

Arctera™ InfoScale for Kubernetes

Enterprise Resiliency and Storage
Management for Kubernetes.

Executive Summary

Containers have become a mainstream solution for building and running IT services that help businesses reduce management overhead and focus on delivering new innovations. Many organizations currently use or are considering Kubernetes to provide the advanced orchestration and management functionality needed to streamline container operations. However, Kubernetes does not provide all the tools needed to manage stateful containerized applications that require enterprise-grade persistent storage and resiliency.

Arctera InfoScale for Kubernetes provides software-defined persistent storage and resiliency functionality designed to deliver an enterprise foundation for stateful applications running on Kubernetes environments. Today, more than half (55 percent) of organizations are actively using containers, and another 18 percent are in the discovery stage¹. It enables you to run your stateful containerized applications in the cloud, with advanced features focused on three key principles:

- **Resiliency:** Protect your containerized applications from unplanned downtime
- **Mobility:** Help move your stateful applications and data across Kubernetes clusters for improved resiliency
- **Efficiency:** InfoScale for Kubernetes delivers an intuitive and simple operating experience that extends Kubernetes functionality

InfoScale for Kubernetes integrates seamlessly with Kubernetes environments, enabling you to easily deploy and manage persistent storage and resiliency for containerized applications and services. This helps ensure that your containerized applications have the foundation they need to be production ready. This document provides an overview of the features and functionality of InfoScale for Kubernetes.

Solution Value

InfoScale for Kubernetes augments and extends Kubernetes by providing persistent storage for stateful containerized applications, and cluster-level resiliency that protects you against local and site-wide outages. InfoScale extends Kubernetes environments with:

- **Resiliency:** Achieve a near-zero recovery time objective (RTO)/recovery point objective (RPO) for mission critical applications with real-time namespace replication between Kubernetes clusters. Integrated fencing technology ensures data integrity within your Kubernetes clusters and can protect your applications against unplanned downtime
- **Mobility:** With replication between Kubernetes clusters, you can easily move metadata and persistent data between environments, which helps eliminate lock-in and mitigates the effects of cloud service outages
- **Efficiency:** Easily configure high performance persistent shared storage with a container-native deployment architecture that optimizes operational efficiency and simplicity

InfoScale for Kubernetes is deployed as a containerized application in a Kubernetes cluster, making it easy to use—with minimum time to value—while also providing a native user experience for Kubernetes administrators and architects.

Solution Overview

As containers become more frequently deployed to serve a wider variety of use cases, Kubernetes environments will often include several types of containerized applications, each with their own storage and resiliency requirements. InfoScale for Kubernetes has several unique features that enable organizations to run important stateful applications in containers. Creating Kubernetes environments capable of supporting critical applications involves several challenges that cannot be resolved with native tools or services. InfoScale for Kubernetes integrates with native Kubernetes components and infrastructure to provide the resiliency and persistent storage your IT services need. Figure 1 shows an overview of InfoScale for Kubernetes.

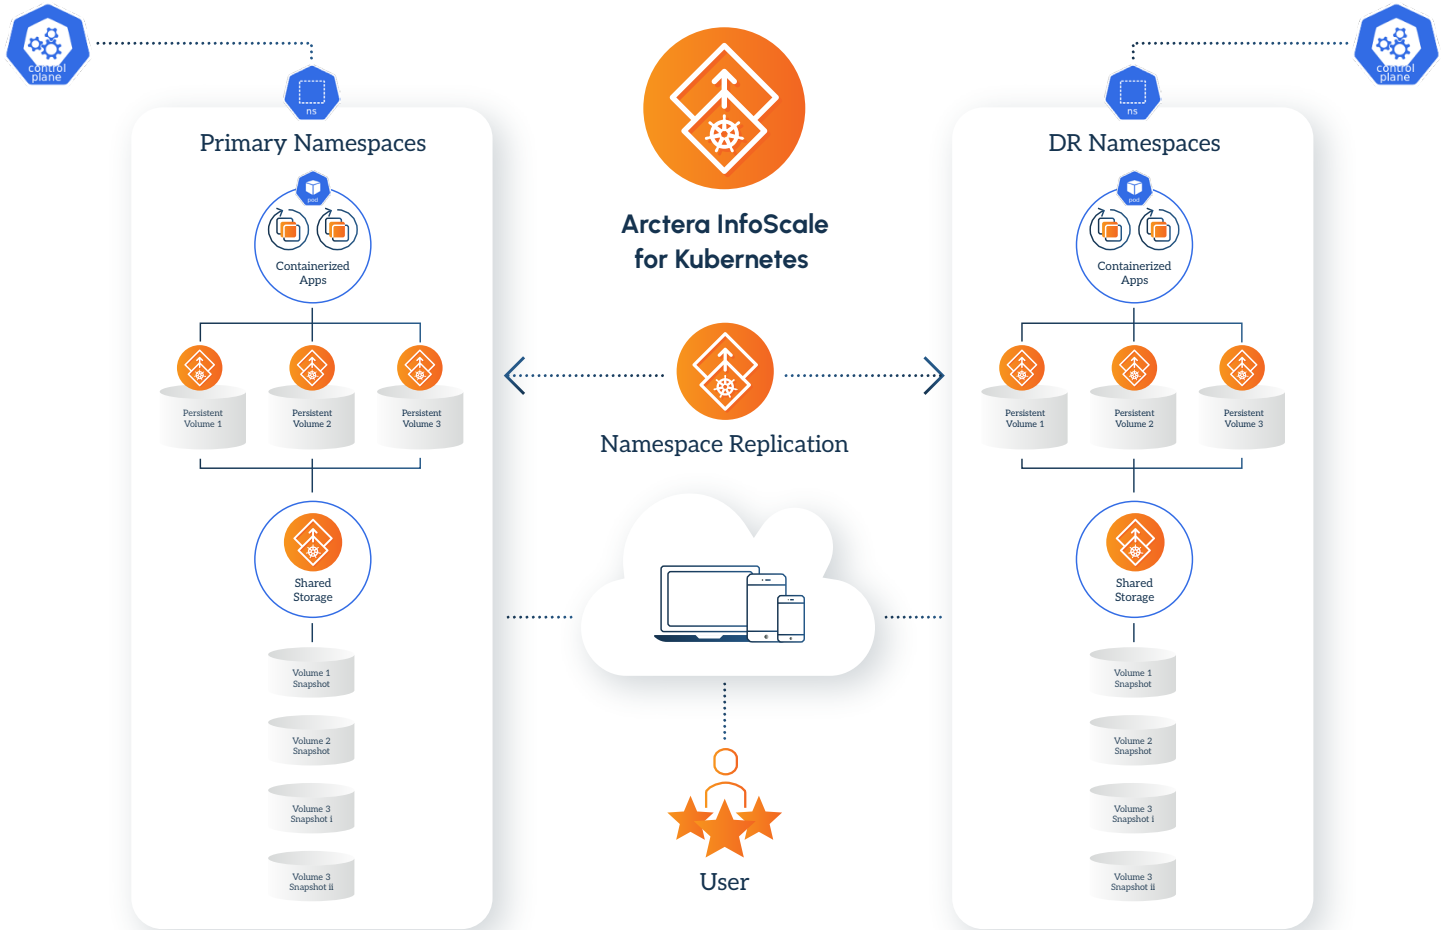


Figure 1. InfoScale for Kubernetes overview

Application Resiliency

Disaster Recovery (DR) for Kubernetes Clusters

InfoScale for Kubernetes includes a disaster recovery management solution that can be deployed as an operator-based feature. The DR manager enables you to migrate and take over Kubernetes clusters for planned and unplanned recovery operations. InfoScale for Kubernetes manages the resiliency process with optimized replication between clusters that maintains write-order fidelity, and replicates application state and cluster metadata.

InfoScale's DR manager for Kubernetes is designed to protect your applications against site-wide outages, where an entire cluster may go offline—in planned and unplanned scenarios. The DR manager enables:

- **Resiliency:** Ensure that your Kubernetes clusters are protected against site-wide outages with a full production-ready DR environment. An advanced feature called Fire Drill allows you to non-disruptively test and validate that the applications in your DR site can be recovered and brought online in the event of a primary site failure
- **Mobility:** Disaster recovery can be configured using plans that allow for namespace granularity, which gives you the flexibility to provide resiliency for one or more applications, or full Kubernetes clusters
- **Efficiency:** Support for synchronous and asynchronous replication with optimized data transfer between sites. You can also optimize both application uptime and operational efficiency, with the ability to perform maintenance and administrative operations non-disruptively if system downtime is required

Figure 2 shows an overview of how the InfoScale for Kubernetes DR manager can be configured to provide resiliency for Kubernetes namespaces and clusters.

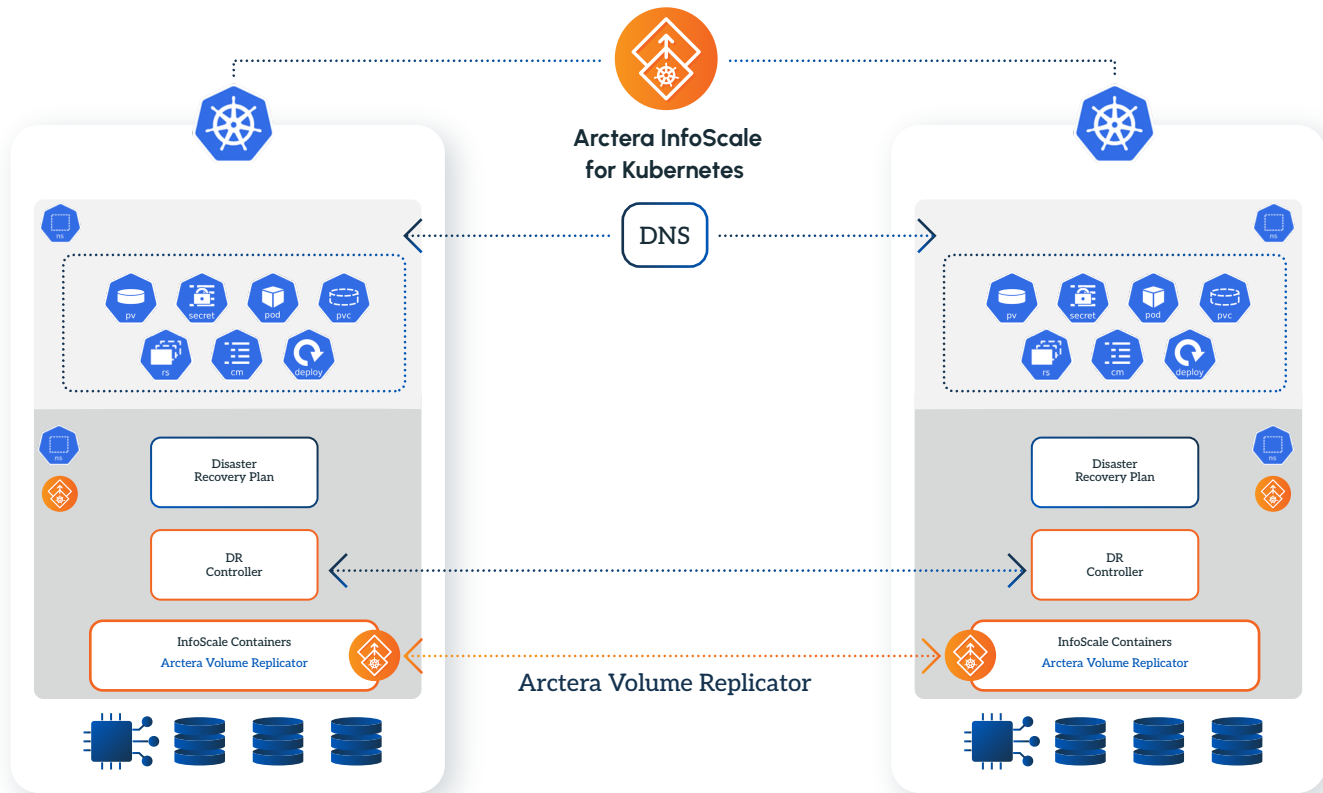


Figure 2. Arctera InfoScale DR manager for Kubernetes

Fencing Controller

With InfoScale for Kubernetes, containerized applications are automatically protected against data corruption due to a split-brain scenario. Split-brain can occur in any clustered environment in the event of a node/hardware failure that disrupts cluster communications and membership. InfoScale for Kubernetes provides I/O fencing capability using a driver container dedicated to managing the fencing process. This prevents data from being written by nodes within the Kubernetes cluster that have failed due to hardware or network communication problems. If a node failure is detected by Kubernetes, the InfoScale fencing driver can ensure the persistent volumes being used by application pods on the failed node are no longer accessible by fencing this node out of the cluster. In the event of a communication loss between cluster nodes (also known as worker nodes), the InfoScale fencing driver relays this information to the Kubernetes master which can then mark the node as failed and move pods to another node.




Figure 3. InfoScale I/O fencing in Kubernetes helps ensure data integrity

Persistent Storage

Container Storage Interface (CSI) Plug-In

InfoScale for Kubernetes provides software-defined storage for Kubernetes environments using a CSI plug-in. This allows you to create persistent InfoScale storage volumes that are managed by the integrated cluster volume manager and cluster file system. This storage is then presented to pods as persistent storage for your stateful containerized applications running in Kubernetes. InfoScale is transparent to the application, as the application pod is only aware that it is writing to a persistent volume claim (PVC) presented by Kubernetes. The CSI plug-in workflow is shown in Figure 4.

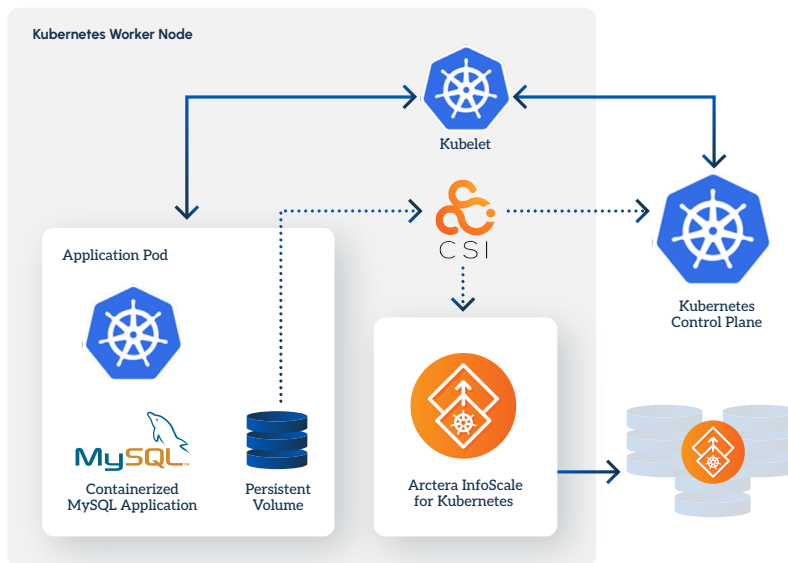


Figure 4. InfoScale for Kubernetes CSI plug-in overview

Persistent Volumes

Kubernetes storage classes are used to manage the attributes of InfoScale persistent volumes that are mounted by Kubernetes inside application pods using the InfoScale CSI plug-in. InfoScale provides several storage class configuration options that can be used to create persistent storage volumes. Storage classes are defined for performance, resiliency, and security, and can be customized to meet application requirements.

InfoScale persistent volumes are provisioned by Kubernetes using the CSI plug-in either dynamically or statically:

- **Dynamic provisioning:** Volumes are created at the same time as the containers and application pod using Kubernetes with the InfoScale CSI plug-in and InfoScale container installed; an InfoScale persistent volume claim binds the storage accessible to the application pods to the InfoScale persistent volume that is available to Kubernetes cluster nodes
- **Static provisioning:** Volumes are created within the Kubernetes cluster by directly accessing the InfoScale containers, creating a volume, and exposing this volume to containerized applications within the cluster; InfoScale statically provisioned volumes can also be used to simplify the process of migrating traditional applications into containers; application data volumes outside a Kubernetes cluster can be migrated to volumes within the cluster and then managed by InfoScale within the cluster using the InfoScale CSI plugin; Arctera support should be contacted to help validate requirements and environmental variables involved in the migration process

Figure 5 shows how InfoScale persistent volumes and persistent volume snapshots are provisioned and used by application pods in a Kubernetes cluster.

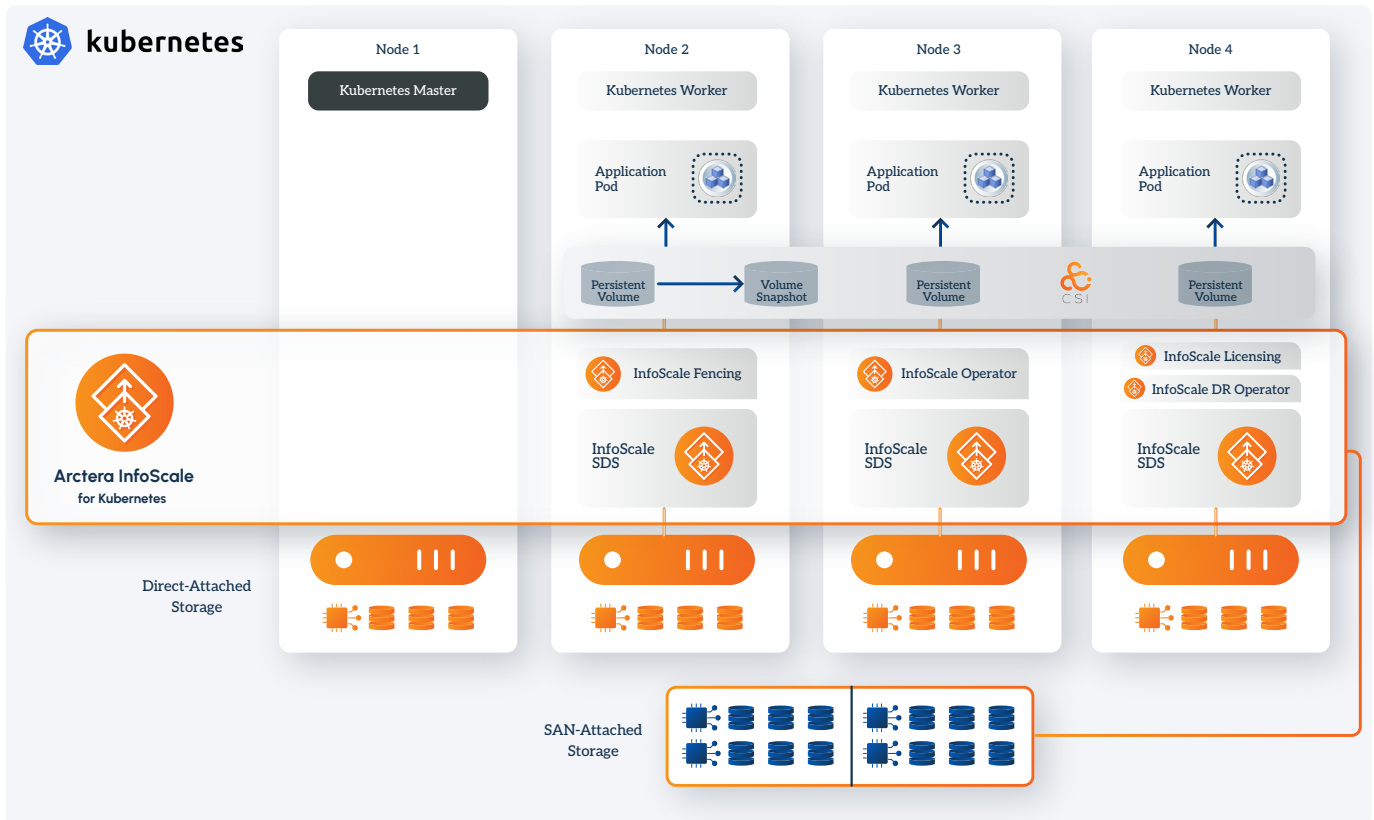


Figure 5. InfoScale persistent volumes and persistent volume snapshots

Storage Management

InfoScale for Kubernetes includes advanced storage management features designed to support stateful enterprise workloads running in Kubernetes environments.

- Data access: InfoScale supports multiple access modes: ReadWriteOnce (RWO), ReadOnceMany (ROX), and ReadWriteMany (RWX). This allows you to provision storage more granularly, based on your application's requirements and operational standards
- Encryption: An encrypted storage class is available for workloads that require additional data security
- Snapshots: Create space-optimized copies of persistent volumes that can be provisioned statically or dynamically. Snapshots can be used to reinstate volume contents on the original volume or on a new persistent volume that you provision
- Volume cloning: Create an exact duplicate of a specified existing persistent volume. Volume clones can be used by any persistent volume claim and can be deleted without affecting the original volume

Figure 6 shows an overview of different use cases where InfoScale persistent volume snapshots are used to support multiple functions and services within an organization running stateful applications in Kubernetes.

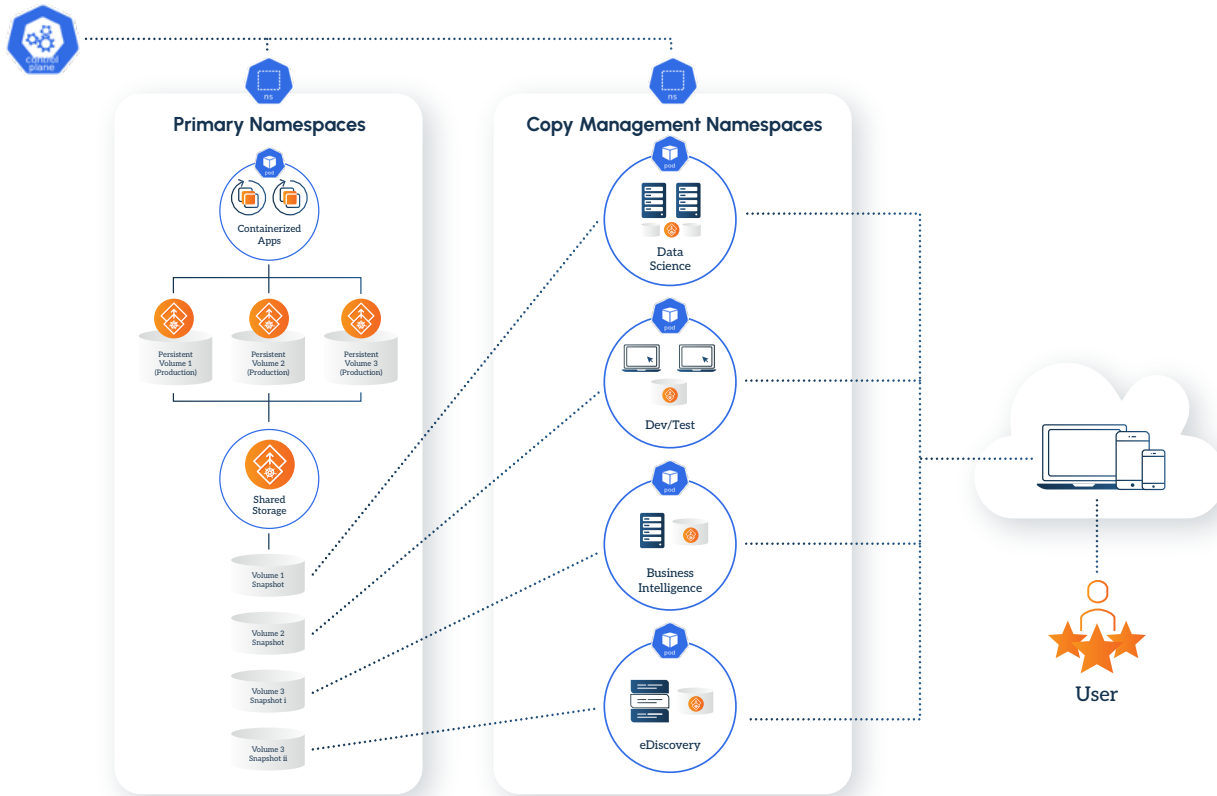


Figure 6. InfoScale for Kubernetes Shared Storage persistent volumes and persistent volume snapshots used for multiple purposes

Conclusion

Digital transformation has driven enterprises toward IT solutions that provide better efficiency and scalability for their IT services. Containerizing applications and managing them with Kubernetes can help businesses deliver more efficient innovation, but this operating model lacks some of the enterprise functionality needed by most applications. Arctera solves this problem by providing an enterprise-focused foundation with storage management and resiliency for Kubernetes that can operate at scale. This unique solution has several key benefits:

- ✓ **Resiliency:** Software-defined persistent storage for Kubernetes with advanced features designed to support enterprise IT service deployments
- ✓ **Mobility:** Enterprise-grade resiliency architecture that provides full disaster recovery and resiliency for Kubernetes clusters
- ✓ **Efficiency:** Optimize operations by eliminating point products and reducing overhead with a storage and resiliency solution that deploys as a containerized application in Kubernetes

Arctera empowers businesses to deploy business-critical stateful applications in Kubernetes. Designed for ease of use, flexibility, and scalability to support large container deployments, InfoScale for Kubernetes delivers an enterprise solution that provides the tools you need to run your applications in containers with maximum confidence.

1. 451 Research's Voice of the Enterprise: DevOps,H1 2020

About Arctera

Arctera helps organizations around the world thrive by ensuring they can trust, access, and illuminate their data from creation to retirement. Created in 2024 from Veritas Technologies, an industry leader in secure multi-cloud data resiliency, Arctera comprises three business units: Data Compliance, Data Protection, and Data Resiliency. Arctera provides more than 75,000 customers worldwide with market-leading solutions that help them to manage one of their most valuable assets: data. Learn more at www.arctera.io. Follow us on X @arcteraio.



arctera.io

For global contact information visit: arctera.io/contact