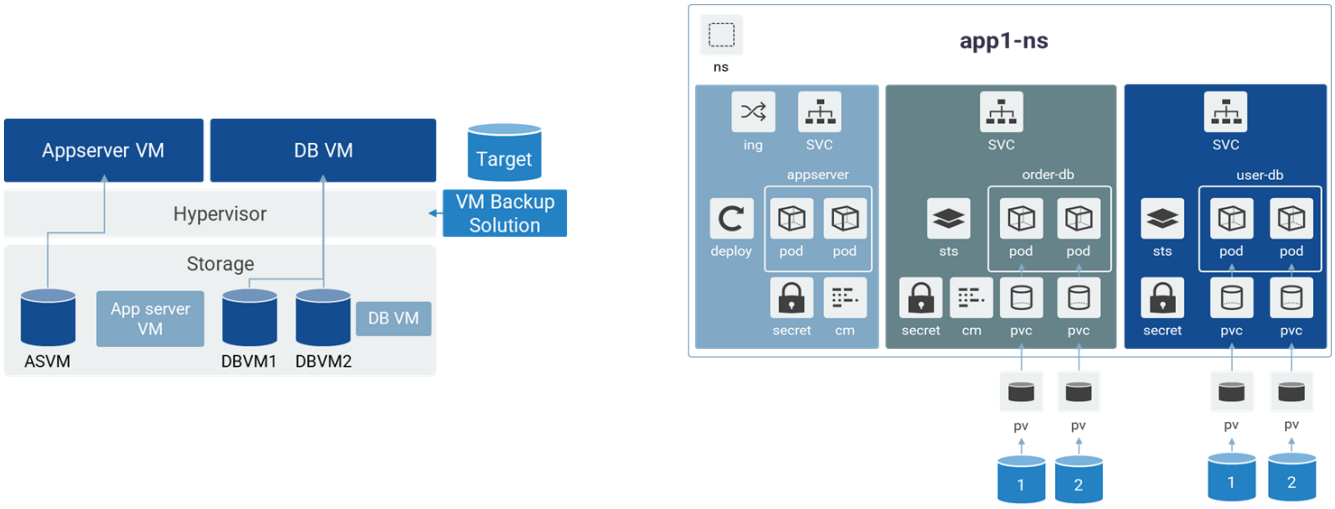

Source: Veeam and Enterprise Strategy Group, a division of TechTarget, Inc.

Given the architectural application differences, as illustrated in Figure 3, it is easy to see why traditional "VM-centric" solutions are insufficient for protecting the data of containerized applications.

- Containers of a single application are in a single or multiple Kubernetes clusters that can be distributed across multiple servers. While backing up applications based on the underlying physical infrastructure, a traditional solution can never be aware of the underlying architecture needed to protect and restore a Kubernetes application and its data store. This lack of visibility translates into unnecessary time and manual effort spent on protecting components and, subsequently, restoring applications.
- Traditional solutions will back up and restore at the VM and storage level. With Kubernetes applications, containers and the data stores are decoupled. For example, it is possible for the containers of a single application to rely on data from multiple database types (MySQL, Cassandra, MongoDB, etc.). When using traditional solutions, there is the risk of not backing up the right servers hosting the application's containers and inadvertently not protecting the right data.
- Because the containerized application and its data are decoupled, backing up data during an incremental production cycle cannot be done efficiently, as the containers and the PVs are not bound by a set of data interfaces.

Why Use Veeam Kasten?

Veeam Kasten has been designed to be "application-centric," offering capabilities such as:

- **Application-specific protection.** Veeam Kasten performs application-specific backups by automatically discovering all components of a Kubernetes application, along with any dependencies and configuration. The platform discovers these components whether they reside in a single or multiple clusters within single, hybrid, or multi-cloud environments.
- **Policy-based automation and monitoring.** Because Veeam Kasten automatically discovers all components of a Kubernetes application, organizations can automate backups according to policies implemented via a user-defined schedule. Organizations no longer need to schedule maintenance windows for backups. Moreover, backups are designed to be application-consistent so that data or transactions are not inadvertently lost and disrupt normal business operations.
- **Fast restoration.** Veeam Kasten enables multiple and granular recovery options using either local snapshots or external copies. More importantly, while traditional solutions restore entire VMs and volumes, Veeam Kasten can restore specific objects of a Kubernetes application (pods, PVs, configurations, etc.). For restoring multiple applications simultaneously, organizations can leverage blueprints to properly orchestrate the restoration of a single or multiple applications so that business operations are not disrupted.

Why This Matters

Relying on traditional server-based backup and recovery solutions for protecting data of containerized applications is not the right approach. Simply put, traditional applications do not properly account for the differences in application architectures.

Enterprise Strategy Group reviewed Veeam Kasten and concluded that the solution is more suited to data protection of containerized applications. Relying on the backup of virtual servers and associated storage to protect containerized applications is inefficient and resource-intensive, as containers can be distributed across multiple virtual servers. Veeam Kasten is designed to automatically discover and back up the containers of any given application, avoiding the need to backup entire virtual servers and their storage volumes.

Myth: 'Kubernetes supports high application availability; thus, my data is protected.'

High application availability and data protection are two different issues. Application availability must factor in several types of disruption, including failure of underlying IT infrastructure, natural disasters, accidental data deletion, and data corruption. In these cases, entire applications must be redeployed to other virtual servers or other clouds. In addition, disaster recovery options must be available so that organizations can achieve their stated recovery time objectives (RTOs) and recovery point objectives (RPOs).

Data protection ensures that any stateful application can be rolled back to a normal operational state from previously taken point-in-time snapshots. While the Kubernetes platform can help an application maintain availability, the platform will not protect the data related to the application and cannot prevent data loss, data corruption, or ransomware. Moreover, should a disruptive event occur, the platform does not ensure business continuity. The most recent backup must be used to restore applications and services to minimize the disruption to the expected user experience.

While applications are protected at the infrastructure level with traditional solutions (i.e., a VM-based application fails over to an isolated and geographically diverse location), these solutions do not account for the Kubernetes application architecture, which is independent of the underlying hardware. Traditional solutions cannot ensure Kubernetes application consistency, such that the application can return to an acceptable operational state. Without the appropriate disaster recovery workflow to account for application awareness and application consistency, meeting RTOs and RPOs can be difficult to achieve, risking loss of revenue and customer goodwill, as well as brand erosion.

Why Use Veeam Kasten?

Because Veeam Kasten is designed to be application-aware, application-consistent, and cloud-native, organizations can implement the required data protection that supports both high application availability and disaster recovery scenarios. Specifically, Veeam Kasten:

- **Maintains multiple restore-ready backups.** Instead of recovering at the infrastructure level, backups are performed at the component level of the Kubernetes application, enabling faster application restoration with minimal or no data loss.
- **Automates import and restoration of the application.** In Kubernetes environments, organizations can relaunch an application on an existing platform (e.g., containers are moved to other VMs), move it either on premises or to an alternative public cloud entirely, or deploy it across a hybrid cloud.
- **Automates disaster recovery workflows.** Veeam Kasten leverages application awareness to ensure that any disaster recovery workflows maintain application consistency. Only the components of the affected applications fail over to designated locations.

Why This Matters

While Kubernetes can restore applications to an operational state, the platform is not designed for data protection. Should events such as data loss, data corruption, or ransomware attacks occur, a data protection solution must be in place so that the application can return to an operational state with minimal business impact.

Enterprise Strategy Group reviewed Veeam Kasten and concluded that the solution is better suited for disaster recovery scenarios for Kubernetes applications. While Kubernetes can redeploy application components automatically should the underlying infrastructure fail, Kubernetes does not protect against data loss. With Veeam Kasten, organizations can use point-in-time snapshots to restore the application to an operational state, ensuring business continuity.

Myth: 'Kubernetes is impervious to ransomware.'

When protecting Kubernetes applications, organizations must consider the underlying architecture and existing security gaps. Unlike VM-based applications, Kubernetes applications consist of many components that are not isolated from a single virtual server. Because a Kubernetes application can be distributed across many virtual servers, more attack vectors exist. Other gaps overlooked include exposed API servers (due to the use of microservices), data mismanagement, inadequate or overly permissive access controls, running configurations on outdated versions of Kubernetes, or using older versions of backup solutions.

Why Use Veeam Kasten?

While Veeam Kasten is not designed to intercept a cyberattack, the platform offers capabilities to minimize any subsequent business impacts, specifically:

- **Early threat detection.** Organizations can view the security status of an entire deployed Kubernetes data protection platform to prevent any threat or attack from further infiltration.
- **Encrypted and immutable backups.** By maintaining secure and application-consistent backups, organizations can redeploy applications (whether within the same cloud or on other public clouds) in disaster recoveries.
- **Accelerated application recovery.** Once the attack has been identified and isolated, applications can be restored to an operational state using specified restore points to maintain business continuity with minimal data loss.

Why This Matters

No application is impervious to ransomware. Kubernetes applications are no exception. Unfortunately, the distributed and disaggregated nature of Kubernetes applications presents more security gaps than VM-based applications.

Enterprise Strategy Group reviewed Veeam Kasten and concluded that the solution, while not designed to intercept attacks, can help in minimizing any long-term damage that results in subsequent disruption of normal business operations, extended downtime, lost revenue, and brand erosion. Veeam Kasten can help organizations proactively monitor for attacks to isolate as early as possible, while enabling encrypted and immutable backups to facilitate service restorations either via redeployment or service restoration, whether onto another public cloud or different infrastructure platform.

Myth: 'I already know how to move applications with my traditional solution.'

Whether migrating applications between different public or private clouds or porting applications from an on-premises data center to the public cloud, application mobility is important. Application mobility can be leveraged include ongoing digital transformation efforts, workload balancing amongst multiple clouds to mitigate costs or risk, and cluster upgrades.

However, organizations cannot afford to tolerate extended downtime while moving applications. Using traditional solutions cannot adequately support application mobility requirements. Migrating services to another platform requires restoring storage volumes so that the data remains intact. Yet, restoring volumes requires establishing new data interfaces between the affected VMs and the underlying storage. Until then, the application is unavailable.

The underlying architecture of Kubernetes applications enables application mobility, as all application elements are independent of the underlying infrastructure. The application can be migrated with minimal code changes and IT effort.

Why Use Veeam Kasten?

With Veeam Kasten, organizations can easily migrate applications while end users experience minimal to no downtime. To accomplish this, Veeam Kasten supports:

- **Automatic import and restore.** By generating an application export, organizations can easily import changes to the application while accounting for any difference in the environment (e.g., new infrastructure or new public cloud) with minimal manual effort spent on recoding or manual configuration.
- **Environment isolation.** Organizations can isolate application migrations without adversely affecting other applications running on the same infrastructure.
- **Encrypted in-flight and at-rest data.** To protect data during application migration and service restoration, data must be encrypted or risk infiltration by bad actors, especially those launching ransomware attacks.

Why This Matters

Migrating Kubernetes applications with traditional solutions will result in unnecessary downtime and waste of IT effort. Entire VMs need to be moved and storage volumes need to be fully restored before the application is fully operational.

Enterprise Strategy Group reviewed Veeam Kasten and concluded that the solution can support Kubernetes application mobility because the application and the data are not tied to the underlying infrastructure. Organizations can move applications to other clusters or public clouds without incurring unnecessary downtime.

Conclusion

Using traditional backup and recovery solutions on Kubernetes applications does not provide adequate data protection. These solutions were built with VM-based applications in mind. However, the architecture of these applications is completely different from that of Kubernetes applications. The distributed and disaggregated nature of the latter, which are rooted in a container architecture, raises issues that cannot adequately be addressed by traditional solutions.

Veeam Kasten is designed to address the unique challenges that Kubernetes applications present. By automatically discovering all architectural components of a Kubernetes application, Veeam Kasten can protect the application and its data by creating application-aware and application-consistent backups, then using those backups to restore the applications with the desired level of granularity. Organizations no longer need to consider traditional solutions for data protection, specifically because these solutions are focused on protecting and restoring applications at the infrastructure level. Since Kubernetes applications are decoupled from the underlying infrastructure, service restorations can occur in a more efficient and effective manner.

Throughout our evaluation, Enterprise Strategy Group determined that Veeam Kasten can deliver the data protection and recovery capabilities required by Kubernetes applications. We examined this by comparing how traditional solutions and Veeam Kasten address four major use cases: backup and restore, disaster recovery, application mobility, and ransomware protection. Veeam Kasten is better suited to these use cases when considering stateful and containerized applications running on Kubernetes, as the platform is application-aware and application-consistent. Since traditional solutions are VM-centric, the data protection offered to Kubernetes applications can be ineffective, inefficient, and incomplete.

When considering data protection solutions for Kubernetes applications, numerous myths exist as to why a traditional solution will suffice. Enterprise Strategy Group encourages you to challenge those myths and examine why Veeam can better serve your needs.

About Enterprise Strategy Group
TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com