INFINIDAT

SOLUTION BRIEF

Kubernetes Data Protection Using the Infinidat CSI Driver

Kubernetes, the fastest-growing infrastructure software, provides high availability and scalability of application services. These benefits do not extend to customer data. As a result, data protection is a critical priority for Kubernetes workloads.

With the increasing use of containers and Kubernetes in the enterprise beyond DevOps, companies are adopting enterprise-class storage for their stateful, container-based workloads. Using the Infinidat Container Storage Interface (CSI) Driver, one can effortlessly weave containerized applications into the same storage framework as other bare metal and virtual workloads. This movement of containerizing applications has brought the need for data protection of these environments to the forefront.

Kubernetes environments continue to experience the following data protection challenges:

- **Application consistency**
- **Volume consistency**
- **Scalability**
- **Performance impact**
- **Disaster recovery**

Infinidat has collaborated with many of the top Kubernetes data protection vendors to ensure that you can safely back up and restore your critical applications and data from the InfiniBox® platform to the Infinidat backup platforms – InfiniBox and InfiniGuard®. The Infinidat CSI Driver supports those using Commvault, IBM Storage Protect Plus, Kasten K10, Trilio for Kubernetes, or Veritas NetBackup.

Whichever platform used to protect Kubernetes applications, it must automatically discover all the components running on the cluster and treat that application as the unit of atomicity. It is crucial for the application to include the state spanning across all storage volumes and databases, as well as the configurable data in Kubernetes objects, such as ConfigMaps and Secrets.

With the Infinidat CSI Driver and the data protection solutions, enterprises can safely backup and restore their:

- Kubernetes-orchestrated clusters, including namespaced and non-namespaced API resources and objects
- Applications, which include supported API resources/objects (such as Secrets, ConfigMaps, Namespaces, and StorageClasses) that can be listed, created, or re-created using the Kubernetes API server
- Annotations on Pods, DaemonSets, Deployments, and StatefulSets
- Helm chart-based applications, including helm configuration and annotations (supported only for on-premises access nodes)
- Configuration-related volumes (configMap, downwardAPI, projected, secret)
- Persistent storage objects (PersistentVolumeClaims, PersistentVolumes)
- CSI-enabled out-of-tree volume plug-ins (recommended)
- Legacy in-tree volumes (VMware vSphere volume plug-in)
- PersistentVolumeClaim volumes created from a VolumeSnapshotClass
- Container image registries (containerized, virtualized)
- etcd Kubernetes backing store and SSL certificates (on-premises environments and self-managed cloud environments only

Data Protection for Kubernetes Workloads

Backup

Restore

Disaster Recovery

A persistent volume can be mounted on a host in any way supported by the resource provider. The Infinidat CSI Driver supports various access modes:

- ReadWriteOnce (RWO) - The volume can be mounted as read-write by a single node. ReadWriteOnce access mode still can allow multiple pods to access the volume when the pods are running on the same node.
- ReadOnlyMany (ROX) - The volume can be mounted as read-only by many nodes.
- ReadWriteMany (RWX)- The volume can be mounted as read-write by many nodes.

The Infinidat CSI Driver with Leading Kubernetes Data Protection Solutions

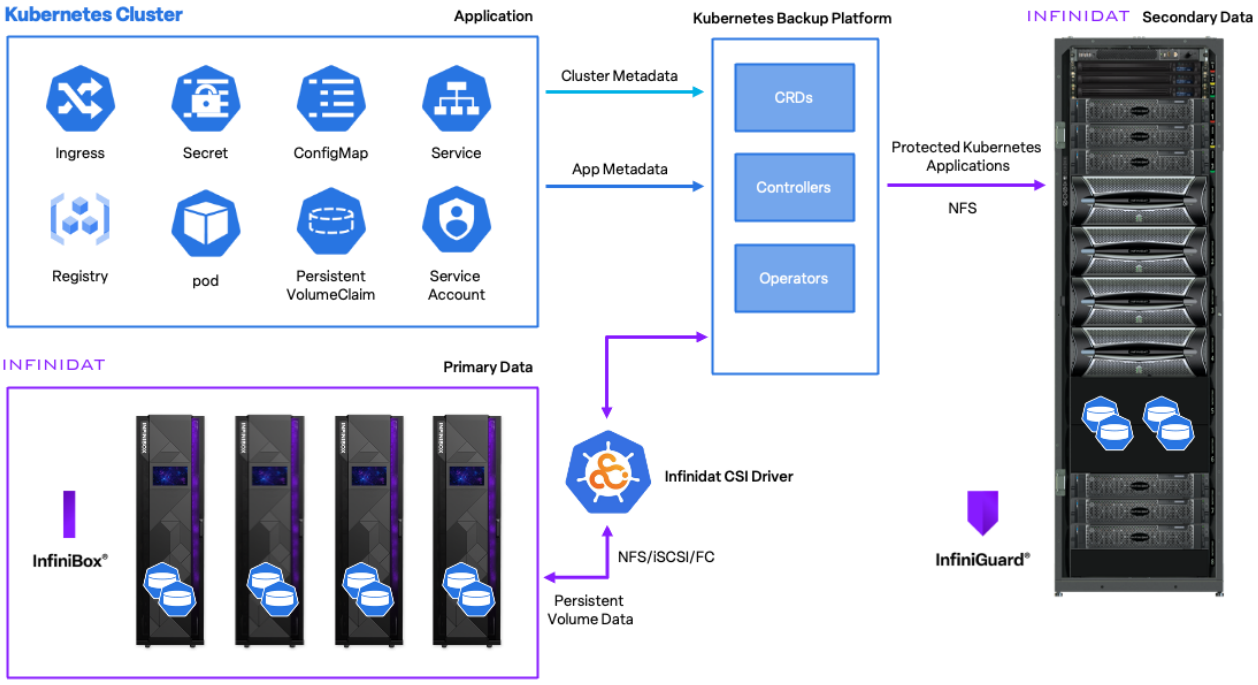

Fig. 1 – Kubernetes data protection configuration for Infinidat primary and backup storage

The Infinidat CSI Driver can be used along with your Kubernetes data protection software for full protection of your environment. By scheduling regular backups and testing your restore capabilities, your can be ready in case of any event, including hardware failures, ransomware attacks or natural disasters.

Backup

The Kubernetes data protection software continuously monitors the cluster for changes and triggers the backup process based on the defined policies. It identifies the relevant Kubernetes objects and prepares them for backup. By creating a VolumeSnapshot resource, one can obtain a snapshot of the data at that point in time. The software interacts with the CSI driver to create snapshots of the required volumes associated with the Kubernetes objects. The CSI driver communicates with the InfiniBox storage system via API calls to create consistent point-in-time snapshots of the volumes. These snapshots capture the data in the volumes at the time of the backup. Along with the volume snapshots, the software captures metadata related to the backup, including labels, annotations, and other relevant information about the Kubernetes objects being backed up. This metadata is crucial for accurate and efficient restores. The software performs verification and integrity checks on the backed-up data to ensure its reliability.

Restore

One can determine the specific backup that you want to restore. The data protection software maintains a catalog of backups, which can be browsed through the UI or queried via the CLI to find the relevant backup. You can then decide on the restore options based on requirements. For example, you might choose to restore the entire application, specific namespaces, or individual volumes.

Once the desired target location and any additional parameters the CSI driver requires are specified, the software communicates with the CSI driver to orchestrate the restore operation.

Provision of new PVs and PVCs: If the original PVs and PVCs associated with the backup are unavailable or cannot be used, the software will provide new ones based on the restore specifications.

The software coordinates with the CSI driver to restore the data from the backup snapshot to the newly provisioned PVs. The CSI driver will create another snapshot from the backup snapshot, then change it from read-only mode to a write-enabled mode and export it to the cluster.

Once the data restore process is complete you can verify the success of the restore operation by checking the logs or using the provided validation mechanisms to ensure the integrity and correctness of the restored data.

Disaster Recovery

Once you have set up your data protection software for DR purposes, you will be exporting your backups to InfiniGuard via NFS.

In the event of a disaster, you will deploy a fresh cluster and configure your CSI driver, recover your data protection software from NFS, and then recover all the applications using the normal restore operations of the data protection software.

Benefits of the Joint Solution

The combination of Infinidat and the referenced data protection solutions provides enterprises who are adopting Kubernetes for production workloads more assurance. As the use of Kubernetes in these environments continues to grow, Infinidat continues to innovate with the ecosystem of data protection vendors to ensure that you are provided with:

- High-performance storage that helps you seamlessly implement backup and recovery, disaster recovery, and stateful application mobility.
- Multi-protocol flexibility so that you can manage Kubernetes Persistent Volumes attached via block and file protocols, including Fibre Channel, iSCSI, NFS, and NFS-Treeq.
- Multi-petabyte scalability to support hundreds of thousands of PVs per InfiniBox system and control multiple InfiniBox arrays within a single Kubernetes cluster.
- Advanced enterprise features to manage native InfiniBox snapshots and clones, including restoring from snapshots, and importing PVs created outside of the InfiniBox CSI Driver.

sales@infinidat.com

INFINIDAT