


Veritas InfoScale for Kubernetes

Storage Management and Data
Services for Containers



Contents

Introduction	3
InfoScale Enterprise Solution Value	3
Veritas InfoScale Storage Architecture	4
Veritas InfoScale Storage	5
Veritas File System	6
Veritas Volume Manager	6
Cluster File System	7
Cluster Volume Manager	7
Kubernetes Container Storage Interface (CSI) Plugin.	7
I/O Fencing	8
SmartIO and FSS.	8
Veritas Volume Replicator (VVR)	9
InfoScale for Containers/Kubernetes (K8s)10
Storage Management11
InfoScale Persistent Volumes11
I/O Fencing12
Red Hat OpenShift Container Platform (OCP) Integration13
Conclusion13
References.14
Table of Figures15
Table of Tables15

Introduction

Veritas Technologies is an industry leader in developing data resiliency and availability solutions that focus on the protection and management of digital assets critical for a company's success and business continuity. One of our flagship products, InfoScale, enables enterprise-class software-defined storage management along with high availability and disaster recovery (HADR) for all data centers, including on-premises, hybrid and multi-cloud. InfoScale works with applications running on physical, virtual and cloud platforms, with support in version 8.0 extended to include applications running in Kubernetes environments. InfoScale is a comprehensive, industry-proven solution that helps users manage enterprise readiness for modern mission-critical applications, focused on 3 key principles:

- **Application Performance:** eliminate overhead and complexity by using InfoScale to build a high-performance software-defined storage environment using any combination of commodity and specialized storage hardware that intelligently manages data so that it is always available on the best storage tier. This improves application performance and maximizes resource utilization while significantly reducing infrastructure overhead, costs and complexity.
- **Application Agility:** enable applications to be managed for the highest performance and uptime across different operating systems and platforms that may span multiple geographic regions. InfoScale helps users avoid vendor lock-in by enabling applications to run in a highly available configuration on any OS and any platform, including public cloud and hybrid cloud configurations. InfoScale also provides automated cloud migration between on-premises systems and between different cloud providers, including both public and private clouds.
- **Application Availability:** integrate directly with native application and infrastructure components to ensure that the underlying infrastructure, applications, and data are managed to provide the highest possible performance and uptime. InfoScale can also provide near-instantaneous automated recovery that minimizes the impact of service disruptions and outages.

This technical overview will assist in understanding the Veritas InfoScale solution for Containers. InfoScale enables organizations to leverage highly performant and scalable software-defined storage to achieve maximum application uptime and performance for containerized environments in a Kubernetes environment.

InfoScale Enterprise Solution Value

InfoScale has several unique features that offer significant value to customers looking to improve application performance and reduce costs by maximizing flexibility of architecture with a software-defined approach to data management in Kubernetes Environments. Creating an environment that's capable of supporting your most mission critical applications presents some challenges that cannot be completely resolved with native tools. InfoScale is designed to integrate with nearly any IT infrastructure to provide resiliency and high performance for containerized applications.

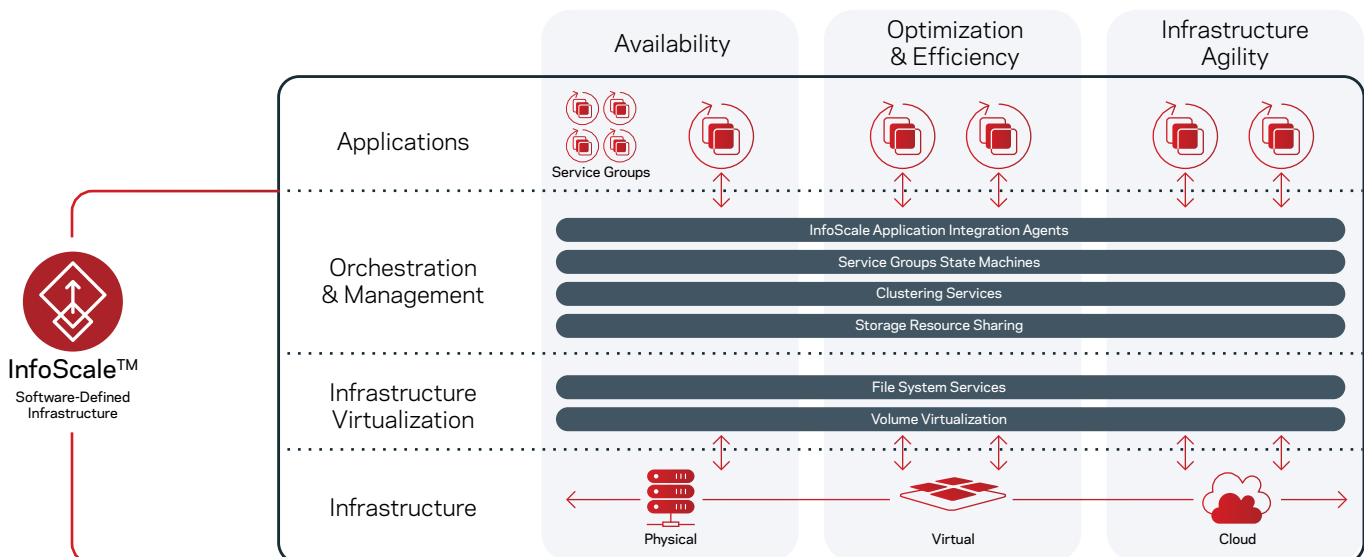


Figure 1. InfoScale Solution Value

InfoScale Enterprise has customized agents designed for mission critical applications that understand and manage the application components and resources not normally monitored by native system tools. This ensures that the application has the highest possible uptime and either meets or exceeds performance expectations. InfoScale can also intelligently manage the HADR process and nearly eliminate the need for manual user intervention for applications to be successfully brought online or recovered in a DR scenario. There are several other benefits of using InfoScale to manage both application storage and the overall HADR solution:

- **Data Availability and Management:** leveraging industry-leading Dynamic MultiPathing (DMP) along with snapshot and checkpoint features ensures that applications can access vital data. Additionally, I/O fencing and I/O shipping technologies protect that critical data from both hardware and software failures in the underlying infrastructure. All of this software-defined storage can be built on a variety of attached storage – be it direct-attach, SAN, or storage on the local nodes themselves.
- **Replication management:** a component of InfoScale, Veritas Volume Replicator (VVR) has the capability of providing a near-zero RTO/RPO for mission critical data that can scale to support the largest workloads. With VVR, there are some additional benefits such as maintaining write order fidelity, multi-target support, heterogeneous system configuration (including public cloud) and zero data loss, and the use of commodity storage hardware.
- **Kubernetes integration:** With the 8.0 release, InfoScale now provides cloud native storage and data services in a Kubernetes environment. By deploying containerized InfoScale alongside a CSI plug-in, InfoScale’s industry-leading data services can be leveraged by containerized applications. And, in Kubernetes environments, VVR supports replication of the container and cluster configuration data as well, enabling disaster recovery of applications.

Veritas InfoScale Storage Architecture

Veritas InfoScale is a fully software-defined solution that has an extensive compatibility matrix and is operating system and platform agnostic. It can be deployed on industry standard hardware and can support a wide range of enterprise applications regardless of their underlying infrastructure. InfoScale is composed of several different products mostly delineated by licensing scheme. At its core, InfoScale has two distinct parts: InfoScale Availability and InfoScale Storage. (Figure 2) As InfoScale Availability is not completely available in containerized environments, we will be discussing that only briefly in this paper. In addition, a new release is InfoScale Kubernetes Edition (or IKE), which is focused on delivering InfoScale features and services for the Kubernetes platform.

InfoScale Storage is the Software Defined Storage group of services and features. It enables users to build an enterprise-class storage solution to manage and protect any type of enterprise data and can be scaled up or down to accommodate any environment or budget. InfoScale Storage consists of the components in table 1.

Features	Developer Edition	Storage Edition	Enterprise Edition
Overview	Free non-production use with a limited number of nodes (3 nodes) for single K8s cluster and all storage edition features	All features included in InfoScale storage SKU for a single K8s cluster (either on-prem or cloud) offered as a subscription model	All features included in InfoScale storage SKU for a single or multi cluster plus advanced features like Disaster Recovery (DR) for on-prem, cloud and hybrid models offered as a subscription model
Cluster Nodes	3	Not limited	Not limited
Resiliency Features	DR features not available	DR features not available	DR features with cluster and application-level resiliency
Number of clusters	1	Any number of single K8s cluster configurations	Multi-cluster with resiliency
Support	Limited support	Full support and software upgrades included with subscription	Full support and software upgrades included with subscription

Component	Description
Veritas File System	A POSIX-compliant extent-based file system known as VxFS that runs on most Unix and Linux variant operating systems
Veritas Volume Manager	A storage management utility known as VxVM that manages physical disks as logical data volumes that are presented to an operating system as a physical device on which file systems can be created
Veritas Volume Replicator	A software-based data replication utility known as VVR that enables consistent block-level data replication between a VxVM managed source data volume and one or more remote data volumes
Cluster File System	A VxFS file system running in parallel access mode which allows for multiple systems (cluster nodes) to access the same file system data simultaneously
Veritas File Replicator	A periodic file level replication utility that tracks and replicates changed files in a VxFS file system. It uses file system checkpoints to read and replicate changed files between a source checkpoint and one or more target checkpoints, with no impact on the application.

Table 1. InfoScale Storage Components

Veritas InfoScale Storage

Veritas InfoScale Storage consists of several main components that provide the basis for building a highly available storage infrastructure for the Red Hat OpenShift platform:

1. Veritas File System (VxFS) is a POSIX-compliant journaling enterprise file system designed to maximize application performance.
2. Veritas Volume Manager (VxVM) is the storage management subsystem that allows you to create, organize, and manage logical data volumes and disk groups used by applications and databases and the underlying physical disks and logical unit numbers (LUNs).
3. Cluster Volume Manager (CVM) is the add-on to VxVM's storage virtualization layer that enables the storage devices to be accessible by multiple nodes in the cluster simultaneously.
4. Cluster File System (CFS) is the extension to VxFS that creates file systems that can be shared by multiple nodes in the cluster concurrently.
5. Flexible Storage Sharing (FSS) enables individual nodes to share direct-attached storage with other nodes in the cluster at the physical disk level. FSS is a feature of CVM and CFS that allows a file system to be built on top of a volume shared with FSS.
6. Veritas Volume Replicator (VVR) enables optimized replication of data between InfoScale managed data volumes.

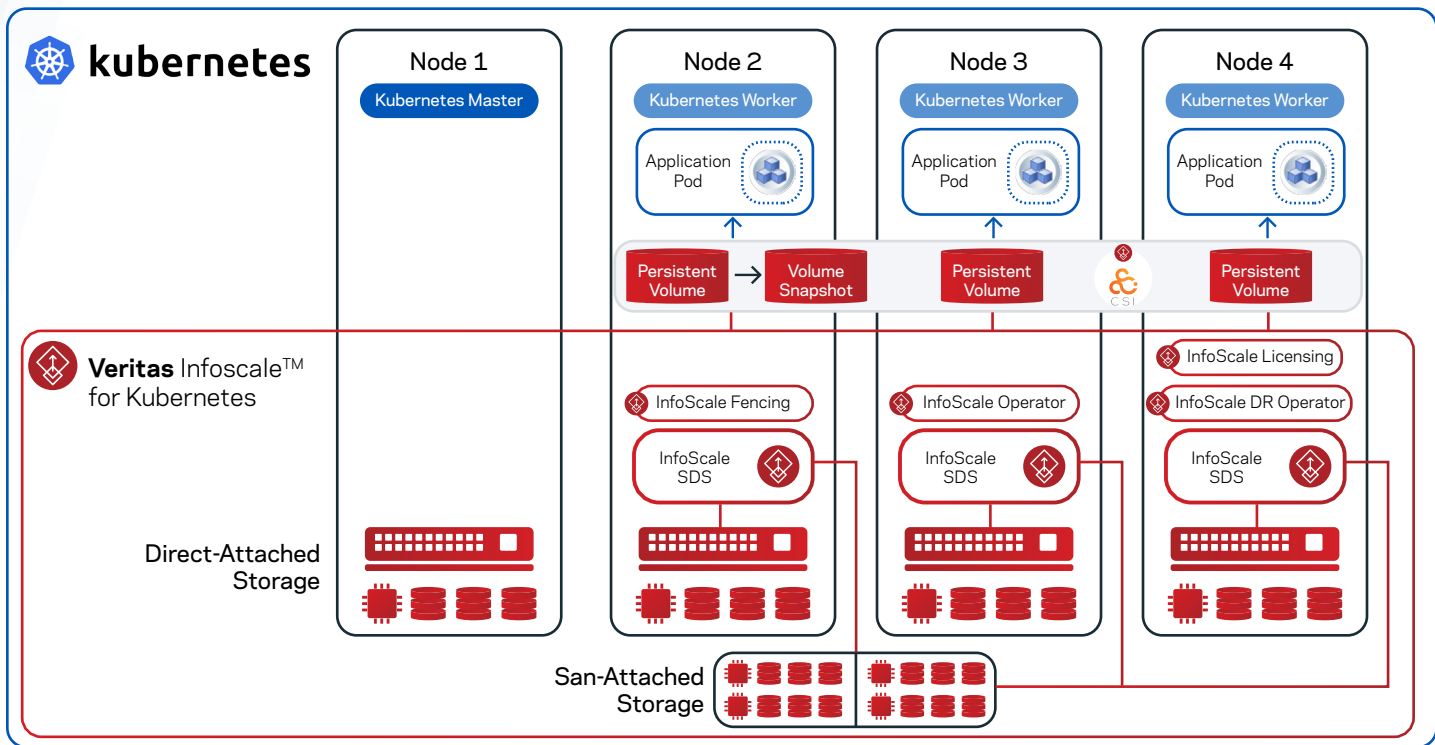


Figure 3. Cluster File System and Flexible Storage Sharing

Figure 3 shows the InfoScale on Kubernetes Architecture, including the operator pods and the varying types of storage that can be used to create the Cluster File System used to deliver storage to the cluster.

Veritas File System

The Veritas File System, known as VxFS, is an extent-based POSIX-compliant journaling file system capable of managing large volumes of data that is designed to provide high performance and availability for applications. VxFS also has several advanced features that maximize application performance while optimizing the data footprint and supports large file systems of up to 16 petabytes.

Veritas Volume Manager

The Veritas Volume Manager, known as VxVM, is a software-based data management utility that provides online disk storage management for compute environments and Storage Area Network (SAN) environments. With support for the Redundant Array of Independent Disks (RAID) model as well as erasure coding, VxVM can be configured to protect against disk and hardware failure, and to increase I/O throughput.

VxVM overcomes restrictions imposed by hardware disk devices and LUNs by providing a logical software-defined volume management layer. After first creating disk groups, which consist of all of the physical disks or LUNs that will be used for a single purpose, you can create volumes that span multiple disks and LUNs. The disk group is the atomic unit in which VxVM instances import (gain access to), deport (relinquish access to), activate (present to VxFS) and deactivate (withdraw accessibility to) disks. VxVM maintains a redundant, persistent record of each disk group's membership, volumes and other underlying structures in dedicated private regions of storage on the disks in a disk group. VxVM provides the tools to improve performance and ensure data availability and integrity. You can also use VxVM to dynamically configure storage while the system is active.

VxVM has several other advanced features for data management, such as FlashSnap for optimized point-in-time data copies and Portable Data Containers that enable the migration of data between different OS platforms.

Cluster File System

Veritas Cluster File System, or CFS, enables you to concurrently mount the same file system on multiple nodes and is an extension of the industry standard VxFS. Unlike other file systems which send data through another node to the underlying storage, CFS is a true Storage Area Network (SAN) file system. All data traffic happens over the SAN, and only the metadata traverses the cluster interconnect.

CFS uses a distributed locking mechanism called Global Lock Manager (GLM) to ensure all nodes have a consistent view of the file system. GLM provides metadata and cache coherency across multiple nodes by coordinating access to file system metadata, such as inodes and free lists. The role of GLM is set on a per-file system basis to enable load balancing.

Cluster Volume Manager

In general terms, a Volume is a unit of storage carved out of a physical disk device. CVM presents a consistent volume state across an InfoScale cluster as nodes import and access volumes concurrently. It also enables all nodes in a cluster to access their underlying storage devices concurrently. CVM transforms the read and write requests that CFS addresses to volume blocks into I/O commands that it issues to the underlying disks.

The primary difference between CVM and VxVM is that CVM allows disk groups to be imported on all the systems in the cluster concurrently, whereas VxVM only allows a disk group to be imported on a single node at a time.

All CVM instances in a cluster must always present the same view of disk group and volume configuration, even in the event of:

- Storage device failure. For example, if a disk is added to or removed from a mirrored volume, all CVM instances must effect the change and adjust their I/O algorithms at the same logical instant
- Cluster node failure. If a cluster node fails while it is updating one or more mirrored volumes, CVM instances on the surviving nodes must become aware of the failure promptly, so that they can cooperate to restore volume integrity

CVM always guarantees that all instances in a cluster have the same view of shared volumes, including their names, capacities, access paths and “geometries”. Most importantly, CVM also manages volume states. For example: whether the volume is online, the number of operational mirrors, whether mirror resynchronization is in progress, and so forth. A volume’s state may change if a device fails or a node fails, or an administrative command is issued.

Kubernetes Container Storage Interface (CSI) Plugin

InfoScale provides software-defined storage for Kubernetes environments through the use of the InfoScale Container Storage Interface (CSI) plugin. This powerful plugin allows you to use InfoScale volumes created on any combination of locally-attached commodity storage and SAN-attached storage, and managed via CVM and CFS as persistent storage for your stateful containerized applications running in Kubernetes. This storage management 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 Plugin is shown in Figure 4.

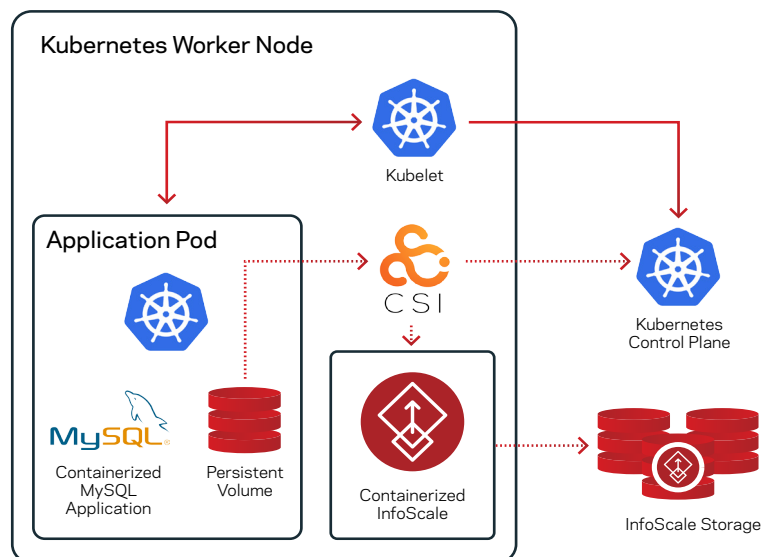


Figure 4 InfoScale CSI Plugin

I/O Fencing

A condition known as the “split-brain” condition occurs when there is communication disruption between cluster nodes. This can result in data corruption due to the fact that InfoScale (and other cluster software) cannot always distinguish between a system failure and an interconnect failure. The “split-brain” condition can also occur if a node within the cluster is so busy that it appears to be hung and pauses communication with the other cluster nodes. The split-brain condition can occur in all clustered storage implementations, including K8s. To mitigate and resolve the split-brain condition, Veritas InfoScale implements an I/O Fencing system that guarantees data integrity by determining which nodes in the cluster should remain in the event of a communication disruption. When a disruption occurs, the node (or nodes) which has failed is ejected from the cluster and prevented from accessing the data disks. One important thing to remember in a containerized environment is that all of this happens at the InfoScale layer, below the containerized infrastructure – the containerized applications are not even aware that this is happening, as their storage continues uninterrupted, unless they’re on the node that suffers a failure of some sort.

The key to protecting data in a shared storage environment is guaranteeing that there is always just a single consistent view of cluster membership. In other words, when one or more systems stop sending heartbeats, the InfoScale software must determine which nodes can continue to participate in the cluster membership and how to handle the other nodes.

There are three I/O Fencing options available in Veritas InfoScale:

1. Disk-based I/O fencing – members of a cluster notify other nodes in the cluster that they are still present by registering themselves to special “coordinator disks.” If there is a failure, the coordinator disks ensure that the surviving nodes permit write operations to data disks. Disk-based I/O fencing requires SCSI-3 Persistent Reservation compatible disk devices.
2. Server-based I/O fencing – A special coordination point server maintains a registry of node membership in a cluster instead of disks, like in disk-based I/O fencing.
3. Majority-based I/O fencing – This is used when coordinator disks or coordination point servers are not available. When a network partition happens, one node in each sub-cluster is elected as the racer node, while the other nodes are designated as spectator nodes. The sub-cluster with the majority number of nodes survives while nodes in the rest of the sub-clusters are taken offline.

For more information on IO Fencing behavior, see “How I/O fencing works in different event scenarios¹.”

Flexible Storage Sharing (FSS)

As an advanced feature of CVM and CFS, Flexible Storage Sharing (FSS) enables the network sharing of local storage across a cluster. The local storage can be in the form of either Direct-Attached Storage (DAS) or internal disk drives and can be either spinning disk or solid-state disk. Network sharing of storage is enabled by using a network interconnect between the nodes of a cluster and leveraging CVM and CFS to allow multiple systems to access the disks simultaneously, and this is a real strength when dealing with containerized environments. You can leverage the local storage on multiple container nodes to provide shared storage for the entire containerized cluster environment without having to spend money on SAN or NAS.

FSS allows network shared storage to co-exist with physically shared storage, and logical volumes can be created by CVM using both types of storage, enabling a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities, without requiring physically shared storage. This is transparent to file systems and applications. In a containerized environment, this enables persistent volumes to be created on this shared storage (which can be spread across multiple nodes, and isn’t necessarily SAN storage) and presented to all the nodes in the cluster simultaneously.

Veritas Volume Replicator (VVR)

VVR is a comprehensive solution for platform-independent replication of data volumes managed by VxVM or CVM. VVR enables cost-effective replication of data over IP networks. This provides organizations with an extremely flexible, storage hardware independent alternative to traditional array-based replication solutions as well as a robust mechanism for moving data into public cloud environments. VVR provides the flexibility of block-based continuous replication as well as file-based periodic replication with Veritas File Replicator (VFR).

VVR also manages and maintains write-order fidelity. This is an important feature that ensures data consistency which is not typically managed by native storage replication solutions. Write-order fidelity means that VVR tracks writes on the primary volume in the order in which they are received and applies them on the secondary volume in the same order. Maintaining write order fidelity ensures that the data on the secondary volume is consistent with the data on the primary volume. Data at the secondary volume can be behind in time (async replication), but VVR ensures that it is a consistent image of the primary RVG at a point in the past.

VVR is a component of InfoScale Storage and is included with InfoScale Enterprise. Its primary use cases are:

- Replication between like or heterogeneous on-premises systems
- Cloud data mobility (data replication from on-prem to cloud, cloud to on-prem, cloud to cloud)
- Data migration from legacy systems to containerized application systems

VVR can work independently or as part of an InfoScale Storage or InfoScale Enterprise cluster. In scenarios where InfoScale is managing storage and HADR for a cloud environment, VVR can manage the data replication. VVR supports replication of data stored in cloud volumes, regardless of the storage type, to other public cloud environments or to a customer managed on-premises storage volume.

VVR has an advanced feature called Adaptive Sync that improves sustained throughput for latency-sensitive applications by automatically switching from synchronous to asynchronous mode and vice versa based on cross-site latency. Adaptive Sync enables the configuration of time-outs for IO so that if any IO duration exceeds the time-out, an acknowledgment that the write operation has completed is returned immediately regardless of whether the write operation is completed at the remote site. When VVR is in asynchronous mode, it detects when latency returns to normal and automatically switches replication back to synchronous mode. The Adaptive Sync functionality enables your applications to withstand higher IO latencies while preserving the ability to maintain remote copies of your application data with a low RPO.

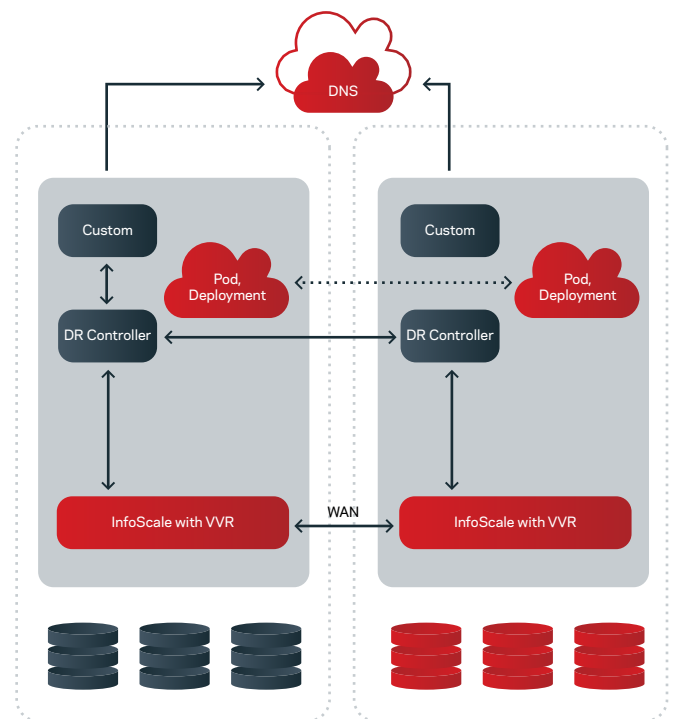


Figure 6

In a container or Kubernetes environment, disaster recovery planning requires that the metadata not just for the containers, but also for the container infrastructure be replicated so that the application comes up cleanly. As a solution for that problem, InfoScale DR operator along with VVR replicates not only the data volumes required for applications to run correctly but also the metadata for the containerized environment itself. This allows for rapid recovery of containerized applications in the event of a site failover. The figure shows this replication mechanism.

InfoScale for Containers/Kubernetes (K8s)

Containers have become a popular solution for efficiently developing and running applications. Containers provide excellent application portability for moving applications between environments and they also help you improve efficiencies by making it easy to standardize the resources your applications require. However, containers and container orchestration engines like Kubernetes do not natively provide all the functionality needed to manage stateful and mission-critical applications.

InfoScale's Container Storage Interface (CSI) plug-in works with Kubernetes to provide advanced storage management for containerized applications. InfoScale storage services provide the functionality needed by stateful applications running in containers that is not available natively in a Kubernetes environment. InfoScale's enterprise functionality integrates with Kubernetes to provide a container management platform suitable for running stateful and mission-critical applications that require:

- **Advanced storage management:** InfoScale's Container Storage Interface (CSI) plug-in allows Kubernetes to provide InfoScale persistent storage volumes to containerized applications being managed within a Kubernetes namespace. InfoScale's Flexible Storage Sharing can be used to provide high performance storage using the disks directly attached to the Kubernetes cluster nodes.
- **Application and system availability:** InfoScale's advanced I/O Fencing capabilities enable organizations to bring failed nodes and/or application pods back online quickly without worrying that data has become corrupted. By quickly removing access to shared data from failed nodes or pods, InfoScale prevents the corruption of data that may be in use by other nodes or pods. This enables the applications to remain online and servicing requests while administrators work to repair the failures.

InfoScale, beginning with version 8.0, is now deployed directly as a cloud-native, containerized application in its own pod on each Kubernetes cluster node. The InfoScale CSI plug-in provides the interface between the application pods that require storage and the InfoScale-managed storage, and there are other services responsible for managing I/O fencing, Disaster Recovery (if configured), licensing of InfoScale itself, and to keep InfoScale up-to-date. Figure 9 provides an overview of how InfoScale integrates with Kubernetes and containers to provide persistent storage for containerized applications.

Enterprise Storage Management for OpenShift powered by InfoScale for Kubernetes

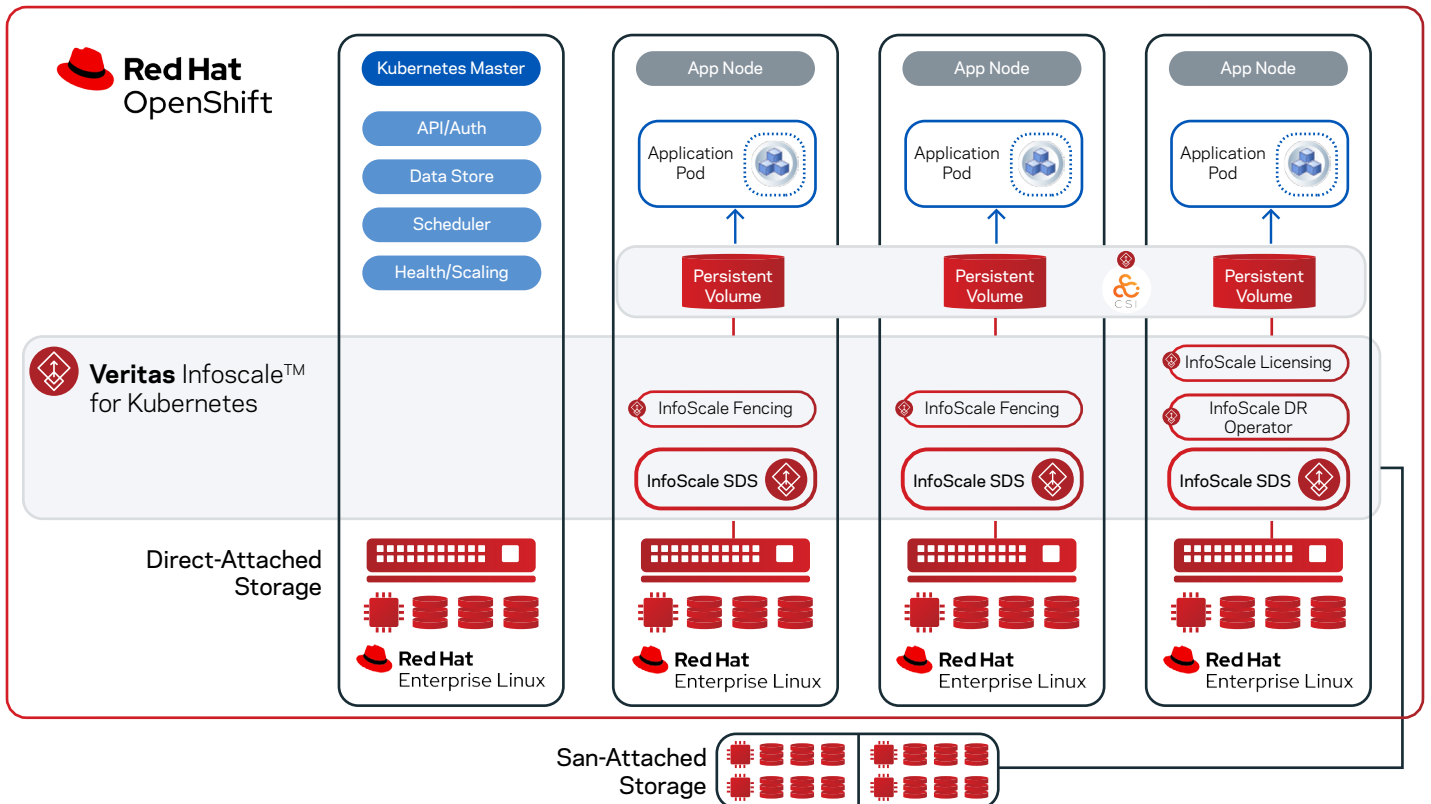


Figure 7 InfoScale running in a Kubernetes environment

Storage Management

The InfoScale CSI plug-in allows Kubernetes to mount InfoScale storage volumes formatted with VxFS inside the containers being managed, which can then be used by stateful applications without the risk of data loss if a container is powered off or removed from the pod. Kubernetes uses the InfoScale CSI plug-in to interface with the InfoScale storage volumes which are created on the Kubernetes cluster nodes. A storage request is made by Kubernetes using a persistent volume claim (PVC) described in a yaml file. Once a persistent volume is created, it is bound to the PVC and made available to the application in the container.

InfoScale Persistent Volumes

Kubernetes Storage Classes are used to manage the attributes of InfoScale storage volumes that are mounted by Kubernetes inside application pods using the InfoScale CSI plug-in. InfoScale provides several different 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 storage volumes are provisioned by Kubernetes using the InfoScale CSI plug-in either dynamically or statically.

- **Dynamic provisioning:** volumes are created at the same time as the container 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 an InfoScale 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. Veritas support should be contacted to help validate requirements and environmental variables involved in the migration process.

The InfoScale CSI plug-in enables Kubernetes to use InfoScale to create persistent volume snapshots that can be used by containerized applications for multiple purposes such as data protection, resiliency and analytics. Snapshots are used as static persistent volumes and the snapshot creation process has no impact on production data. In addition to persistent volume snapshots, you can also create persistent volume clones that help protect containerized applications against hardware failures.

I/O Fencing

Containerized applications with shared storage provided by InfoScale 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, after installing a container dedicated to managing the fencing process, provides advanced I/O fencing by preventing data from being written by nodes within the Kubernetes cluster that have failed due to hardware or network communication failures. 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. This prevents data corruption by allowing only the working nodes to continue normal operations. In the event of a communication loss between cluster nodes (aka worker nodes), InfoScale's fencing driver relays this information to the Kubernetes master which can then mark the node as failed and move pods to another node.

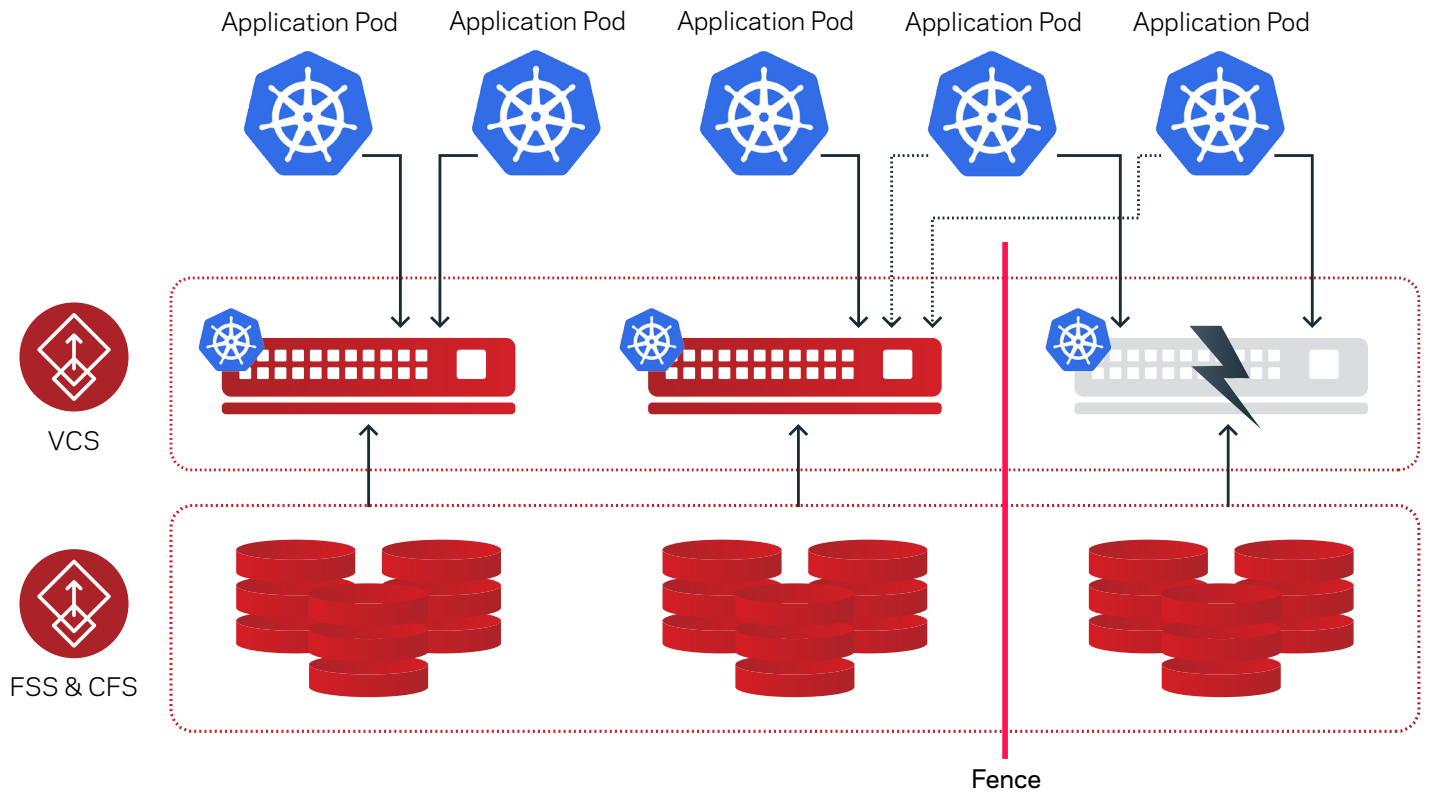


Figure 8. I/O fencing in Kubernetes using InfoScale

Red Hat OpenShift Container Platform (OCP) Integration

Starting with version 8.0, InfoScale can be used within the Red Hat OpenShift Container Platform (OCP) to provide persistent software-defined storage using local disk or SAN. Leveraging the OpenShift platform allows customers to take advantage of a fully-integrated, enterprise-ready Kubernetes deployment architecture backed by Red Hat's customer support and experience.

InfoScale is a certified RedHat operator and available as an Operator bundle on Red Hat Catalog and OperatorHub.io . This will enable developers to deploy InfoScale with a single-click native deployment and provision static or dynamic storage based on application needs from the Red Hat OpenShift console.

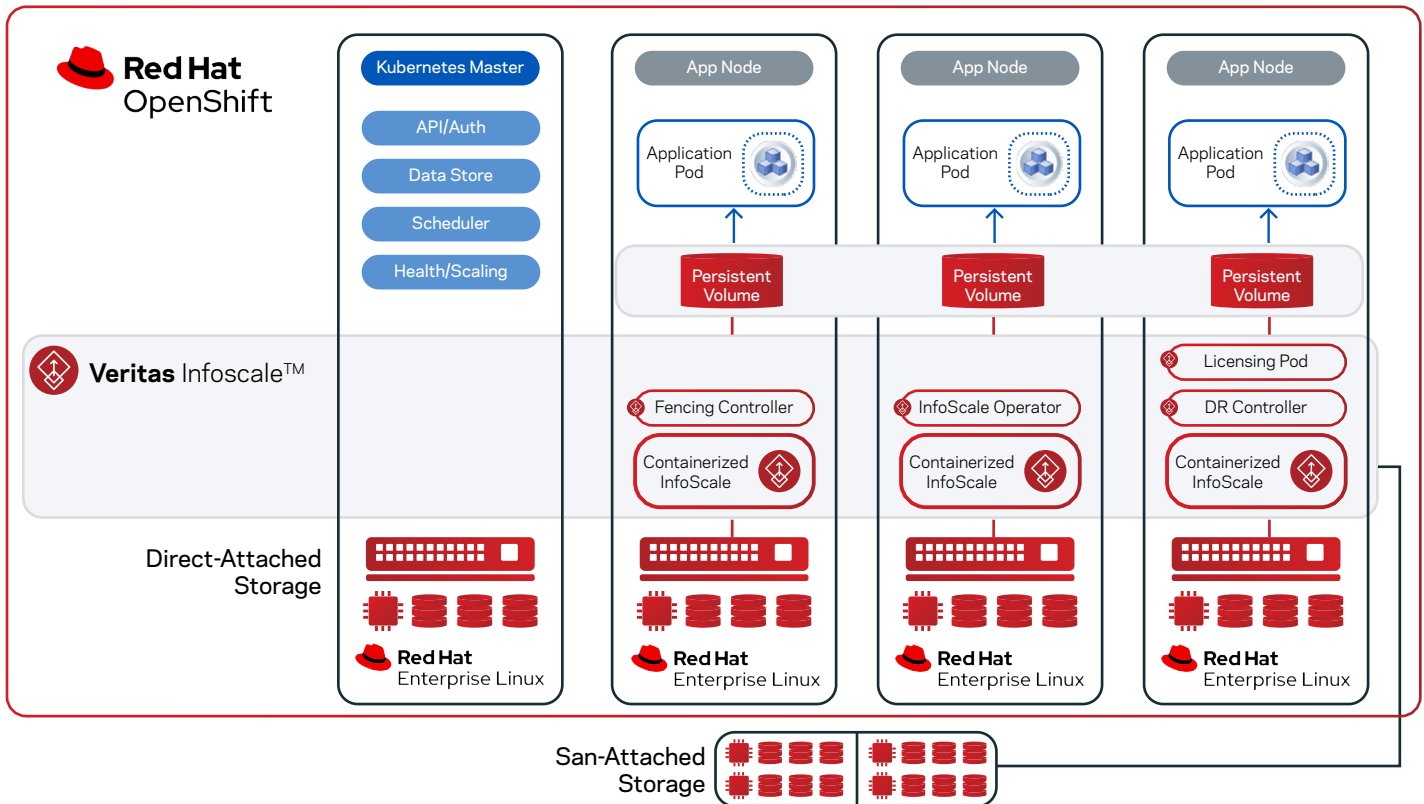


Figure9. InfoScale in a Red Hat OpenShift environment

Conclusion

Digital transformation has driven enterprises toward IT solutions that provide better efficiency and scalability for their IT services. Containerizing applications and managing operations with orchestration engines like Kubernetes can help businesses deliver more efficient innovation, but this operating model has several gaps in the enterprise functionality needed by most applications. Veritas solves this problem by providing an enterprise-focused foundation for platform-agnostic storage management for containers and Kubernetes that can operate at scale. This unique solution lets you run your containerized applications on any storage infrastructure.

With a focus on usability and functionality, Veritas enables businesses to take advantage of the benefits of containerization with advanced features that integrate seamlessly with containerized environments. Designed for flexibility and scalability to support the largest container environments, Veritas delivers an enterprise software-defined storage solution that provides the tools you need to run your applications in containers with maximum availability and protection.

References

Veritas InfoScale™ Intelligent Monitoring Framework:

https://www.veritas.com/content/support/en_US/article.100013636

Veritas InfoScale™ Availability 8.0 Administrator's Guide – Linux:

https://www.veritas.com/content/support/en_US/doc/79798461-79798510-0/v54435307-79798510

Veritas InfoScale™ Storage 8.0 Administrator's Guide – Linux:

https://www.veritas.com/content/support/en_US/doc/79798461-79798506-0/id-SF170661273-79798506

Veritas InfoScale™ Enterprise Data Sheet:

https://www.veritas.com/content/dam/Veritas/docs/data-sheets/veritas-infoscale-enterprise_21350331.pdf

Veritas InfoScale™ Main Site:

<https://www.veritas.com/availability/infoscale/resources>

Veritas InfoScale™ Solutions for Cloud Environments:

https://www.veritas.com/content/support/en_US/doc/130803809-130803829-1

Veritas Volume Replicator Administrator's Guide – Windows & Linux:

(Windows): https://www.veritas.com/content/support/en_US/doc/106632433-106632437-0/wxrt-tot_v20407449-106632437

(Linux): https://www.veritas.com/content/support/en_US/doc-viewer.79604030-79604037-0.index

Veritas InfoScale™ 8.0 SmartIO for Solid-State Drives Solutions Guide – Linux:

https://www.veritas.com/content/support/en_US/doc/83133640-83133643-0/v101741999-83133643

Veritas Flexible Storage Sharing Data Sheet:

https://www.veritas.com/content/dam/Veritas/docs/data-sheets/21327034_GA_ENT_DS-Veritas-Flexible-Storage-Sharing-EN.pdf

Storage Foundation Cluster File System High Availability 8.0 Administrator's Guide – Linux:

https://www.veritas.com/content/support/en_US/doc/79798461-79798510-0/v54435307-79798510

Cluster File System in Veritas InfoScale Data Sheet:

https://www.veritas.com/content/dam/Veritas/docs/data-sheets/V0628_GA_ENT_DS_Cluster-File-System-InfoScale-EN.pdf

Virtual Business Service – Availability User Guide:

https://www.veritas.com/support/en_US/doc/133157225-133157228-0/br74_viom_tot_v87485226-133157228

Table of Figure

Figure 1. InfoScale Solution Value	3
Figure 2. InfoScale Product Suite	4
Figure 3. HADR Configuration Scenarios.	6
Figure 4. VIOM Management Integration Points	7
Figure 5. Sample VBS Configuration	9
Figure 6.	10
Figure 7 InfoScale running in a Kubernetes environment	11
Figure 8. I/O fencing in Kubernetes using InfoScale	13
Figure 9. InfoScale in a Red Hat OpenShift environment	14

Table of Tables

Table 1. InfoScale Storage Components	5
---	---

1. https://sort.veritas.com/public/documents/vie/7.0/linux/productguides/html/vcs_admin/ch09s06s04.htm

About Veritas

Veritas Technologies is a global leader in data protection and availability. Over 80,000 customers—including 87 percent of the Fortune Global 500—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

VERITAS™

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact
information visit:
veritas.com/company/contact