

Veritas InfoScale™ Enterprise for Oracle® Real Application Clusters (RAC)



Manageability and availability for Oracle® RAC databases

Overview

Veritas InfoScale™ Enterprise for Oracle® Real Application Clusters (RAC) offers a proven solution to help customers implement and manage highly available RAC databases. InfoScale Enterprise for Oracle RAC enables IT organizations to select the most appropriate operating system and storage hardware for their environment, all without compromising management capabilities. Being hardware platform agnostic, this solution allows IT teams to have a single view for database, clustering, and storage management tasks while facilitating installation, configuration, and central management. Moreover, it increases flexibility, scalability, and performance while reducing system downtime.

Key Benefits

- **Scalable database performance**—Utilize database accelerators and multiple physical paths to disks
 - **Identify and remove I/O bottlenecks**—Map database objects down the storage hierarchy to the disks
 - **Enable stretch RAC environments**—Achieve high availability and disaster recovery by utilizing a campus cluster configuration for Oracle RAC
 - **Reduce database storage costs**—Automatically move less frequently used data to slower, less expensive disks
- **Simplify management of Oracle RAC**—Centralize multi-node management and make RAC almost as easy to manage as a single-node non-RAC database
 - **Support for Oracle RAC 12c**—Supports installation of and upgrade to Oracle RAC 12c and use of multi-tenant (container and pluggable database) architecture
 - **Flexible Storage Sharing (FSS)**—Combine shared and direct-attached storage (DAS) on Linux®, IBM® AIX®, and Oracle® Solaris™ platforms
 - **SmartIO**—Granular and intelligent application-level caching; enables storage QoS for mission-critical applications on Linux, AIX, and Solaris platforms
 - **Facilitate off-host processing**—Create easy-to-use database clones to enable data analysis and backups
 - **Increase efficiency of database backup and recovery**—Protect the database from logical errors by providing point-in-time copies
 - **Ensure data integrity**—Eliminate the risk of data corruption in the event of a “split brain” condition

Key Benefits

While most application cluster implementations are intended to increase availability, Oracle RAC also attempts to improve scalability by using multiple servers for the same workload. This scale-out approach to clustering introduces management complexity to server, storage, and database administration. InfoScale Enterprise for Oracle RAC minimizes this complexity by enhancing the native capabilities of Oracle RAC with a highly available, scalable, non-disruptive server and storage management solution that is independent of operating system and storage hardware. The centralized management capability enables users to add and remove nodes and storage capacity without impacting application availability. InfoScale Enterprise enables Oracle RAC tablespaces to grow online without the need to pre-allocate storage capacity. A single cluster file system and volume management tool facilitates creation of a shared Oracle home that simplifies ongoing maintenance and patch management.

Support for Oracle RAC 12c

InfoScale Enterprise for Oracle RAC supports installation of and upgrade to Oracle RAC 12c. This includes support for existing features of Oracle RAC in 12c and the use of multi-tenant (container and pluggable database) architecture.

SmartIO

Enable granular and intelligent caching at the application level. With SmartIO, application architects, server administrators, and database administrators can move reads inside the server and enable storage Quality of Service (QoS) for mission-critical applications. This feature is supported on Linux, AIX, and Solaris platforms.

Flexible Storage Sharing (FSS)

Simplify SSD/Flash adoption and realize the performance benefits without compromising on high-availability and flexibility. Using FSS, scale-out servers with Direct-Attached Storage (DAS) and experience near-local read and write performance to and from remote disks. Logical volumes can be created using both types of storage creating a common storage namespace. Logical volumes using network shared storage provide data redundancy, high availability, and disaster recovery capabilities—without requiring physically shared storage—transparently to file systems and applications.

FSS use cases include support for current InfoScale Enterprise for Oracle RAC use cases, off-host processing, DAS benefits leveraged with existing InfoScale Enterprise for Oracle RAC features, FSS with File System level caching, and campus cluster configuration. Installing InfoScale Enterprise for Oracle RAC automatically enables the FSS feature and no separate license is required. This feature is supported on Linux, AIX and Solaris platforms.

Facilitate off-host processing

Relational databases provide a single view of the data to all applications. While this dramatically improves the quality of information available to users and managers, it constrains an enterprise's ability to use the same data for data analysis or backup.

Both require point-in-time images, which cannot be made while a database is being updated by business applications. The FlashSnap feature enables administrators to set up reusable point-in-time copy policies, verify storage resources, and create full online database volume clones and space-saving file system checkpoints with minimal impact on production applications and users. Database snapshots can be migrated to secondary servers and used for resource-intensive processes such as backups, decision support, reporting, and testing— independent of the storage hardware being used.




Figure 2. Migrating database snapshots to secondary servers




Figure 1. Flexible Storage Sharing (FSS) delivers performance without compromising on flexibility

Increase efficiency of database backup and recovery

InfoScale Enterprise for Oracle RAC enables efficient backup and recovery of Oracle RAC databases using Storage Checkpoint technology. Storage Checkpoint is a disk- and I/O-efficient snapshot technology for creating a “clone” of a currently mounted file system (the primary file system). Unlike a full file system copy that uses separate disk space, all Storage Checkpoints share the same free space pool where the primary file system resides, greatly reducing the need for extra storage. A direct application of the Storage Checkpoint facility is Storage Rollback. Because each Storage Checkpoint is a consistent, point-in-time image of a file system, Storage Rollback is the restore facility for these on-disk backups. Storage Rollback simply rolls back blocks contained in a Storage Checkpoint into the primary file system for very fast database recovery.

Ensure data integrity

When multiple systems/nodes have access to data via shared storage, the integrity of the data depends on internode communication ensuring that each node is aware when other nodes are writing data. When the coordination between the nodes fails, it results in a “split brain” condition—a situation in which two servers try to independently control the storage, potentially resulting in application failure or even corruption of critical data, which can then require days to recover, if recovery is even possible.

I/O fencing is the Veritas method of choice for ensuring the integrity of critical information by preventing data corruption. InfoScale Enterprise for Oracle RAC has implemented I/O fencing using the industry standard SCSI-3 persistent group reservation technology, as well as alternate non SCSI-3 fencing technology allowing a set of systems to have temporary registrations with the disk and coordinate a write-exclusive reservation with the disk containing the data.

With I/O fencing, Veritas ensures that errant nodes are “fenced” and do not have access to the shared storage, while the eligible node(s) continue to have access to the data, virtually eliminating the risk of data corruption.




Figure 3. Implementing I/O fencing

Scalable database performance

There is a strong movement towards the consolidation of disparate database systems onto even larger RAC clusters. The major concern in any consolidation effort is maintaining respectable performance and/or meeting committed performance service-level agreements (SLAs). InfoScale Enterprise for Oracle RAC improves the overall performance of database environments by providing extensions to Oracle Disk Manager (ODM), a database accelerator technology that enables Online Transaction Processing (OLTP) performance equal to raw disk partitions, but with the manageability benefits of a file system. It delivers the same performance benefits as Quick I/O, but also provides tight database integration for easier manageability. Moreover, with the Dynamic Multi-pathing feature of InfoScale Enterprise for Oracle RAC, performance is maximized by load-balancing I/O activity across all available paths, from the server to all major hardware RAID array products. Further, with this feature, there is no need for third-party multi-pathing software, reducing the total cost of ownership.

Identify and remove I/O bottlenecks

The challenge with maximizing database performance is having visibility from the data file to the storage spindle. The storage mapping feature in InfoScale Enterprise for Oracle RAC provides organizations with a proven approach to that visibility, when used with Veritas InfoScale™ Operations Manager. When organizations have detailed database mapping information, a detailed understanding of the storage hierarchy, and knowledge of where each data file resides, performance bottlenecks can be eliminated. Armed with this information, IT organizations can minimize I/O performance bottlenecks by dynamically moving data files to different logical units on different physical spindles, or to another array altogether. The performance tuning capabilities of InfoScale Enterprise for Oracle RAC gives organizations the flexibility to be as sophisticated as they choose, or let InfoScale Enterprise manage data architecture with hot relocation, a feature that automatically detects and replaces failed disks using a free disk pool.

Enable stretch RAC environments

A campus cluster configuration provides local high availability and disaster recovery capability in a single InfoScale Enterprise for Oracle RAC cluster. This configuration uses data mirroring to duplicate data at different sites. No host or array replication is involved in the process. InfoScale Enterprise for Oracle RAC supports campus clusters that employ shared disk groups mirrored with Cluster Volume Manager (CVM).

A InfoScale Enterprise for Oracle RAC campus cluster is similar to a basic InfoScale Enterprise for Oracle RAC cluster except that the data is mirrored across multiple sites. When a site fails, the InfoScale Enterprise for Oracle RAC cluster in the secondary site continues running and as parallel applications are already running on the secondary nodes, they are not affected. The ability to run RAC in a campus cluster configuration results in a more highly available disaster recovery environment. Using a Coordination Point Server (CPS) in a campus cluster environment enables the use of a third site as an arbitration point without requiring SAN connectivity to the third site.

Reduce database storage costs

As the size of a relational database grows, so does the size of the inactive data. The result: a steep rise in storage costs and an abysmal drop in database performance. This solution matches data storage with the data's usage requirements so that data is relocated to less expensive storage based on policies defined by the administrator. SmartTier enables organizations to define data movement via policies based on partition name, log files or database files. Data movement can be defined for file objects as well as entire files on individual volumes.

The key benefit of this capability is that less-frequently used data is created on or automatically moved to slower, less expensive disks, allowing frequently-accessed data to be stored on the faster disks for quicker retrieval. Moreover, data file creation or movement can occur without taking the database offline and can be completely transparent to users and applications.

Other Product Highlights

- **Storage capacity planning**—Simulate various Storage Checkpoint creation and retention models in a production environment
- **Flexible management**—Offer intuitive Web or command line interface options for local or remote management
- **Intelligent workload management**—Increase automation of cluster administration, maximize application uptime, and improve utilization of server resources
- **Cluster-wide logical device naming**—Simplify management of SAN-based storage
- **Support for Oracle® Enterprise Manager Grid Control**—Monitor Veritas InfoScale™ Availability, raise alerts and violations based on resource stack and map database objects on the Veritas storage stack
- **Support for Oracle Enterprise Manager 12c**—Monitoring clustering and administering SmartIO and snapshots

Supported operating systems

- IBM AIX
- Oracle Solaris
- Red Hat® Linux
- SUSE® Linux
- Oracle® Enterprise Linux

¹ SmartTier includes and extends capabilities previously offered by Dynamic Storage Tiering (DST)

More Information

Visit our website
www.veritas.com

About Veritas Technologies LLC

Veritas Technologies LLC enables organizations to harness the power of their information, with solutions designed to serve the world's largest and most complex heterogeneous environments. Veritas works with 86 percent of Fortune 500 companies today, improving data availability and revealing insights to drive competitive advantage.

Veritas World Headquarters

500 East Middlefield Road
Mountain View, CA 94043
+1 (650) 933 1000
www.veritas.com