

Oracle's Technologies for High Availability

Sasidhar Pendyala, Oracle Software India Ltd., India Development Centre
(sasidhar.pendyala@oracle.com)

Introduction :

The world never stops, Neither should one's business, servers, or applications. The consequences of downtime, whether planned or unplanned, can be disastrous for mission critical applications like e-commerce, defense, aviation systems etc. Generally downtime of mission critical systems can be minimized through reliability and redundancy of components. Disaster recovery measures include geographic separation, isolation of dependencies between data centers, and redundancy of entire systems across data centers.

Oracle database offers a wide range of functionality for combinations of high availability, scalability and manageability . Oracle® is one of several companies working to run their product on Clusters.

Oracle's Technologies provide solutions for :

- High Availability for minimizing of unscheduled downtime
- Performance Scalability, the ability of the system to handle greater workloads and the ability to continue to add capacity to the system
- Cost and Complexity of the system
- Mean Time To Repair (MTTR), the average speed with which service is restored to end-users following a service outage.
- Disaster recovery
- Minimizing scheduled downtimes due to upgrades, patches and maintenance.

This paper talks about Oracle's technologies for High Availability, which refer to minimizing downtime and Scalability.

Oracle's Technologies for High Availability

The alternatives commonly associated with High Availability and Performance Scalability of Oracle databases are:

1. Oracle Distributed Databases

A distributed database is an application database in which parts of the data are located on multiple

databases that are connected together using *database links*. *Synonyms* can be used to mask the actual location of different tables from the application code. Distributed database is a solution for scalability problems in particular with the advent of the *two-phase commit*, which guarantees that transactions commit or rollback properly across multiple Oracle databases.

2. Oracle Standby Database

A standby database is a mirror of a primary database where the archived redo log files from the primary database are applied over the standby database, continuously. When the primary database goes down the standby database will be available to become the primary database and database service can resume with minimal interruption. The standby database can be opened for use within 2-30 minutes.

Oracle8i™ Data Guard provides a complete Automated Standby Database solution for Oracle8i databases. Features like automatic shipping and applying the archived redo log files by the standby database makes this simple and proven feature even more attractive. Planned downtime for maintenance can be reduced because production processing can be moved to the standby and back, easily and quickly. With the new Oracle8i ability to open the standby database for queries only, it might be possible to use the standby as a "reporting server".

Oracle9i™ Data Guard introduces a new standby database mechanism, the Logical Standby Database. Instead of applying the log files directly, the standby database extrapolates SQL statements from the log files and applies those statements to the server as any normal set of SQL commands. As a result of this SQL-based maintenance method, a standby database is available for normal read write operations.

3. Replication

Replication is the process of copying and maintaining database objects in multiple databases that make up a distributed database system. Oracle Advanced Replication allows an application to update any replicas of a database, and have those changes automatically propagate to other databases,

while ensuring global transactional consistency and data integrity. In the event of a disaster at one of the sites, the surviving databases will remain online. Replication is available in the following configurations:

- Read-only Slave (one *master* copy of data which can be updated, and one or more *slave* copies which are read-only)
- Updateable Slave (one *master* copy of data which can be updated, and one or more *slave* copies which can also be updated)
- N-way Masters (multiple *master* copies of data, with rules for *conflict resolution*, should conflicts arise from updates)

Replication is a viable alternative for high-availability, as each database is equipped to continue operating normally should either the network or other database cease functioning. It is also a viable alternative for scaling performance, as each database can handle transactions without synchronous *server-to-server* traffic occurring for each transaction.

4. Oracle® Parallel Server

In concept, *Oracle Parallel Server* uses clustered hardware to permit multiple, independent database instances to operate against a single database. It is complex, but if implemented correctly, applications can tolerate failures with minimal or no downtime. If any node should fail, then the database instances on the surviving nodes of the cluster are instantly available. Oracle Parallel Server makes failover transparent to users.

In some cases of clustering technology, high availability systems deliver the added benefit of scalability because there are redundant computers that can share the workload. For example in case of Oracle Parallel Server, if implemented correctly, applications can scale performance across all the nodes in the cluster. It is right to say that database engines work best on a single node where individual processes need to work closely together. But there is more to a database application than the engine at the core. The whole client side of the database system, which handles the requests from users, can be spread across different nodes, thereby spreading the workload across the clustered system. If more users come on the system, an extra node can be added.

Oracle® Fail Safe gives workgroups and departments, running Oracle® on Windows NT and 2000 clusters with a Microsoft® Cluster Server (MSCS), an easy way to minimize downtime. It

ensures fast automatic database failover during both planned and unplanned outages. With Oracle Fail Safe it is possible to reduce downtime for an Oracle database as well as any application that can be configured as a Microsoft Windows service.

Oracle9i™ Real Application Clusters introduce a breakthrough in clustered database technology, called Cache Fusion. The system provides both transparent application scaling, as well as super-fast fail over for high availability of data and applications. Oracle delivers a fully integrated high availability e-business solution built on Oracle9i Real Application Clusters. *Oracle9i Real Application Clusters* is the culmination of many years of Oracle's development on the Cache Fusion architecture for Oracle Parallel Server. In this environment multiple server nodes have simultaneous access to a common database, and lock management is handled by the Oracle Database via high speed server-cache to server-cache data transfer. The resulting architecture is linearly scalable, highly available, easier-to-manage, and has a significantly lower Total Cost of Ownership than comparable proprietary solutions in the market today.

Oracle® Real Application Clusters Guard is an enhanced configuration and an integral part of Oracle Real Application Clusters. It tightly integrates Oracle's enhanced recovery features with the cluster framework of the platform to provide a configuration that leverages the best high availability technology each platform has to offer. The Oracle Real Application Clusters Guard architecture is designed to build on the strengths of traditional high availability solutions and provide the following functions:

- Automated, fast recovery and bounded recovery time from failures that crash the Oracle instance
- Automatic capture of diagnostic data when certain types of failures occur
- Enforced primary/secondary configuration.
- Elimination of delays that clients experience when reestablishing connections after a failure

Conclusion :

Each of these alternatives will score differently depending on the system reliability and availability requirements of an Enterprise. Also these alternatives are not mutually exclusive. It is possible to mix and match these alternatives and get the best possible solution.