



Focus On: Oracle and High Availability

What is High Availability?

Quite simply, 'High Availability' systems must deliver virtually full-time availability. A highly available system must avoid single points of failure by encompassing hardware and/or software redundancy. In the event of a failure, the tasks that are being performed by the failed component are automatically taken over by the backup component.

Oracle offers a number of 'high availability' options:

- **RAC**
- **Data Guard**
- **Failsafe**
- **Replication**

Some of these can be used in conjunction with each other.

This document incorporates information on high availability products and processes, as well cost information and availability required. Our aim is to help DBAs decide which high availability solution best fits their systems.

1) Oracle Real Application Clusters (RAC)

RAC relies on clustered hardware and permits multiple instances to share a single database. It offers high availability, deploying Transparent Application Failover where, in the event of a failure, the active sessions on a failed node or instance are automatically failed over to enduring instances on other nodes. Scalability is also delivered through the ability to spread the database workload across all active nodes.

Based on Oracle's Cache Fusion architecture, Oracle Real Application Clusters enables sharing of frequently accessed data across all the servers in a cluster. If the query request is served by a remote cache, then the block is transferred across the high speed cluster interconnect from one node's cache to another. This 'cache fusing' happens automatically and is transparent to the application. This transparency is the key technology that provides the fast, efficient scaling of Oracle9i Real Application Clusters.

RAC can be combined with Data Guard and Advanced Replication.

2) Oracle Failsafe

Oracle Fail Safe is high availability software that allows users to configure and verify Windows clusters and to automatically fail over Oracle databases and applications. It is included with every Oracle 11g and Oracle 10g license for Windows 2003 Release 2 and Windows 2008.

In the event of a system failure, Oracle Fail Safe works with Microsoft Cluster Server to restart Oracle databases and applications on a surviving cluster node. It also enables flexible fail back when nodes are returned to service; planned failover for rolling cluster upgrades and workload balancing; and scripting of high availability tasks through its command line interface.

Oracle Failsafe is optimised for Windows customers with database and application workloads that can be handled by a single system. Oracle Failsafe solutions can be deployed on Windows NT and Windows 2000 clusters. It supports up to 4 servers in a cluster with Windows 2000.

Supported products include:

- Oracle Databases (Standard and Enterprise Editions of Oracle9i and Oracle8i)
- Oracle Applications release 11i
- Oracle iAS components, including:
 - Oracle Forms Services
 - Oracle Reports Services
 - Oracle HTTP Server
 - Oracle Intelligent Agent
 - Oracle Service for Microsoft Transaction Server (for Oracle8i release 8.1.7)

Oracle Failsafe can be combined with Oracle Data Guard and Advanced Replication.

3) Oracle Data Guard

Oracle Data Guard protects systems from mistakes, corruptions and other disasters that may threaten a database, by providing physical and logical standby database protection and disaster recovery features. It can be used alongside traditional backup, restoration and cluster techniques to deliver data protection and availability.

With Data Guard, standby databases are maintained as consistent copies of the production database. If the production database fails and becomes unavailable, Data Guard can switch any of the standby databases to take on the production role, thereby minimising the impact of the database outage.

Oracle Data Guard can be combined with Advanced Replication and Oracle RAC

4) Oracle Replication

Replication is the process of copying and maintaining replica versions of database objects (e.g. tables) in a distributed database system. It is suitable for applications, which usually need data to be periodically synchronized between central systems and large numbers of smaller, remote sites that often operate disparately.

Replication can also improve performance and increase availability of applications because alternate data access options becomes available. Even if one site becomes unavailable, users can continue to query or even update the remaining locations. For example, if network traffic is high, users can access a local database rather than a remote server to minimise that network traffic.

Data can be replicated in three key ways:

- **Read-only materialized views**, where one master table and multiple remote copies are refreshed at intervals by pulling changes from the primary site to the backup site.
- **Updateable materialized views**, where partial or entire snapshots are taken of the environment and are updated at regular intervals.
- **N-way master replication**, where multiple master tables are updateable from all sites, i.e. changes that are applied at one site are captured and stored locally before being forwarded and applied at each of the remote locations.

Oracle Replication falls in to the realm of 'high availability' because, if parts of the distributed database are down, function can continue as parts of the data will still be accessible. It detects and attempts to automatically resolve conflicts in the data to ensure convergence throughout the replicated environment. It also allows information to be shared across multiple database versions and hardware platforms.

Oracle Replication can be combined with Data Guard and Oracle RAC.