Essential Capabilities of Disaster Recovery for Kubernetes

INTRODUCTION TO KUBERNETES DISASTER RECOVERY

You are responsible for the teams implementing new business applications, entering new markets, and attracting a new set of customers through innovative software. As a result, you have chosen to take a cloud-native approach to software development and operations with the adoption of Kubernetes in your organization.

The new applications built and run on Kubernetes are crucial to your organization's success in maintaining or expanding your market lead. But with the adoption of any new technology, there are many factors to pay attention to: new development methodologies, new teams, new staff, new technologies, new partners, new vendors, and new challenges.

As increasingly mission-critical applications move to Kubernetes, disaster recovery tops the list of concerns for enterprise IT leaders.

Figure 1: RTO/RPO demands by application type

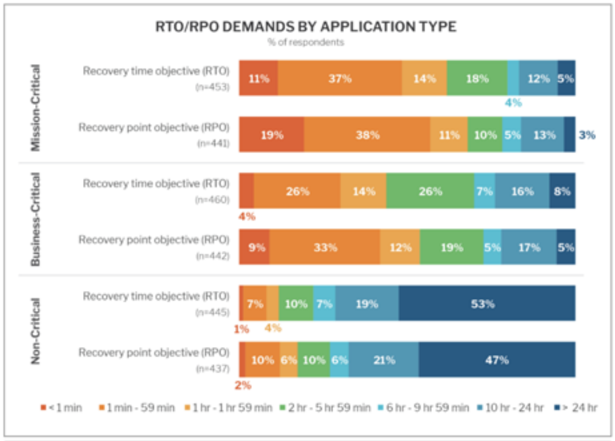

Source: 451 Research, Voice of the Enterprise: Storage, Workloads and Key Projects 2019

After making significant investments in Kubernetes, the last thing you want to happen is for applications to be unavailable because of some disaster out of your control—cloud providers go down, data centers lose power, services are unavailable, connectivity is disrupted, and customers are unhappy.

The Uptime Institute reports that outages are increasingly being blamed on third-party service providers such as co-location facilities or public cloud vendors. When 31% of outages are caused by an environment you rely on but don't control or have insight into and these outages are combined with network failures (30%) and IT/Software errors (28%)[1], you need a reliable disaster recovery solution for applications running on Kubernetes.

According to 451 Research, for mission-critical apps, 57% require an RPO of <1 hour and 48% require an RTO of <1 hour. Even non-critical applications have DR requirements that can stress already overburdened IT teams.

KEY TAKEAWAYS

Portworx solves the three major challenges to ensuring an enterprise-grade disaster recovery solution for your cloud-native Kubernetes applications. Built from the ground up specifically for Kubernetes, Portworx PX-DR is:

1 Container granular and understands Kubernetes' primitives such as namespaces for Kubernetes-native protection and recovery.

2 Application aware for a single enterprise-grade solution without the need for multiple point products.

3 Push-button easy and "just works," based on automation across all clouds, delivering zero data loss and fast recovery.

[1] "Is 99.99 An Industry Myth? Uptime Institute Study Shows Outages Are Common And Costly", Uptime Institute, August 2018

Yet, even with these skill gaps, enterprise IT teams must still deliver robust DR solutions for many enterprise applications.

Your disaster recovery solution not only needs to be portable and easy to use, but it also must be aware of the specific containerized applications and their individual technology components.

As a result, an easy-to-use and portable disaster recovery solution built specifically for Kubernetes is more important than ever.

The looming problem is that

60%

of large organizations face a **growing skills gap** in their DevOps and infrastructure based teams needed to implement these solutions[2]

UNIQUE CHALLENGES ORGANIZATIONS FACE PROTECTING KUBERNETES APPLICATIONS

Containers Are Fundamentally Different From Virtual Machines

Containerized applications are different from applications running in virtual machines. To successfully protect and then restore a containerized application, you must orchestrate a complex series of synchronized actions across a distributed system. This is because the application is likely running in multiple containers, and those containers are on different nodes within a Kubernetes cluster.

This is a vastly different architecture than the traditional single application in a virtual machine paradigm we've seen for the past 15 years.

Your existing backup and recovery solution may work very well for your current virtual machine based applications, but it won't work with containers or Kubernetes. Trying to use a traditional enterprise backup system that understands virtual machines will, at best, leave you with a false sense of security and, at worst, unavailable applications with corrupted data. The first rule of data protection is never to lose data. With a virtual machine-based solution, you're at risk.

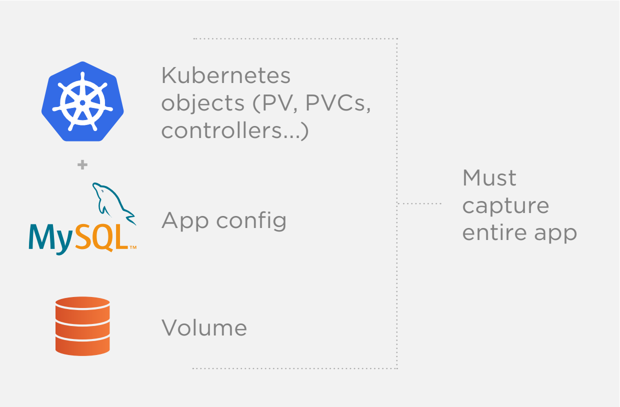

Seamless migrations, faster recovery

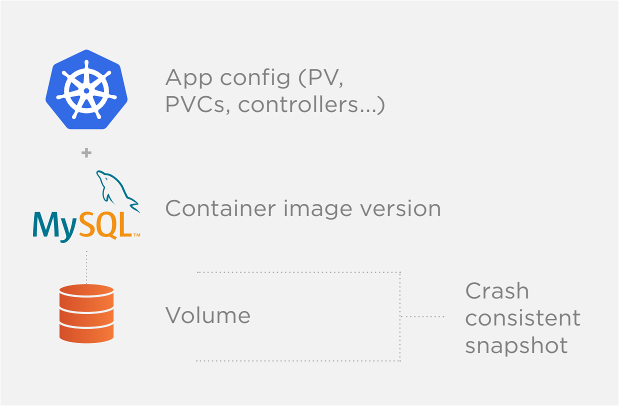
App config (PV, PVCs, controllers...)
+
MySQL
Container image version
Volume
Crash consistent snapshot

Other recovery solutions

[2] "IT infrastructure teams face a DevOps and cloud-native technology skills gap", 451 Research, November 2019

Applications Require Automated Intelligence For Recovery

Recovering a Kubernetes application after a disaster is not as simple as starting a new container in another location. Simple snapshots are no longer sufficient to ensure data consistency.

Containerized application components are deployed and scaled individually, each with their own container image, deployment configuration, state rules, extensions, life cycle operations, dependencies, and data. The additional configuration data and application business rules are often stored as metadata in your Kubernetes cluster and need to be protected and recovered to allow for application failover to work properly.

Furthermore, today's containerized applications are more than just a single component with their container image and data. Applications in a microservices architecture are comprised of front-end services combined with numerous middleware layers implementing business logic that is, most importantly, connected to persistent data services. This entire application stack must be backed up and recovered as a group.

> "
>
> There is a DevOps team. There's conversations. They've been around to say, 'How can we help you?' And we haven't found that they have enough technical knowledge yet to help us on the storage side but those discussions exist and there's discussions that maybe eventually we'll take a storage person and move them onto the DevOps team to see how they can help."
>
> - IT/Engineering Managers and Staff, Financial Services
> 50,000-99,999 Employees, $10bn+ Revenue[3]

46%

of organizations declared that **application portability without significant refactoring** is very important,[4] and the need for your disaster recovery strategy to "just work" is real.

Finally, the most popular data services (Kafka, Cassandra, Elastic, MySQL, and MongoDB) are all built by different communities with very little in common with respect to their operational life cycle and the skills required to operate them.

The increasing popularity of distributed data services demands a rethink of application protection, one that takes into account both the data and declarative operational metadata.

Fast Recovery with Zero Data Loss Across Every Cloud

Backup and data protection have always been the worst mix of the mundane and critically important. You want it to be boring because when it's exciting, it means something has gone wrong. You want it to "just work" and be part of your operational procedures as you scale your applications and environments.

Most critically, recovering your application with zero data loss in a timely fashion is the most important aspect of your disaster recovery strategy.

[3] "IT infrastructure teams face a DevOps and cloud-native technology skills gap", 451 Research, November 2019
[4] "Hybrid and Multi-Cloud: Economically, Often the Best Choice", 451 Research, December 2019
[5] "Capacity growth and disaster recovery are leading storage pain points for enterprises", 451 Research, March 2019

Over 53% of businesses report a RTO (Recovery Time Objective) of less than one hour for mission critical applications[5], but more than 50% take between one and four hours to recover[6]. There is still a long way to go for most organizations to achieve their availability requirements.

Finally, the need for your protection and recovery strategy to "just work" across all of your environments is more crucial than ever before. With over 65% of large organizations having already implemented a hybrid multi-cloud strategy across on-premises and at least one public cloud, 31% of whom are using more than two public clouds, you can't afford to have bespoke solutions for each provider.[10]

PORTWORX DELIVERS DATA PROTECTION FOR KUBERNETES

Portworx understands Kubernetes and containerized applications. We built a series of storage and data management solutions specifically to solve the challenges faced by organizations as you modernize your IT.

Portworx PX-DR is purpose-built to protect your Kubernetes applications, enable fast recovery with zero loss of data, and ensure your teams scale without requiring specialist skills for each new containerized technology you rely on.

Further increasing demand for full **end-to-end automation** of all aspects of the application lifecycle.

>41%

of outages cost over **USD $1,000,000**[7]

30%

outages take longer than **5 hours** to recover[8]

80%

of outages were **deemed to be avoidable**[9]

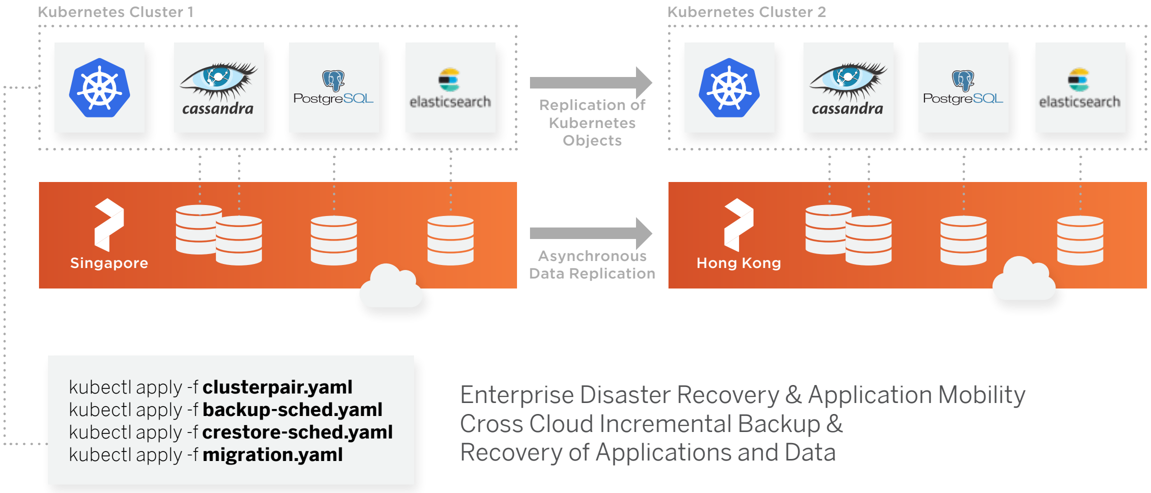

```
kubectl apply -f clusterpair.yaml
kubectl apply -f backup-sched.yaml
kubectl apply -f crestore-sched.yaml
kubectl apply -f migration.yaml
```

Enterprise Disaster Recovery & Application Mobility
Cross Cloud Incremental Backup &
Recovery of Applications and Data

[7,9] "Uptime Institute data shows outages are common, costly, and preventable", Uptime Institute, June 2018
[6,8] "Uptime Institute's Datacenter Survey, Part 1: Outages, PUE, Edge Trends", Uptime Institute, June 2018
[10] "Vote Studies Show Hybrid, Multi-cloud Are Now Standard Cloud Strategies," 451 Research, July 2019

Built for Kubernetes

Containerized applications typically run in multiple containers across multiple hosts. The delineators are Pods and Namespaces (or Project by some Kubernetes distributions). Portworx PX-DR understands both the Pod and Namespace constructs, enabling you to protect an entire application at the either container-granular or Namespace level.

By protecting the Pod or entire Namespace, you get peace of mind that, regardless of your application configuration or the placement across machines in your cluster, you can simply and easily select applications to protect.

Protecting your application is more than simply orchestrating snapshots, although that is complex enough. Portworx's data protection solutions also make it simple to restart your application quickly in another Kubernetes cluster—regardless of cloud provider or location.

By ensuring the protection of your application, its configuration, and, most importantly, its data, Portworx delivers true Kubernetes native disaster recovery.

Application Aware

Today's containerized applications are increasingly built with a diverse set of technologies—presentation, message streaming, analysis, and data storage—that have vastly different operational models and communities.

Portworx PX-DR speaks to **Kubernetes** natively and has pre-built integrations with many of the leading data services **used by many of the Global 2000** today.

The only way to scale your application and data protection strategy is either to mandate a restricted set of technologies that limit your agility or employ a solution that handles the intricacies natively, enabling you to accelerate your transformation.

App-consistent backups means understanding apps

1. Flush & Lock tables in the background
2. Status complete and return to CRD

A. Flush memory
B. Status complete and return to CRD

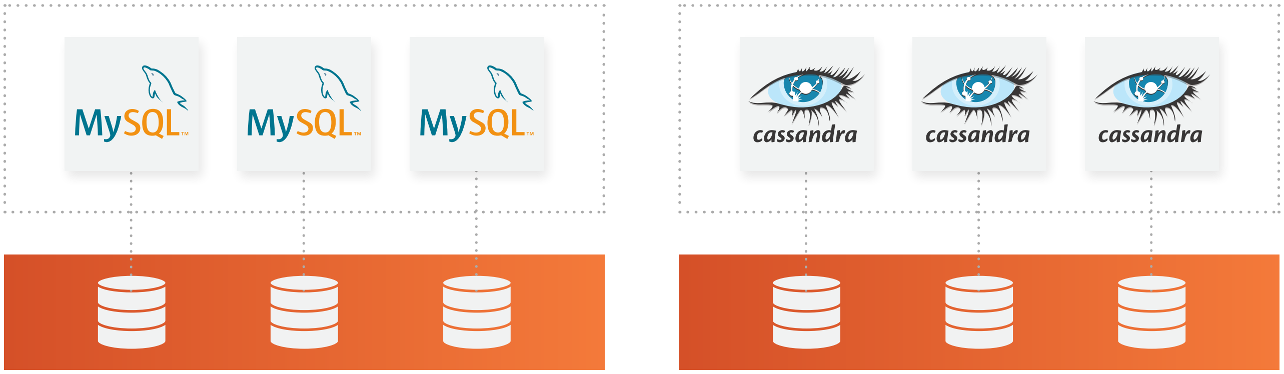

3. Freeze filesystems and snapshot
4. Unfreeze filesystems
5. Release table Lock

C. Freeze filesystems and snapshot
D. Unfreeze filesystems

This means that you can continue to innovate and deploy new and transformative applications, safe in the knowledge that you don't need individual specialists for every technology. Portworx takes care of the underlying integration for you—freeing you up to simply set protection schedules that meet your availability requirements.

Consistently Reliable DR in a Multi/Hybrid-Cloud World

Your applications don't live in a single controlled environment. Even if you're only using a single cloud vendor, you've deployed applications across multiple regions. As you continue to adopt a multi/hybrid cloud approach, this complexity will increase.

Push-button easy that *just works*, based on automation across all clouds delivering **zero data loss** and **fast recovery.**

Portworx PX-DR is built specifically with multiple and heterogeneous environments in mind. Delivering zero data loss (Zero RPO) and fast recovery times (Low RTO) is a key requirement for application resilience—and Portworx's solutions deliver this.

When deploying your applications across very low latency connected clouds—for example, with high speed metropolitan private links—Portworx enables you to achieve Zero RPO and Zero RTO. This means no data loss and instant recovery in the event of a site failure.

In more traditional configurations where your applications are deployed across diverse geographic locations, with Portworx, you can experience fast recovery in seconds to minutes without any data loss.

Portworx PX-DR data protection empowers organizations to far exceed the expectations of most businesses' one-hour recovery objective.[11]

When your applications work even when your clouds don't, your customers keep coming back.

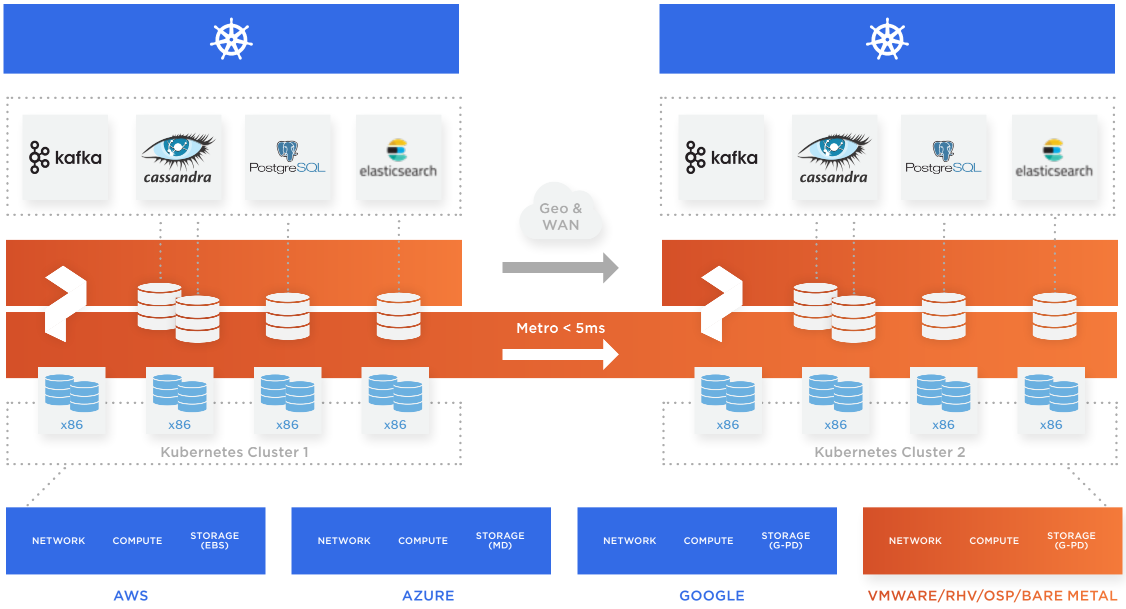

[11] "Capacity growth and disaster recovery are leading storage pain points for enterprises", 451 Research, March 2019

Delivering On Our Promise For Royal Bank of Canada

Royal Bank of Canada wanted to take advantage of the benefits they saw in Kubernetes with Red Hat OpenShift©. However, they were not able to meet strict availability requirements with just Kubernetes. With Portworx Enterprise Storage Platform for Kubernetes and PX-DR, Royal Bank of Canada was able to achieve Zero RPO and an RTO of less than two minutes. That means they are able to recover their application in another data center in under two minutes with no data loss.

Without Portworx's disaster recovery capabilities, RBC would not have been able to ensure availability and thus would not have been able to deploy their application on OpenShift©.

CONCLUSION

Your business is moving fast; your competitors are, too. In this highly-connected world where customer experience is king, failing your customers is not an option.

Mistakes occur, outages happen, and clouds fail—don't get caught without a robust, agile, and fast recovery solution that ensures your applications are up and running again quickly and that you don't lose any customer transactions.

The containerized applications you are deploying on Kubernetes are your growth engine, but they are also built differently than your traditional applications. You need to protect them differently with a solution built specifically for Kubernetes that understands containers.

Portworx PX-DR solves these challenges and enables you **to build a platform for success.**

To truly scale and take advantage of DevOps, requiring specialized staff for each technology is a limiting factor you can't afford. Your disaster recovery solution needs to be application-aware and to enable your highly-skilled generalists to scale without limits.

Hybrid multi-cloud is the new normal, which means constraints that limit functionality across providers is untenable. Your applications need to use whatever services they require in whichever cloud you choose.

WHAT'S NEXT

Portworx has been helping organizations in the Global 2000 and more for over 5 years. We understand what you're trying to achieve and where most of the challenges are. We are ready to help you.

As a next step, we recommend referring your lead architect to the Portworx Disaster Recovery for Kubernetes overview, or you are welcome to speak directly with a specialist.

To see Portworx PX-DR in action, check out this 5 minute explainer video.