InfoWorld

GET TECHNOLOGY RIGHT

IT White Paper

Data Protection and Availability with Oracle Active Data Guard 11g

Compliments of:

ORACLE

Data Protection and Availability with Oracle Active Data Guard 11g

An integrated new high availability-disaster recovery solution enhances quality of service while reducing cost and complexity. IDG tests the new Oracle Active Data Guard option to assess its value to DBAs.

Improving Quality of Service for Critical Business Transactions

Efficient business operations, high quality customer service, conformance with government regulations, and safeguarding corporate information assets all depend upon achieving the highest possible level of data protection and availability. Thus, it is no surprise that data protection and availability are among the top priorities of business continuity initiatives for companies of all sizes across all industries.

To meet these needs, Oracle Data Guard offers a comprehensive data protection, disaster recovery (DR) and high availability (HA) solution for the enterprise. Oracle Data Guard is a central component of an integrated Oracle Database HA solution set that helps organizations ensure business continuity by minimizing the various kinds of planned and unplanned downtime that can affect their businesses.

As an extension of this solution set, Oracle now offers Oracle Active Data Guard, an option of Oracle Database 11g Enterprise Edition. This new option helps enterprises that want to improve the performance of mission-critical applications by offloading resource-intensive activities such as reporting and backups to a replica of their production database.

Oracle Active Data Guard enables one or more physical standby databases to be open for read-only access while changes from the production database continue to be applied. All queries reading from a physical replica execute in real-time, and return current results. As a result, any operation that requires up-to-date read-only access can be offloaded to the replica, enhancing and protecting the performance of the production database. Oracle Active Data Guard also enables the use of fast incremental backups when offloading backups to a standby database, and it can provide additional benefits of HA and DR protection against planned or unplanned outages at the production site.

DBAs can use Oracle Active Data Guard to:

- **Increase performance** - Offload unpredictable workloads to an up-to-date replica of the production database
- **Simplify operations** - Eliminate management complexity that accompanies traditional replication solutions
- **Eliminate compromise** - Reporting replica is up-to-date and online at all times, which is not possible with traditional storage mirroring technology
- **Reduce cost** - An Oracle Active Data Guard physical standby database can also provide DR services and/or serve as a test database, with no additional storage or servers required

IDG Evaluates Oracle Active Data Guard

Oracle Data Guard is designed to provide a transactionally complete standby database in case of disaster. Essentially, this feature provides the same protection as cold failover clustering, but it goes further. Clustering is meant to protect against server-level failures: If a CPU, motherboard, or other component of the server fails, then the clustering mechanism will fail over to the inactive clustering partner in a cold failover cluster. A DBA can concentrate on getting the primary server working again while business continues.

However, any cluster solution does not protect against storage array failure, logical corruption, or network or site

failure. Because clusters share the same set of disks, if a shared disk fails past the point of repair, then all nodes of the cluster will be made unavailable. Because cluster nodes share the same location and the same network, if the network between the client community and the database server goes down (due to a faulty router, firewall, or switch, or even a rat chewing through a cable), clustering also gives no protection.

As a result, clustering is limited in its capabilities to provide a complete HA solution. Moreover, in a cold failover cluster (active-passive scenario), the standby node of the cluster isn't available while the primary is online.

Remote-mirroring addresses the key limitations of local clustering. Comprised of two completely different servers with separate storage resources at different locations, remote mirroring not only gets past the shared storage limitation, but it also addresses network and site failures. If a network outage occurs, the administrator disables remote-mirroring, mounts the remote volumes, and starts Oracle at the recovery site; the administrator then fails over the application and resumes business processing on the standby node. However, just as with local clustering, the standby systems remain idle, waiting for a failover to occur.

Oracle Active Data Guard, introduced in Oracle Database 11g, lets you have your standby and use it, too. This option lets DBAs set a standby instance into read mode to support real-time queries, solving one of the biggest issues businesses have with their OLTP databases: how to optimize performance by separating OLTP and read-only activity while providing near real-time reporting. This is just one of the capabilities Oracle provides to enable users to utilize standby systems. In addition, there is the previously mentioned ability to offload fast incremental backups Finally, a new feature included in Oracle Data Guard at no extra charge is called Snapshot Standby. This feature lets DBAs put the standby database into a temporary read/write mode, letting them test database changes while still providing the original HA/DR protection.

Test Bed

The primary database was stored on a RAID-5 array with 14 SATA spindles, directly attached to an Intel-based Win¬dows Server 2003 Enterprise server (32-bit with Service Pack 2 installed) running Oracle Database 11g. The serv¬er was configured with four 3GHz Intel Xeon CPUs and 4GB of RAM. Testers used the Quest Benchmark Factory 5.5 load generator to simulate activity of 50 users against the database, running Benchmark Factory on the same server as Oracle Database.

The standby database was stored on a RAID-5 array with 10 SCSI spindles, directly attached to an Intel-based Windows Server 2003 Enterprise server (32-bit with Service Pack 2 installed) running Oracle Database 11g. The standby server was configured with two 1.4GHz Intel Xeon CPUs and 2GB of RAM. The same standby database was used to test both Oracle Active Data Guard and Oracle Data Guard 11g Snapshot Standby.

Test Results

Testing Oracle Active Data Guard was straightforward. Using Benchmark Factory, testers generated a small workload of 10 users, writing to a couple of tables in the primary, and watched the data move across to the standby server. The InfoWorld team tested queries of the standby, and failovers with the standby in read mode and with the standby in Snapshot Standby (read/write) mode.

Under these light test conditions (including a relatively small database with no major changes made to the standby), there were no problems. The data moved over quickly, queries performed well and produced expected results, and failovers were successful and fast. There was no delay failing over from read-only mode, and failovers took less than 5 minutes even from Snapshot Standby (read/write) mode. After rolling back from Snapshot Standby, testers verified that the standby was consistent with the primary database by counting rows in each of the tables on each side. Oracle Active Data Guard functioned as expected throughout all tests.

The tests also demonstrated an additional advantage of Oracle Data Guard over remote-mirroring. With Data Guard, Oracle is always mounted and running with a "hot" copy of the production database on the standby server. Compared to remote mirroring, this ability makes failover simpler and faster. Moreover, using Oracle Active Data Guard eliminates any uncertainty regarding the state of the standby database while it is in standby role.

InfoWorld Test Center Evaluation

Oracle Active Data Guard is a huge leap forward. It not only allows DBAs to make the standby readable, but to do so with minimal configuration changes on the client. DBAs can create database services that correspond to specific workloads. For example, DBAs can create a service associated with the OLTP workload that is available on the primary database as well as a read-only service for query workload that is available only on the standby database. When the application connects to the OLTP service, it is automatically connected to the primary database; when it connects to the read-only service, the queries are routed to the standby to fulfill the reporting requests. This is powerful functionality that makes it easy for shops to address their reporting needs using current OLTP data.

In most systems, testing changes requires DBAs to restore production to a test server, understanding that its data will be a few days behind most of the time. In addition, these systems make it difficult to simulate the production workload, even with benchmarking software; as a result, most companies dismiss this type of testing, assuming it cannot be a viable part of their test process. With Snapshot Standby (and an associated Oracle Database 11g feature, Database Replay, also known as Real Application Testing), DBAs could initiate a resurgence of this type of application testing.

The bottom line: DBAs will love what Oracle Active Data Guard can do and the freedom it gives their organization in putting otherwise idle standby databases to good use. Although testers found that setting up Oracle Data Guard is challenging for DBAs new to Oracle, once it was up and running it was easy to use. Switching between modes takes just a single statement at the command line. No doubt many Oracle shops will have great success with Oracle Active Data Guard.

Invest in Quality of Service

Oracle Active Data Guard provides a common infrastructure for data protection, availability, scaling performance by utilizing standby systems, and reducing business risk by thoroughly testing system changes – all while using a common management interface. These activities can be supported on a single replica of the production database; Oracle Data Guard does not increase cost or complexity by requiring multiple replicas, where each uses a different technology to address different requirements. Oracle Active Data Guard is easier, less expensive, and more functional. It is an option that is hard to beat.

IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

I D C E X E C U T I V E B R I E F

Two Birds, One Stone: A Single Approach to Ensure Continuous Database Availability and Better Performance

March 2008

Adapted from *Worldwide Database Management Systems 2007–2011 Forecast and 2006 Vendor Shares* by Carl Olofson, IDC #209611

Sponsored by Oracle Corp.

Introduction

Standby database servers have traditionally been deployed to protect enterprise data from disruptive business events or failures and prevent any data loss. They enable data processing to continue with little or no interruption in situations where the production database fails or becomes unavailable. Such an outage could be a scheduled shutdown for maintenance or an unplanned outage due to data corruption, human error, or natural disasters.

Database administrators deploy standby databases in the datacenter or at remote locations. They function as exact transaction-by-transaction mirrors of the primary database and thereby ensure continuous availability of the primary database. Normally, the standby database server functions entirely in passive standby mode, receiving a stream of replicated database transactions until that moment when the primary server goes offline. When this happens, the database load is immediately switched to the standby database without loss of time or data.

Until recently, the standby database only mimicked the primary database and was unusable for any other purpose. To those paying the bills, it seems eminently unfair supporting a standby database server for a system that is to be used only in case of emergency.

Consequently, there have been many attempts in recent years to create a standby database server with capabilities that would allow it to not only serve as a replica of the primary database but also handle other types of workloads such as reporting and querying. Third-party and database management systems (DBMS) vendors have used various approaches to solve this problem — but at a cost in terms of performance or possible loss of data. Within the past few years, some key vendors have developed new technology advancements that make it possible to turn a necessary insurance investment into a valuable and more productive resource.

The Challenges of Stealing Standby Database Cycles

Most replication techniques for databases depend on technologies that lock up the entire database, making it unavailable for updates. A few products use their own techniques to get around this problem and thereby make the standby (or secondary) database available at least for read-only access. The goal is that any techniques used must enable the standby database to match the live database with as little latency as possible.

The most common techniques involve reading update transactions as they are being written into the log and streaming them into the log for the standby database, which is running a constant roll forward operation against the log entries. This approach not only ensures that the standby database matches the primary as each primary database update is committed but also causes the standby database to operate as if it were in perpetual recovery mode, making it unavailable for any other type of processing.

Other techniques involve capturing database calls and mirroring them to the standby database so that it performs every operation exactly as the primary database does. This approach also precludes any other activity because the standby database needs to process transactions at the same rate as the live database because any delay in performing the mirrored operations would result in a corresponding delay in failing over to the standby database. Also, the standby database is often unable to process select statements due to the time and resources required to handle locking and data currency issues.

Technologies Enabling Reading of Standby Database Servers

Recently, vendors have developed products and features that allow standby databases to operate in a read-only mode. The goal is that any techniques used must enable the standby database to match the live database with as little latency as possible, yet still permit query and reporting functions. These approaches include technologies that do one or more of the following:

- Stream-optimized transactions that enable the standby database to process updates faster than the primary database, opening up the extra capacity for query processing
- Optimizing lock and buffer management in the standby database to ensure that query processing has as little impact on the execution of the replicated transactions as possible
- Other "secret sauce" that DBMS vendors have built in, but will not divulge, that accomplishes this goal

Benefits of Offloading Workloads to the Standby Database Servers

Typical database workloads perform about 85–90% reads to each write operation; therefore, users can get tangible value from being able to use their backup servers, and the backup servers can be used to relieve the query and reporting workload from the primary database. This allows the primary database more processing capacity for update transactions. As a result, the primary database can process a larger or more mission-critical workload faster.

For example, if the primary database is heavily loaded performing time-critical update operations, it is essential that those updates be recorded in a timely manner. Any other operations, such as database queries or report generation, just add an additional burden to the primary database's workload. This ability to channel simple queries and reporting to the standby database enables the primary database to run faster. Furthermore, with this capability offloaded, the primary database can be tuned to perform a specific type of operation.

IT and database administrators are tired of having to pay full price for the cost of systems, software licenses, additional storage, network connections, and so forth for a standby database server that will be used only for maintenance or as outage prevention insurance. The ability to use these systems for other purposes that deliver information and value to the users justifies their initial and ongoing costs.

Considerations

Decisions about the appropriate solution that can ensure continuous database availability for organizational needs should be based on existing business requirements and criteria for overall database performance and processing requirements, as well as data integrity and accuracy. Due consideration must be given to the answers to the following questions:

- Does the organization have a standby database such as is described in this document that is unavailable for query and reporting but that could be used for that purpose if the technology is available to do it?
- If so, is the primary database challenged in delivering required performance while servicing a blended update and query workload, and would the ability to offload the query portion of that workload significantly reduce such a challenge?
- Does the organization consider the added value derived from the ability to use the standby database server for query and reporting worth the additional expense in terms of additional fees for licensing of the product or feature that enables it from the vendor?

If the additional cost of such functionality represents a small increase of the overall cost, and the value would be great (due to demand for more database access and better throughput on the primary database), then the extra cost may be well justified.

Organizations must also carefully evaluate their needs and these trade-offs to ensure that their database is able to perform other reporting and query capabilities if required without jeopardizing the security, performance, and integrity of the primary database when outages occur.

Careful consideration should also be given to the following characteristics of the environment to determine if the read and query load on the standby database server will jeopardize its ability to fail over successfully when required:

- Size of the database
- Number of concurrent users
- Transaction speed and security requirements
- Consistency of network performance
- Complexity and location of systems and network architecture
- Amount of data backed up on a regular basis and the retention policies
- Administrative resources available to manage the environment

Conclusion

The development of a standby database server that can be successfully deployed as a single solution to both ensure continuous database availability and be used to offload query and read functionality from the primary database has been an arduous and extraordinarily difficult journey. New technology advancements now make it possible for organizations to purchase and deploy this useful technology as part of their own database solutions and investments.

The ability to utilize standby database servers in new ways to increase workload sharing delivers real benefits in terms of measurable value for improving the performance of the primary database as well as getting more value and productivity from the backup database investment.