VERITAS

Veritas NetBackup for Kubernetes Data Protection

Optimize your Kubernetes environment with unified, enterprise-grade resiliency, choice and flexibility.

White Paper | September 2022

Contents

Executive Summary
Why NetBackup for Kubernetes?
Operational Simplicity
Kubernetes-Native and API Integration
Efficient Management
Simplified Administration
Enterprise-Grade Resiliency
Application Centric
Snapshots
Flexible Recovery
Security
Choice and Flexibility
Adopt Any Cloud
Support Any Storage
Infrastructure As Code
NetBackup for Kubernetes Architecture and Components
Summary
Appendix
Versions

Executive Summary

Digital transformation has accelerated the need for enterprises to offer their services to customers easily over the web. This transformation is a culture change and a survival need that mandates more agility to improve the customer experience, reduce friction, increase productivity, and elevate profitability. Kubernetes (K8s) is an open source container orchestration platform that automates deploying, managing, and scaling containerized applications' processes. Kubernetes clusters can span hosts across on-premises, hybrid, and multi-clouds and is an ideal platform for hosting cloud-native applications to provide the consistency teams need to work with multiple clouds by creating abstractions that bring all deployments into one environment.

As more Kubernetes workloads go into production, enterprises need to maintain compliance. The flexibility and scalability Kubernetes offers also introduces levels of data protection challenges. Solutions not designed to be Kubernetes-native cannot tie into the CI/CD pipeline and can negatively impact the cluster. In Kubernetes, applications are made up of many components, so application protection and recovery need to be well orchestrated, or the applications may not be able to recover. With various distributions of Kubernetes running on-premises and in the cloud, organizations need Kubernetes native APIs to support any on-premises or in-cloud distribution and be able to back up anywhere and recover anywhere.

An Enterprise Data Protection solution for Kubernetes must be application-centric. One application workload could comprise hundreds of components and cross many containers. Unlike Virtual machines where you back up the VM, containers are ephemeral and merely serve as a holding place. That's why Kubernetes data protection should focus on the application, not the container.

With the following design principles, Veritas NetBackup offers scalability, simplicity, flexibility, and platform- and storage-agnostic and application-centric protection.



Figure 1. An overview of the features an effective Kubernetes data protection solution should provide.

NetBackup Kubernetes data protection achieves operational simplicity, enterprise-grade resiliency, and choice and flexibility.

- Use native tools to deploy, communicate, and interact with the Kubernetes cluster.
- Tie into enterprise data protection frameworks for SLA and Security management.
- Automatically detect the creation of containerized applications and apply the right backup policy to each container, if necessary.
- Support any distribution of Kubernetes in a unified way for the greatest choice of infrastructure for the customer.
- Automatically scale up or down to provide protection for the largest and most demanding environments.
- Recover the application from any type of failure or ransomware scenario.

Why NetBackup for Kubernetes?

Data is an integral part of an organization, and it must be protected and recoverable from the edge to the core to the cloud. Veritas designed NetBackup[™] for Kubernetes to offer operational simplicity, enterprise-grade resiliency, and choice and flexibility for Kubernetes workload protection (see Figure 2).



Figure 2. A summary of the features the NetBackup for Kubernetes solution provides.

Operational Simplicity

Veritas enables backup administrators to protect their K8s environment without becoming Kubernetes experts and for Kubernetes Admins to protect their applications without learning a backup tool. Our backup and recovery solution is optimized by using Kubernetesnative constructs for the Kubernetes administrators and a simple web UI for backup administrators.

Kubernetes-Native and API Integration

NetBackup uses Helm charts to enable customers to deploy all other Kubernetes resources. In addition, we provide a full suite of RESTful application programming interfaces (APIs) with role-based access control (RBAC) that enables self-service for users to align with their CI/CD pipeline and native tools. In NetBackup 8.2 or later, you can test out NetBackup APIs in your own environment with the Swagger interface at https://<master-server-name>/api-docs/index.html.

Efficient Management

NetBackup's self-service agentless management with RBAC and API-driven workflows lets you delegate tasks, reducing the time spent on administrative activities and free up resources for higher-impact functions.

Simplified Administration

NetBackup provides a simple web UI to back up and recover Kubernetes workloads. With a few simple clicks, you can manage Kubernetes protection jobs, and you can automate the protection of your application namespaces with dynamic templates to ensure new namespaces are protected as soon as they come online.

Enterprise-Grade Resiliency

With NetBackup for Kubernetes, you can protect your environment without impacting or disrupting mission-critical applications. NetBackup is built with resilient Kubernetes constructs that deploy and work seamlessly with your environment.

Application Centric

Kubernetes consists of many abstractions that wrap applications and their data and provide interfaces for container orchestration services. For example, the underlying storage is provided to Pods via Kubernetes Persistent Volumes (PVs). They allow allocating a specific amount of storage to an app and configuring write/ read access permissions, I/O limits, and storage security. Similarly, there are many other objects—such as Secrets, Service Accounts and Jobs—that control how containers within a Pod communicate and how the data is accessed by various microservices. Traditional data protection tools do not know how to interact with these abstractions to make true application backups, including Kubernetes objects, application configuration, and data. NetBackup is Kubernetes-aware and protects your applications seamlessly. NetBackup discovers and coordinates the snapshot and recovery of all components that comprise an application, including all PVs, config files, and custom resources.

Snapshots

The Container Storage Interface (CSI) is a standard for exposing arbitrary block and file storage systems to containerized workloads on Container Orchestration Systems (COs) like Kubernetes. Most storage vendors have their own CSI plug-ins to work exclusively with their storage arrays.

NetBackup leverages Kubernetes-native snapshots with CSI plug-ins. The native snapshots do not impact performance and provide the availability of always-on cloud operations.

Flexible Recovery

NetBackup provides the recovery flexibility to back up once and recover from any level of disaster. For example, with NetBackup, you can backup and recover the ETCD configuration of the cluster, roll back instantly from a snapshot of the namespace or restore an entire workload within a namespace, individual resources, or just one or all of the associated PVs to the same or an alternate Kubernetes cluster or distribution.

Security

NetBackup for Kubernetes ties into the overall NetBackup Ransomware capabilities. With this additional level of security, you have the confidence of knowing you have encryption in flight and at rest, and your backups are checked for any potential anomalies to keep you safe.

Choice and Flexibility

NetBackup spans any CNCF-certified Kubernetes distribution—on-premise or in the cloud—to enable unified backup and recovery.

Adopt Any Cloud

NetBackup is designed to use Kubernetes-native APIs to support any on-premise or in cloud distribution. With NetBackup, you can backup your Kubernetes data anywhere and recover anywhere.

Support Any Storage

NetBackup leverages CSI plug-ins to support any storage with production-ready CSI plug-ins. This gives Kubernetes users more options for storage and makes the system more secure and reliable. Customers can choose disk or cloud-based storage or leverage NetBackup Appliances for scale-out storage, storage, and network optimization using our deduplication technologies or immutable storage for even greater security.

Infrastructure As Code

NetBackup includes a full suite of APIs, providing the choice and flexibility you need to build custom workflows to fit into your CI/CD pipelines. In addition, you can use Swagger to check the format and examples of these APIs.

NetBackup for Kubernetes Architecture and Components

A Kubernetes cluster is a set of nodes that run containerized applications. Containerizing applications package an app with its dependencies and some necessary services. Kubernetes clusters allow containers to run across multiple machines and environments: virtual, physical, cloud-based, and on-premises.

There are two main components in the NetBackup for Kubernetes solution: the NetBackup Kubernetes Operator and the NetBackup Elastic Data Mover that get deployed into a NetBackup dedicated namespace.



Figure 3. An overview of the main components of the NetBackup for Kubernetes solution.

Kubernetes namespaces include all the custom resources and configuration that comprise that namespace. Persistent Volume (PV) and Persistent Volume Claim (PVC) are two important concepts in Kubernetes storage.

The **controller node** contains all the cluster information and access. Communication to and within the cluster is handled by an API server. For example, any communication to the NetBackup Custom Operator or calling of snapshots goes through the API server to make the request.

NetBackup dedicated namespaces are deployed with the NetBackup Kubernetes operator (KOps). This operator is a custom resource Veritas developed to go into Kubernetes environment and facilitate all the calls by Velero through the API server. It communicates information and metadata back to the NetBackup Primary Server, defines the schedules, and facilitates all communication between NetBackup and Kubernetes. The NetBackup KOps will communicate with the CSI plug-ins through the API server to take any snapshots and storage of those snapshots to any cloud-compliant storage and communicate that status back to the Primary Server. Once the snapshot is taken, if desired, a backup copy of that snapshot is written to longer term storage as a second copy to meet customer SLA's. This movement is optimized using a data mover written by NetBackup to be cloud-optimized to scale up and down resources as needed to provide the best performance with the least impact on the cluster.

Summary

Kubernetes is a portable, extensible, open-source platform for managing containerized workloads and services that facilitates both configuration and automation. NetBackup for Kubernetes is designed to optimize the Kubernetes environment for the backup admin. Its application-centric design provides unified, enterprise-grade resiliency for all Kubernetes platforms and delivers the choice, and flexibility organizations have come to expect from Veritas.

Appendix

Kubernetes Terms

Key Term	Description		
API Server	The Kubernetes API server validates and configures data for the API objects, which include pods, services, replication controllers, and others. The API server services REST operations provide the front end to the cluster's shared state through which all other components interact.		
Controller Manager	The Kubernetes controller manager is a daemon that embeds the core control loops shipped with Kubernetes. In Kubernetes, a controller is a control loop that watches the cluster's shared state through the API server and makes changes, attempting to move the current state towards the desired state.		
CSI /CSI Plug-in	The Container Storage Interface (CSI) is a standard for exposing arbitrary block and file storage systems to containerized workloads on Container Orchestration Systems (COs) like Kubernetes. Most storage vendors have written their own CSI plug-ins to work exclusively with their storage arrays.		
Jobs	A job in Kubernetes is a supervisor for pods carrying out batch processes, that is, a process that runs for a certain time to completion, such as a calculation or a backup operation.		
ETCD	Kubernetes uses ETCD as the back end for service discovery and stores the cluster's state and its configuration.		
Persistent Volume (PV)	Kubernetes makes physical storage devices like SSDs, NVMe disks, NAS and NFS servers available to a cluster in the form of objects called Persistent Volumes (PVs). A PV is an abstraction for the physical storage device you have attached to the cluster. If you are using Kubernetes on Google's or Amazon's cloud, you can make your Google SSDs or EBS volumes available to your containers in the form of PVs. Each PV is consumed by a Kubernetes pod (or pods) by issuing a PersistentVolumeClaim object, a PVC. The easiest way to create the PV/PVC pair for your Pod is to use a StorageClass object and then use the storageclass to create your PV/PVC pair dynamically whenever you need to use it.		
Pods	Pods are the smallest deployable units in Kubernetes. As the official documentation states, "A pod (as in a pod of whales or pea pod) is a group of one or more containers, with shared storage/network resources, and a specification for how to run the containers." So in simple terms, a pod is the mechanism for how a container actually gets turned "on" in Kubernetes.		
Storage Class	A storage class is a Kubernetes object that stores information about creating a PV for your pod. With a storageclass, you do not need to create a PV separately before claiming it.		
CI/CD	Continuous integration (CI) and continuous delivery (CD) embody a culture, a set of operating principles, and a collection of practices that enable application development teams to deliver code changes more frequently and reliably. The implementation is also known as the CI/CD pipeline.		

Versions

Version	Date	Author	Key Updates
1.0	Sep 2021	Rachel Zhu	Original document.
2.0	Sep 2022	Rachel Zhu	Updated for NetBackup 10

About Veritas

Veritas Technologies is a leader in multi-cloud data management. Over 80,000 customers including 95% of the Fortune 100 -rely on us to help ensure the protection, recoverability, and compliance of their data. Veritas has a reputation for reliability at scale, which delivers the resilience its customers need against the disruptions threatened by cyberattacks, like ransomware. No other vendor is able to match Veritas' ability to execute, with support for 800+ data sources, 100+ operating systems, 1,400+ storage targets and 60+ clouds through a single, unified approach. Powered by our Cloud Scale Technology, Veritas is delivering today on its strategy for Autonomous Data Management that reduces operational overhead while delivering greater value. Learn more at www.veritas.com. Follow us on Twitter at @veritastechllc.

VERITAS

2625 Augustine Drive Santa Clara, CA 95054 +1 (866) 837 4827 veritas.com

For global contact information visit: veritas.com/company/contact