



kubernetes

MAPR DATA FABRIC FOR KUBERNETES BENEFITS

Persist data for containerized applications

Scale data and performance as containers grow

Benefit from MapR tickets, for end-to-end security

Multi-tenant deployment and access

CONTACT US

Info@mapr.com

TRY MAPR

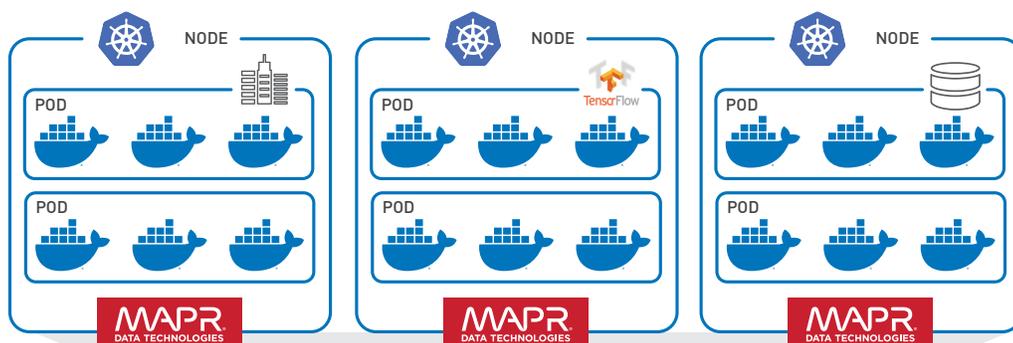
[Download](#)



MAPR DATA FABRIC FOR KUBERNETES

BUILD STATEFUL CONTAINER APPLICATIONS BY EXTENDING DATA FABRIC TO KUBERNETES AND CONTAINERS

MapR Data Fabric for Kubernetes provides persistent storage for containers and enables the deployment of stateful containers. It addresses the limitations of container use by providing easy and full data access from within and across clouds and on-premises deployments. Now stateful applications can easily be deployed in containers for production use cases, machine learning pipelines, and multi-tenant use cases.



MAPR CONVERGED DATA PLATFORM

- High Availability
- Real-Time
- Unified Security
- Multi-Tenancy
- Disaster Recovery
- Global Namespace

ON-PREMISES, MULTI-CLOUD, EDGE

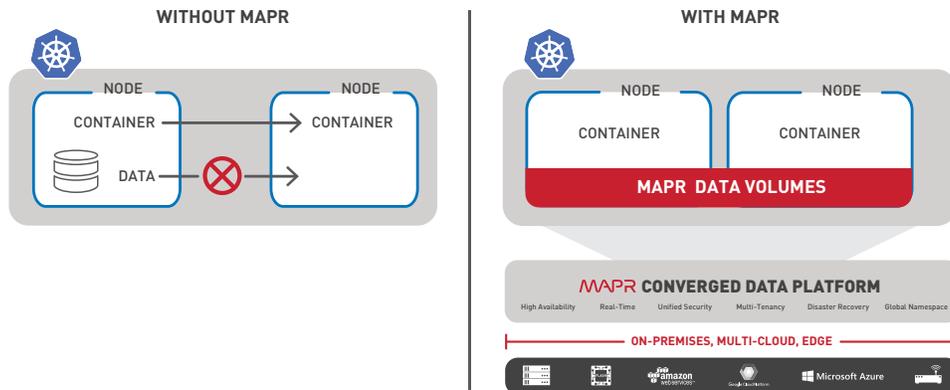


SOLUTION BENEFITS

As organizations adopt containers at a larger scale, and as they move containers as a deployment model to production, other aspects such as storage, monitoring, and performance becomes essential as well. MapR Data Fabric for Kubernetes provides several features that will assist organizations in their journey with containers.

PERSIST DATA

MapR integration with the Kubernetes storage plugin allows for MapR volumes to be mounted for access by containers. As new containers are deployed, data volumes can be created and retained, even when containers are deleted. If containers are moved across environments, the MapR global namespace provides access to the data, independent of where the containers reside.



Applications can be synchronized and updated with a unified view and without disruption. Data stored in MapR can benefit from the full-fledged advantages of high availability, security, data protection, and disaster recovery.

SCALE DATA AS YOU SCALE CONTAINERS

By using Kubernetes DaemonSet on MapR, containers can scale seamlessly without any disruption to the existing environment. With its “grow as you go” feature, MapR handles growth in data without it having to be moved to a separate, dedicated environment. Containers deployed in cloud or in edge clusters can easily gain access to data in MapR, and containers can be seamlessly added and moved across clusters.

HIGH PERFORMANCE

Containers running on client hosts can achieve high bandwidth access to data hosted on MapR. Container clusters managed by Kubernetes can provision and access data on any type of hardware infrastructure: SSD, NVMe, or HDD. As containers and data scale, performance scales linearly.

SECURE YOUR CONTAINER DATA

MapR tickets can be used to set up user authentication and authorization methods to access the persistent data. As more containers are added, each gets its own secure, persistent sandbox of compute, storage, and memory resources. This is an essential benefit, particularly for organizations that have multi-tenant environments.

HIGH AVAILABILITY

MapR distributed architecture ensures critical data is never lost, through tunable levels of replications and distributing copies across

the cluster. Automatic failover ensures data is always available without any disruption to containerized applications.

DATA PROTECTION

Enterprise data services, like mirroring and consistent point-in-time snapshots, protect critical data. Applications are becoming fluid because of containers, since containers can be moved or recreated across clusters. Recovering from major outages or having planned outages is made easier by migrating containers and mirroring/snapshotting data for containers.

ENABLING WORKLOADS WITH MAPR DATA FABRIC FOR KUBERNETES

MapR Data Fabric for Kubernetes, combined with the capabilities of the MapR Converged Data Platform, is beneficial to organizations desiring to run workloads in containers.

- **All enterprise application.** Essentially all enterprise applications can be run as containers on MapR. An exhaustive set of MapR capabilities, such as high availability, security, and high performance, enables organizations to deploy quickly.
- **Applications shared across multiple tenants.** Sharing applications and adding new users have less infrastructure implication on MapR.
- **Edge workloads.** The MapR small footprint edition enables applications running on edge clusters to be deployed in containers.
- **CI/CD.** MapR offers a consistent, reliable platform, integrated with Kubernetes for continuous development and integration.
- **Machine learning applications.** Portable and smart applications can be deployed using MapR Streams for real-time feed of data and machine learning pipelines.



For more information visit mapr.com

MapR and the MapR logo are registered trademarks of MapR and its subsidiaries in the United States and other countries. Other marks and brands may be claimed as the property of others. The product plans, specifications, and descriptions herein are provided for information only and subject to change without notice, and are provided without warranty of any kind, express or implied. Copyright © 2018 MapR Technologies, Inc.