



StorPool
DISTRIBUTED STORAGE

Persistent Storage for Kubernetes

High-performance persistent volumes
for large Kubernetes environments

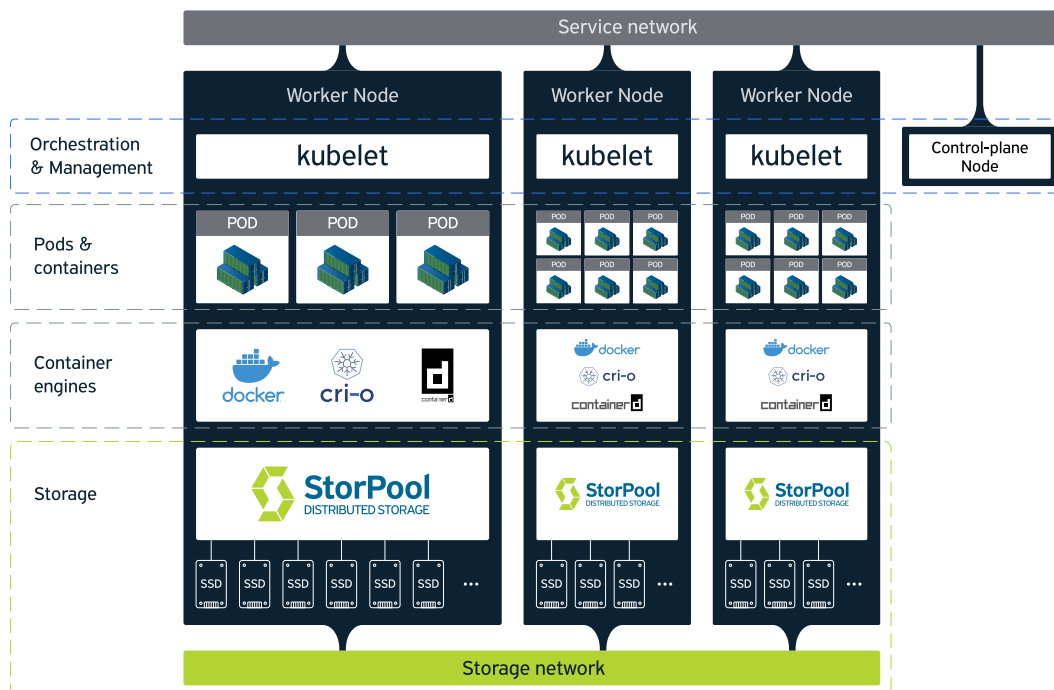


High-Performance Storage for Kubernetes

Kubernetes is emerging as the industry standard for deploying containers at a scale. It provides the developers with an API-centric way to manage containerized applications, which self-manage, heal and scale. Kubernetes has changed the way in which we build, ship and run applications. Now, it has become the foundation for modern infrastructure and New IT.

Containers need to work together and share data. Modern business applications require persistent storage and enterprise features like high availability, security, and performance.

StorPool's Integration with Kubernetes



StorPool has a native integration with Kubernetes through which persistent volumes are provided to the pods. StorPool's CSI plugin allows Kubernetes to create persistent volumes backed by storage, provided by the StorPool cluster. This fast and reliable storage can be used as the default for all data, not only databases or stateful applications/microservices.

The setup consists of a StorPool storage cluster, redundant network layer, and a number of bare metal Kubernetes nodes with StorPool client (initiator) installed. Each Kubernetes persistent volume is backed by a StorPool volume attached as a block device to the node where the pod, which requested it, is running.

StorPool's CSI plugin allows on-premise Kubernetes clusters to use StorPool as persistent storage. It supports dynamic provisioning. The integration is designed for bare-metal nodes. Further, StorPool's integration with Kubernetes provides persistent volumes that are stored in the StorPool cluster and can be dynamically attached/detached to different Kubernetes nodes as needed, and it is block-level storage. StorPool supports persistent volume claims in ReadWriteOnce or ReadOnlyMany, like all modern block device provisioners.

Features

The storage for containers must prevent data loss and ensure data is accessible wherever the containers are. StorPool provides a modern storage architecture designed for containers and a wide set of data management features.

- End-to-end data integrity
- Efficient Copy-on-write snapshots and clones
- High availability
- Multi-stack support
- Scalability
- Backups and Disaster Recovery
- Advanced Monitoring and Statistics

Kubernetes Persistent Storage Across Many IT Stacks

Managing new-age IT and modern infrastructure requires multi-stack support. While providing persistent volumes for Kubernetes, StorPool supports multiple technologies and cloud management systems. It can provide persistent shared storage from one storage system to several IT stacks. The list of supported cloud orchestration systems is one of the widest in the industry and includes OpenStack, VMware, OnApp, OpenNebula, CloudStack, and proprietary Cloud Management Systems.




Running modern applications and databases on high-performance storage with Kubernetes





StorPool
DISTRIBUTED STORAGE

Get in Touch

 +1 415 670 9320

+44 (0) 20 7097 8536

 info@storpool.com

sales@storpool.com

[Get Started](#)

[Book a Demo](#)

About StorPool

StorPool is the fastest software-defined block storage on the market, used by public and private cloud builders, enterprises, MSPs, SaaS, hosting and cloud providers. It comes as software, plus a fully-managed data storage service that transforms commodity hardware into a fast, highly available and scalable shared-storage system.

StorPool is the superior alternative to mid- and high-end SANs and All-Flash Arrays (AFA) for mid- and large-scale deployments (hundreds of terabytes to petabytes of storage).