

A horizontal decorative bar consisting of a solid red segment on the left, followed by four segments with abstract, colorful patterns in shades of blue, orange, green, and purple.

## Oracle 9i RAC and the Red Hat Global File System

### **Abstract**

This application brief addresses how Red Hat Enterprise Linux and the Red Hat Global File System can simplify shared storage management for Oracle 9i RAC deployments.

June 2004



## **The Challenge: Providing ease of management for Oracle 9i RAC deployments on shared storage clusters**

Oracle 9i RAC is a scalable, highly-available, clustered database that runs efficiently on Red Hat Enterprise Linux and low-cost servers and storage. In December 2003, Red Hat Enterprise Linux and Oracle RAC set the world record for the TPC-C benchmark, achieving over 1 million transactions per minute on a shared storage cluster of 16 industry servers. Users are scaling Oracle databases without incurring the equipment or support costs of large SMP servers.

Managing a shared storage hardware environment required by Oracle9i RAC can be complex. Management issues include

1. A separate Oracle installation must be performed on each 9i RAC node.
  - More than 100,000 files must be installed on each node (this includes the Oracle Home and root directories)
  - Home and root directories must be separately managed for each node since these directories cannot be shared
2. Unless 9i RAC is deployed on top of a file system, a minimum of nine raw devices (one for each table space) must be managed, plus additional raw devices for user space.
  - Raw devices cannot be resized
  - Raw devices appear as “unused” space and can be inadvertently written over by other applications if proper care is not taken
  - Certain advanced Oracle features are not supported on raw disk (ONF, ELT, and I/O fencing for example)


## **The Solution: Simplified Shared Storage Management with Red Hat Enterprise Linux and GFS**

With Red Hat Enterprise Linux and GFS, 9i RAC management is similar to a single instance of the Oracle database (as opposed to managing multiple Oracle Home directories and node specific information associated with each node). With a Red Hat Enterprise Linux and GFS solution, you can:

1. Install Oracle once with a single Oracle shared home and shared root.
2. Expand and manage shared storage for 9i RAC with the ease of a file system but with high-performance of direct I/O to disk.
3. Run POSIX-compliant applications on 9i RAC nodes.
4. Avoid data corruption and downtime with the leading Linux operating system — Red Hat Enterprise Linux — and the most mature and stable cluster file system for Linux available today.

## The Raw Devices Approach: Oracle 9i RAC Configuration

A typical Oracle 9i RAC installation is shown below.



Oracle 9i RAC installation requires the following:

### 1. Multiple Oracle instances.

- An Oracle Home for each node and node-specific information (NSI) associated with each node. An eight-node configuration will have:
- Eight local disks attached to each node (or similarly, eight physical or virtual storage devices on a storage network)
- Eight Oracle Home directories
- Node-specific information on each of the eight local disks – requiring each local disk to be managed and protected

### 2. A raw storage device for each table space.

- Results in a minimum of 10 raw devices to manage when required Oracle tables (system, undo, redo, temp, users, tools, index) and user space is considered
- Limited to 255 raw devices


*Note: Other servers that may share the same storage device (e.g. over a SAN) cannot see the raw devices that belong to 9i RAC. Therefore, Oracle data can be inadvertently overwritten by other servers that share the storage device.*

### 3. Unnecessary downtime and administrative effort.

- Raw devices cannot be dynamically expanded
- Updates to the Oracle home directories on each node are required
- Raw devices can be much more difficult to backup, restore, and manage than file-based storage

## The Red Hat Solution: Red Hat Enterprise Linux, Red Hat GFS and Oracle 9i RAC

A typical Oracle 9i RAC solution with Red Hat GFS is shown below.



Oracle 9i RAC is made to access data through Red Hat GFS. Red Hat GFS clusters multiple nodes so that all nodes can have equal access to the same files. Red Hat GFS supports direct I/O access to disks for maximum performance. Configuration and manageability is vastly improved. Note the following benefits for the Red Hat Enterprise Linux, Red Hat GFS, and Oracle 9i RAC solution:

1. A Single Oracle Instance – Support for a shared Oracle Home and centralized NSI, 9i RAC is much easier to install and manage. Over 100,000 Oracle files and executables are installed in a single Oracle Home directory. The Oracle Home is then shared by all 9i RAC nodes. Unlike OCFS and certain other cluster file systems, node-specific information (NSI) accessed by 9i RAC can be centrally stored under Red Hat GFS – eliminating the need to separately manage NSI for each node.
2. High-performance (`direct_io`) datafile storage. Create as many data files as are required by the database.
3. Cluster nodes do not have to be exactly the same and can run any POSIX-compliant application.
4. The functionality and stability expected in a mature cluster file system, including quotas and dynamic volume growth, is available to 9i RAC through Red Hat Enterprise Linux and GFS. Downtime to reconfigure and/or format disks for 9i RAC is eliminated and data integrity is assured.
5. Red Hat Enterprise Linux and GFS supports over 300 nodes in a single shared storage cluster – well beyond the supported cluster size in 9i RAC.

*Note: When OCFS is used to provide file system functionality for Oracle 9i RAC, there is no support for POSIX-compliant applications. In addition, all nodes must be identical, resulting in limited scalability.*

## **Results:**

1. Easier to manage shared storage environments to support Oracle 9i RAC
2. Red Hat Enterprise Linux and GFS make administrators more productive because all administrative tasks can be accomplished on a single cluster node.
3. Proven stability, data integrity, performance, compatibility, and Oracle 9i RAC management cost savings.

**[www.redhat.com](http://www.redhat.com)**