VERITAS™

Veritas InfoScale

Technical Overview:
Managing Mission-Critical applications
in a Software-Defined Data Center

Contents

Contents

Introduction

Veritas Technologies is a 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, Veritas InfoScale™, is designed to enable enterprise software-defined storage management as well as high availability and disaster recovery (HADR) for all data centers, including on-premises, hybrid and multicloud. InfoScale is a comprehensive, industry-proven solution that helps organizations manage enterprise readiness for modern mission-critical applications, focused on three key principals:

1. **Application availability**—InfoScale integrates directly with application components to ensure the underlying infrastructure and the application itself are managed to provide the highest possible performance and uptime. InfoScale also provides instantaneous automated recovery that minimizes the impact of service disruptions and outages.

2. **Performance**—Eliminate overhead and complexity by using InfoScale to build a high-performance, software-defined storage environment using commodity hardware that intelligently manages data so it's always available on the fastest storage tier. This approach improves application performance and maximizes resource utilization while significantly reducing infrastructure overhead, costs and complexity.

3. **Agility**—InfoScale's platform-independent architecture enables organizations to manage applications for the highest performance and uptime across different operating systems and platforms that may span multiple geographic regions. InfoScale helps avoid vendor lock-in by enabling applications to run in a highly available configuration on any operating system (OS) and any platform, including public cloud and hybrid-cloud configurations. InfoScale also provides automated cloud migration from on-premises systems and between different cloud providers.

This technical overview will explain the Veritas InfoScale Enterprise solution, which includes InfoScale Availability and InfoScale Storage. InfoScale Enterprise enables organizations to combine the HADR requirements of IT applications with highly performant and scalable software-defined storage to achieve maximum application uptime and performance using any OS or platform.

InfoScale Solution Value

InfoScale has several unique features that offer significant value to organizations looking to improve application performance and reduce costs by maximizing architectural flexibility without being locked into a specific technology. 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 the world's most demanding 'always-on' applications (see Figure 1).

InfoScale has customized agents designed for mission-critical applications that understand and manage the application components and resources not monitored by native system tools. This functionality ensures the application has the highest possible uptime and either meets or exceeds performance expectations. InfoScale can also intelligently
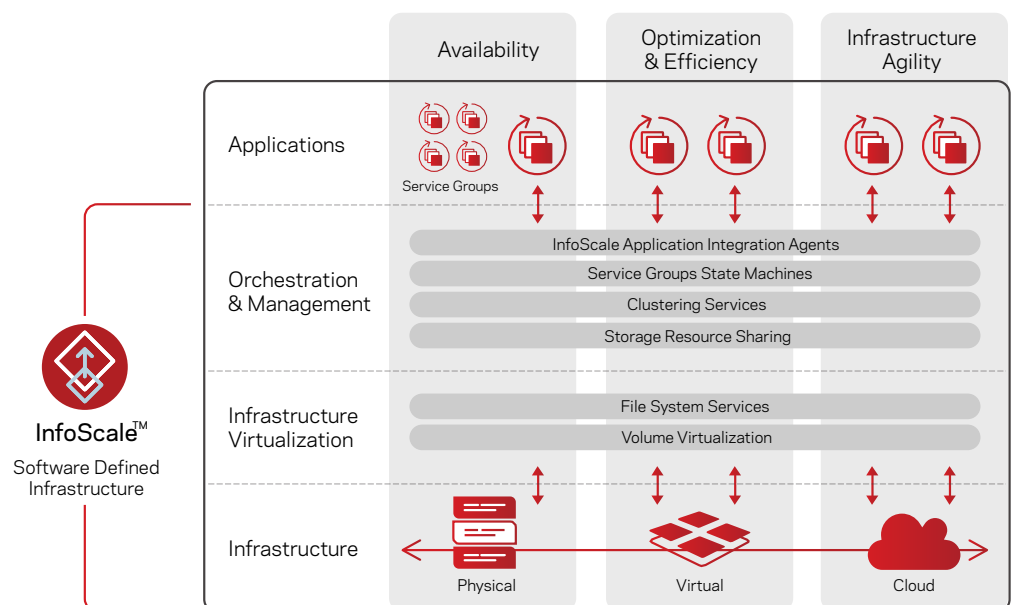

Figure 1. A summary of InfoScale's solution value.

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:

- **Near-instant fault detection**—By introducing processes that run at the kernel level, InfoScale Availability can respond to application failures almost instantly when they occur and can take action to maximize application uptime or if required, failover to another site. This design also significantly reduces compute resources required for monitoring and can help prevent data corruption by reducing the time to action in the event of a failure.

- **Maximized performance**—Regardless of the underlying infrastructure, InfoScale can intelligently cache application data so it's always available on the highest performance storage tier. Going one step further, InfoScale can manage the intelligent data caching process on highly available storage clusters that can be built using direct-attached storage on commodity hardware. This software-defined approach provides enterprise storage features and better application performance at a significantly reduced cost, with no vendor lock-in.

- **Replication management**—InfoScale provides the flexibility to integrate with either third-party storage replication technology or Veritas Volume Replicator. Both options can provide a near-zero recovery time objective/recovery point objective (RTO/RPO) for mission-critical data and can scale to support the largest workloads. With Volume Replicator, there are some additional benefits such as maintaining write order fidelity, multi-target support, heterogeneous system configuration (including public cloud) and zero data loss.

- **Cloud integration**—InfoScale can manage HADR functionality for both hybrid-cloud and native cloud environments. InfoScale Availability can also manage data transfer between on-premises systems and public cloud environments, making it easy to move applications between platforms. Organizations can also move applications between public cloud platforms such as Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform to support a multicloud HADR strategy.

- **HADR Firedrill**—InfoScale Availability can use the Firedrill feature to manage and run a simulated test on an isolated, non-production network segment to ensure systems at the secondary site are working properly prior to a full failover event. It does so by using snapshots of production data that are then attached to temporarily provisioned systems used for testing purposes. InfoScale also manages the cleanup of the Firedrill environment when it's no longer needed.

- **Management simplicity**—Veritas InfoScale Operations Manager (VIOM) is a single-pane-of-glass management console that lets you manage any InfoScale component from one place. In addition to management of InfoScale solutions, VIOM also provides visibility into non-InfoScale infrastructure, so you can use it to identify potential issues within an environment that could lead to unexpected downtime.

InfoScale can further facilitate an enterprise HADR solution with an integrated feature where application tiers can be grouped together in a way that represents the entire business service the application provides. This group is known as a **Virtual Business Service (VBS).** A VBS represents a multi-tier application as a single, consolidated entity that augments the HADR provided for the individual application tiers. Using a VBS, you can completely automate the recovery or migration of a complex, multi-tier application, making it easy to provide HADR for an entire business service.

InfoScale Architecture

InfoScale is a fully software-defined solution that has an extensive compatibility matrix and is OS- and platform-agnostic. You can deploy it on industry-standard hardware, and it can support a wide range of enterprise applications, regardless of their underlying infrastructure. InfoScale is composed of several products, mostly delineated by licensing scheme. At its core, InfoScale has two distinct parts: InfoScale Availability and InfoScale Storage. (See Figure 2.)

InfoScale Availability consists of the components shown in Table 1. It encompasses all the services and features designed to ensure applications are highly available. There are several application-specific agents available with InfoScale Availability that ensure optimal application performance and resiliency.
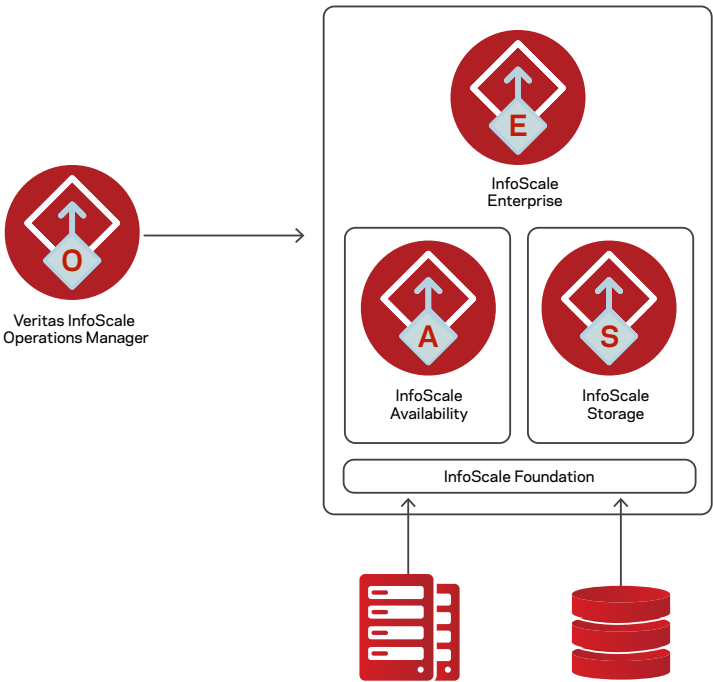
.

Figure 2. The components of the InfoScale product suite.

Component	Description
Service Groups	A collection of hardware and software resources and their dependencies that are required for an application to run.
High Availability Daemon	A process that collects information about resource states from the application agent on the local system and forwards it to all cluster members.
Application Agents	A software package for systems managed by InfoScale that enables the management of hardware and software resources.
Intelligent Monitoring Framework (IMF)	An event-based notification system integrated with certain application agents that provides near-instantaneous notification of resource state changes.
Veritas InfoScale Operations Manager (VIOM)	A web-based centralized management console for InfoScale Enterprise that provides monitoring, visualization and management of resources.

Table 1. InfoScale Availability Components.

InfoScale Storage is the software-defined storage group of services and features. It enables you 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 shown in Table 2.

Component	Description
Veritas File System (VxFS)	A POSIX-compliant extent-based file system known as VxFS that runs on most Unix and Linux variant operating systems.
Veritas Volume Manager (VxVM)	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 you can create file systems.
Veritas Volume Replicator	A software-based data replication utility that enables consistent block-level data replication between a VxVM-managed source data volume and one or more remote data volumes.
Cluster File System (CFS)	A VxFS file system running in parallel access mode, which allows 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 2. InfoScale Storage Components.

InfoScale Availability

InfoScale Availability consists of five main components that work together to create a highly available clustered system infrastructure:

1. Service Groups

2. High Availability Daemon (HAD)

3. Application Agents

4. Monitoring – Intelligent Monitoring Framework (IMF)
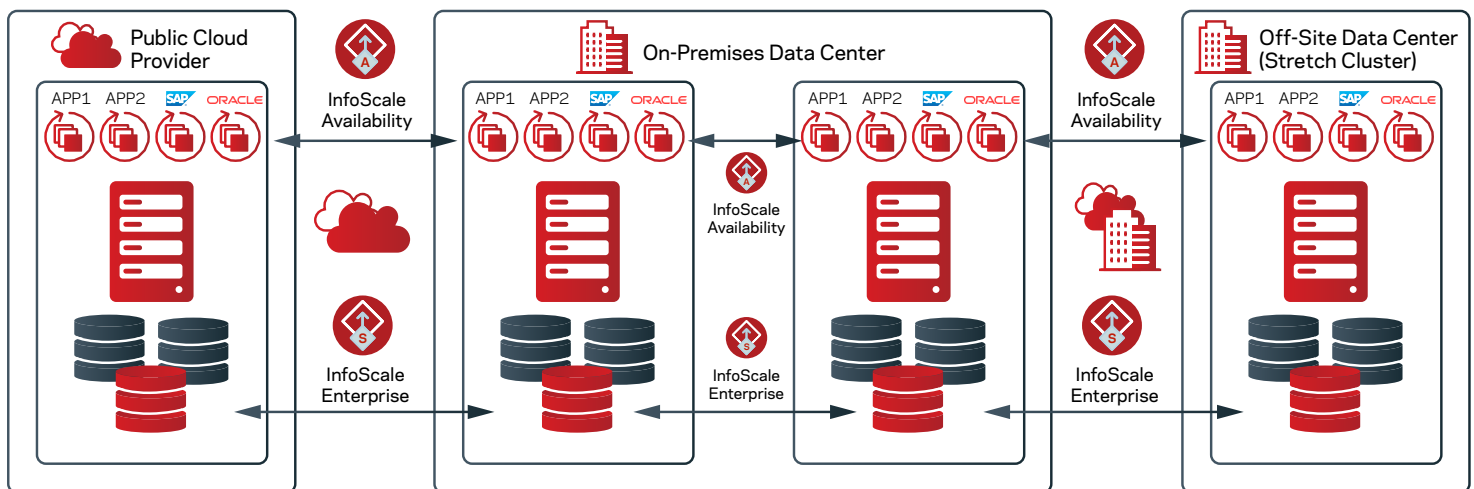
5. Veritas InfoScale Operations Manager (VIOM)

Figure 3. An example of different HADR configuration scenarios.

Figure 3 shows the different HADR configurations that are possible with InfoScale Availability. These range from HA within a local cluster or even a single node, single site cluster to a multi-site clustered configuration that spans geographic regions. Replication is managed by InfoScale Availability using either Volume Replicator or direct integration with supported storage subsystems that provide native replication functionality.

Service Groups

A service group is a virtual container of resources managed by InfoScale that are required for an application to run. The service group defines the individual resources as well as the dependencies between them and the order in which they need to be managed for an application to be online. Service groups allow InfoScale to manage all the hardware and software resources of a managed application as a single unit. Resources include components such as Network Interface Cards (NICs), IP addresses, disk groups, volumes and mount points. InfoScale manages both the resources and their dependencies to ensure the application is online and can immediately act in the event of a resource failure.

Service groups have multiple attributes that define how it will manage resources for an application. Attributes each have a definition and a value. The definition describes the scope of the attribute and the value contains the input. For example, the SystemList attribute is a user-defined attribute that contains the list of systems on which the service group is configured to run as well as their priorities.

There are three types of service groups:

1. **Failover service groups**—Service groups that run on one system in the cluster at a time. Failover service groups are used for most applications that do not support multiple systems simultaneously accessing the application's data.

2. **Parallel service groups**—Service groups that run simultaneously on more than one system in the cluster. A parallel service group is more complex than a failover group and is appropriate for applications that manage multiple application instances running simultaneously without data corruption.

3. **Hybrid service groups**—Service groups used for replicated data clusters that combine both failover and parallel service groups. The hybrid service group behaves like a failover service group within a system zone and a parallel service group across system zones.

Service groups can be dependent on each other. For example, a managed application might be a web application that is dependent on a database application. Because the managed application consists of all components that are required to provide the service, service group dependencies create more complex managed applications. When you use service group dependencies, the managed application is the entire dependency tree.

High Availability Daemon

The Veritas Cluster Server (VCS) High Availability Daemon (had) is the cluster "engine" and is also known as the Veritas High Availability Engine. An instance of had runs on each cluster node and dynamically maintains a replicated state machine that always provides all nodes with the same view of the cluster state.

The had manages the following functions:

- Building the cluster configuration using the configuration files.
- Managing information distribution when new nodes are added to the cluster.
- Responding to user input.
- Acting when something fails.

The *had* uses agents to monitor and manage resources. It collects information about resource states from the agents on the local system and forwards it to all cluster members. The engine that runs on each node has a completely synchronized view of the resource status on each node.

Application Agents

In general terms, an agent is a software package that runs on the systems being managed as part of an InfoScale HADR solution. Agents are multi-threaded processes that provide the logic to manage resources. Agents manage the system components (or "resources") required for an application to be online. Data from the agents is also fed back to had for analysis and actionability and to VIOM for reporting.

InfoScale agents provide HA for specific resources and applications. Each agent manages resources of a certain type. For example, the SAP agent manages SAP components such as NetWeaver and HANA databases. Typically, agents start, stop and monitor resources and report state changes.

InfoScale Availability has agents for most tier 1 applications such as SAP, Oracle, Tibco and Microsoft applications. In situations where an application agent is not available, a generic resource group option is available that enables functionality like that provided by the InfoScale Availability agents.

There are three types of agents available with InfoScale Availability:

1. Application agents for managing enterprise applications such as SAP, Tibco and WebSphere as well as custom applications.

2. Database agents for managing database applications such as Oracle, MySQL, SAP HANA and Microsoft SQL Server.

3. Replication agents that manage hardware and software data replication technologies such as EMC SRDF, NetApp SnapMirror and Hitachi TrueCopy.

InfoScale Availability also supports third-party custom agents for use in situations where a specific agent is not available. These third-party custom agents are developed and supported by the third parties.

Intelligent Monitoring Framework

The IMF is an extension of the application agent. It provides near-instant fault detection and is essentially a kernel-level process that is installed as part of the InfoScale Availability agent software.

IMF is an integrated part of the InfoScale Availability agent framework. IMF allows the agents to register the resources to be monitored with a notification module that enables immediate (event-based) notification of resource state changes without having to periodically poll the resources to find the resource current state. This process enables InfoScale Availability agents to act immediately in the event of a system fault.

The IMF is a key component of InfoScale Availability that helps eliminate downtime for enterprise applications with near-zero RPO/RTO requirements. Not all InfoScale Availability agents are IMF-aware. You can find a list of agents that support IMF here.

Veritas InfoScale Operations Manager

VIOM is a platform and vendor-agnostic centralized management console for InfoScale Availability, InfoScale Storage and other third-party infrastructure. VIOM is used for monitoring, visualization and management of system and storage resources (see Figure 4). VIOM is also a reporting engine and can generate multiple reports, including a risk analysis report that can summarize issues that may arise within an environment that could reduce HADR readiness.

A typical VIOM deployment consists of two main components: a management server and managed hosts. The management server hosts the web-based user interface (default URL: https://<hostname>:14161/vom). Depending on the usage scenario, VIOM may also discover virtualization environments and SAN/NAS infrastructure as well as SAN fabrics.
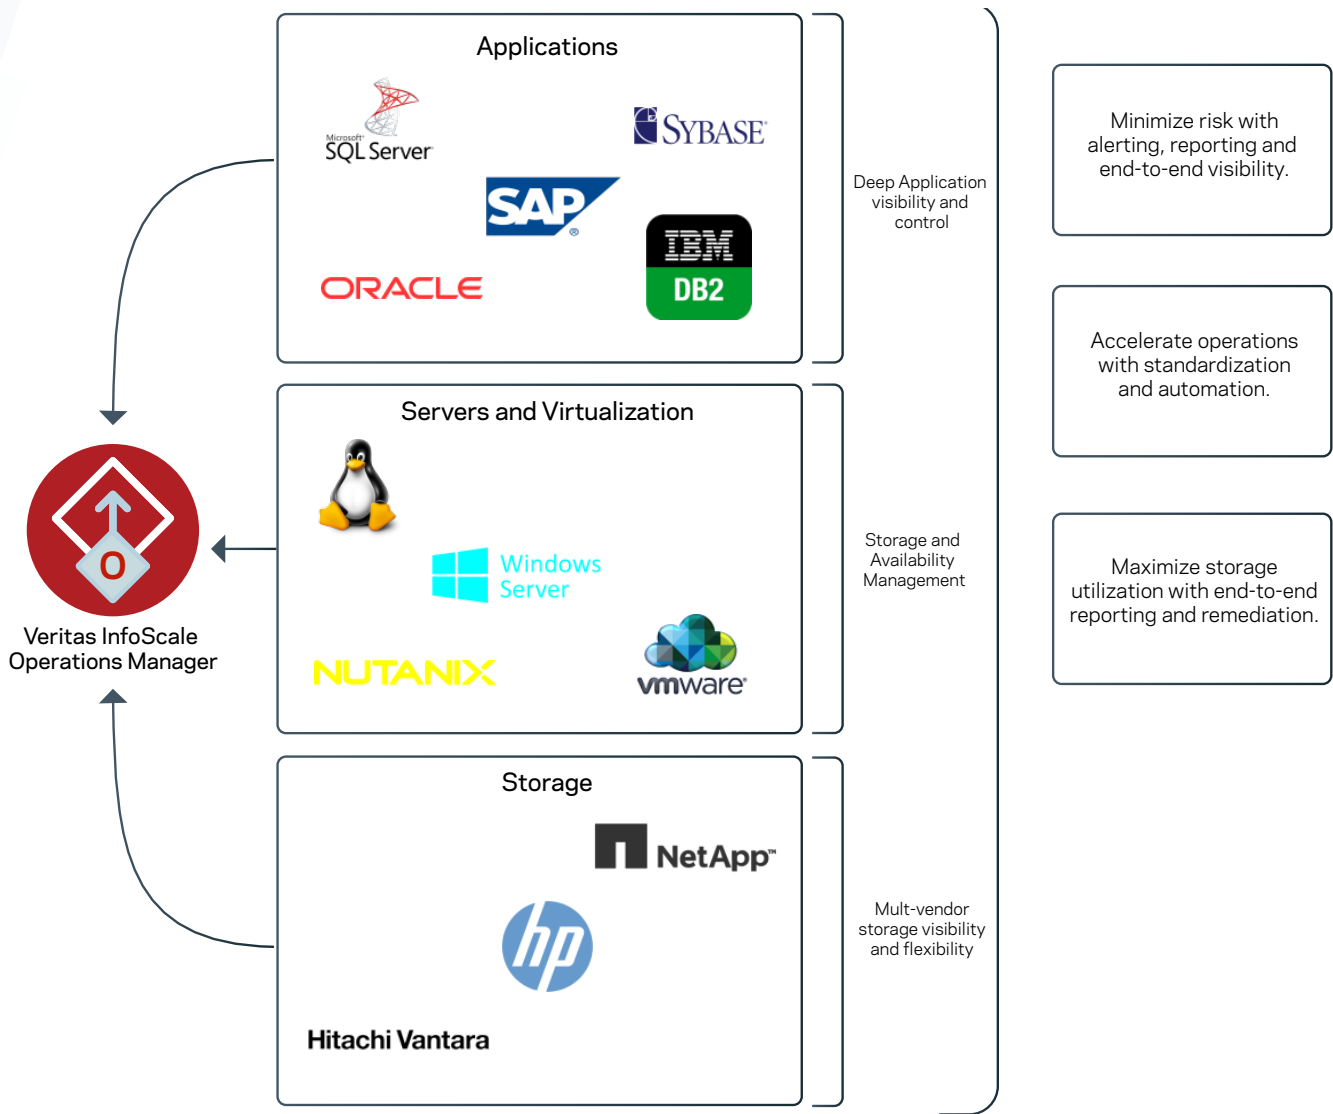
Figure 4. An overview of VIOM management integration points.

You can add managed hosts into VIOM either using agents or as agentless hosts in situations where there is no InfoScale software installed on the target hosts. VIOM manages agentless hosts using SSH or WMI. The level of visibility within VIOM for agentless hosts is the infrastructure only; there is no application visibility for agentless hosts.

Virtual Business Services

Enterprise applications typically consist of multiple systems deployed in tiers that work together to provide an overall business service. A VBS is designed to manage complex, multi-tier applications as a single entity that represents the overall business service the application tiers are providing.

With the VBS, you can work across a heterogeneous environment, which enables IT to ensure the availability of multi-tiered applications across almost any platform or infrastructure. It doesn't matter if the web server sits in a VMware virtual machine (VM), the application server in KVM on Linux and the database on physical Big Iron. If the platform falls within the InfoScale support matrix, it can be included as a tier in a VBS.

A VBS manages dependencies between the service tiers by allowing you to configure the order in which the service groups in the VBS are brought online in a start operation and taken offline in a stop operation. The VBS does not alter the dependencies that are configured for clusters included in the VBS tiers.

VIOM provides the management framework and is required to create and manage a VBS with InfoScale Availability (see Figure 5).
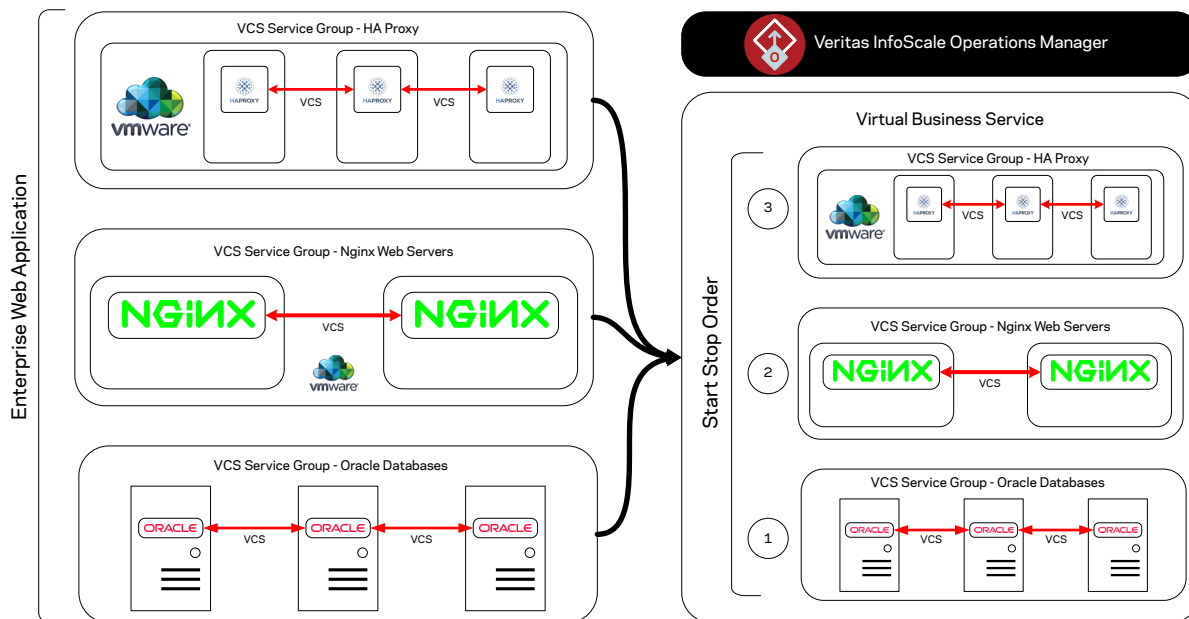
Figure 5. A sample VBS configuration managed by Veritas InfoScale Operations Manager.

InfoScale Storage

InfoScale Storage consists of six major components that provide the basis for building a highly available clustered storage infrastructure:

1. Veritas File System (VxFS) is a POSIX-compliant enterprise file system designed to maximize application performance.

2. Veritas Volume Manager (VxVM) is the storage management subsystem that allows you to create logical data volumes used by applications and databases from physical disks and logical unit numbers (LUNs).

3. Cluster Volume Manager (CVM) is the storage virtualization layer that enables the mapping of volume block addresses on to storage devices.

4. Cluster File System (CFS) creates a file system that can be shared by multiple nodes in the cluster.

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. CFS can be built on top of a volume shared with FSS.

6. Veritas Volume Replicator enables optimized replication of data between InfoScale-managed data volumes.

Figure 6 shows a four-node InfoScale cluster mounted to a CFS consisting of storage exported from local disks with FSS. The CFS allows all nodes in the cluster to read and write data to a shared namespace. Data stored on the CFS can be modified by any node on the cluster.
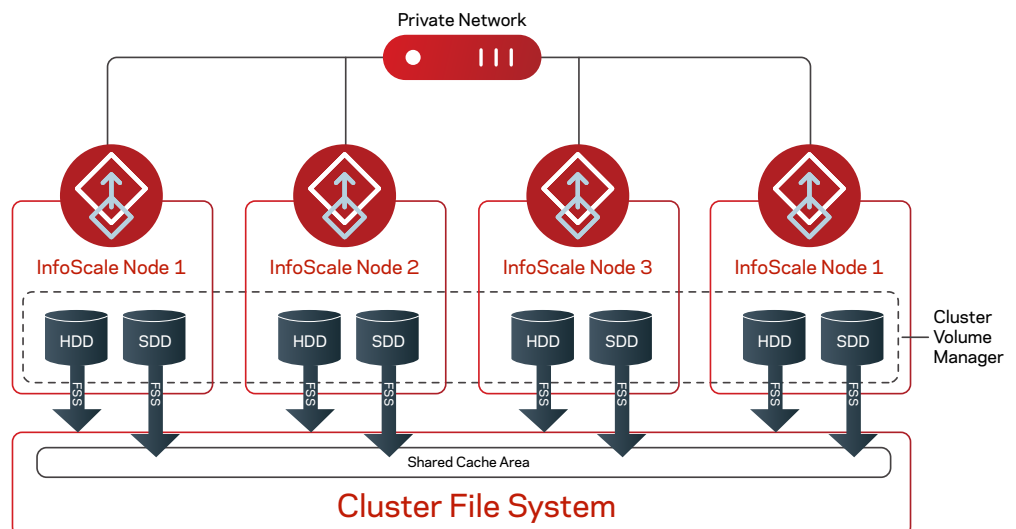

Figure 6. A depiction of how the Cluster File System (CFS) and Flexible Storage Sharing (FSS) enable node sharing.

Veritas File System

The VxFS is an extent-based, POSIX-compliant journaling file system capable of managing large volumes of data and is designed to provide high performance and availability for applications. The VxFS has several advanced features that maximize application performance and optimize the data footprint. The VxFS supports online growing and shrinking of the file system, compression, encryption and thin reclamation, which allows you to release free data blocks of a VxFS file system to the free storage pool of a thin storage LUN.

The VxFS allocates disk space to files in groups of one or more adjacent blocks called extents. The VxFS defines an application interface that allows programs to control various aspects of the extent allocation for a given file.

The VxFS supports large file systems of up to 16 petabytes.

Veritas Volume Manager

The 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 increase I/O throughput. The VxVM also provides features that enhance fault tolerance and fast recovery from disk failure or storage array failure.

The VxVM overcomes restrictions imposed by hardware disk devices and LUNs by providing a logical, software-defined volume management layer, allowing volumes to span multiple disks and LUNs. The VxVM provides the tools to improve performance and ensure data availability and integrity. You can also use the VxVM to dynamically configure storage while the system is active.

The 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 platforms.

Cluster File System

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

The CFS uses a distributed locking mechanism called Global Lock Manager (GLM) to ensure all nodes have a consistent view of the file system. The 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 the 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. The 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. The 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 CVM organizes disks into disk groups whose membership is specified by administrators. The disk group is the atomic unit in which CVM instances import (gain access to), deport (relinquish access to), activate (present to CFS) and deactivate (withdraw accessibility to) disks. The CVM 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.

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 they can cooperate to restore volume integrity.

The 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 important, the CVM also manages volume states, including 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.

I/O Fencing

The "split-brain" condition occurs when communication disruption between InfoScale cluster nodes can cause data corruption when InfoScale cannot 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 occurs in all clustered storage implementations. To mitigate and resolve the split-brain condition, 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 that has failed is ejected from the cluster and prevented from accessing the data disks.

The key to protecting data in a shared-storage cluster environment is to guarantee there is always 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 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 the surviving nodes permit write operations to data disks. Disk-based I/O fencing requires SCSI-3 Persistent Reservation—compatible disk devices. (See Figure 7.)

2. Server-based I/O fencing—A special coordination point server maintains a registry of node membership in a cluster instead of disks, as in disk-based I/O fencing.

3. Majority-based I/O fencing—This option 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, and the other nodes are designated as spectator nodes. The sub-cluster with the majority number of nodes survives and nodes in the rest of the sub-clusters are taken offline.

For more information on I/O fencing behavior, see "How I/O fencing works in different event scenarios."

The CFS involves a primary/secondary architecture. One of the nodes in the cluster is the primary node for a file system. Although any node can initiate an operation to create, delete or resize data, the GLM master node carries out the actual operation. After creating a file, the GLM master node grants locks for data coherency across cluster nodes. For example, if a node tries to modify a block in a file, it must obtain an exclusive lock to ensure other nodes that may have the same file cached invalidate this cached copy.

Flexible Storage Sharing

FSS enables network sharing of local storage across a cluster. The local storage can be in the form of Direct-Attached Storage (DAS) or internal disk drives. Network shared storage is enabled by using a network interconnect between the nodes of a cluster.

FSS allows network shared storage to coexist with physically shared storage, and you can create logical volumes using both types of storage, enabling a common storage namespace. Logical volumes using network shared storage provide data redundancy, HA and DR capabilities without requiring physically shared storage. This process is transparent to file systems and applications.

You can use FSS with SmartIO technology for remote caching to service nodes that may not have locally attached solid-state drives (SSDs).

SmartIO

The SmartIO feature of InfoScale Storage enables data efficiency on higher-performance storage (such as SSDs) through intelligent I/O caching. Using SmartIO to improve efficiency, you can optimize the storage cost per IOPS. SmartIO does not require in-depth knowledge of the underlying hardware technologies. SmartIO uses advanced heuristics to determine what data to cache and how that data is removed from the cache. The heuristics take advantage of InfoScale's visibility into the characteristics of the workload.

Figure 7. A depiction of SCSI-3 I/O fencing in an InfoScale cluster.

SmartIO uses a cache area on the target device or devices. The cache area is the storage space SmartIO uses to store the cached data along with the metadata for the cached data. SmartIO supports different types of read and write caching. The type of cache area used determines whether it supports file system caching or volume caching. To start using SmartIO, you can create a cache area with a single command while the application is online. A tool called SmartAssist is available to help estimate the optimal SmartIO cache size based on analysis of system components and data.

When the application issues an I/O request, SmartIO checks to see if the I/O can be serviced from the cache. As applications access data from the underlying volumes or file systems, certain data is moved to the cache based on the internal heuristics. Subsequent I/Os are processed from the cache.
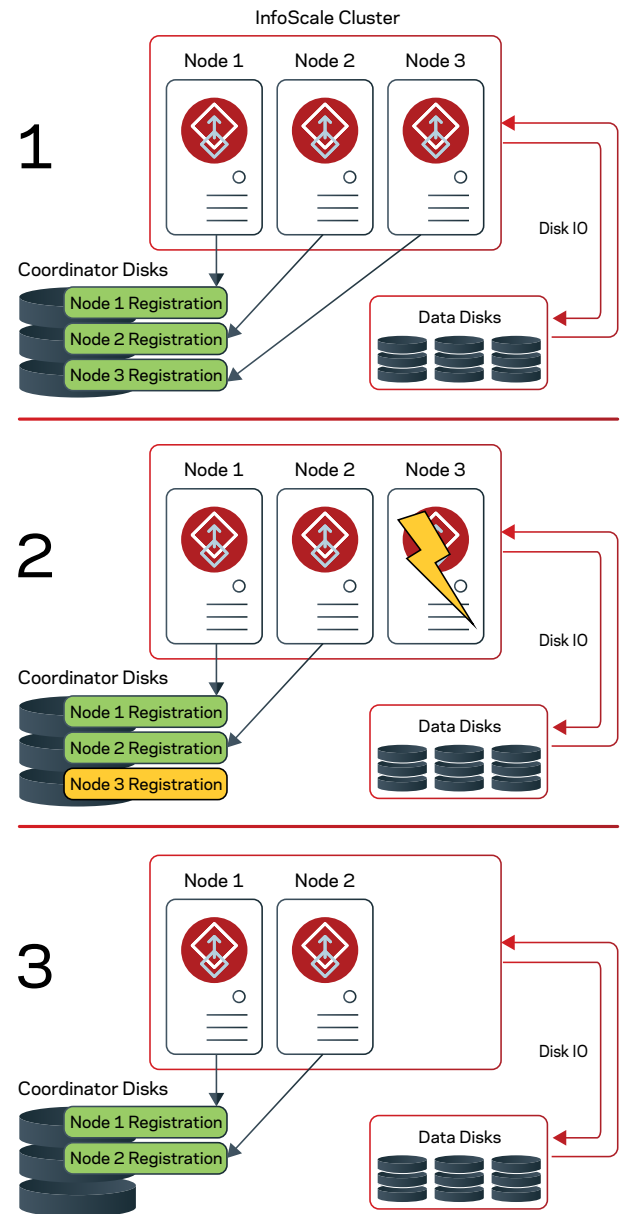
SmartIO and FSS

SmartIO supports the use of SSDs exported by FSS to provide caching services for applications running on Veritas Volume Manager (VxVM) and Veritas File System (VxFS). In this scenario, FSS exports SSDs from nodes that have a local SSD. FSS then creates a pool of the exported SSDs in the cluster. From this shared pool, FSS creates a cache area for any or all nodes in the cluster. Each cache area is accessible only to the node for which it is created. (See Figure 8.)

Veritas Volume Replicator

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

Figure 8. An overview of how InfoScale uses FSS and SmartIO to provide caching services.

Volume Replicator includes the following components, as shown Figure 9:

- **Replicated Volume Group (RVG)**—A group of related volumes within a given VxVM disk group configured for replication. A related volume is a set of volumes where writes must be replicated in order on a secondary site.

- **Storage Replicator Log (SRL)**—A buffer of writes for an RVG. Each RVG contains one SRL and writes to data volumes in the RVG are first queued in the SRL on the primary host before they are sent to the secondary.

- **Replication Link (RLINK)**—The link between a primary and secondary RVG. A primary RVG can have up to 32 associated RLINKs.

- **Data Change Map (DCM)**—Tracks writes when the SRL overflows. The DCM is also used for the initial synchronization of volumes with a secondary site when starting replication. The DCM becomes active only when the SRL no longer has space to hold accumulated updates.

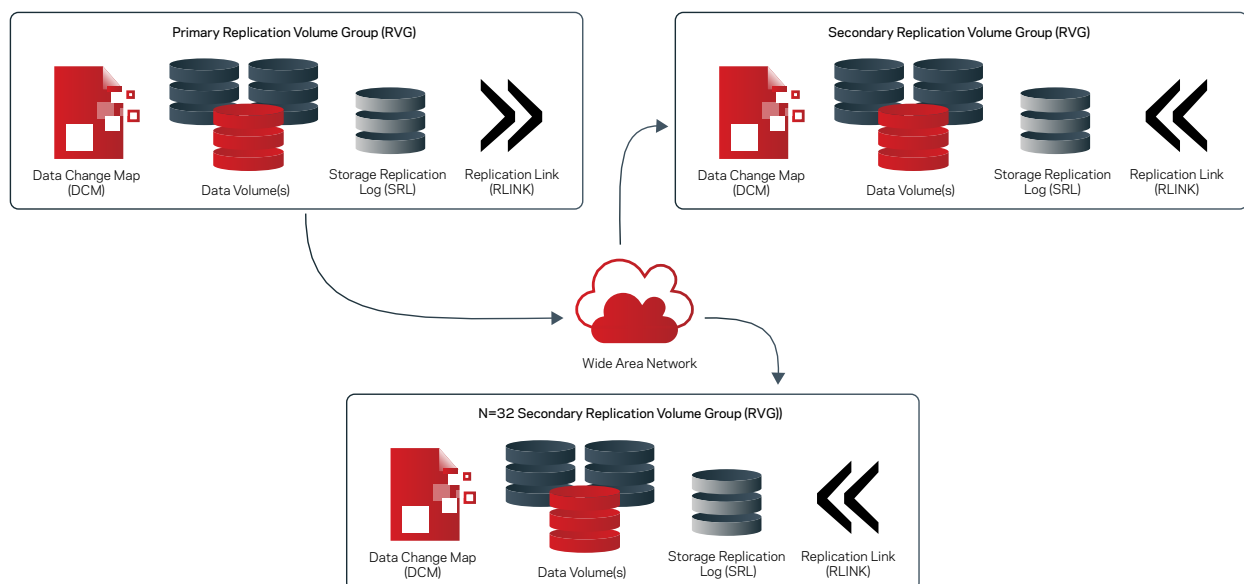- **Replicated Data Set (RDS)** —A grouping of a primary RVG and one or more secondary RVGs.

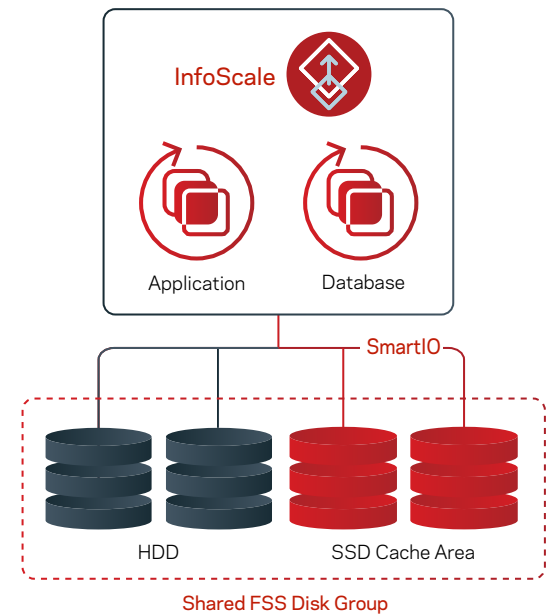*Figure 9. A summary of Volume Replicator components.*

Volume Replicator 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 Volume Replicator tracks writes on the primary volume in the order in which they are received and then applies them on the secondary volume in the same order. Maintaining write-order fidelity ensures 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 Volume Replicator ensures it is a consistent image of the primary RVG at a point in the past.

Volume Replicator is a component of InfoScale Storage and is included with InfoScale Enterprise. Its primary use cases include:

- Replication between like or heterogeneous on-premises systems.
- Cloud data mobility (data replication from on-prem to cloud, cloud to on-prem and cloud to cloud).

Volume Replicator 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, Volume Replicator manages the data replication. Volume Replicator 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.

Volume Replicator 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 I/O so that if any I/O 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 Volume Replicator is in asynchronous mode, it detects when latency returns to normal and automatically switches replication back to synchronous mode. The Adaptive Sync functionality allows your applications to withstand higher I/O latencies while preserving the ability to maintain remote copies of your application data with a low RPO.

InfoScale Networking

InfoScale clusters are formed from multiple discrete compute nodes connected by private networks, enabling inter-cluster communication. This communication is required for all the nodes to perform functions such as acquiring exclusive "locks" to storage and file system resources, updating cluster node membership (preventing split-brain errors) and transferring data from one node to another to fulfill I/O requests for data not stored on the node local to the request.

Figure 10 shows two private networks configured for the InfoScale cluster. This approach is required to provide redundancy in case one of the private networks fails.
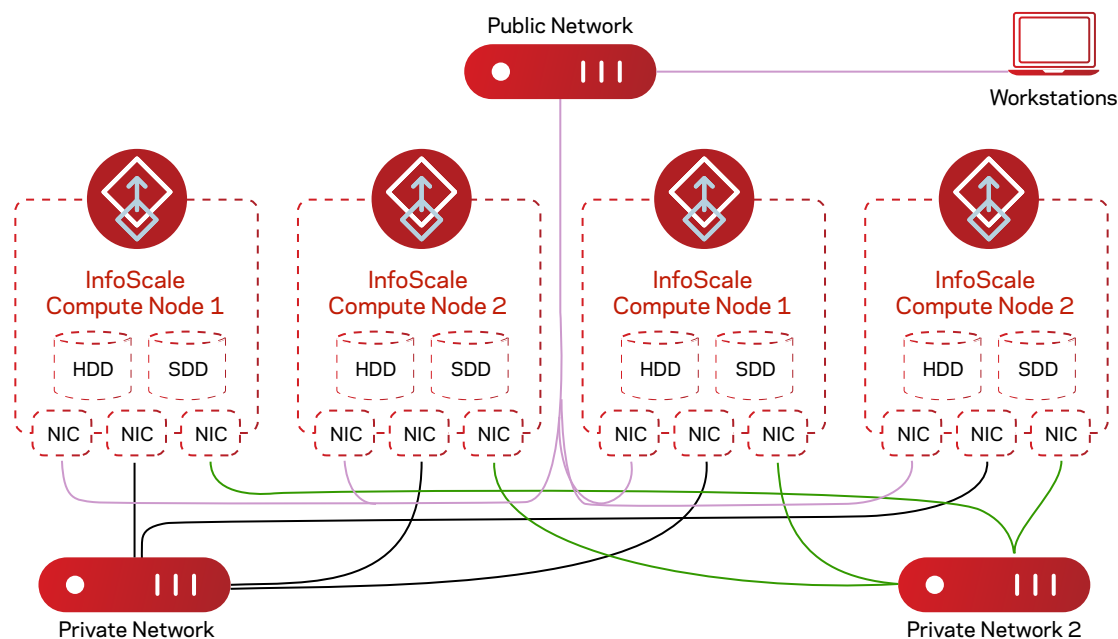

Figure 10. An overview of InfoScale networking between two private networks.

Cluster communication is conducted over the special high-performance Low Latency Transport (LLT) protocol. LLT enables each node to discover other nodes in the cluster by means of a broadcast. Once the nodes have been discovered, heartbeats are sent between all the nodes so each node is aware of any changes in the status of other nodes in the cluster.

InfoScale Enterprise

Figure 11 shows an example of an InfoScale Enterprise solution that combines both the HADR functionality of InfoScale Availability as well as the software-defined storage features and functionality offered by InfoScale Storage. Together, the two products work to create a highly available, highly scalable and high-performance environment that maximizes application performance and availability.

Each InfoScale Availability had instance monitors the resources connected to its node and takes appropriate action (for example, initiates application failover) if it detects a critical resource failure.

All had instances obtain cluster configuration and policy information from a single cluster configuration file normally called main.cf. The configuration file specifies cluster resources and their organization into service groups.
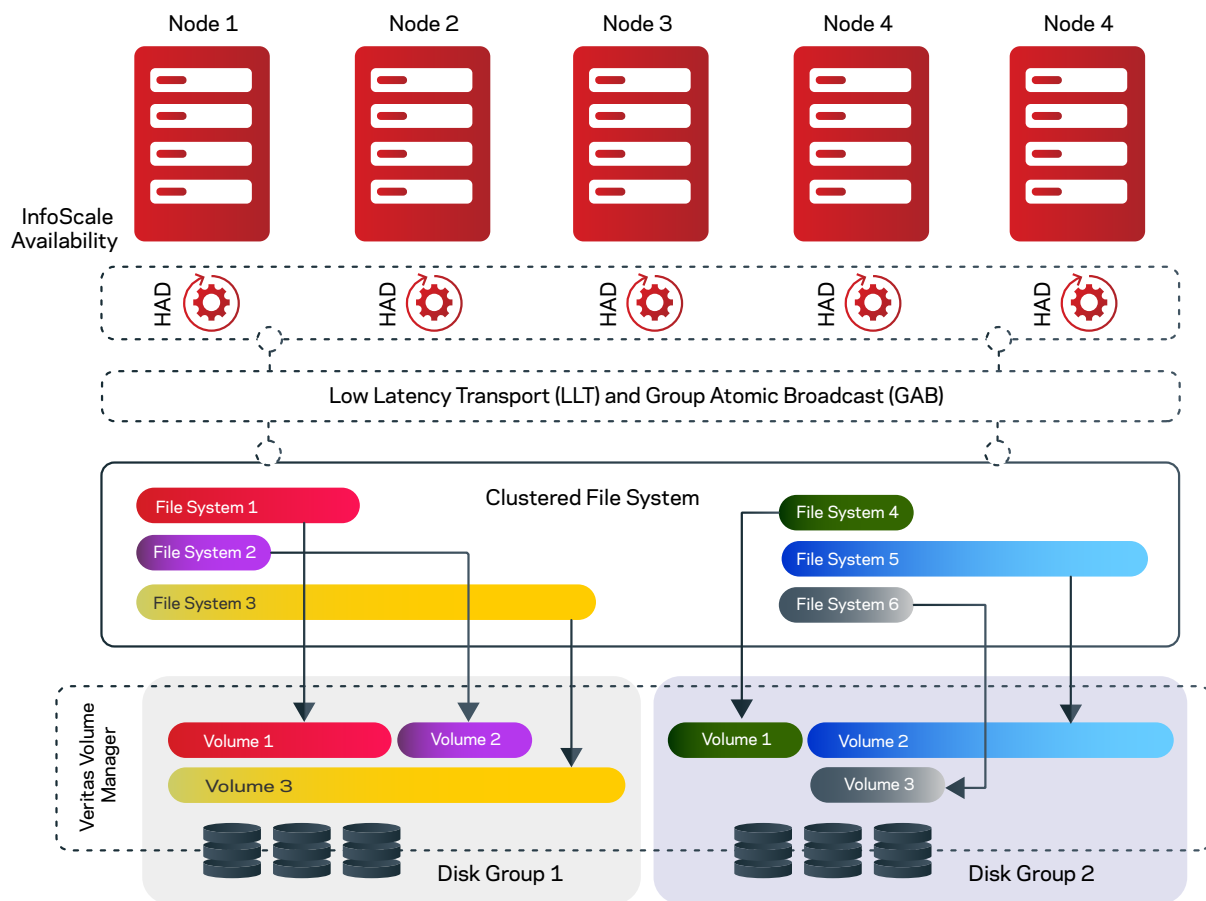

Figure 11. An example of an InfoScale Enterprise solution that includes VCS,CFS and VxVM.

InfoScale Availability encapsulates applications and the resources they require to run as service groups. Examples of components typically included in service groups are disk groups, CVM volumes, CFS file system mounts, network interfaces and IP addresses. The InfoScale Availability framework manages service groups by monitoring their resources while they are operating and by starting and stopping them in response to changes in the cluster state and/or administrative commands.

CVM instances are structured as VCS parallel service groups, with an instance running on each cluster node. All instances access storage devices directly through a storage network using Fibre Channel, iSCSI, SAS technology or FSS. Each CVM instance transforms CFS read and write requests into I/O operations on the underlying devices and executes the corresponding commands.

Cloud Integration

HADR in the Cloud

InfoScale supports several private and public cloud providers and can be used in these cloud environments to provide HADR capability that is not commonly available with cloud-native tools. Using InfoScale in the cloud provides several benefits:

- **Application focus**—Although public cloud infrastructure is designed to provide excellent availability and durability for compute and storage systems, InfoScale is focused on the applications that run on top of this infrastructure. Providing all the same benefits as an on-premises deployment, InfoScale Availability cloud agents are developed specifically for cloud services and can manage the cloud compute, network and storage resources required for both your infrastructure and application to be online in the cloud.

- **Quality of Service**—With InfoScale, your applications can be deployed in a highly available configuration that spans across cloud zones, regions or even cloud service providers to ensure you're protected against cloud service disruptions and outages.

- **Architectural flexibility**—InfoScale supports multiple HADR configurations that are not limited to specific cloud providers. InfoScale can manage applications regardless of the underlying infrastructure, including public cloud infrastructure, so you can easily implement a multicloud strategy with a single solution.

- **Performance**—Like on-premises InfoScale environments, instant fault detection for applications running in cloud environments ensures action is taken instantly in the event of an application failure. InfoScale cloud agents integrate with native cloud monitoring and reporting tools to provide increased operational visibility into the application.

With InfoScale, applications can be replicated to supported public and private cloud environments. InfoScale supports operations such as migrating from physical UNIX systems on-premises to Linux systems in the cloud with minimal configuration. Volume Replicator also provides the functionality to extract your data from the cloud to either bring data back on-premises or to move data to another cloud provider. This capability gives end users the architectural flexibility required to implement a multicloud strategy without needing multiple tools or professional services to manage data transfer between cloud providers.

Cloud-Based Software-Defined Storage

With InfoScale Availability providing HADR capability for cloud environments, InfoScale Storage can significantly improve the performance and efficiency of the underlying cloud-native storage services with features like FSS and SmartIO. InfoScale Storage provides enterprise functionality for cloud environments beyond what's available with native cloud tools, offering some key benefits:

- **Performance**—Although public cloud infrastructure offers higher performance storage options, there are limitations at the system level that minimize overall performance (IOPS). With InfoScale SmartIO intelligent caching, application reads can be served from faster volumes using SSD storage and writes can be served from a cheaper storage tier. This approach significantly improves application performance with minimal additional cost.

- **Scalability**—With FSS, you can create the resilient shared storage volumes needed to horizontally scale enterprise applications using public cloud infrastructure. InfoScale also enables granular resource scaling. When an application needs additional compute or storage resources, they can be scaled dynamically and independently, which reduces operating costs and provides infrastructure flexibility for your applications.

Enterprise Storage and Availability in the Cloud

InfoScale Enterprise combines the HADR features of InfoScale Availability and the storage management and performance benefits of InfoScale Storage to provide the enterprise functionality needed to confidently run a tier 1 application in a public cloud environment. With InfoScale Enterprise, you have the tools needed to manage your high-priority applications in the cloud:

- Application-specific integration for near-instant fault detection and failover.

- Highly performant block-level shared storage using cloud-native infrastructure.

- Simplified multi-tier application management from a single console.

- Flexibility to architect your applications to run on any or across multiple cloud platforms.

InfoScale currently has solution templates available in public cloud marketplaces to simplify the purchasing and deployment experience. Solution templates are available in the AWS Marketplace (AWS CloudFormation Template), the Azure Marketplace (Azure ARM Template) and the Google Cloud Platform Marketplace (Deployment Manager Template).

Hybrid Cloud

Managing storage and high availability for environments that consist of both on-premises and public cloud infrastructure can be challenging and may require multiple point tools. InfoScale can fully support a hybrid approach to public cloud consumption, offering bidirectional HADR and storage management between on-premises environments and supported public cloud providers. InfoScale manages the application components both on-premises and in the cloud environment, and Volume Replicator manages the data replication between on-premises and cloud data volumes. (See Figure 12.)

With an agnostic approach to operating systems and platforms, InfoScale is well suited for deployment in hybrid-cloud HADR configurations and can be tailored to support varying RPO and RTO requirements. InfoScale can also support hybrid-cloud HADR configurations with multiple cloud providers, enabling a resilient and performant multi-cloud strategy.
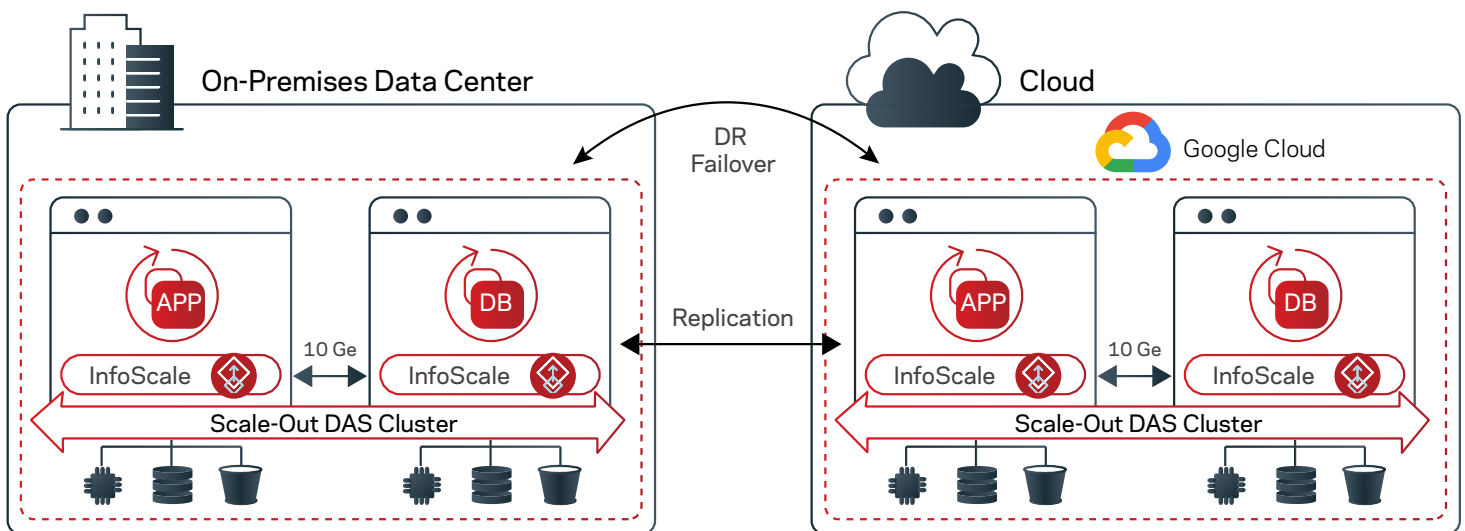

Figure 12. A sample InfoScale hybrid cloud design concept.

Public Cloud

Public cloud services have limited options for providing HADR and enterprise storage functionality for cloud-native applications deployed in an infrastructure as a service (IaaS) model. In most cases, cloud-native applications have the same RPO and RTO requirements as on-premises applications running in a clustered or other high-availability configuration. InfoScale offers the same benefits, functionality and configuration options for the HADR and storage management of cloud-native applications with the added benefit of providing data mobility between cloud providers. This capability helps eliminate being locked in to any specific cloud service provider and enables a high-availability configuration that can protect against cloud provider outages. (See Figure 13.)
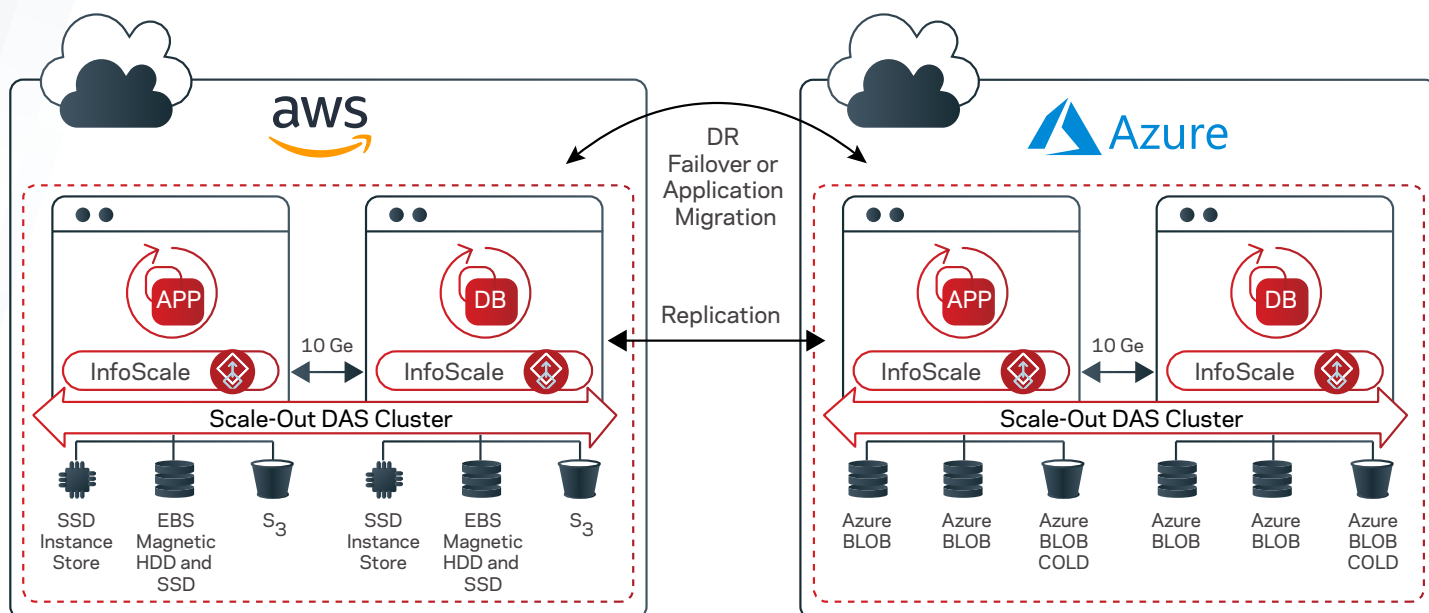
Figure 13. A sample InfoScale public cloud design concept.

InfoScale provides flexible configuration options and is a certified solution for tier 1 applications such as SAP and Oracle that run in public cloud environments. For more information on using InfoScale in public cloud environments, see the following documents:

- SAP NetWeaver in AWS
- SAP HANA in AWS
- Oracle in AWS
- InfoScale in Microsoft Azure

Cloud Migration

There are several tools available for cloud migrations, including those offered by most public cloud service providers. These tools are designed for migration purposes only, however, and don't offer any additional benefit beyond a one-time migration.

InfoScale can support the migration of nearly any application to the cloud while also offering some added benefits over cloud provider migration tools:

- **Rehearsal**—Using the Firedrill feature, InfoScale manages the testing of an application on a non-production network segment in the cloud using temporarily provisioned cloud compute instances and snapshots of the production data volumes. It can do so on-demand using cloud resources, minimizing cost and operational overhead.

- **Enterprise readiness**—In addition to managing the application migration to the cloud, InfoScale also provides high availability for your applications once they're migrated to the cloud environment.

- **Failback**—With full bidirectional operations support for public cloud services, InfoScale can also move applications back on-premises for any reason once they're migrated and online in the cloud environment.

InfoScale supports cloud migrations for on-premises environments being migrated to the cloud and for cloud-native environments being migrated within the cloud or to a different cloud service provider.
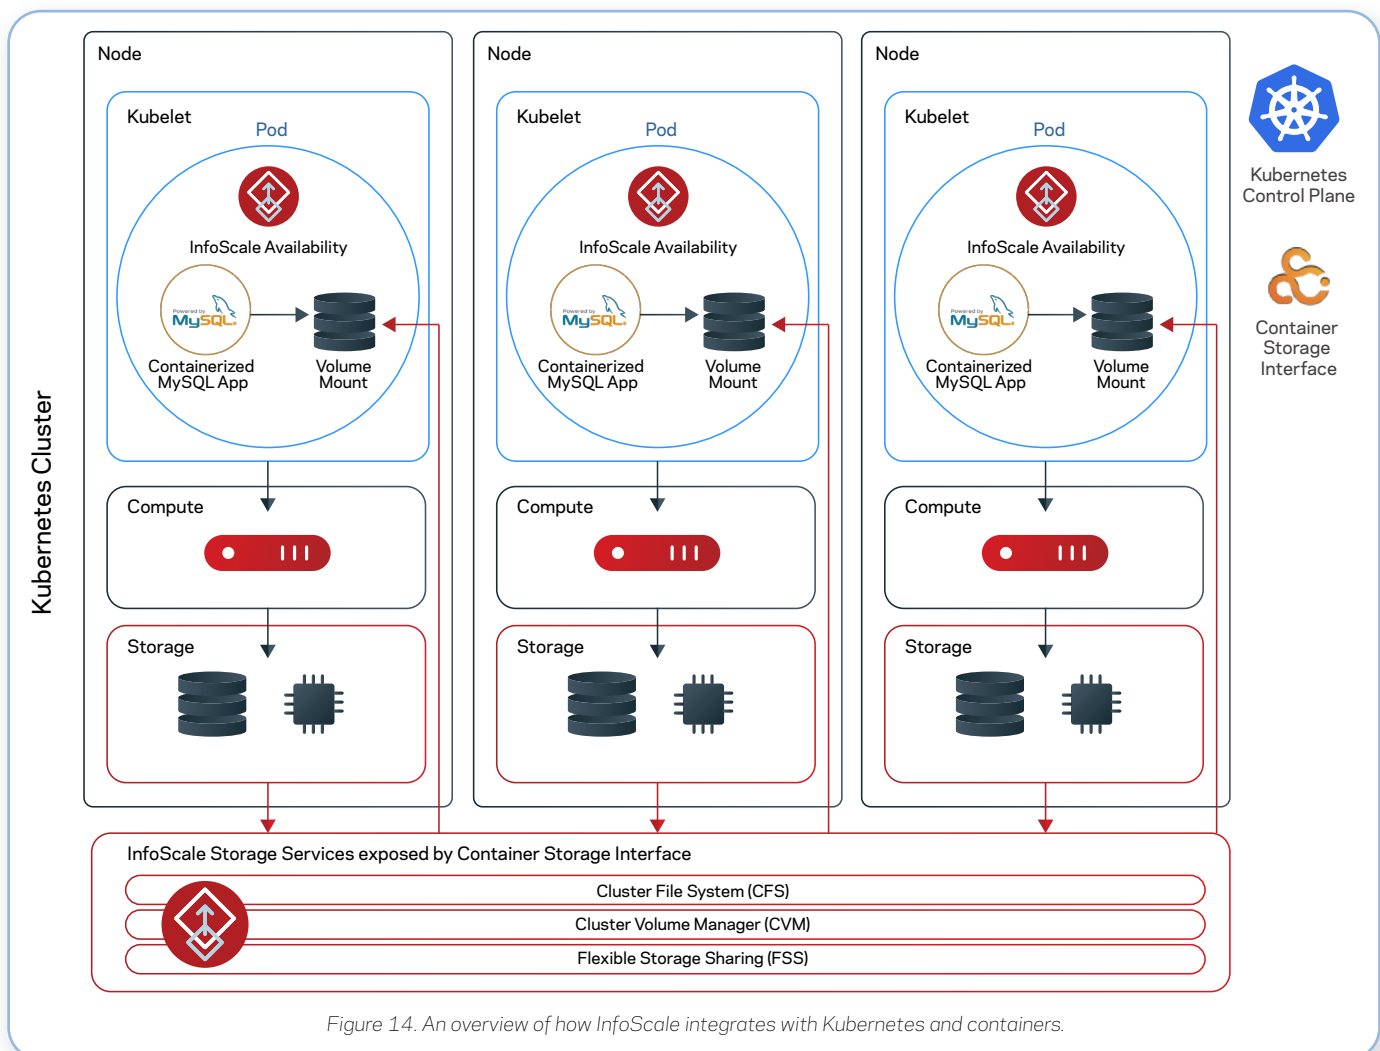
InfoScale for Containers

Containers have become a popular solution for efficiently developing and running applications. Containers provide excellent portability for moving applications between environments and they also help 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 and Veritas Cluster Server (VCS) high-availability agents work with containers and Kubernetes to provide advanced storage management and application-aware high availability for containerized applications. InfoScale availability and 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 the following:

- **Advanced storage management**—InfoScale's CSI plug-in allows Kubernetes to provide InfoScale persistent storage volumes to containerized applications being managed within a Kubernetes namespace. You can use InfoScale's FSS to provide high-performance storage using the disks directly attached to the Kubernetes cluster nodes.
- **Application high availability**—InfoScale's VCS agents are customized to manage high availability for containerized applications by integrating with Kubernetes liveness probes. VCS agents provide visibility into the status of applications running in containers and can also monitor the infrastructure resources required for the applications to be online.

InfoScale is deployed directly on Kubernetes cluster nodes and the InfoScale CSI plug-in provides the interface between Kubernetes and InfoScale. Figure 14 provides an overview of how InfoScale integrates with Kubernetes and containers to provide high availability and persistent storage for containerized applications.

Figure 14. An overview of how InfoScale integrates with Kubernetes and containers.

High Availability

InfoScale VCS agents provide high availability and recovery automation for containerized applications by monitoring critical application processes and resources. VCS works with Kubernetes by returning liveness probe requests. Kubernetes liveness probes are used to query VCS agents to find out if the application and its resources are running. If an application process fails, it can be restarted without restarting the entire container. InfoScale manages application high availability in a containerized environment using the same approach it uses for applications running on traditional virtualized or physical infrastructure. There are two deployment options available for VCS agents to manage containerized applications:

- **VCS deployed as a sidecar container**—This option involves provisioning an extra container that is deployed within a Kubernetes pod with applications that need to be highly available. The VCS sidecar implementation includes the application monitoring agent, the mount agent to manage availability of the persistent volume claims and the network agent to ensure the network resources are available. When the application pod is deployed, the application container and InfoScale sidecar container are launched concurrently in a shared Kubernetes process namespace. The shared process namespace enables the InfoScale sidecar container to monitor application processes.

- **VCS deployed inside the container**—This option involves creating the container and installing VCS within the container in the same manner VCS would be installed on a traditional server. VCS agents can also be included in the application container image. High availability is provided with the same functionality as the sidecar container deployment. The Kubernetes Kubelet conducts liveness probes on the VCS agent and can kill and restart the container if required, based in the information provided by the VCS agent.

Both deployment options provide health checks and in-depth monitoring for containerized applications. Figure 15 shows both VCS agent deployment options for managing application high availability in a Kubernetes environment.
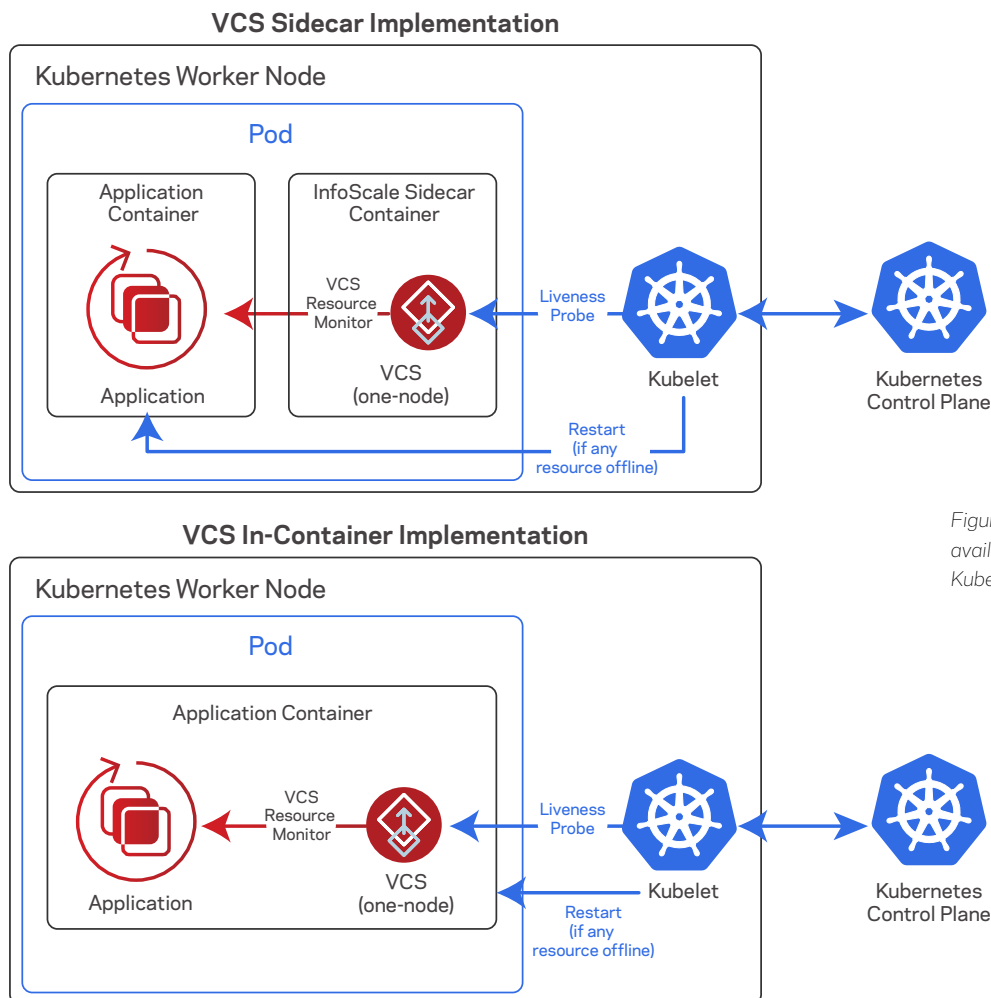

Figure 15. An overview of how to achieve high availability for containerized applications in Kubernetes using InfoScale VCS agents.

Storage Management

InfoScale's CSI plug-in allows Kubernetes to mount InfoScale storage volumes formatted with VxFS inside the containers it's managing, which stateful applications can then use 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 that 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. 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 directly on the InfoScale cluster prior to the creation of Kubernetes application pods and containers. InfoScale statically provisioned volumes can also be used to simplify the process of migrating traditional applications into containers by allowing traditional applications and containerized applications to coexist within the same Kubernetes cluster. Application data in InfoScale persistent volumes can be unmounted from a traditional application and then mounted by Kubernetes using the InfoScale CSI plug-in into an application pod hosting a containerized version of the same application.

Figure 16 shows the Kubernetes workflow for mounting an InfoScale persistent storage volume inside a container running a MySQL database.
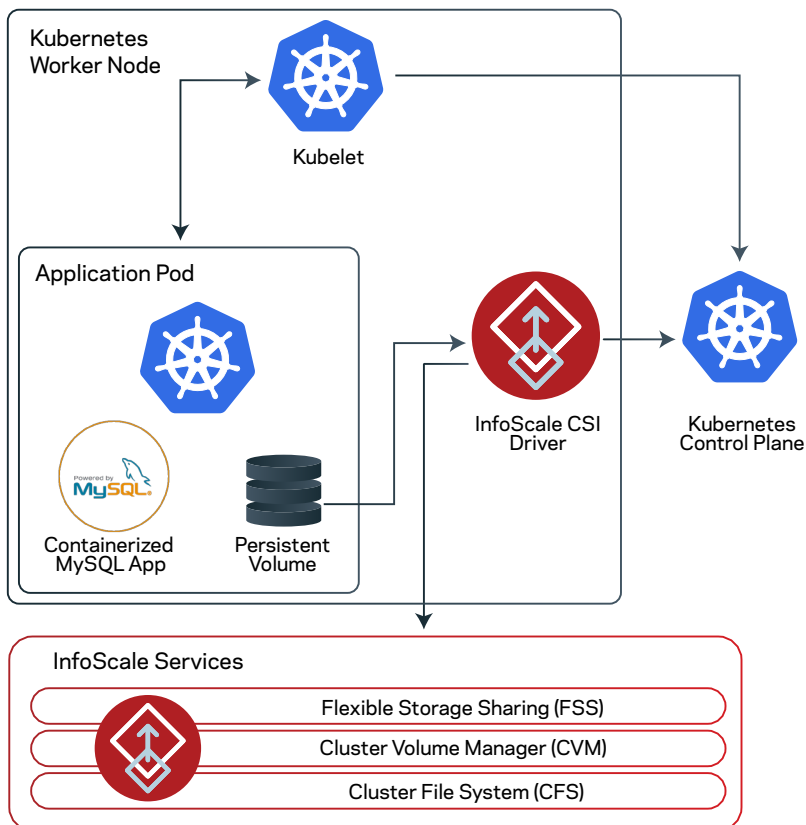

Figure 16. The Kubernetes workflow for mounting an InfoScale persistent storage volue in a container.

The InfoScale CSI plug-in enables Kubernetes 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, which can occur in a clustered environment in the event of a node/hardware failure that disrupts cluster communications and membership. InfoScale provides advanced I/O fencing by preventing data from being written to nodes within the Kubernetes cluster that have failed due to hardware and 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 process prevents data corruption by allowing only the working nodes to continue normal operations. In the event of a communication loss between cluster nodes (or "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. (See Figure 17.)
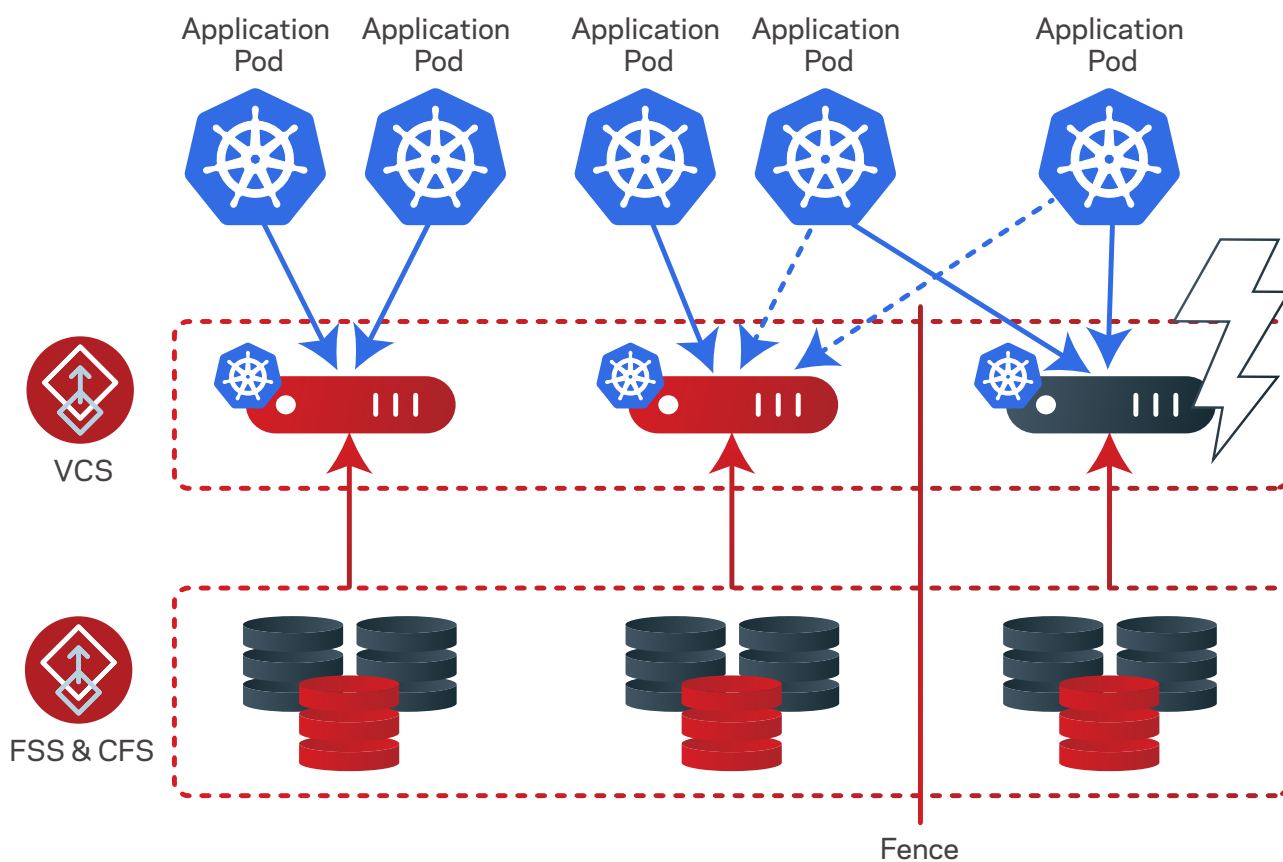

Figure 17. The I/O fencing process in Kubernetes using the InfoScale CSI plug-in.

Conclusion

The InfoScale Enterprise architecture provides a highly available, highly scalable, high-performance storage and application management suite that enables organizations to deliver their business services with the confidence they'll always be available and running in an optimal state. InfoScale is a proven solution that offers several key benefits for enterprises with demanding SLAs:

- **Flexibility**—Enterprise-level HADR and storage management for any application, any platform, anywhere. InfoScale's agnostic approach to high availability and storage management lets you architect the best solution for your applications without being locked into a specific technology or service provider.

- **Simplify management**—Use a single console and reduce the need for manual intervention in the HADR process, which can be time-consuming and error prone. Virtual Business Services further simplify managing system availability by logically representing complex, multi-tier applications as a single entity you can manage with a single click.

- **Enterprise-level availability**—With instant application fault detection and the ability to manage several HADR configurations across operating systems, platforms and even cloud services, InfoScale provides best-in-class options for managing HADR for your most important applications.

- **Increased confidence**—InfoScale's integrated, non-disruptive DR Firedrill feature allows you to maximize production uptime and increase your confidence in the HADR and cloud migration processes.

- **Best-in-class performance**—As a software-defined solution, InfoScale can maximize resource utilization and significantly improve application performance with intelligent data management capability that works with any application, on any type of infrastructure. This approach significantly reduces operational costs, eliminates vendor lock-in and provides complete freedom of architecture to build and run your applications on the infrastructure that best fits your business requirements.

With the ability to provide exceptional application performance and high availability as well as architectural flexibility, InfoScale enables businesses to improve their application SLAs while reducing infrastructure footprints by integrating cloud into their IT strategy. InfoScale can help ensure maximum application performance and uptime in heterogeneous environments. Whether running on-premises, in a hybrid cloud configuration or entirely within a cloud environment, InfoScale provides an enterprise software-defined storage and application availability solution for any platform.

References

- Cluster File System in Veritas InfoScale Data Sheet

- Storage Foundation Cluster File System High Availability 7.4.1 Administrator's Guide – Linux

- Veritas Flexible Storage Sharing Data Sheet

- Veritas InfoScale 7.4.1 SmartIO for Solid-State Drives Solutions Guide – Linux

- Veritas InfoScale Availability 7.4.1 Administrator's Guide – Linux

- Veritas InfoScale Enterprise Data Sheet

- Veritas InfoScale Intelligent Monitoring Framework

- Veritas InfoScale Solutions for Cloud Environments

- Veritas InfoScale Storage 7.4.1 Administrator's Guide – Linux

- Veritas InfoScale Website

- Veritas Volume Replicator Administrator's Guide – Windows

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.

VERITAS™

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

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