

# Symcloud Storage

## For Kubernetes in any cloud

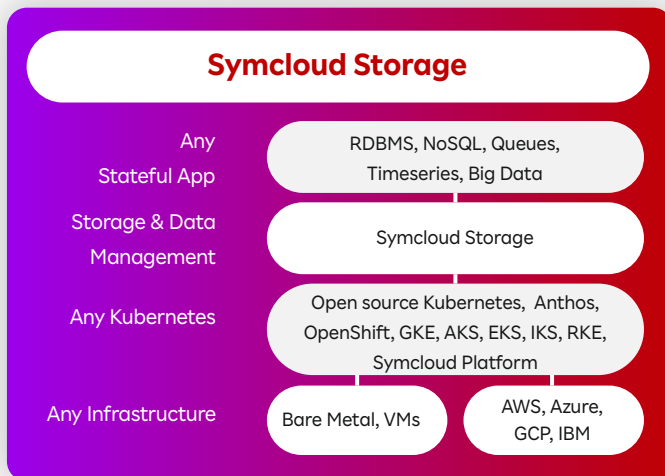
Symcloud Storage brings advanced storage and data management to any Kubernetes running distribution on-premises or in any cloud. Symcloud Storage seamlessly integrates with Kubernetes-native administrative tooling such as Kubectl, Helm Charts, and Operators through standard APIs, CLI and an easy-to-use GUI. Developers and DevOps teams can deploy Symcloud Storage as a standard Kubernetes operator. The entire solution can be fully automated based on user-defined policies that automatically track all application and storage constructs over the lifetime of the solution, drastically reducing time to

outcome and human error.

Symcloud Storage delivers unprecedented ease of use and automation, while delivering bare-metal performance at scale for your stateful, cloud-native applications (data + metadata + Kubernetes configuration). Symcloud Storage is built from the ground up with cloud-native constructs designed to provide a unique, automated cloud-native experience.

With Symcloud Storage, all lifecycle and data protection operations are performed in the context of the entire application-aware state (data + metadata + Kubernetes configuration), not just storage volumes. As a pioneer of application-aware cloud-native storage, Symcloud Storage understands your application's lifecycles, as well as the many functions and roles that make them up, including the in-depth interaction between storage and its ever-changing Kubernetes application constructs.

Symcloud Storage makes it easy for developers, QA and Ops teams; they no longer need storage expertise to manage complex applications, and can instead focus on high-value projects.



### The challenge of cloud-native Kubernetes and containerization

Kubernetes, although essential for next-generation cloud scale and automation, adds a layer of complexity when it comes to application data protection and disaster recovery. In the past, storage could easily be mapped to a host or a virtual machine (VM). Hence, data protection and recovery mechanisms were centered around storage only. But, with container-based solutions, like Kubernetes, the application is deconstructed into multiple components that will scale independently and have different relationships with your data. Therefore, in the case of application cloning, storage scaling, migration, replication, environment changes and recovery, the relationship between applications and their data becomes critical. Both must be managed together as a whole to achieve accurate and immediate outcomes.

### Why stateful applications need Kubernetes-aware storage

Kubernetes was initially designed for stateless workloads. Stateful workloads need a storage layer that is Kubernetes- and application-aware to accelerate data management, provisioning and data protection. Kubernetes does not actively manage the underlying storage for stateful workloads, and relies on an external interface called the Container Storage Interface (CSI). Symcloud Storage uses the CSI to deliver the industry's most robust stateful application support across any cloud and environment. It is fully automated to drastically reduce recovery time, migration time, and human error, while introducing industry-leading functionality.

## Symcloud Storage Benefits

Reduce time to outcome and human error with superior automation, over multiple tasks and data protection lifecycle events:

- Automate data lifecycle management such as backup/recovery, snapshots, rollback, and cloning of the entire application state, providing enterprise-grade availability
- Automate tracking of your container-to-storage relationships as they change over time; ultra-reliable, eliminating stealth breaks after application upgrades and reconfigs

Improve productivity for your developers, QA and Ops teams, freeing up valuable time:

- Collaborate across teams by cloning entire applications in minutes
- Confidently move application and storage to new environments, Cloud, Core, Edge, MEC, enterprise
- Get started quickly with an easy-to-use, no-storage-expertise-needed solution that simplifies application and data management

### Protect your existing investments

- Scale with linear bare-metal performance as your applications and data grow over time
- Manage any workload, big data, enterprise, 5G, MEC, databases, content delivery, gaming, network functions, stateless and stateful
- Utilize your existing disks and storage arrays
- Deploy and manage reliable storage on every popular Kubernetes distribution

## Symcloud Storage Highlights

### Availability and data protection

- Fully application-aware data protection of not only your storage but your applications, metadata, configurations and secrets – the full package or specific modules can be easily scaled, migrated to new environments and rolled back at the press of a button or automated event
- Automated discovery of nodes and disks
- Auto-reconnect to disks on node failure, for storage arrays and cloud drives
- Automatic data resync on node recovery
- Advanced multi-data center, multi-cloud and stretch cluster disaster recovery, with synchronous and asynchronous replication
- Full quiescence of complex databases before snapshotting, with commits on disk writes, not memory cache writes
- Advanced encryption and key management support
- Backup by application namespaces or other custom groupings

### Performance management

- Built from the ground up to scale with bare-metal performance, for both network- and storage-intensive workloads
- IOPs-based QoS enforcement
- Live data rebalancing to eliminate hotspots
- Automated, policy-driven, data locality with affinity/anti-affinity based on compute and storage, as well as individual storage components

### Capacity management

- Application to spindle visibility with advanced monitoring and an easy-to-use GUI, API and CLI
- Selective replication and storage for all or individual volumes on a per application basis
- Policy-driven, always-live expansion of data volumes
- Thin, differential snapshots and clones, where only the incremental changes need storage
- Data compression
- Policy-driven, by storage type, auto-provisioning based on drive type (NVMe, HDD, SSD), rack, cluster and custom location
- Supports block and file storage as well as ReadWriteMany, ReadWriteOnce and ReadOnlyMany access modes

### Deploy anywhere

- Runs on-prem with your existing disks and arrays
- Is fully integrated into your favorite cloud vendor, with hybrid cloud capabilities
- Easy onboarding of any Kubernetes application using Helm, YAML and existing application operators
- Broad Kubernetes distribution support with custom integration into all major distributions including, open source Kubernetes, Anthos, OpenShift, GKE, AKS, EKS, IKS, RKE, and Symcloud Platform

