SPOTLIGHT

Sponsored by: Intel

Organizations must be able to guarantee zero RPO and low RTO for their Kubernetes environments before they can run mission-critical applications in them.

Mission-Critical Applications in Kubernetes Need Performant Disaster Recovery

December 2023

Written by: Johnny Yu, Research Manager, Storage and Computing

Introduction

One of the greatest difficulties of running mission-critical applications in Kubernetes is ensuring they are highly available and resilient against outages.

Companies lean on strong business continuity and disaster recovery (BCDR) technology and practices to ensure their critical applications are always available. For applications running on premises, in virtual machines (VMs), or in clouds, it isn't difficult to create a stable, reliable foundation for building applications. Most IT organizations know how to provide BCDR for applications in these "traditional" environments, and there are plenty of enterprise BCDR offerings available on the market.

Cloud-native container and VM environments running on Kubernetes work differently. Generally, in traditional environments, applications are linked to the storage they live on, so protecting the storage is as good as protecting the application. For Kubernetes, any given storage environment could be home to clusters containing the resources and components for multiple different applications. Businesses adopting containers and VMs on Kubernetes can't simply use the same BCDR tools they use for traditional environments because of this fundamental difference in application architecture. To ensure availability and resiliency, a solution must be able to recover at the application level.

An IDC survey illustrated this difficulty succinctly (see Figure 1). The survey asked organizations what they found to be the most difficult aspects of refactoring legacy applications for cloud-native environments. Adapting adjacent processes (such as data protection and data security) to work with the newly containerized applications was the most common response, beating out factors such as cost, complexity, and time spent.

AT A GLANCE

KEY TAKEAWAY

To run containerized applications in any serious capacity, organizations must first build a performant container BCDR system. This involves container orchestration software that is tightly integrated with the backup software responsible for preserving the state of Kubernetes clusters and the applications and data within. In addition, the hardware powering the system must be fast enough to keep recovery time objectives low and recovery point objectives at zero.

FIGURE 1: *Adapting Existing Processes to Support the Newly Containerized Enterprise Apps Proved Challenging*

Q *What is the most challenging aspect of refactoring legacy applications for cloud-native environments?*

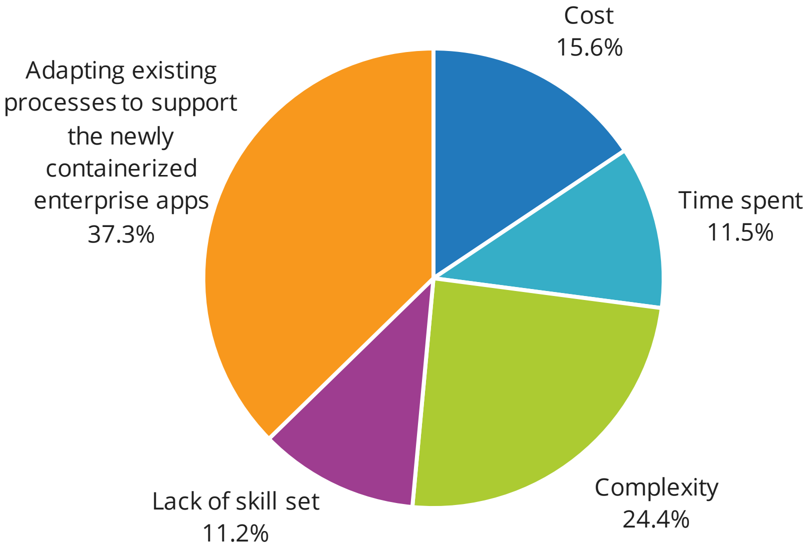

n = 410

Source: IDC's IT Infrastructure for Storage and Data Management Survey, October 2021

Since traditional BCDR solutions don't always translate well to Kubernetes, companies will have trouble protecting VMs and container applications running on Kubernetes to the same extent as applications running in traditional environments. This creates a roadblock for running mission-critical applications in Kubernetes, since those are applications for which there should be no compromising on reliable, performant BCDR.

To run these applications in any serious capacity, organizations must first build a performant container BCDR system. This involves container orchestration software that is tightly integrated with the replication software responsible for preserving the state of Kubernetes clusters and the applications and data within. In addition, the hardware powering the system must be fast enough to keep recovery time objectives (RTOs) low and recovery point objectives (RPOs) at zero.

Hallmarks of Performant Container BCDR

BCDR for mission-critical applications must possess the following:

» **Zero RPO and low RTO.** These are fundamental to BCDR for critical applications, regardless of the environment they are running in. Since these applications are the lifeblood of the business, they must not lose any data and must be available within minutes after an outage.

- **Flexibility in BCDR policies.** A BCDR solution shouldn't solely focus on the most critical applications. The ability to use asynchronous replication or snapshots for non-mission-critical applications allows for greater coverage, enabling organizations to create a spectrum of availability options and map their applications accordingly.
- **Integration between orchestration software and replication software.** For critical applications, software must not be a bottleneck for BCDR. Integration with container orchestrators allows replication software to quickly capture not just data but also application artifacts such as configurations and secrets to rebuild everything in a secondary environment.
- **Integration between software and hardware.** A performant BCDR solution will consist of an ecosystem of products that are tightly integrated and certified to work well together. Hardware-software vendor partnerships in this space usually involve co-development and joint support.
- **Automation.** The bulk of performant BCDR is preparation, such as setting different recovery policies for applications depending on their criticality, discovering workloads that need BCDR policies applied, and running disaster recovery (DR) tests. Automation not only minimizes the administrative burden for this important work but also can speed up recovery by automating tasks such as following runbooks and bringing applications and microservices up in the correct order.

Benefits

Organizations that can build a high-performance BCDR system for containers and VMs on Kubernetes stand to gain the following benefits:

- **Run critical applications in Kubernetes.** This is the most important outcome. Organizations can't run mission-critical applications in containers and VMs on Kubernetes if they don't have a way to ensure the application never loses data and never goes down for more than a minute or less.
- **Maintain compliance.** Ensuring critical applications stay running isn't simply a matter of reducing downtime and its associated costs. Some businesses have applications that are critical enough that a prolonged outage puts them out of regulatory compliance.
- **Achieve scalability.** Since containers are lightweight, running applications in them instead of on servers or VMs allows for greater scalability. This is enhanced further by automation, which allows for large-scale deployment without increasing the administrative burden.

Trends

Businesses have been moving away from treating BCDR as an infrastructure task and instead have taken an application-centric perspective. Rather than looking at how quickly certain environments and storage systems can be brought back online, the focus has shifted toward application availability. Organizations see a spectrum of availability options and then map their applications based on criticality. This perspective affects how organizations interact with containers. If businesses find they cannot achieve the best levels of availability using containers, they will not deploy their most critical applications on them.

It is also important to note how nascent container technology is. Companies that have deployed containerized applications in the past have been repatriating them into traditional environments after realizing how much more complexity is involved in running and supporting a refactored legacy application. IDC research found more than half of the respondents had decided to reverse about half of their containerization work (see Figure 2). As the market transitions out of these early days, organizations will take a more strategic approach to containerizing — and protecting — their legacy applications.

FIGURE 2: ***Containerized Applications Getting Repatriated***

Q *Within two years, what percentage of workload functions that you have running via containers do you expect to be repatriated into traditional environments?*

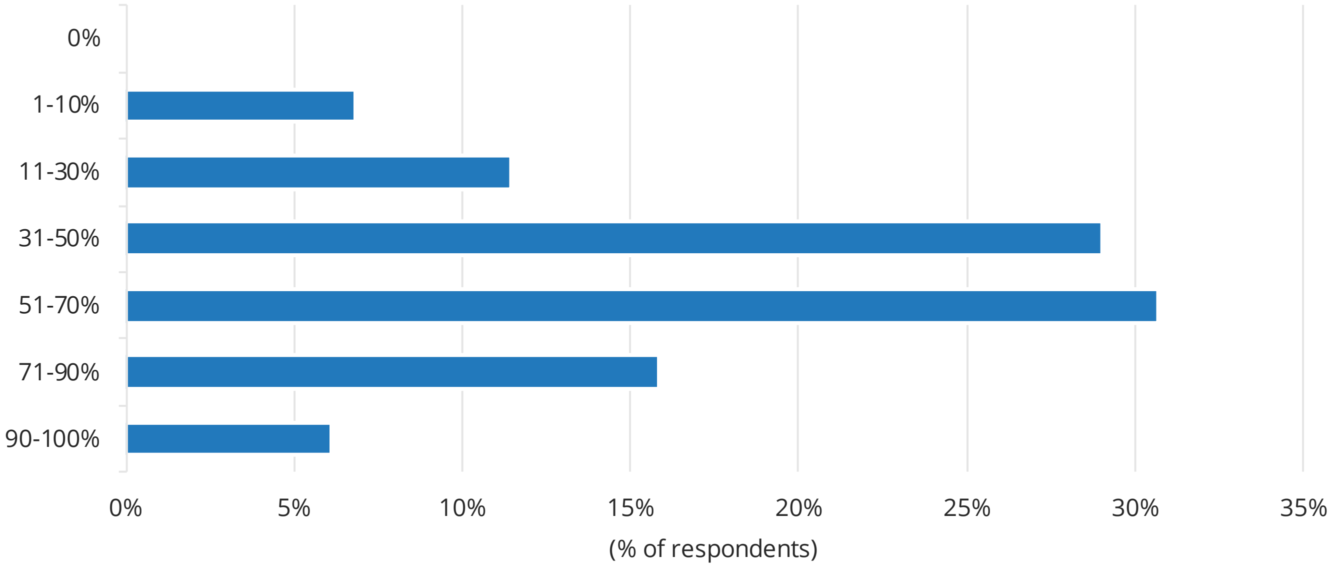

n = 410

Base= all respondents

Notes:

This survey is managed by IDC's Quantitative Research Group.

Data is not weighted.

Use caution when interpreting small sample sizes.

Source: IDC's IT Infrastructure for Storage and Data Management Survey, October 2021

Ransomware attacks have become a fact of IT life. Recovering from cyberincidents has become a much more critical use case for BCDR, and organizations are looking to shore up their defenses against cybercriminals. Although attacks specifically targeting container environments aren't as common, it is only a matter of time before organizations run enough critical applications on container platforms that it becomes lucrative for criminals to extort them.

As the containers and Kubernetes market matures and as organizations refactor their VM-hosted applications to work in containers, VM experts are being forced to become Kubernetes experts. Luckily, there are tools available that help this modernization process by allowing VMs and Kubernetes applications to run side by side on the same platform. BCDR that can support these tools will be in great demand, as it would allow organizations to make this transition at their own pace and without losing BCDR functionality.

Considering Intel

To enable businesses to run their mission-critical applications in Kubernetes, Intel worked together with Red Hat and Pure Storage to build a high-performance container BCDR solution. Intel and Red Hat codeveloped a reference architecture using Red Hat OpenShift and 4th Gen Intel Xeon Scalable processors. Portworx DR by Pure Storage provides DR capabilities for the solution.

Intel's Xeon processors provide the processing power for the BCDR solution. Compared with the previous generation, the 4th Gen Xeon chip added more accelerators to increase computing performance, with the aim of supporting heavy workloads such as artificial intelligence, high-performance computing, and data analytics. For example, Intel's 4th Xeon Scalable processors are available on AWS in the following instance types: M7i, M7i-flex, C7i, and R7i and R7iz.

Enterprise storage data movement is one of the latest-generation Xeon's targeted use cases as well, making it fitting for performant BCDR. This solution needs to move large amounts of data very quickly to keep data between primary storage and secondary storage in sync.

The container orchestration platform used for Intel's reference architecture is provided by Red Hat OpenShift, a cloud-native application platform for secure development life-cycle management. It features standardized workflows, support for multiple environments, continuous integration, and release management. Together with Intel's Xeon processors, these two form the software-hardware backbone of the BCDR solution.

Red Hat OpenShift is a family of products that includes Red Hat OpenShift Container Platform, which provides platform, developer, data, and application services. OpenShift Container Platform has a feature called OpenShift Virtualization, which allows users to run traditional applications in VMs alongside modern, cloud-native applications in containers. This provides management consistency across applications on a common hybrid cloud platform.

Red Hat OpenShift also includes several advanced products. Red Hat Advanced Cluster Management for Kubernetes allows for large-scale, multicluster management and policy enforcement across Red Hat OpenShift and other Kubernetes distributions. Red Hat Advanced Cluster Security for Kubernetes allows organizations to securely build, deploy, and run cloud-native applications. Red Hat Quay provides a security-focused and scalable private registry platform for managing content across globally distributed datacenter and cloud environments. Red Hat also has an extensive list of technology partners with validated integrations, including Portworx by Pure Storage.

Portworx Enterprise is Portworx's container storage and data management platform, and its PX-DR component provides BCDR for Intel and Red Hat's solution. PX-DR is capable of synchronous replication for data in storage between clusters in the same metro region, so changes made to the primary cluster are automatically reflected in the DR cluster. The tool can also perform asynchronous replication for other Kubernetes objects defined by scheduled policies for any two clusters, regardless of geographical location. As it's designed for the Kubernetes BCDR use case, Portworx PX-DR can capture all the components necessary to recover at the application level.

Intel's reference architecture takes advantage of the deep integration between Red Hat and Portworx. All components of Portworx Enterprise, including PX-DR, are visible from OpenShift's console. This enables customers to view their Kubernetes clusters, monitor the health of the clusters' underlying storage, and check data consumption without using

Portworx Enterprise's interface. They would only need to switch over to Portworx to execute storage and BCDR functions.

These three components of the reference architecture combine into a high-performance BCDR solution that aims to deliver zero RPO and RTOs of less than two minutes for organizations' most critical Kubernetes applications. The solution also provides higher-RTO asynchronous DR for less critical applications, offering businesses a spectrum of availability to which they can map their applications by criticality.

Challenges

The container market is still nascent, to the point that not all organizations are running a meaningful amount of their Kubernetes workloads in production and the workloads that are typically run aren't critical. Intel's solution addresses the needs of companies that have hit the performant BCDR roadblock for Kubernetes, but many aren't yet at the stage of considering what mission-critical applications they should move into their Kubernetes environments.

In addition, Red Hat OpenShift, while popular, isn't the only Kubernetes distribution or container orchestrator available. Organizations could be running other Kubernetes distributions such as Rancher, Docker, or HashiCorp Nomad, or they could be using cloud-based managed Kubernetes services such as Google's GKE, Azure AKS, or AWS EKS. Once they have a platform of choice, it is often difficult to get organizations to switch.

Conclusion

Companies can't reap the benefits of Kubernetes for their most mission-critical applications if they don't build a performant container BCDR first. They need to make sure these applications can be recovered with zero RPO and low RTO.

Intel, in partnership with Red Hat and Pure Storage, offers high-performance BCDR through a reference architecture that leverages powerful computing from Intel's Xeon processors, a features-rich Kubernetes distribution through Red Hat OpenShift, and robust BCDR software in Portworx-DR. Between the processing power of the Xeon chips and the integration between OpenShift and Portworx, the solution removes many of the hardware and software bottlenecks that can interfere with data and application availability.

> Companies can't reap the benefits of Kubernetes for their most mission-critical applications if they don't build performant container BCDR first.

As container and Kubernetes technology matures, more companies will be looking to deploy it, and they'll inevitably hit the BCDR roadblock. Intel's solution addresses this problem. By putting BCDR in IT teams' minds at the start of container application deployment, organizations will be able to start planning out ways to make their Kubernetes environments ready to support mission-critical applications — and avoid treating performant BCDR as an afterthought.

About the Analyst

Johnny Yu, *Research Manager, Storage and Computing*

Johnny Yu is a research manager within IDC's Infrastructure Software Platforms research group. He covers storage controller software; data replication, protection, and archiving; storage device management; and container data management, with a focus on how businesses optimize costs and secure their storage environments as their infrastructure expands beyond their datacenters.

MESSAGE FROM THE SPONSOR

The adoption and growth of Kubernetes from development and testing environments to supporting stateful, mission-critical applications has exposed the gap in the storage and data services taken for granted on legacy platforms. While Kubernetes has proven to be capable with regards to creation, orchestration, and deployment of containerized applications, the Achilles' heel is container storage and data management. Portworx is partnering with Intel and Red Hat to provide container storage and data management, built for Kubernetes, that delivers the performance, security, disaster recovery, and automation required by platform engineering teams to scale mission-critical, containerized applications in production. Learn more about our partnership around BCDR for Kubernetes by attending this webinar: https://portworx.com/webinar/disaster-proof-your-containerized-apps/.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.