HIGH AVAILABILITY OPTIONS FOR ORACLE DATABASE

Dan Norris, dnorris(at)piocon.com, Piocon Technologies, et. al.

INTRODUCTION

Today more than ever, businesses need uninterrupted access to mission-critical data. To remain competitive in the global economy, decision makers must have reliable access to timely data so that they can react to changing circumstances. Oracle Database is one of the most widely used relational databases in today's business world, with companies using Oracle Database to store their most sensitive and vital data.

Because of the importance of the applications it supports, ensuring the availability of the Oracle Database has been a subject of a tremendous amount of examination and effort within the community of Oracle Database users. Oracle Database administrators have developed several techniques to improve the availability of the database. Each technique targets a certain availability level and, accordingly, each technique has varying costs, implementation times and complexity associated with it. Logically, techniques that yield high availability tend to be more complex, take longer to implement and are more costly.

This document outlines the primary methods for improving the availability of an Oracle Database. The most basic options are discussed first, followed by a series of increasingly complex options. This organized progression of options is known as the Oracle Database High Availability Spectrum. This spectrum is represented by the diagram below:

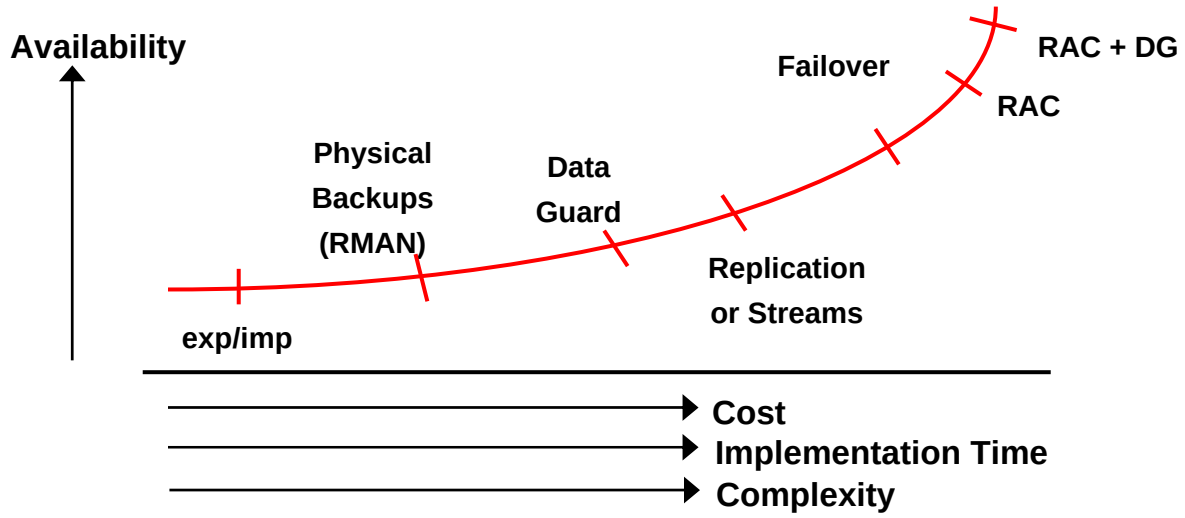

This diagram depicts the relationship between cost, implementation time, and complexity. The availability achieved for a particular option is non-linear, but it does bear a direct relationship to these costs.

The sections that follow describe each of the options shown in the diagram with a discussion about the advantages and disadvantages for each option.

Oracle has also created the Maximum Availability Architecture (MAA). This architecture provides an architectural "blueprint" for maximizing availability for your Oracle environment and includes some of the options outlined in this paper.

HIGH AVAILABILITY OPTION: EXPORT/IMPORT

Oracle Export (exp) and Import (imp) are complementary utilities used to perform logical database backup and recovery. They can also be used to move Oracle data from one machine, database or schema to another.

The export/import utilities write data in an Oracle proprietary binary format from the database into operating system files called dump files. These dump files are transportable across operating systems, platforms and most database versions. The dump files contain data and all the DDL commands needed to re-create the database objects requested.

Exports can be created at many different levels:

- A table-level export can export one or more tables or partitions of a table.
- A schema-level export can export one or more schemas.
- A full export will export the entire database.

ADVANTAGES

- Oracle Export/Import provides a backup and recovery strategy that is simple to implement. Certified Oracle Database administrators will need no extra training or specialized expertise to use the Export/Import utility.
- The export/import approach is relatively stable. These tools have been a part of the database release since at least Oracle version 5 and have remained a very reliable way to back up the database. The export and import utilities are available with every Oracle database and are widely known for their compatibility between releases.
- The export utility will detect physical corruption and ensure that all data can be read during the export process.
- Table structures and data can be transferred between Oracle databases, even if they reside on platforms with different hardware and software configurations.
- Table structures and data can be transferred between databases running different versions of Oracle. This method is often used as a mechanism to upgrade smaller databases.
- Performing an export/import reorganizes your data and can help to eliminate database fragmentation.
- Tables can be imported into a different schema owner than the one from which they were exported.

DISADVANTAGES

- Performing an export/import process on a large database can be very time consuming. The time required to export and import data is one reason that only small databases are upgraded via export/import.
- An export/import process is very CPU intensive and will reduce overall system performance while the export/import process is running.
- Dump files for large databases are very large, so significant amounts of disk space can be required to conduct an export for large databases.
- The dump files have to be transferred from one machine to another if transporting data. Because of their large size, the dump files can be difficult to transport in a timely manner.
- Large dump files may become corrupted and become unusable.
- Operations can be difficult to restart. If an export fails, it usually has to be restarted from the beginning.
- The time to recover can be long if there is a complete loss of the database.
- Some of the newest data types are not supported by export and import.
- Export and import utilities will likely not be enhanced in future releases of Oracle Database. Customers are being encouraged to migrate to Oracle Data Pump starting with Oracle Database 10g.

Oracle export and import utilities are common and have an established track record. The export and import utilities are commonly used in conjunction with other options since they provide additional, unique recovery options. As an option for high availability, export and import are the most basic and almost every customer will require higher availability than what export and import can provide.

High Availability Option: Data Pump

Oracle Data Pump is a new feature introduced with Oracle 10g and is an enhanced version of the export/import utilities. It is a server-side utility that can be used to unload and load data as well as move data and metadata. Data Pump is designed as the eventual replacement for the original export and import utilities. The original import utility will very likely be supported forever as a means to import dump files from earlier versions of Oracle. However, the original export and import utilities will eventually be deprecated and only the Data Pump utilities will be supported for logical export.

The Data Pump infrastructure is accessible though the PL/SQL package DBMS_DATAPUMP which enables database administrators to create custom data movement utilities using Data Pump.

Oracle claims that a Data Pump export is about twice as fast as an original export and that a Data Pump import is 15 to 45 times faster than an original import. Basic tests conducted by various Oracle user-community groups largely confirm that Data Pump is faster, but with varied levels of improvement.

Advantages

- Operations performed using Data Pump export and import utilities are much faster than operations performed using the traditional export and import utilities.
- Data Pump supports restarting jobs and allows the monitoring of progress for loads and unloads.
- Data Pump automatically manages multiple, parallel streams of export and import for maximum throughput.
- Data Pump implements self-tuning mechanisms to ensure that tables are exported and imported in the most efficient manner.
- Data Pump allows the database administrators to perform an export and import over the network without any data movement.
- It is possible to detach from and re-attach to long-running jobs.
- Data Pump has the ability to estimate how much space an export job would consume, without actually performing the export.
- Oracle Data Pump is included in the Oracle Database license starting with version 10g.

Disadvantages

- Data Pump import can only read Data Pump export dump files. It is not compatible with the traditional export or import utilities.
- Parameter files created for the original export and import utilities are not compatible with the new Data Pump utilities.
- If disk space is at a premium, this option might not be viable as there is no compression of the dump file offered by Oracle. It is possible to use a third-party compression tool if available to you.
- Startup time is longer for a Data Pump job since Data Pump is specifically designed for large jobs with lots of data.
- Data Pump can be very resource intensive—it will consume as much CPU, memory, and I/O bandwidth as your setting of the PARALLEL parameter will allow.
- All export files created and all import files to be read must be located on the database server. Unlike the client-server architecture utilized by the traditional export and import utilities, there is currently no way to create a Data Pump export file on an Oracle client installation.
- Data Pump is a relatively new technology, so adoption is limited at the present time. As a result, community knowledge levels are relatively low. This will obviously change as time passes and more sites adopt Data Pump as their standard logical backup utility.

Data Pump is a great utility that has gained significant momentum due to its performance advantages and programmatic interfaces. As a high availability option, Data Pump has similar characteristics to the traditional export and import utilities. Most sites will require higher availability than what is afforded by Data Pump, but it will be common for Data Pump to be used in conjunction with other options outlined later in this paper.

HIGH AVAILABILITY OPTION: ORACLE RECOVERY MANAGER

Oracle Recovery Manager (RMAN) offers a fast and reliable backup solution. It features useful functionality such as online backups, incremental backups, block media recovery, automation of backup management tasks and integration with third-party media management systems. RMAN backups can be seamlessly offloaded to a physical standby database, allowing customers to gain additional benefit from their disaster recovery environment.

ADVANTAGES

- RMAN automatically executes all underlying database procedures before and after backup or restore, eliminating the need to depend on OS and SQL*Plus scripts. It provides a common interface for backup tasks across different host operating systems, and it offers features not available through user-managed methods, such as parallelization of backup/recovery data streams, backup files retention policy and detailed history of all backups.
- Incremental backups capture changes to your database on a block-by-block basis with respect to the most recent incremental backup. Incremental backups are generally smaller and can be created more quickly than full database backups. Recovery from an incremental backup is faster than recovery that relies solely on redo logs alone. Recovering with incremental backups requires no additional effort on the part of the database administrator. If incremental backups are available, RMAN will use them during recovery.
- If disk space is limited, then you can still reduce database recovery time by keeping RMAN backups compressed on disk. Uncompressing the backed up files is not necessary since recovery operations can directly use the compressed backup files. This can provide significant time savings when compared to restoring from tape media.
- The Block Media Recovery functionality allows RMAN to fix a corrupted block (detected on backup) while the data file remains online, and non-affected data continues to be available for selecting and updating. This increases data availability and reduces the mean time to recovery (MTTR) by selectively restoring and recovering the damaged blocks. Minimal I/O is needed because redo is only applied to damaged blocks.
- RMAN takes advantage of intimate knowledge of Oracle block structures to provide high backup and restore performance and efficient file compression. By default, when creating backup sets, RMAN backs up only blocks that are in use (or have ever been used) and saves disk space by merging blocks into as few backup pieces as necessary.
- New in Oracle Database 10g, the flash recovery area allows administrators to setup notifications on disk space usage and automate obsolescence of expired backup sets via RMAN client command-line or Oracle Enterprise Manager (OEM) interfaces.
- Administrators can take advantage of the RMAN client command-line interface, whose commands are written in an OS-independent scripting language, or administrators can use the OEM Backup and Recovery Management console. With OEM, performing backup job scheduling and recovery operations is completely automated via step-by-step wizards.
- RMAN in Oracle Database 10g Release 1 eliminates the need for database administrators to manually create an auxiliary instance for tablespace point-in-time recovery (TSPITR). RMAN takes care of instance creation on the same server as the target database and removal upon completion of the tablespace recovery.
- Through one standard Media Management Layer (MML) API, third-party media management vendors can integrate their media managers with RMAN. Database administrators make a minor change in the RMAN backup commands to utilize tape instead of disk-based backups.

DISADVANTAGES

- In the event that an organization needs to take full advantage of all of RMAN's features, database administrators need to create a separate database to store the RMAN catalog. This separate database will contain all the information related to your backup and restore operations.
- RMAN uses its own language to perform the backup and restore operations.
- RMAN integration with third-party media managers via the MML API may require purchasing an additional license from the backup vendor. This license may be expensive.
- For sites that have relied on their own scripts for backups in the past, getting familiar with RMAN can be a significant effort. RMAN has its own scripting language, syntax, and vocabulary for backup components.
- There is no scheduler included with RMAN, so scheduling still has to be handled outside of the backup utility.
- Compatibility between RMAN and database releases can sometimes create problems. The RMAN utility, the RMAN catalog database, and the backup target may all be different releases and compatibility must be ensured.

Oracle RMAN is a backup and recovery utility. As a high availability option, physical backup is generally considered the baseline lowest-level of availability allowed. While not always considered a true "high" availability option, RMAN provides a simple way for even novice database administrators to conduct physical backups that are complete and valid. Most customers that look for true high availability require more than physical backups because of the amount of time needed to recover a physical backup.

HIGH AVAILABILITY OPTION: ORACLE DATA GUARD

Oracle Data Guard builds on the standby database technology that was first introduced with Oracle 7 and has improved greatly since then. Through Oracle 9i Release 1, the standby database always was a physical standby database in recovery mode, applying archive logs received from a primary database. The logical standby database, which applies the changes on an open database using Oracle LogMiner, was introduced with Oracle 9i Release 2 (9.2).

First available as an optional feature of Oracle 8i, Data Guard provided methods to automate physical standby database maintenance and a process to gracefully activate a standby database environment without losing any data. These features were just the beginning of a rich set of capabilities introduced in Oracle 9i to provide complete automation for creation, monitoring, maintenance, and activation of a standby database environment.

Oracle Data Guard ensures data protection and disaster recovery capability for enterprise data. By providing a comprehensive set of services that create, maintain, manage and monitor one or more standby databases, Oracle Data Guard enables databases to survive disasters and data corruption. Data Guard maintains these standby databases as transactionally consistent copies of the production database. If the production database becomes unavailable due to a planned or unplanned outage, Data Guard can switch any standby database to the production role, minimizing the downtime associated with the outage. With Data Guard, administrators can optionally improve production database performance by diverting resource-intensive backup and reporting operations to standby systems.

A Data Guard configuration consists of one production database and one or more physical or logical standby databases. The databases in a Data Guard configuration are connected by Oracle Net and may be dispersed geographically. There are no restrictions on where the databases are located if they can communicate with each other.

ADVANTAGES

- Data Guard maintains transactionally consistent database copies to provide protection against unplanned downtime and disaster. Standby sites can be configured to receive updates immediately or lag behind the primary site to allow data recovery for logical corruptions.
- Data Guard can be employed to reduce planned downtime for hardware upgrades, Oracle patch set installation and database upgrades.

- Built-in features allow automatic detection and resolution of missing recovery information following temporary loss of connectivity between the primary and standby database(s).
- Multiple levels of data protection and performance can be employed to balance data availability against system performance requirements.
- Reporting and backup operations can be diverted from the production database to standby database(s) to make efficient use of system resources.
- Managed, automatic role transition and application notification minimize planned and unplanned downtime.
- Following a failover, automatic resynchronization of a failed primary database restores the environment to a valid configuration.

DISADVANTAGES

- All servers in a Data Guard configuration (even logical standby) must utilize the same platform and operating system.
- An equal or greater amount of disk space must be used for the standby database. This is because the standby (physical or logical) is an exact copy of the production database, thus their sizes are the same.
- Standby databases are additional databases that require the administrator's time to maintain, upgrade, patch, and monitor.
- In order to effectively test the standby database, regular testing must be conducted to utilize the standby database.
- Data Guard requires specific knowledge that is not necessarily common to all database administrators.

Oracle Data Guard is an excellent option for disaster recovery since it has few restrictions and does not require any specialized hardware. Most sites use Data Guard as a method for disaster recovery, but they generally do not have automated procedures in place to activate a Data Guard environment in the event of a failure. Instead, most organizations choose to activate the standby site manually after verifying that the primary site is not viable. This results in a recovery time that is generally faster than a physical recovery, but not as fast as options described later in this document.

HIGH AVAILABILITY OPTION: ADVANCED REPLICATION

Oracle Advanced Replication is one method for replicating data to another database. Advanced Replication has also been known by other names over its many years. Since it was introduced in Oracle 7, this capability has also been known as Advanced Symmetric Replication or Multi-master replication. Some organizations have elected to use replication as a means to provide high availability. Most locations use replication to provide an active-passive solution where only one site in the replication agreement is used at any given time.

Oracle Advanced Replication has two primary forms: synchronous and asynchronous. When synchronous replication is used, Oracle Database will make the changes on all remote sites at the very same time that they occur on the source system. This behavior requires all sites to be available in order for any changes to occur. If one database is unavailable for any reason (network error, password reset, or database failure), then all sites are unable to make changes to the replicated object(s). For this reason, synchronous replication is not used very often, but it is the best way to ensure that all sites in the replication agreement converge. Asynchronous replication is implemented by placing queues on each node in the replication agreement. These queues are populated with all changes that are made to all replicated objects. Database jobs are created to process these queued transactions. If a failure occurs and contact is lost with the remote site(s), the queue holds all transactions until the remote site(s) is available.

ADVANTAGES

- Oracle Advanced Replication has a long history and has been widely deployed. Because of this, it has few bugs and its capabilities are well-defined.

- There is no additional cost and this feature is not a licensed option. It is available with the Personal Edition and Enterprise Edition licenses for Oracle Database beginning with Oracle 9i.
- The master sites for a replication environment only have to be network accessible to one another. Therefore, these master sites can be located anywhere in the world. The master sites must follow compatibility rules for database version, but there are no requirements for platform, operating system, or other infrastructure requirements.
- Architecturally, Advanced Replication is relatively simple to comprehend and configure. Additionally, OEM has tools available to facilitate the configuration and management of a replication environment.

DISADVANTAGES

- Triggers on tables are used to populate the deferred transaction queue. Sometimes, these triggers can introduce significant overhead for DML on replicated objects.
- LONG and LONG RAW data types are not supported for replication. These types should be converted to LOB columns, which can be replicated.
- Most database administrators have not worked with Oracle Advanced Replication, so additional training may be required to configure and manage a replicated environment. In order to properly manage and troubleshoot a replication environment, database administrators will need to dedicate some time and effort on learning replication concepts as well as syntax and replication-related database views.
- There is no easy, automated way to ensure that users will be able to connect to the remote site when the primary site fails.
- Unless you can guarantee that changes will only be made to one site at a time, conflict resolution routines must be implemented in order to resolve data convergence issues.
- Network bandwidth requirements can be substantial, depending on the volume of data being replicated. This can also contribute significantly to the overhead on the system and may affect user performance unless a dedicated network is configured for replication network traffic.
- While Advanced Replication has not been desupported or deprecated by Oracle Corporation, other, superior replication technologies have been introduced by Oracle and other vendors.
- The delay for a replicated change to propagate to a remote site can often be 60 seconds or more. There are settings to control the delay time for propagation, but increasing the propagation frequency (thereby reducing the delay) will create significant overhead (especially CPU) for the originating site.

Oracle Advanced Replication has a long history, and it has certainly has been perfected since its release. Although Oracle is likely to continue to support Oracle Advanced Replication for several more releases, other Oracle replication technology, like Oracle Streams, may be the better choice for new implementations. Customers that need a time-tested, stable environment to move data from one database to another find Advanced Replication very effective. For large numbers of replicated objects, the overall cost of maintenance and overhead on the system will be higher, so most organizations use Advanced Replication for relatively small numbers of objects—not typically a whole database.

HIGH AVAILABILITY OPTION: STREAMS

Introduced in Oracle 9i Release 2 (9.2), Streams builds on the functionality provided by Oracle LogMiner to capture database changes by reading the redo stream. With this architecture, Oracle can avoid the costly overhead of triggers on the database objects to capture changes as they happen. In most databases, Oracle Streams can read directly from the online redo stream, providing a very fast replication time.

Oracle Streams operates three primary processes: capture, propagate, and apply. Based on a set of rules that serve as filters, the capture process scans the redo stream for matching transactions. When a change matches, it is converted (or "captured") into a

logical change record (LCR) that is enqueued for further processing. The queue is propagated by a dedicated process that moves LCRs from a capture queue to an apply queue (these queues are usually in different databases). On the receiving end, an apply process reads LCRs from the queue and applies those changes appropriately (again, based on a set of rules).

Some of the most powerful applications for Streams are extract, transform and load (ETL) processes. These ETL processes are commonly utilized to populate data warehouses from online transaction processing (OLTP) databases. By making use of some customization of the Streams rules, OLTP changes can be captured and transformed into one or more changes for target objects in the data warehouse. This feature warrants further investigation for those requiring ETL processing, but goes beyond the scope of this paper.

While Streams is a relatively new technology, the first two releases of Oracle Database 10g have significantly enhanced and improved this feature. Oracle Support often recommends Streams when users encounter replication issues with Advanced Replication.

ADVANTAGES

- Since it uses redo log information instead of triggers to gather the changes for replication, overhead on the source system can be significantly lower than other replication options. In fact, it is possible to configure the capture process to run on a node other than the source system, creating even less overhead on the original source system. This is a very significant advantage.
- This is relatively new technology and there is significant momentum behind it so, when compared to Advanced Replication, support is easier to obtain for any troublesome issues.
- OEM has tools available to configure and manage a Streams environment.
- There is no additional cost and this feature is not a licensed option. It is available with the Enterprise Edition license for Oracle Database starting with Oracle 9i.
- The source site can continue to function even if the destination site is not available.
- Due to the speedy propagation time that is often characteristic of Streams, most applications will not need to be aware that replication is taking place. Also, rapid propagation can also reduce the number of update conflicts that occur (and therefore will often negate the need for conflict resolution programming).

DISADVANTAGES

- Especially in Oracle 9i Release 2 (9.2), it is difficult to remove a single object from Streams without removing the entire configuration and starting over.
- Documentation for using Oracle Streams in version 9.2 is not very comprehensive. Presumably, because of its new status in 9.2, few Metalink notes were published on its shortcomings and problem issues. With version 10g, documentation was improved significantly and Metalink notes are now available for most common tasks.
- The architecture for Oracle Streams is relatively complex and there are multiple additional processes and parameter settings that must be implemented in order to be successful. This architectural complexity can lead to challenging troubleshooting situations when issues are encountered.
- Learning how to implement and manage a Streams environment generally requires significant time for database administrators. Since it relies on technology that relatively few database administrators have used, it is difficult to leverage past experience to speed up the learning process.

Oracle Streams is one of the best ways to replicate data from one database to another. It is a flexible and adaptable subsystem that will serve many needs involving moving data from one database to another with low overhead costs. Customers looking for a way to replicate data quickly using the latest technology will use Oracle Streams. Streams provides faster failover capability and can also make data available at multiple sites, but has maintenance and configuration overhead higher than some other options.

HIGH AVAILABILITY OPTION: FAILOVER CLUSTERS

Failover clusters (also called "cold failover clusters") have been used with Oracle databases for many years. Originally, they were an environment that only existed on Unix systems, but they have expanded to operate on Windows servers as well.

A failover cluster is a group of two or more servers that both have access to the same storage and are on the same network. Similar to Oracle Real Application Clusters (RAC) clusters, clusterware software is employed to monitor and maintain cluster membership. Processes, network addresses and storage are configured within the clusterware framework so that the cluster knows what components are required in order to make an application (such as a database) available for use.

All application components run on a single node in the cluster (just as they would on a stand alone system) and when that node fails or needs to be taken offline, the clusterware software is instructed to stop the application's processes, deconfigure the application's network addresses, unmount the application's storage and then make those items available on a surviving cluster member. This does result in a short outage to the users, but since these steps are completely handled by the clusterware in an automated way, the outage is much shorter than if the database had to be restored or the system repaired. Since the process of "moving" an application and its dependencies to another node is automated, if the failure is unplanned (like a system crash in the middle of the night), the clusterware can bring up the application on a surviving cluster member without any interaction required from the administrators.

The clusterware software and an extra network interface on each cluster member (for a private interconnect network) are the only additional costs for an environment like this one. There is a significant amount of planning required to build a failover cluster that effectively supports multiple applications or databases. Often, reconfiguration of networks and storage is also required, but that rarely requires additional hardware purchases in today's typical data centers that include VLANs and SAN storage. Clusterware software is available from most Unix vendors (HP, IBM, Sun, Red Hat, Novell) as well as other software companies (Oracle, Veritas).

ADVANTAGES

- Failover clusters are typically easy to implement and require no application changes at all. Applications (including Oracle databases) do not need to be cluster-aware or have any special adaptations in order to operate as they normally do on a standalone host.
- Many clusterware vendors include scripts and modules to properly monitor and handle normal start and stop functions for Oracle components (database and listener). This translates into better support for Oracle databases in these cluster environments as well as shorter implementation cycles since tested and certified software can be employed.
- Oracle Clusterware can be utilized without additional license (as of September 2006) if some relatively simple requirements are met. It provides the capability to handle monitoring and failover for any process and can handle dependencies between processes (like start the database before the app server).
- The systems used in a failover cluster do not need to have the same hardware configuration. They may have different numbers of CPUs and different amounts of memory. However, they must use the same operating system and processor architecture.

DISADVANTAGES

- Often, failover clusters for databases will require additional testing to determine the application behavior when the database is restarted without restarting the application server.
- Some sites are challenged to find a storage layout and management scheme that allows them to scale the number of databases or applications to a significant number.
- The system administrators and database administrators need some additional training to understand necessary changes in management practices to maintain the cluster applications.

- Some failover cluster software licenses can be expensive.
- Troubleshooting problems can be more challenging since there are more moving parts in a multi-node cluster than a standalone system.
- A failover event causes an outage for the application. The duration of that outage is commonly advertised as very short but, for most Oracle Database environments, the outage is usually between 2 and 5 minutes. Testing this outage duration and setting appropriate expectations are important points for a new implementation.

Failover clusters are a common choice for organizations that need automated high availability, but do not want to utilize Oracle Real Application Clusters for cost or complexity reasons. The biggest benefit of a failover cluster environment is that the Oracle Database has no knowledge or configuration differences from a normal single-instance environment. Outage periods for a failover cluster are generally less than 5 minutes with no intervention required. For organizations that require higher availability, Oracle Real Application Clusters is the most common option.

HIGH AVAILABILITY OPTION: ORACLE REAL APPLICATION CLUSTERS

Oracle Real Application Clusters (RAC) is a database option that enables multiple systems to run instances of a single database. The database is placed on shared storage and is accessed directly from each node that runs an instance. The set of nodes are bound together using clusterware.

Clusterware software manages cluster membership, maintains heartbeats between all nodes via multiple methods to ensure that no node suffers a partial failure and corrupts the database. In the case that some node has a partial failure, the clusterware software will fence (eliminate) that node from the cluster via a method that ensures it cannot corrupt the shared storage or database. Oracle Clusterware (previously known as Oracle Cluster Ready Services) is a mandatory part of every 10g RAC installation.

Oracle Real Application Clusters also leverages many features of Automatic Storage Management (ASM) quite well. ASM makes the traditional problem of managing shared storage much easier by providing an easy way to manage all database storage without relying on a file system at all. While ASM is not a feature exclusive to RAC, it is very useful for RAC environments.

With Cache Fusion, Oracle RAC can be used for almost any application without changes or special adaptations to the application. When configured properly, users can connect to any node in the cluster and their sessions will sustain an instance failure by being reconnected to a surviving instance on another cluster member. This type of configuration provides a very fast failover and users may not even know that a failure has occurred. For high-availability configurations, two-node clusters are very common. Sometimes, a RAC environment may require more than two nodes to handle normal workload. In that case, an N+1 configuration is the most common scenario.

ADVANTAGES

- Very fast failovers so that users are able to resume work activity quickly and often are not aware that an outage has happened.
- This product is undergoing active development at Oracle and the future is very bright.
- There is a well-defined, active community of support (www.oracleracsig.org).
- Hardware vendors and software vendors recognize Oracle RAC as an option and perform certifications against RAC configurations.
- Current Oracle database administrators can easily extend their knowledge to include RAC administration since it shares the same foundation as single-instance databases.
- Sites using failover clusters will be familiar with most of the concepts in a RAC cluster.
- No new or special hardware is usually required, only reconfiguration of existing environments.

- The architecture employed is "shared everything." This means that loss of one node does not affect availability of data. All data are still immediately available to all other nodes in the cluster.
- Training and skilled professional services are available from multiple vendors.
- For small clusters (less then 4 CPUs in the cluster), the Oracle Standard Edition database includes RAC capability with no additional license needed.
- No third party software is required. Oracle owns the whole software stack on top of the operating system.
- RAC environments are just as stable as other Oracle database environments when implemented properly.

DISADVANTAGES

- For most environments, the Enterprise Edition database license is used. To add RAC to that environment, the RAC Option license must be purchased and that can be expensive.
- There is special knowledge required to properly plan and implement a cluster of any real size. RAC is not a "set it and forget it," environment.
- For some applications, reconfiguration and occasionally redevelopment activity is necessary.
- Additional network setup is required to configure a private interconnect network for all cluster nodes. Gigabit Ethernet is common for most clusters built today.

Oracle RAC is the ultimate option for Oracle Database availability. No other option can provide the immediate, non-stop failover that Oracle RAC provides. While cost and complexity are typically higher for RAC environments, most sites that deploy RAC are willing to accept those challenges in exchange for superior technology.

CONCLUSION

Depending on your needs and your budget, one of the availability options or a combination of multiple options examined in this paper should provide a viable solution for your organization. Many sites already employ the basic availability options, but most are focused on maximizing system availability while minimizing cost. This whitepaper has provided the necessary information to help you identify your best choice given your uptime and budgetary requirements.

CREDITS

- Export/Import & Data Pump: Padraig O'Sullivan, University of Maryland
- RMAN: Diego D. Morales, IT Convergence
- Data Guard: Ariel Gustavo Paez, IT Convergence
- Replication and Streams: Diego Cutrone, IT Convergence
- Advanced Replication, Streams, Failover Clusters, and RAC: Dan Norris, Piocon
- Editorial assistance: Keith Thomas, IT Convergence, and Beth Norris, freelance editor

REFERENCES

- Oracle Technology Network (http://www.oracle.com/technology)
- Oracle Database documentation: Oracle Database Utilities book, Oracle New Features Guide book
- Oracle Database documentation: Oracle RMAN Concepts and Administration
- Oracle Database documentation: Oracle Data Guard Concepts and Administration book

- Oracle Database documentation: Oracle 9i Advanced Replication Release 2 (9.2) (A96567-01)
- Oracle Database documentation: Oracle 9i Database New Features Release 2 (9.2) (A96531-02)