

High Availability and Resiliency for Open Source Applications and Databases

Contents

Executive Summary	3
Solution Value	3
High Availability and Resiliency	4
High Availability Management	5
Intelligent Monitoring	5
Enterprise Storage Management	5
InfoScale Application Agents.	6
PostgreSQL Agent	7
MySQL Agent	8
High Availability Topologies	9
Centralized Operations.10
Database Lifecycle Management11
Database Cloning11
Replication12
Data Warehousing.13
Summary13

Executive Summary

Open source solutions are rapidly gaining traction in the IT industry and have become a popular option to run modern IT services. These solutions rely on robust and efficient database management systems to effectively manage application data. As users expect applications to be available and easily accessible at any time, modern applications have an ever-shrinking tolerance for downtime. This has driven a distinct need to architect and deploy applications in a highly available and resilient manner that requires an enterprise solution designed to provide high availability and resiliency that can deliver the always-on experience that end users expect.

Veritas InfoScale™ is an enterprise solution that ensures high availability and resiliency for IT applications. InfoScale has a proven track record of delivering application high availability in environments of all types and sizes — including open source applications and database management systems. InfoScale enhances your applications and databases with advanced features and functionality focused on:

- Achieving high availability for open source databases and the applications they support
- Resiliency that eliminates unplanned downtime with platform-agnostic disaster recovery
- Automation and visibility that simplifies operations with the flexibility to support nearly any application

As a software-defined solution, InfoScale is the ideal foundation to manage high availability and resiliency for open source applications and databases with any uptime requirement. Whether you're operating in an on-premises environment, in the cloud, or in a hybrid cloud environment, InfoScale can help ensure that your applications are always on and delivering a premier user experience. This white paper discusses how InfoScale provides high availability and storage management for open source applications and database management systems.

Solution Value

InfoScale is a multi-faceted and platform agnostic foundation for any type of application that delivers extensive functionality to enhance application and database security, availability, and resiliency.

InfoScale adds significant value for open-source applications databases by providing high availability, resiliency, and enterprise storage management as a single software-defined solution.

InfoScale provides a full spectrum of functionality for open source applications and databases that help ensure your IT services deliver a smooth and predictable end user experience by giving you:

- **Flexibility:** InfoScale can provide high availability and resiliency for nearly any kind of application, including open source databases and the applications they support.
- **Availability:** Automated local high availability ensures that your applications and databases are highly available and performing as expected, or better.
- **Resiliency:** Protect your applications against unforeseen events, disaster scenarios, and other circumstances that may result in unplanned downtime with automation that helps you deliver an always-on experience for your end users.

Modern applications demand availability and uptime — a requirement with direct implications for the underlying databases, operating systems, and storage management systems. InfoScale manages high availability, resiliency, and storage for the entire application stack as a single software-defined solution.

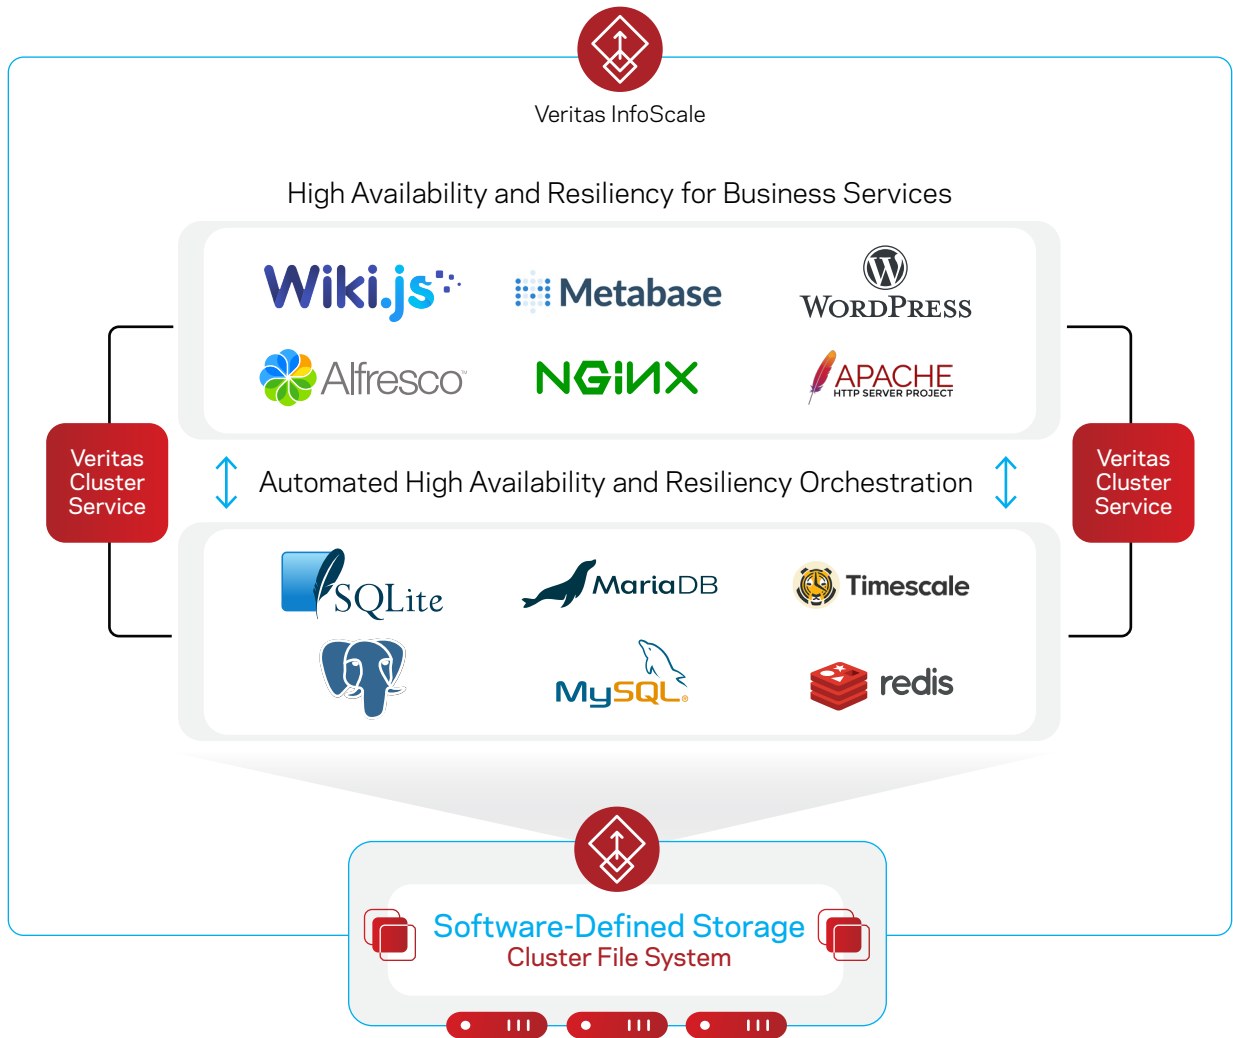


Figure 1. InfoScale provides high availability and resiliency for open-source applications and databases

High Availability and Resiliency

InfoScale Cluster Server, also known as Veritas Cluster Server (VCS), is a proven solution to manage high availability and resiliency for any type of application. InfoScale's customizable management framework is ideal for open source applications and databases, as it gives you the flexibility to easily manage resources that are often tailored to meet specific application and business requirements.

VCS works by organizing systems into a management framework for increased availability. Each system (node) is independent and is managed by the VCS software to form a cluster. By linking the system hardware with intelligent software, VCS can provide and automate application failover and control. When a system or an application under VCS management fails, other systems can take over based on a predefined resiliency plan to bring up services elsewhere in the cluster or in a different cluster.

InfoScale provides high availability and resiliency for systems and applications with a focus on two overarching principles:

- Application management: VCS is a flexible high availability and resiliency management solution that allows you to deploy your applications and databases in a variety of highly available and resilient topologies.
- Enterprise storage management: With InfoScale software-defined storage, you can run your applications and databases with several advanced features — including a clustered file system — that are not provided with native system tools and resources.

By offering application high availability and storage management as a single solution, InfoScale can help reduce your costs, improve your application performance, increase operational flexibility, and eliminate unplanned downtime for open source applications and databases.

High Availability Management

InfoScale VCS agents manage system components (resources) and application processes required for an application to be online. InfoScale agents provide high availability for specific resources and applications, and each agent manages resources of a certain type. For example, the SAP agent manages SAP components such as NetWeaver and HANA databases, and the PostgreSQL agent manages database processes and replication. Typically, agents start, stop, monitor resources, and report state changes. InfoScale has agents for most tier 1 applications, and has a generic agent that can be used to manage nearly any type of application — including open source databases.

Intelligent Monitoring

Application and system monitoring is a key part of effectively managing application high availability. While most solutions use poll-based monitoring that periodically poll resources to monitor application status, InfoScale has an advanced monitoring capability known as the Intelligent Monitoring Framework (IMF), based on InfoScale’s asynchronous monitoring framework kernel driver. With instant notifications from the kernel driver when an event occurs, IMF significantly reduces the time required to detect application failures. IMF allows the VCS agents to register managed resources 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 enables InfoScale availability agents to act immediately in the event of a system fault, and reduces system overhead.

Figure 2 shows an overview of how InfoScale’s IMF provides more advanced monitoring for a PostgreSQL database versus a traditional poll-based monitoring solution.

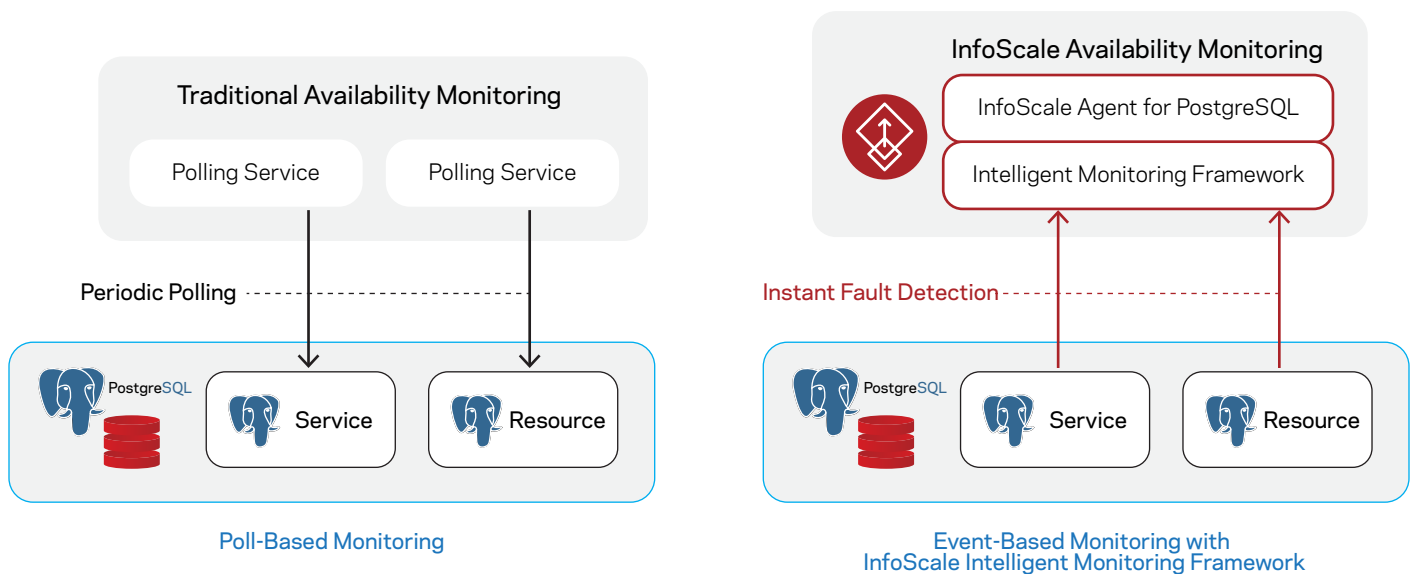


Figure 2. InfoScale’s Intelligent Monitoring Framework provides instant fault detection

Enterprise Storage Management

Ensuring that your applications and databases are deployed with robust storage management is essential to optimize performance and eliminate unplanned downtime. InfoScale’s storage management functionality has a long-standing history of delivering enterprise-grade storage for database management systems. InfoScale’s storage management functionality is based on:

- Veritas File System (vxfs): An advanced file system capable of managing large volumes of data, designed to provide high performance and availability for applications and databases.

- Cluster File System (CFS): A feature of vxfs that provides highly available, parallel access for applications and databases deployed in a clustered configuration. CFS provides better lock management than NFS-based storage services and makes failover more reliable.
- Flexible Storage Sharing (FSS): A feature of CFS that allows logical volumes to be created using block storage, enabling a common storage namespace without requiring physically shared storage. FSS is transparent to file systems and applications, and can be implemented using on-premises block storage (SAN or DAS) and most cloud-native block storage services.
- Veritas Volume Replicator (VVR): Enables platform-independent disaster recovery and resiliency by intelligently managing replication for open source databases. VVR has several advanced features including Adaptive Sync, which improves sustained throughput for latency-sensitive applications by automatically switching from synchronous to asynchronous mode and vice versa based on latency.

InfoScale also includes a feature called Smart I/O that can help improve your database performance by intelligently caching frequently accessed data on fast storage devices for faster reads. By deploying your open source databases on InfoScale-managed storage, you can improve performance, reduce costs, and eliminate unplanned downtime. Figure 3 is an example of how PostgreSQL can be made highly available using the PostgreSQL VCS agent, CFS, FSS, and Smart I/O for performance optimization.

InfoScale Application Agents

InfoScale provides several application-specific agents designed to simplify the management of applications by including application-specific parameters to be monitored and managed directly in the agent configuration management interface.

InfoScale also includes a general use agent — known as the VCS application agent — that enables you to work with nearly any application, including open source applications and databases. For some databases such as PostgreSQL, a specific agent can be used. In other situations, open source application and database high availability and resiliency can be managed using the VCS application agent. The VCS application agent is customizable and provides several benefits for open source applications and databases:

- Flexibility: Manage availability and resiliency for nearly any application or database with minimal configuration.
- Simplicity: The VCS application agent automates manual processes and is included by default with the InfoScale software for use with any InfoScale-managed system.
- Monitoring: The VCS application agent works with the Intelligent Monitoring Framework and offers more advanced and efficient monitoring capability compared to traditional application monitoring.

InfoScale's VCS application agent performs online, offline, and monitor operations. The VCS application agent can also detect discrepancies between the VCS configuration and the underlying infrastructure on a system. This includes things such as the availability

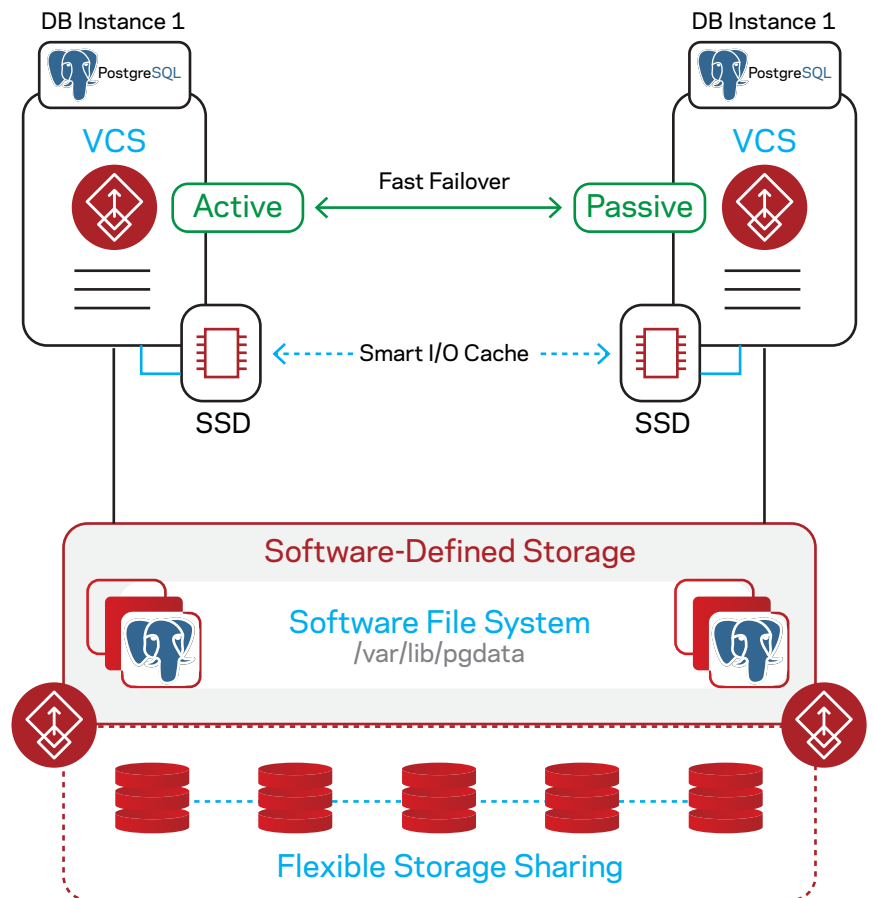


Figure 3. PostgreSQL deployed in a highly available configuration with InfoScale availability and storage management

of a program, execution permissions, and binaries that are needed by applications being managed.

InfoScale also provides a feature called Virtual Business Services (VBS) that allows you to manage multi-tier applications and all their dependencies as a single logical entity (IT business service). This helps maximize application availability by automating operations and reducing the need for manual intervention in the event of an application or system fault. Figure 4 is an example of how a VBS can be configured to manage availability for a two-tier application with an open source database deployed in a highly available configuration on InfoScale's cluster file system.

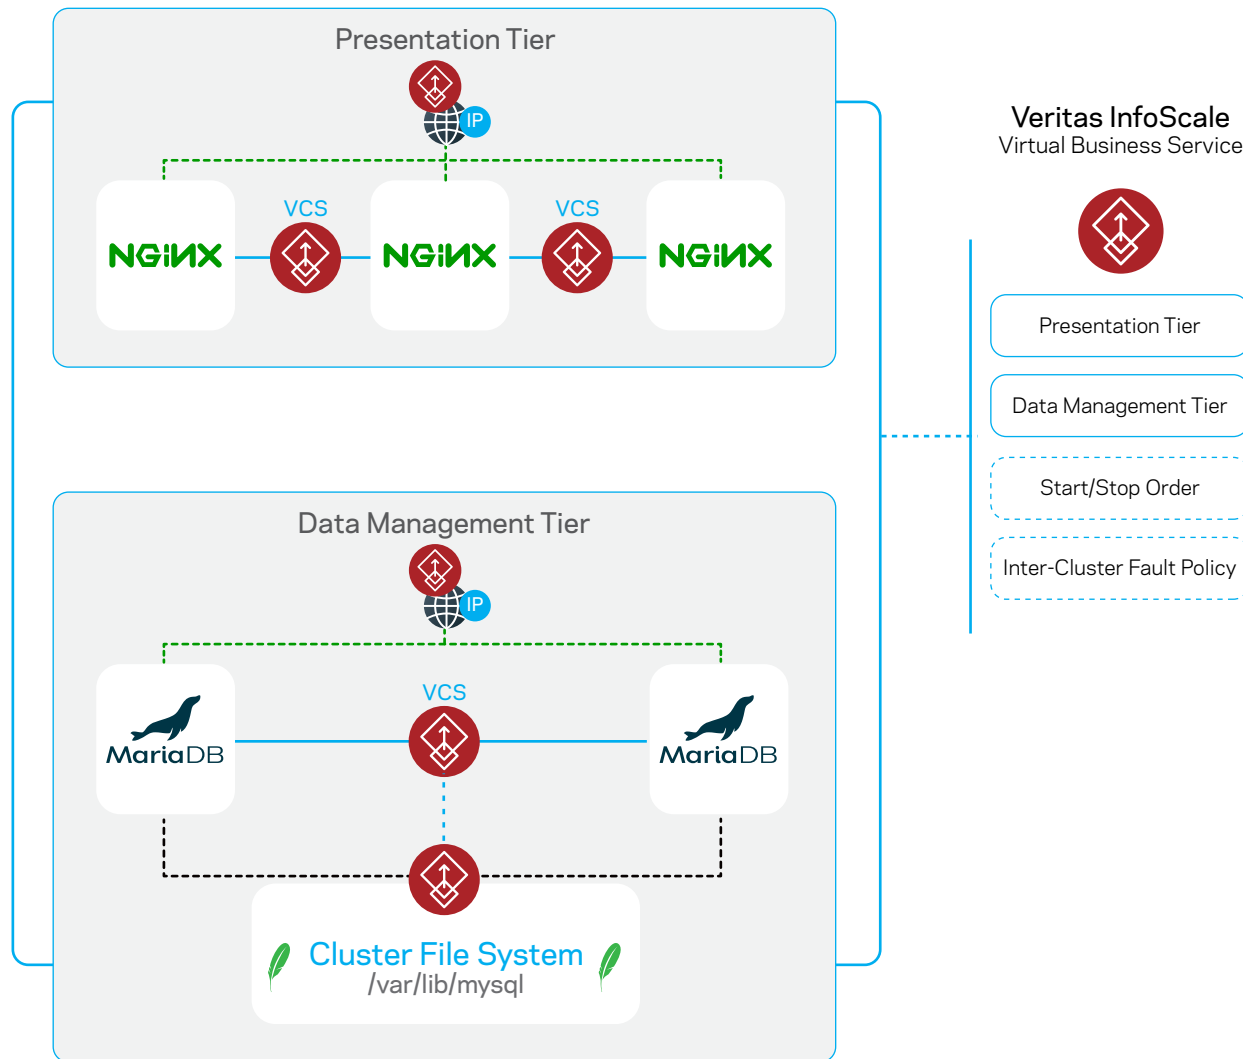


Figure 4. InfoScale automates and simplifies high availability management for open source applications and databases

PostgreSQL Agent

InfoScale natively includes an agent custom developed to manage high availability for PostgreSQL databases. The InfoScale agent for PostgreSQL can manage and provide high availability for PostgreSQL servers and EDB Postgres Advanced Server clustered environments. It can bring a PostgreSQL instance online, monitor its state, detect failures, and shut down the instance as required. There is also an agent for PostgreSQL replication that can provide high availability for the PostgreSQL Streaming replication feature. The PostgreSQL agent can manage several Streaming replication operations:

- Monitor the replication link on the master node
- Automatic takeover on the standby node
- Promote the standby (failover) node to the primary in case there is a database failure on the master node

The InfoScale agent for PostgreSQL works with the InfoScale IMF, which can provide instant fault detection for PostgreSQL processes. In this situation, the InfoScale agent for PostgreSQL registers the process IDs of the PostgreSQL processes, and the agent for PostgreSQL replication registers the Postgres WalSender process with the InfoScale kernel driver. In the event of a state change for any registered resource, the InfoScale agents take the appropriate action based on the state of the resource(s).

In addition to high availability monitoring, InfoScale can provide shared storage services for PostgreSQL databases. This can be done for single instance databases or more complex configurations where the database is configured as a cluster using InfoScale's storage management functionality with a cluster file system. The following graphic represents an example of a PostgreSQL database configured in a highly available topology. It also includes a geographically dispersed recovery site and InfoScale storage management providing advanced shared storage services.

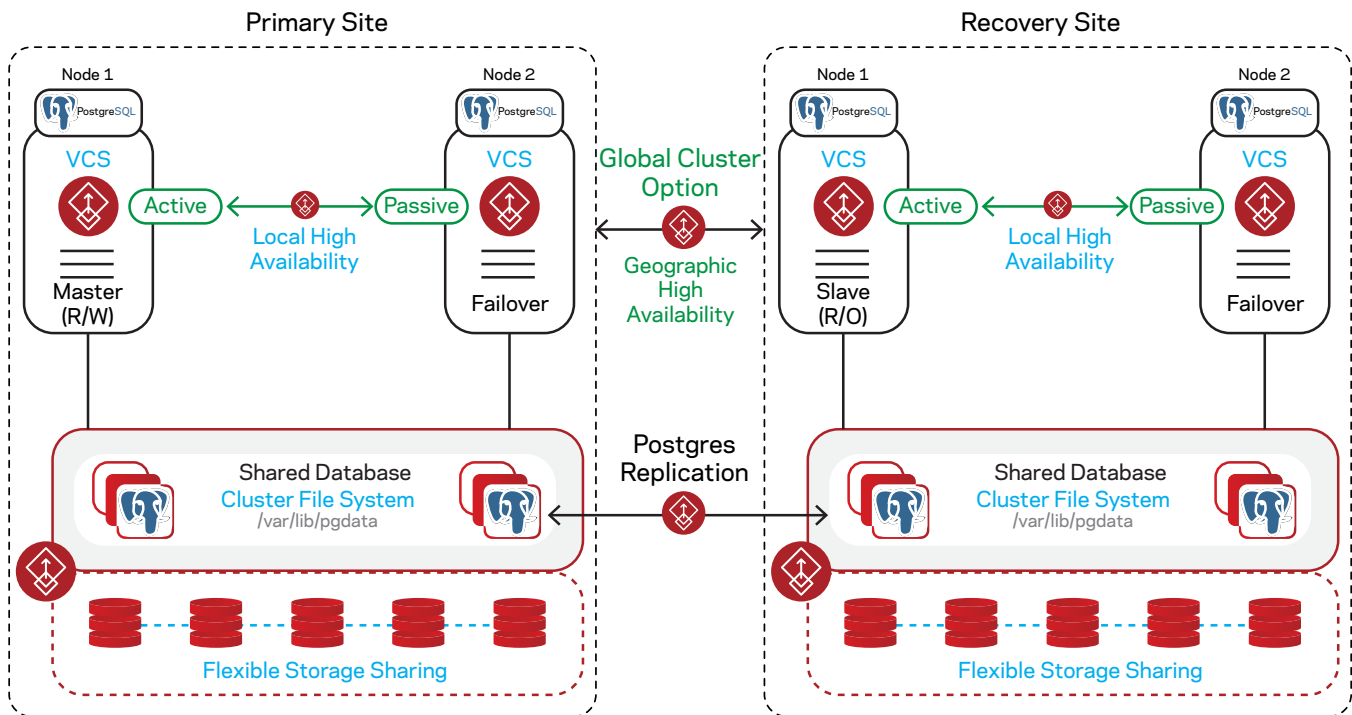


Figure 5. PostgreSQL deployed in a highly resilient architecture with InfoScale VCS and Cluster File System

MySQL Agent

In addition to the PostgreSQL agents, InfoScale includes an agent for MySQL that provides high availability for MySQL databases. The MySQL agent provides the following functionality for MySQL databases:

- Online and offline operations that verify the current state of the MySQL server instance(s) to ensure that a start or stop operation is necessary and completes successfully when executed
- Monitoring with the InfoScale IMF with instant fault detection for MySQL resources and processes
- A clean function that will attempt a graceful shutdown of the MySQL server instance(s) and any processes pertaining to this instance in case there is a failure

Figure 6 shows an example of how the InfoScale MySQL agent can automate failover for a MySQL database running in the cloud.

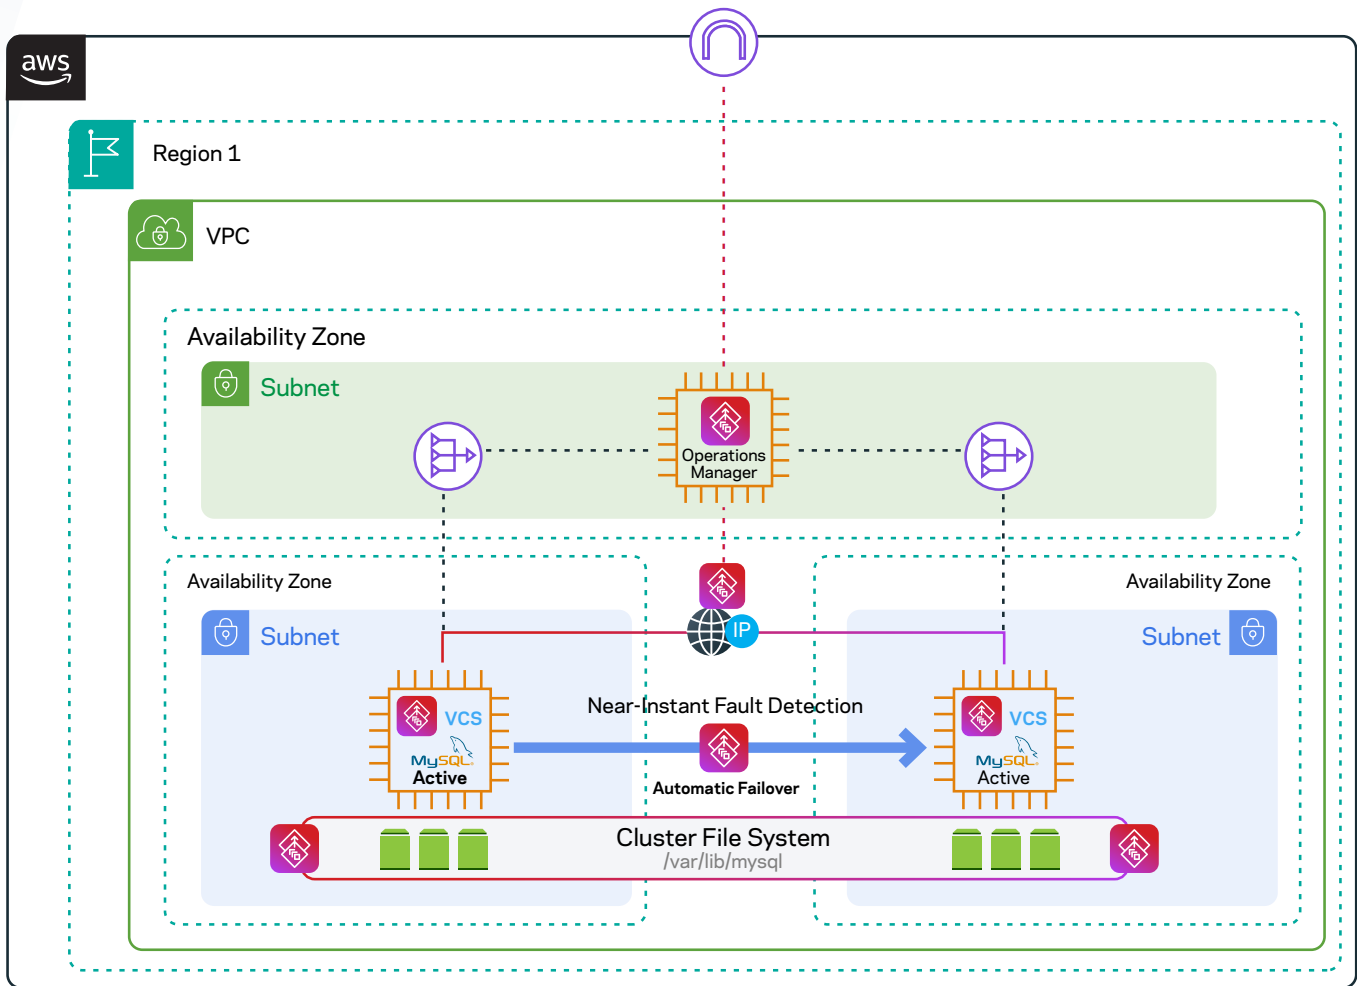


Figure 6. MySQL deployed in a highly available configuration in the cloud using the InfoScale VCS MySQL agent

High Availability Topologies

InfoScale allows you to manage applications and databases using a variety of topologies that can be customized to suit your SLAs and business requirements. This includes basic application and system management configurations for non-critical applications, as well as advanced configurations that can provide near-zero recovery point objectives (RPO) and recovery time objectives (RTO) for business-critical applications. With InfoScale, you can automate high availability management and confidently run open source databases in a variety of topologies focused on:

- High Availability: Deploy your database with automated local high availability to help deliver an always-on end user experience.
- Disaster Recovery: Configure databases using an active-passive topology, where you can recover applications in the event of an outage.
- Performance and Resiliency: Using an active-active topology, multiple nodes work together to manage and process I/O, which can deliver better application performance while providing resiliency and downtime in the event of system failures.

Figure 7 shows an example of InfoScale — known as Veritas Alta™ Enterprise Resiliency when deployed in the cloud — managing an advanced resiliency topology in a public cloud environment that improves performance; and achieves a near-zero RPO over any distance, and a near-zero RTO for an open source database.

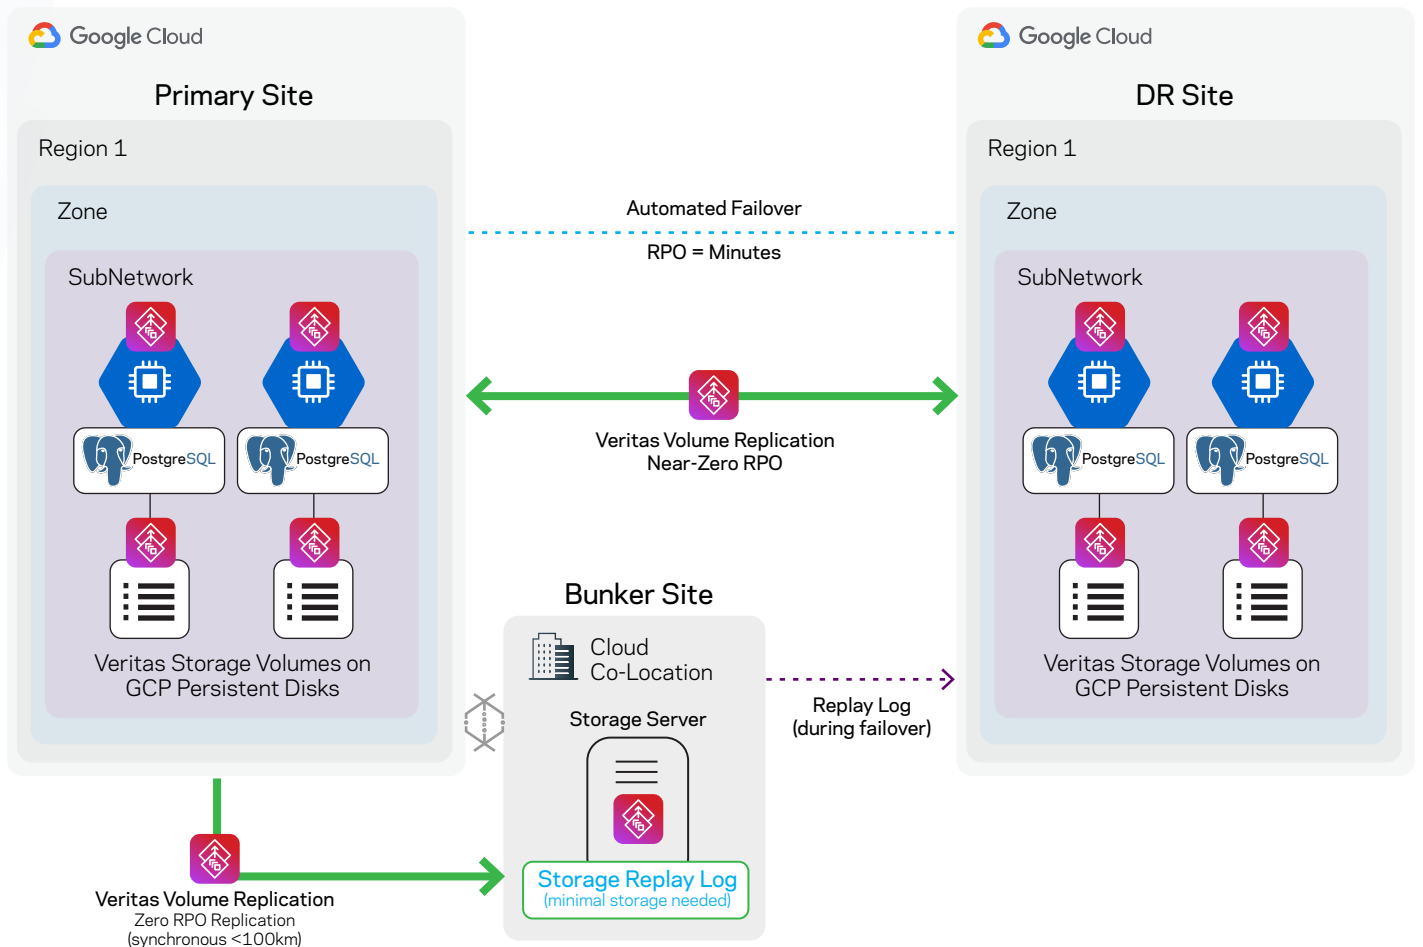


Figure 7. InfoScale high availability and resiliency with a near-zero RPO and RTO for a PostgreSQL database in the cloud

Centralized Operations

InfoScale includes an intuitive web-based console designed to help centralize operations. The 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 provides visibility into non-InfoScale infrastructure, so you can use it to identify potential issues in most environments that could lead to unexpected downtime.

Data from the VCS agents and infrastructure discovery operations is presented in VIOM and can be used for various purposes including analysis and system management. This can provide significant value for open source databases by helping you visualize generic system information that may not be natively managed or visible within the database software. VIOM can help you:

- Increase operational efficiency with integrated patch management and reports that provide information about HADR readiness for your applications and databases
- Optimize application performance and availability with a full application to system view of your operating environment
- Centralize visibility and control to reduce downtime, maintain compliance, and easily visualize operational status with system level reporting and proactive risk analysis

Figure 8 provides a high-level overview of how VIOM works with all components in the operating stack to help ensure full visibility into your environment, while helping to improve efficiency and reduce risks.

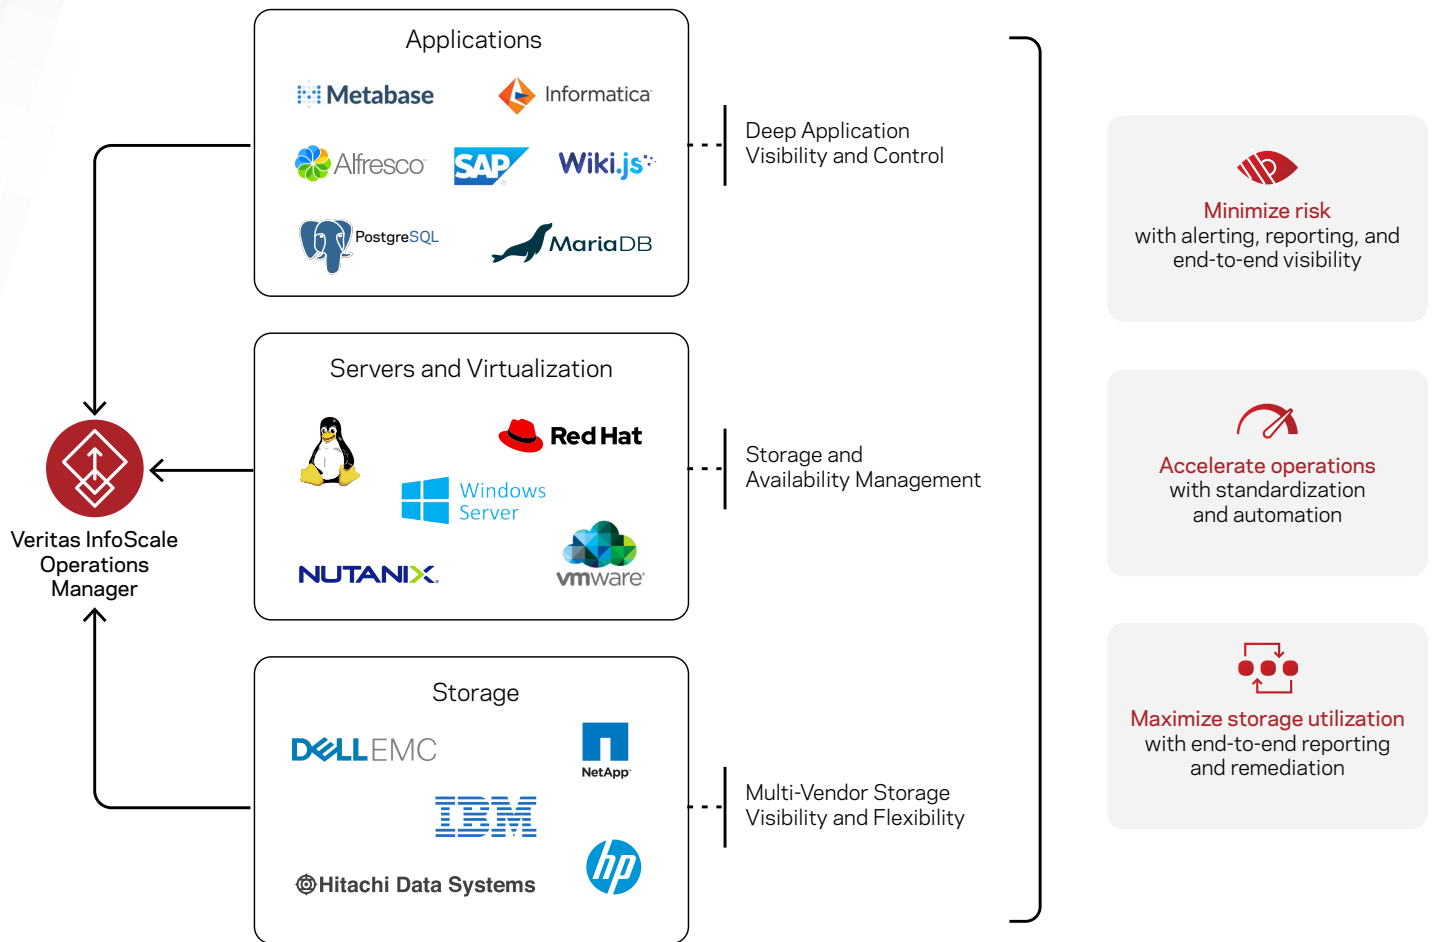


Figure 8. VIOM centralizes management and visibility for applications and infrastructure

VIOM provides an API that can be accessed using HTTPS with any standard HTTPS client. The interface provides the ability to query discovered data and manage user-defined attributes for InfoScale components — including the VCS agents used to monitor open source applications and databases.

Database Lifecycle Management

InfoScale has several advanced features that make it an ideal solution to manage database lifecycles, disaster recovery, and archival. Features such as Database FlashSnap, Volume Replicator, and snapshots can help you manage your databases more efficiently and contribute to a resilient IT environment.

Database Cloning

InfoScale's Database FlashSnap enables you to create a point-in-time copy of a database for backup and off-host processing. With Database FlashSnap, you can quickly and easily make backup copies of your volumes online and with minimal interruption to your end users.

Database FlashSnap also lets you capture an online image of an actively changing database at a given instant — known as a database snapshot. You can use database snapshots on the same host as the production database or on a secondary host sharing the same storage. A database snapshot can be used for off-host processing applications, such as backup, data warehousing, and decision-support queries. When the snapshot is no longer needed, a database administrator can import the original snapshot back to the primary host and resynchronize the snapshot to the original database volumes.

With Database FlashSnap, you can mirror the volumes used by the database to a separate set of disks, and those mirrors can be used to create a snapshot of the database. These snapshot volumes can be split and placed in a different disk group. This snapshot disk group can be imported on a separate host, which shares the same storage with the primary host. The snapshot volumes can be resynchronized periodically with the primary volumes to capture recent changes in the datafiles. If the primary datafiles become corrupted, you can quickly restore them from the snapshot volumes. Snapshot volumes can be backed up and recovered in addition to database cloning.

Replication

Veritas Volume Replicator (VVR) is a fully integrated component of the Veritas Volume Manager (VxVM) that enables optimized replication for data volumes. With VVR, you can:

- Replicate existing VxVM configurations
- Transparently configure replication while an application/database is active
- Replicate application writes on volumes at a source location to one or more remote areas across any distance
- Provide a consistent copy of application data at remote locations

VVR provides a consistent copy of application data at the remote sites. If a disaster occurs at the source location, you can use a copy of the application data at the remote location and restart the application at the remote site. The host at the source location on which the application is running is known as the Primary host, and the host at the target location is known as the Secondary host. You can have up to 32 Secondary hosts in a VVR configuration.

VVR must initially synchronize the Primary host volumes with the volumes on the Secondary host. VVR provides several methods to initialize the application data between the primary location and the remote location, such as using the network, using tape backup, and moving disks physically.

VVR replication operates synchronously or asynchronously and can automatically switch between synchronous and asynchronous replication based on the network bandwidth available for replication. Synchronous replication occurs when the Primary host must wait for the Secondary host to complete the write operation to disk. Asynchronous replication does not require the Secondary write to complete.

The Secondary host's role is easily converted to a Primary role to maintain IT operations when a disaster occurs. A single command can be issued to enable the migration of roles with little to no data loss. Likewise, when the original host has recovered, the Primary role can be transferred back. VxVM features the ability to fail back to a Primary host without performing a complete resynchronization, minimizing recovery time.

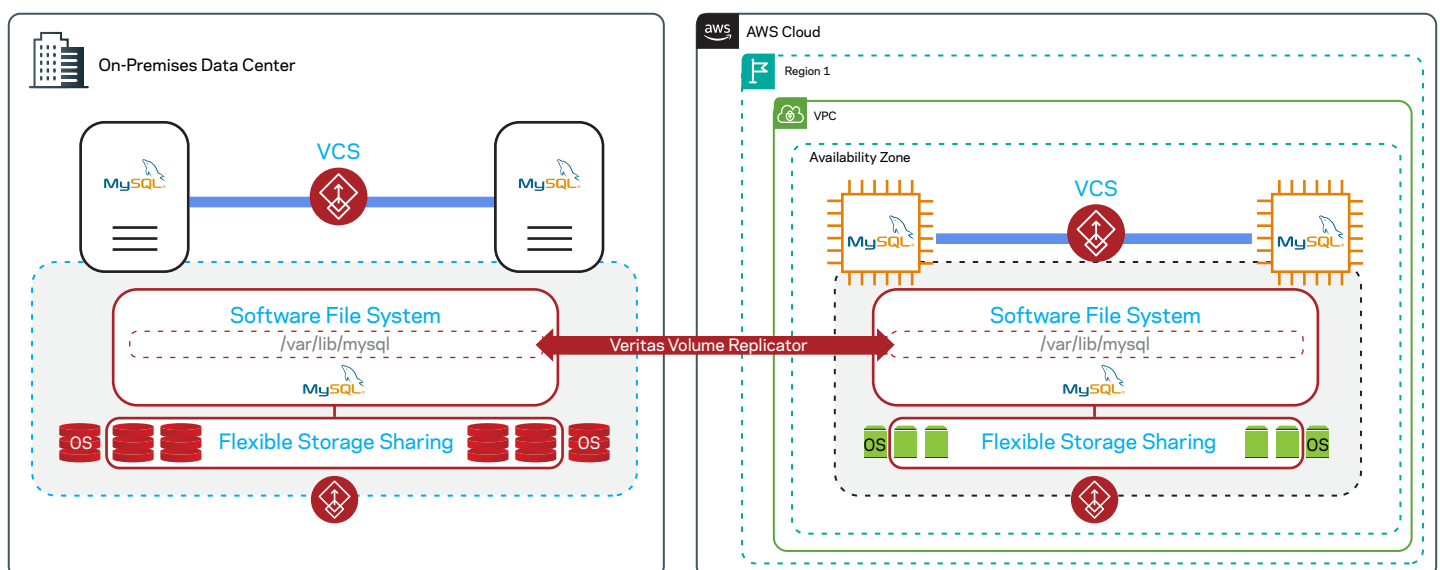


Figure 9. Veritas volume replication for an open source database in a hybrid cloud architecture

Data Warehousing

Using a strategy that combines VVR replication and cloning, InfoScale can help you manage a data warehousing solution. At a high level, you can automate the process of replicating data to a secondary site, which can then be cloned and fed into a data warehouse. If the primary data is no longer needed, it can be truncated, allowing storage space to be reclaimed. In Figure 10, data is extracted from a clone of the secondary copy of the database. When the data is replicated to the secondary server, it cannot be used until a clone is created. Once the clone is created, the data can be read and loaded into the data warehouse. Once the process is complete, the clone can be destroyed. The archived records can be safely purged from the primary site database.

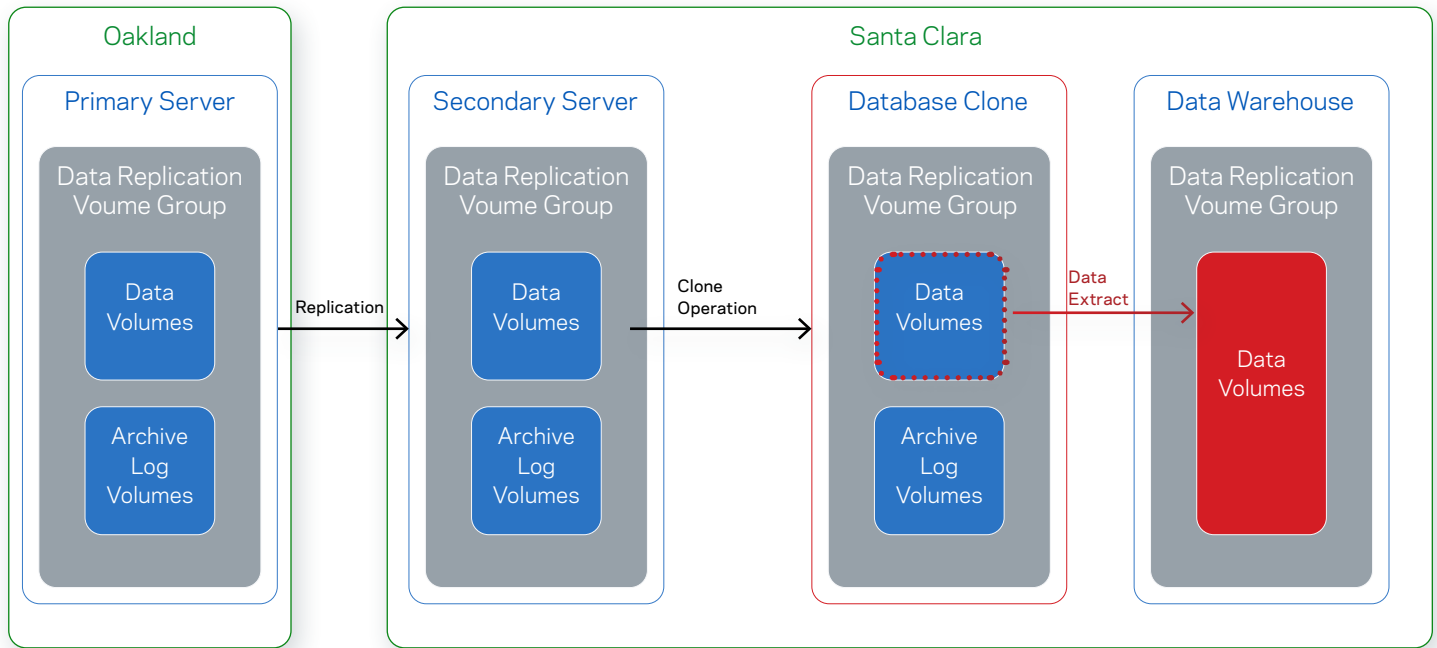


Figure 10. InfoScale's replication and cloning functionality can be used for data warehousing

Summary

Open source applications and databases have become a mainstream and robust option for managing IT services of all types. Ensuring that these IT services deliver a premier user experience is always top of mind, and this requires a proven enterprise platform that can effectively deliver high availability and resiliency for your IT services. InfoScale is the ideal foundation for open source applications and databases that enables you to operate on nearly any platform with confidence — while helping to reduce cost and complexity. InfoScale provides several key benefits for open source applications and databases:

- **Flexibility:** Delivers a single solution for HADR and storage management that works with nearly any type of application and database
- **Availability:** Ensures your applications and databases are highly available and performing as expected, to achieve a premier end user experience
- **Resiliency:** Protects your IT services against unplanned downtime with a solution that delivers zero data loss and near-zero recovery times; and can be used in on-premises, cloud, hybrid cloud, and multi-cloud environments

InfoScale is a comprehensive, multi-faceted solution with a proven track record of improving availability and resiliency for open source applications and databases. It helps avoid unplanned downtime and provides the foundation you need to run your open source solutions with maximum flexibility, availability, and confidence.

About Veritas

Veritas Technologies is the leader in secure multi-cloud data management. Over 80,000 customers—including 91% of the Fortune 100—rely on Veritas to help ensure the protection, recoverability and compliance of their data. Veritas has a reputation for reliability at scale, which delivers the resilience its customers need against the disruptions threatened by cyberattacks, like ransomware. No other vendor is able to match the ability of Veritas to execute, with support for 800+ data sources, 100+ operating systems and 1,400+ storage targets through a single, unified approach. Powered by Cloud Scale Technology, Veritas is delivering today on its strategy for Autonomous Data Management that reduces operational overhead while delivering greater value. Learn more at www.veritas.com. Follow us on X 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