

Optimizing Disaster Recovery

Using Oracle Data Guard on Dell PowerEdge Servers

The high cost of computing system downtime has prompted organizations to view business continuity and high availability as two critical IT concerns. Using Oracle9i™ Real Application Clusters and Oracle® Data Guard on Dell™ PowerEdge™ servers and Dell storage can help administrators cost-effectively achieve data protection, high availability, and resilience for IT infrastructures.

BY PAUL RAD, ZAFAR MAHMOOD, IBRAHIM FASHHO, RAYMOND DUTCHER, LAWRENCE TO, AND ASHISH RAY

Although many formulas exist for calculating the cost of downtime, most IT managers would agree that downtime is simply not acceptable for business today. Enterprises expect their systems to be up and running without interruption, in many cases 24/7. Downtime, whether planned or unplanned, translates into lost opportunities and increased costs.

Dell and Oracle have partnered to offer a high-availability architecture that helps minimize scheduled and unscheduled downtime caused by numerous events, including system failures, site disasters, user errors, data corruption, and maintenance activities.

Dell produces entry-level, midrange, and high-end server and storage clusters built from standards-based components. These clusters are designed to help improve availability by removing all single points of failure within the cluster. At each cluster level, Dell also provides the capability to recover from additional failures, thereby protecting against multiple component failures. Low-cost Intel® processor-based Dell™ PowerEdge™ servers can help businesses implement the degree of availability that best meets their service level objectives.

Oracle® Maximum Availability Architecture (MAA), a High Availability (HA) best-practices blueprint from Oracle, aims to maximize system availability while reducing the

design complexity of an optimal HA architecture. MAA, which uses Oracle HA technologies such as Oracle9i™ Real Application Clusters (RAC) and Oracle Data Guard, provides recommendations that encompass the database tier, the application server tier, network and storage infrastructures, and operational principles. By adopting the MAA methodology, IT organizations can build a simple, robust architecture that helps prevent, detect, and recover from outages with a fast mean time to recovery (MTTR).

Dell and Oracle recommend adopting the MAA methodology on Dell platforms to address requirements for high availability, data protection, and disaster recovery. This article—the result of a joint project between Dell and Oracle engineering teams—explains how the RAC and Data Guard components of the Oracle MAA can be used on Dell PowerEdge servers and storage to build the foundation of an end-to-end, high-availability architecture.

Creating a highly available infrastructure using Oracle9i RAC

In an Oracle9i RAC environment, each node in a cluster runs a separate Oracle instance, and these instances can concurrently access a single, shared database. Although it spans multiple hardware systems, the database appears to applications as a single, unified database system. This

configuration helps provide a very high degree of scalability and availability to enterprise applications:

- The capability to flexibly, transparently, and cost-effectively scale capacity as business needs change
- Fault tolerance to failures within the cluster—particularly node failures

A typical Dell and Oracle9i configuration includes a storage area network (SAN). A Dell/EMC Fibre Channel-based SAN fabric supports multipath routing between SAN switches, helping ensure that no single points of failure exist in the configuration. In a typical topology, a node has multiple Fibre Channel host bus adapters (HBAs), each connected to the same SAN, resulting in multiple paths to the same device. SAN storage devices also can accept multiple Fibre Channel connections.

Although Oracle9i RAC addresses local system failures and provides rapid recovery from node failures or instance crashes, it does not offer protection from site disasters or user errors such as an accidental drop of critical user tables in the database. Such protection is provided by Oracle Data Guard, an integrated feature of the Oracle9i Database Enterprise Edition.

Enabling disaster recovery using Oracle Data Guard

Oracle Data Guard is software that creates, maintains, and monitors one or more standby databases to help protect enterprise data from failures, disasters, user errors, and data corruption. Data Guard maintains standby databases as transactionally consistent copies of the primary database. These standby databases can be located at remote disaster recovery sites thousands of miles from the production data center or they may be located in the same building. If the primary database becomes unavailable because of a planned or unplanned outage, Data Guard can be used to switch any standby database to the primary role, thus minimizing the downtime associated with the outage and helping to prevent data loss.

A standby database is initially created from a backup copy of the primary database. Once the standby database is created, Data Guard automatically maintains it by transmitting primary database redo data to the standby system over standard TCP/IP networks and then applying the redo data to the standby database.

Data Guard supports two types of standby database, each of which uses a different method to apply redo data to the standby




Figure 1. Data Guard Redo Apply (physical standby database)




Figure 2. Data Guard SQL Apply (logical standby database)

database and keep it transactionally consistent with the primary database: Redo Apply, used for physical standby databases, and SQL Apply, used for logical standby databases.

- **Redo Apply process on a physical standby database:** A physical standby database is kept synchronized with the primary database by applying the redo data received from the primary database using Oracle media recovery (see Figure 1). The standby database is physically identical to the primary database on a block-for-block basis, and thus the database schemas, including indexes, are the same. The physical standby database can be opened in read-only mode, and queries can be run on it at that time; however, it cannot run recovery at the same time it is opened as read-only.
- **SQL Apply process on a logical standby database:** A logical standby database contains the same logical information as the primary database, but the physical organization and structure of the data may be different. The SQL Apply process keeps the logical standby database synchronized with the primary database by transforming the redo data received from the primary database into SQL statements and then executing the SQL statements on the standby database (see Figure 2). This enables the logical standby database to be accessed for queries and reporting purposes at the same time the SQL statements are being applied to it.

Figure 3 shows two sites with identical configurations. Each site consists of redundant components so that requests can always be serviced, even if a failure occurs. Each site also contains a set of application servers or mid-tier servers.

The primary site with the primary database uses Oracle9i RAC to protect the database from host and instance outages. The secondary site contains a physical standby database that is maintained by the Data Guard Redo Apply process. The secondary site uses Oracle9i RAC to protect it from local host and instance outages.

Oracle Data Guard offers two simple methods—switchover and failover—to handle both planned and unplanned outages of the primary database. Administrators can initiate both methods directly through simple SQL statements or the Data Guard Manager graphical user interface (GUI), which is the Data Guard administrative interface integrated with Oracle Enterprise Manager.

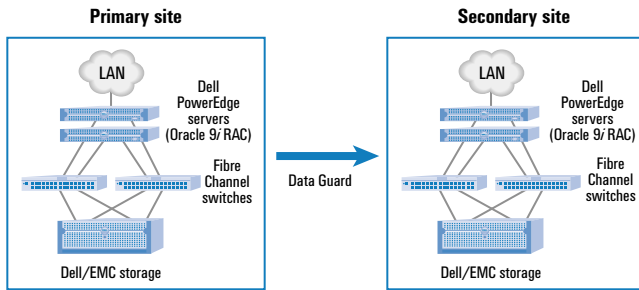


Figure 3. Maximum Availability Architecture using Oracle9i RAC and Oracle Data Guard

Data Guard switchover

Switchover is a planned role reversal of the primary and standby databases to manage scheduled maintenance on the primary database. A switchover operation does not require re-instantiation of the Oracle database, so the primary database can assume the role of a standby database almost immediately. As a result, administrators can perform scheduled maintenance more easily and frequently. For example, administrators can use switchover to perform a system upgrade on the primary site by switching over all the database clients to the standby site as they upgrade hardware on the primary site.

For steps on how to invoke a switchover in a Data Guard configuration, refer to “Physical Standby Database Switchover” in the Oracle Maximum Availability Architecture white paper at http://otn.oracle.com/deploy/availability/pdf/MAA_WP.pdf. For detailed information on how to configure an Oracle RAC Data Guard physical standby database, visit *Dell Power Solutions* online at http://www.dell.com/magazines_extras.

Following a successful switchover, the standby database assumes the primary role and the former primary database becomes a new standby database. In a RAC environment, a switchover requires only one active Oracle instance for each database. If required, administrators may perform a *switchback* operation by doing a subsequent switchover to return the databases to their original roles.


Data Guard failover

Administrators may invoke a *failover* operation when an unplanned catastrophic failure occurs on the primary database or when the primary database cannot be recovered in a timely manner. A Data Guard failover may be accompanied by a site failover to move end users to the new site and database. Once the failover is completed, the primary database can be accessed from the secondary site. Following MAA guidelines, the former primary database must be re-created as a new standby database to restore resiliency.

Typically, little or no data loss is experienced during a failover operation. For detailed information about Data Guard failover, refer to “Physical Standby Database Failover” in the Oracle Maximum Availability Architecture white paper.

Enabling business continuity with Oracle9i RAC and Data Guard

The combination of Oracle MAA, Dell PowerEdge servers, and Dell storage can offer enterprises an easy, low-cost means to implement business continuity and disaster recovery for IT infrastructures. Oracle9i RAC running on Dell PowerEdge servers helps provide the reliability and scalability of a redundant cluster environment. RAC enables high availability by helping provide continuous data access when a node or instance fails, or when performing scheduled system maintenance on a subset of nodes in the cluster.

Oracle Data Guard facilitates data protection and disaster recovery by automating the maintenance of geographically distant standby databases as transactionally consistent copies of the primary database. Data Guard enables easy switchover or failover of a primary database to a standby database if planned or unplanned outages occur at the primary site. Because it is an integrated feature of the Oracle database, Data Guard can be more cost-effective and better optimized to protect Oracle data than host-based or storage-based remote mirroring.¹ For continuous data availability and a resilient, high-availability system architecture, organizations may consider implementing Oracle MAA best practices in combination with Oracle Data Guard and Oracle9i RAC on Dell servers and storage. 

Paul Rad (paul_rad@dell.com) is a senior software engineer in the Dell Database and Application Engineering Department of the Dell Product Group.

Zafar Mahmood (zafar_mahmood@dell.com) is a software engineer in the Dell Database and Application Engineering Department of the Dell Product Group.

Ibrahim Fashho (ibrahim_fashho@dell.com) is the development manager for the Database and Application Engineering Department of the Dell Product Group.

Raymond Dutcher (raymond.dutcher@oracle.com) is a principal member of the technical staff in the Oracle High Availability Systems Group.

Lawrence To (lawrence.to@oracle.com) is a principal member of the technical staff in the Oracle High Availability Systems Group.

Ashish Ray (ashish.ray@oracle.com) is a senior product manager in the Oracle Database High Availability Group.

FOR MORE INFORMATION

Dell and Oracle partnership:

<http://www.dell.com/oracle>

Oracle9i RAC:

<http://otn.oracle.com/products/database/clustering>

Oracle Data Guard:

<http://otn.oracle.com/deploy/availability/htdocs/DataGuardOverview.html>

Oracle Maximum Availability Architecture:

<http://otn.oracle.com/deploy/availability/htdocs/maa.htm>

Oracle Database High Availability:

<http://otn.oracle.com/deploy/availability>

¹For a comparison of Oracle Data Guard and third-party remote mirroring options, visit <http://otn.oracle.com/deploy/availability/htdocs/DGCompTech.html#RemoteMirror>.