Commvault®

eBOOK

Save your apps with Kubernetes backup

5 key data protection considerations for DevOps

KUBERNETES DATA PROTECTION

A growing need

While Kubernetes (K8s) adoption is on the rise, more than 38% of companies report they do not proactively consider data protection for their container-based application environments[1]. However, given that stateful applications are rapidly migrating to containers and that containers are poised to overtake VMs and bare metal servers as the production platforms of choice in the next 24 months[1], ITOps and DevOps pros are increasingly recognizing the importance of container data protection – or experiencing the pain of what happens without it.

Bottom line: if your business is among those rapidly adopting Kubernetes, but you're not yet backing up your entire container ecosystem, you need a plan to integrate data protection into your existing DevOps workflows—fast.

This eBook reveals why your business needs data protection for container-based data, how high-availability alone won't cut it, and 5 key considerations as you determine your strategy for container data backup.

We get it—containers are different than anything you manage today—and you may have containers spread across on-prem and multi-cloud. You may be facing the challenge of constrained resources when it comes to managing data backups, while upskilling your existing teams in containers. You may also be one of many who trust to high availability of key data stores to ensure access to your data and apps. Unfortunately, that's not enough to protect the underlying data that's powering the cluster and application—and the cost of doing nothing is far too high.

PROTECTING CONTAINERS AND THE FUTURE OF YOUR DATA MANAGEMENT

The rise in the popularity of microservices, containers, and K8s points to developers' and users' coalescence around a primary platform for modern cloud app deployments. Enterprises are increasingly using K8s in production deployments signaling a strategic commitment to the platform.

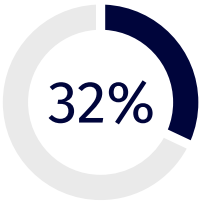

of companies have containerized more than half of their applications[2]

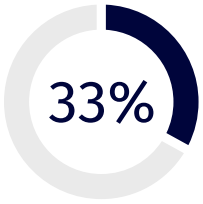

of companies are running more than half of their containerized workloads in production environments[2]

Containers were initially designed to be stateless, meaning that data was not stored (or persisted) in the container after it stopped running. Stateless containers allow for applications to be quickly scaled for their specific task. This enabled DevOps teams to design and build web-scale applications that could adapt at the speed of cloud growth.

Once DevOps engineers had the ability to create their own containers, the migration of stateful applications into containers began, and so did the protection problem. Stateful applications in Kubernetes require proper storage and data management throughout the entire application lifecycle. And while containers are multi-cloud by nature, it can be challenging for the IT and DevOps teams to manage and migrate container persistent data in a cohesive manner.

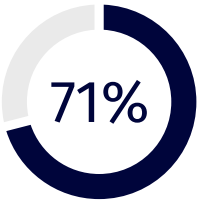

of respondents said container-based applications will be deployed in a combination of public cloud and private environments[3]

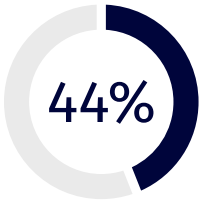

of respondents also noted managing backup and recovery of container-based applications in a hybrid cloud environment as their biggest data protection challenge[4]

Businesses urgently need a consistent way to easily migrate/replicate and protect data in their container ecosystem to both build and recover container-based applications across their hybrid cloud environments.

AVOID THIS COMMON MISCONCEPTION

High availability isn't backup. Here's why you should differentiate between the two

High availability vs. backup
Kubernetes supports setting up high availability clusters along with supporting components such as etcd, which is a consistent and highly available key value store for cluster configuration state. High availability enables Kubernetes to be operational during unplanned infrastructure failures, thereby minimizing application downtime.

However, as we know from years of protecting applications within VMs—high availability alone is not enough and there is a critical need for backing up the Kubernetes application and its associated data. Additionally, in a Kubernetes landscape, we must also consider container image registries, and the cluster state, etc. While high availability can indeed restart the containers during unplanned outages, the application cannot recover and be fully operational if the underlying data that's powering the cluster and application is corrupted or lost. In fact, if the outage spans an entire site you will need to recover to an alternate cluster, requiring all your container application configuration and data.

That's exactly why it's critical to back up all the application and cluster state—to ensure full recovery and to rapidly bring back operational applications and minimize disruption to your business.

Hybrid deployments remain the most common approach with:

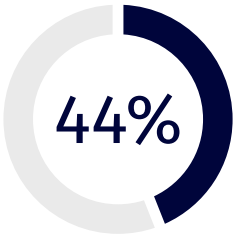

of companies deploying containers in hybrid cloud mode[2]

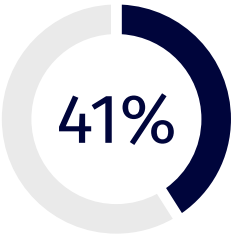

of enterprises adopted cloud-only deployment strategy[2]

Kubernetes is the most adopted and de-facto container orchestration platform.

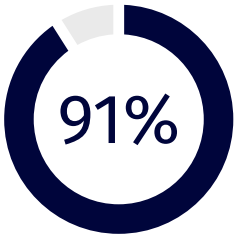

of organizations use Kubernetes to orchestrate containers[2]

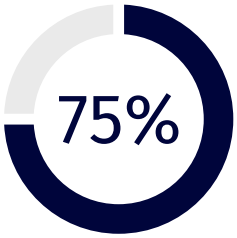

use it in production environments[2]

WHY PROTECT KUBERNETES APPLICATIONS?

With all this in context, let's look at some of the main drivers and scenarios that make Kubernetes backup essential.

Cybersecurity and ransomware attacks are on the rise and can cripple organizations by corrupting business data or leaking sensitive data and credentials. Backing up all the application data and periodically testing for recovery is critical to recovering and protecting against these types of attacks.

Accidental data loss can be accelerated with automation and DevOps propagation of changes, including misconfigurations. Kubernetes provides application development and deployment agility that is transforming organizations. However, any unintentional errors could propagate quickly and expose organizations to risks. Backups provide the ability to roll back to known good states.

Environmental failures can also lead to disaster scenarios and most organizations can recover from these events with a sound data protection and backup strategy in place.

Compliance and regulatory needs are often overlooked and are also a key driver for backing up Kubernetes applications. However, maintaining compliance can be a full-time job that some companies don't have the in-house resources to manage.

Application mobility and portability is an important criterion for adopting Kubernetes containers. Backups and snapshots facilitate and make it easy to port and move applications across clouds.

Keep in mind that:

- Human error contributed to misconfiguration of container workloads **67%**[2] of the time
- **44%** of companies have delayed application deployment into production due to containers or Kubernetes security concerns[2]
- **91%** of surveyed enterprises experienced a security incident in their Kubernetes and container environment[2]

With the recognition that a backup solution is critical with Kubernetes, there are important considerations to keep in mind, including why a cloud-based Backup as a Service (BaaS) solution should be an important part of your strategy.

5 MUST-HAVES FOR KUBERNETES BACKUP SOLUTIONS

As the data in containers becomes increasingly critical to DevOps and ITOps functions—and to enterprise business continuity

There are 5 simple keys to successful data backup and recovery adoption that can save your data (and save you headaches down the road).

1 Seamless hybrid cloud operations

Choose a solution that will run across cloud and on-prem environments through a SaaS delivery model, while also providing storage flexibility. This will help with data migration and application mobility. By allowing for data recovery to occur both through on-premises or cloud locations, your business can enjoy the agility to adopt the Kubernetes distribution and operational model that works for your developers.

2 Ease of operation

Go easy on yourself and your developers, allowing them to work in the most efficient and effective manner. Automatic application discovery means hands-off protection; and Container Storage Interface (CSI) integration means backup is a snap. The best solution is one that's easy to adopt with seamless integration and cohesive management features.

3 Comprehensive protection

Containers are not an island and must have comprehensive oversight however they're being managed. The optimal protection solution will cover your current network infrastructure while being able to extend to Kubernetes without any extra effort required by your IT team.

4 Security prioritization

Any proper solution must incorporate layered security, airgapped defenses, anomaly detection, encryption in flight, at rest and much more. Don't sell your business short when it comes to keeping bad actors and virtual attacks at bay.

5 Central management

An optimal SaaS solution should provide simplified management that reduces overall complexity, not require any infrastructure management, fill in existing skills gaps and be built with a cloud-forward perspective.

OVERCOME THE RISKS

With the proper solution in place, you can be positioned to avoid many of the significant risks and threat of downtime that can come with adopting Kubernetes without a container data protection strategy in place. Your applications can benefit from consistent protection and simplified migration of Kubernetes data, all while maintaining native API-based integration with the container ecosystem.

At the same time, your DevOps and IT teams should not be required to handle a big learning curve when it comes to adopting a backup solution or easily extending data protection to Kubernetes as just another workload. You need to be sure you are working with trusted leaders to have the best data backup and recovery plan in place.

... and what's the one place you can find all this?

THE INDUSTRY-LEADING SOLUTION

Commvault® Cloud for Backup & Recovery for VM & Kubernetes is a cloud-native data protection solution that can help with your Kubernetes protection needs while meeting every single one of the requirements and considerations we've covered up to now. This includes comprehensive protection, a central management interface, cloud-native development, application mobility and deployment flexibility. Because it's cyber resilient, this means reduced management overhead and lower TCO (total cost of ownership), with no backup infrastructure to manage, automatic updates, and a simple subscription model. DevOps teams can stop worrying about data loss and focus on driving application modernization initiatives into production.

When you want a simple solution that provides ultimate flexibility with reliable pricing, Commvault is the only choice.

> Commvault is demonstrating a strong commitment to Kubernetes and its operational models while continuing its robust support for traditional platforms and operations. Commvault can provide data protection for Kubernetes across a broad range of organizations. At the end of the day, Commvault is a market leader for enterprise data protection."
>
> **Enrico Signoretti**
> GigaOm Radar Report

Reach out to a Commvault representative today to discover how our solutions can meet your hybrid cloud needs wherever you're at in the journey.

1 ESG Container Data Protection Nov 2020
2 StackRox Fall 2020. "State of Container and Kubernetes Security."
3 ESG Research, 2020. "ESG Master Survey Results, Data Protection Considerations for Containers"
4 Ibid

To learn more, visit **commvault.com**