

# White Paper

## Disaster Recovery Solutions for Oracle<sup>®</sup> Database Standard Edition RAC

v1 Sept 2011

# Table of contents

Executive Summary	1
RAC Architecture	2
RAC DR Challenges	3
Physical Standby for RAC	5
RAC to RAC	7
RAC to Single Instance	8
Log Extract	9
Transport	10
Log Apply	10
Exception Management	11
More than Technology	13
ROI	17
Conclusion	20

## Executive Summary

*A true DR solution requires physical replication of the primary database, not just logical replication.*

The Oracle database management system is a favorite platform for organizations of all sizes, from small and medium organizations to large enterprises. While these organizations have a range of budgets available to them, they all need to ensure the availability of their systems and the integrity of their data. Irrespective of the organization size, start-up to large enterprise, it is critical that these systems are available to customers at all times.

These systems are often e-commerce or self-service systems that must be available to clients on a 24x7 basis in order for the organizations to provide service and value to their clients, and to remain competitive. Historic forms of data backup such as nightly tape backups and file system replication are not viable options for these round-the-clock systems.

To deliver this level of availability, systems need a combination of high availability (HA) and disaster recovery (DR). While HA ensures that the system continues to operate subsequent to the failure of individual components, DR ensures continued availability in the event of a “disaster”, or complete failure of the production site.

When Oracle is implemented using the Real Application Clusters (RAC) option, it provides aspects of high availability and scalability, but does not in itself deliver DR; DR requires a remote copy of the database. Several options exist for adding DR to Oracle RAC.

When using Oracle’s product suite, remote replication of the database is delivered with Data Guard, necessitating the use of Oracle Enterprise Edition.

To reduce the added cost and complexity of deploying Oracle Enterprise Edition (EE), organizations look for alternatives that will allow them to remain with Oracle Standard Edition (SE). The options are third party tools or home-built solutions. Home-built solutions are risky, costly to maintain, often unproven and difficult to support. Third party tools however can provide a cost-effective, reliable, supported and proven solution to DR.

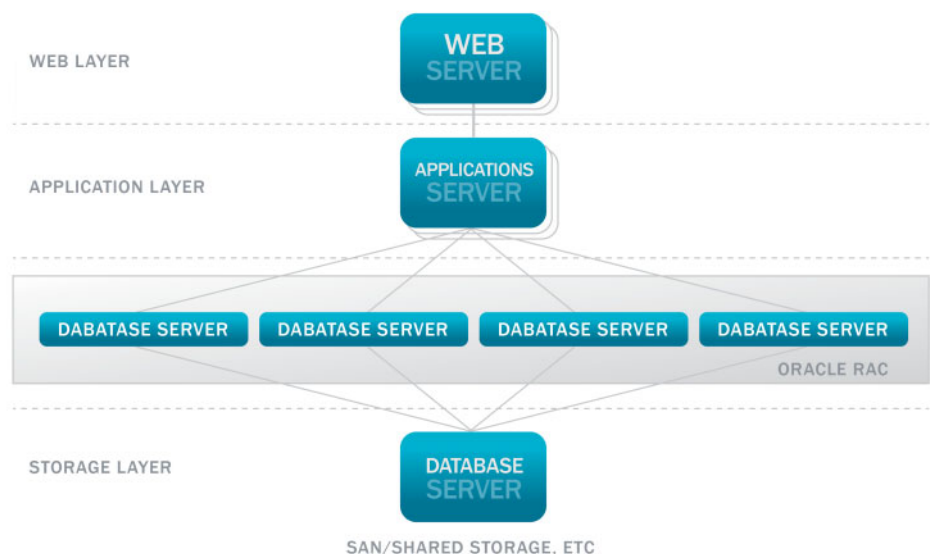
This whitepaper examines the architecture of Oracle RAC and what it provides in terms of high availability and scalability. It then looks at the requirement for DR, the options for delivering disaster recovery with Oracle SE, and reviews the return on investment of various approaches. It discusses Dbvisit Standby as an example of a cost effective and proven fully featured product that delivers DR for any Oracle SE database, including RAC.

# RAC Architecture

*Oracle RAC protects the system against failure of one or more discrete components within a single site.*

Oracle Real Application Cluster (RAC) allows multiple servers to run Oracle RDBMS software simultaneously while accessing a single shared database, thus providing a clustered database. This delivers high availability to Oracle databases so the database remains available to the application (and to users), even in the event where one or more servers fail.

The diagram below illustrates the various infrastructure layers in a modern web application environment, and the method for providing HA and scalability at each layer. Oracle RAC provides HA and scalability at the database server layer. Redundant servers and additional products such as load balancers can be used to provide HA and scalability at the application and web server layers. While all servers are dependent on the physical database, deploying redundant storage solutions, such as a SAN, can provide high availability at the storage level. Although fundamentally different, other architectures such as client-server all have similar multi-tier structures, and achieve the same high availability and scalability benefits at the database server and storage levels from Oracle RAC. Oracle RAC also provides a scalability solution, allowing organizations to deploy additional database servers as required to support additional load. Adding database servers in this way allows the workload to be shared across a larger pool of resources, thereby providing an increase in processing capacity. Oracle RAC protects the system against failure of one or more discrete components within a single site. While this provides high availability and scalability, RAC does not provide disaster recovery.



# RAC DR Challenges

*The sole purpose of disaster recovery solutions is to provide certainty in the event of a site failure. If this cannot be assured, then the approach needs to be questioned.*

Disaster recovery provides redundancy at the site level, ensuring continuity of service in the event that the entire primary (production) site is no longer available. To achieve disaster recovery requires a remote site that can be brought online and into production in the event of a failure of the primary site.

In order to make the standby site useful it is critical that up-to-date and accurate data is maintained within the standby database. In the event of a failover from the primary to the standby site, any data not transferred from the primary site's production database to the standby site ahead of the failure will be lost. To minimize the potential for data loss, software is used to periodically transfer data changes to the standby site, thus keeping the primary and standby databases in step. Data Guard is the Oracle product that provides data replication and delivers disaster recovery. Data Guard is a feature rich and proven solution that can be configured to provide various levels of data synchronization, each providing a different balance of resilience and performance. Data Guard does however require Oracle Enterprise Edition. For Oracle customers with Standard Edition who are otherwise satisfied with the features it offers, this presents a significant and often uneconomic increase in license fees.

Thankfully there are alternatives to Oracle Data Guard, including non-Oracle products developed and supported by third-party vendors, and home-built systems developed and maintained by in house database support teams.

In assessing available options for DR, we need to assess the total cost of ownership of each solution, including implementation, operation and support. We also need to determine the level to which each will deliver an effective and reliable mechanism for:

- (i) Synchronization between the primary and standby sites*
- (ii) Providing a failover process to move operations from a failed primary site to the standby site in a timeframe that meets the agreed service levels*
- (iii) Restoring the full dual-site state once the primary site is operational again*

While a system developed in-house may initially seem like a cost-effective solution, this may not be the case overall. Once developed these systems need to be supported, placing a critical dependency on the staff member or members who originally developed them. These systems also need to be maintained and tested with each change to the underlying environments, including updated operating systems and database management systems, as well as infrastructural changes.



# White Paper

*The option of third party tools provides an ideal balance between the in-house development and Data Guard approaches.*

This can add significant overhead to the total cost of the solution. The final part of the assessment is the reliability of the solution in the areas outlined above. Without extensive testing of exception cases it is difficult to be confident that the solution will operate as required. This level of testing can be difficult to perform in a customer environment with minimal test environments and infrastructure.

The sole purpose of these Data Recovery solutions is to provide certainty in the event of a site failure, and if this cannot be assured, then the approach needs to be questioned.

The risks of an in-house solution can be summarized as follows:

- (i) Development cost unknown – may or may not be cheaper to build than licenses for a third party product.*
- (ii) Reliability of maintaining synchronization, failover management and recovery is all unproven*
- (iii) The need to maintain staff skills in order to provide the ongoing support, maintenance, enhancement and operation of the solution*

The option of third party tools provides an ideal balance between the in-house development and Data Guard approaches, delivering a proven, reliable and supported product in a cost effective manner.

The question then becomes selecting the right third-party tool. Dbvisit Standby is a proven, reliable and stable product used by customers in over 60 countries round the globe. It is used here as an example of a cost-effective third-party tool for providing DR for Oracle RAC on Standard Edition, via physical data replication.

# Physical Standby for RAC

*With a DR implementation on the primary site, if a RAC node fails Dbvisit Standby on another node will automatically take over the log shipment from the failed node*

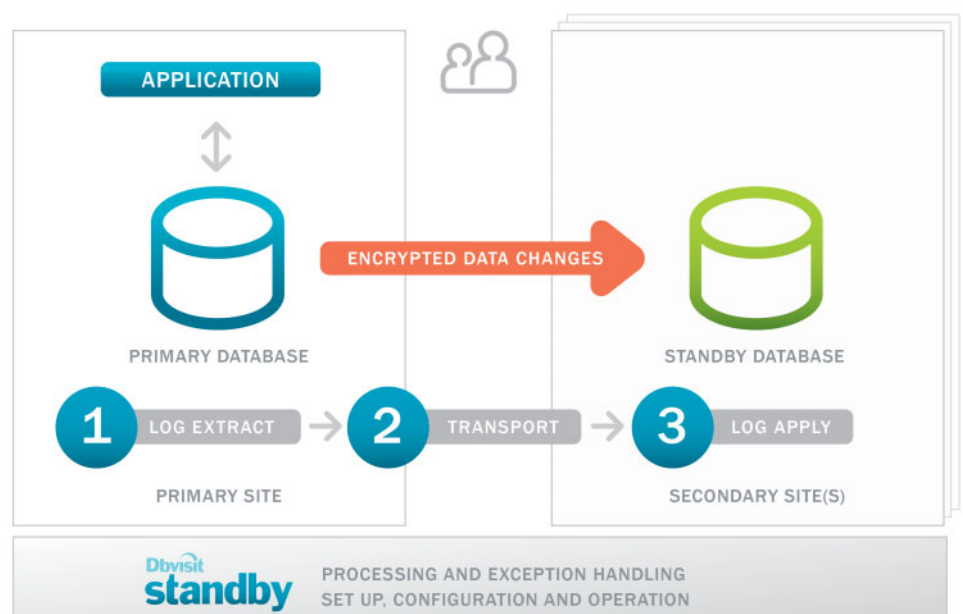
Dbvisit Standby provides a proven, reliable and cost effective solution for disaster recovery for Oracle Standard Edition databases, including those running on RAC implementations.

A Dbvisit Standby configuration consists of a primary (production) database and one or more standby (secondary) databases. The primary and secondary databases connect using a secure mechanism (SSH) over TCP/IP. Databases can exist anywhere provided they can communicate with one another via TCP/IP.

The standby database is initially created from the primary database using the Create Standby Database (CSD) command. Once the standby database is available Dbvisit automatically synchronizes the databases by transmitting the archived redo logs from the primary database and applying them to all of the secondary databases.

The following diagram provides an overview of the data replication process performed by Dbvisit Standby. It illustrates the three key steps in the replication process

- 1) *The extraction of logs on the primary site*
- 2) *The transportation of logs to the secondary site*
- 3) *The application of logs on the secondary site:*





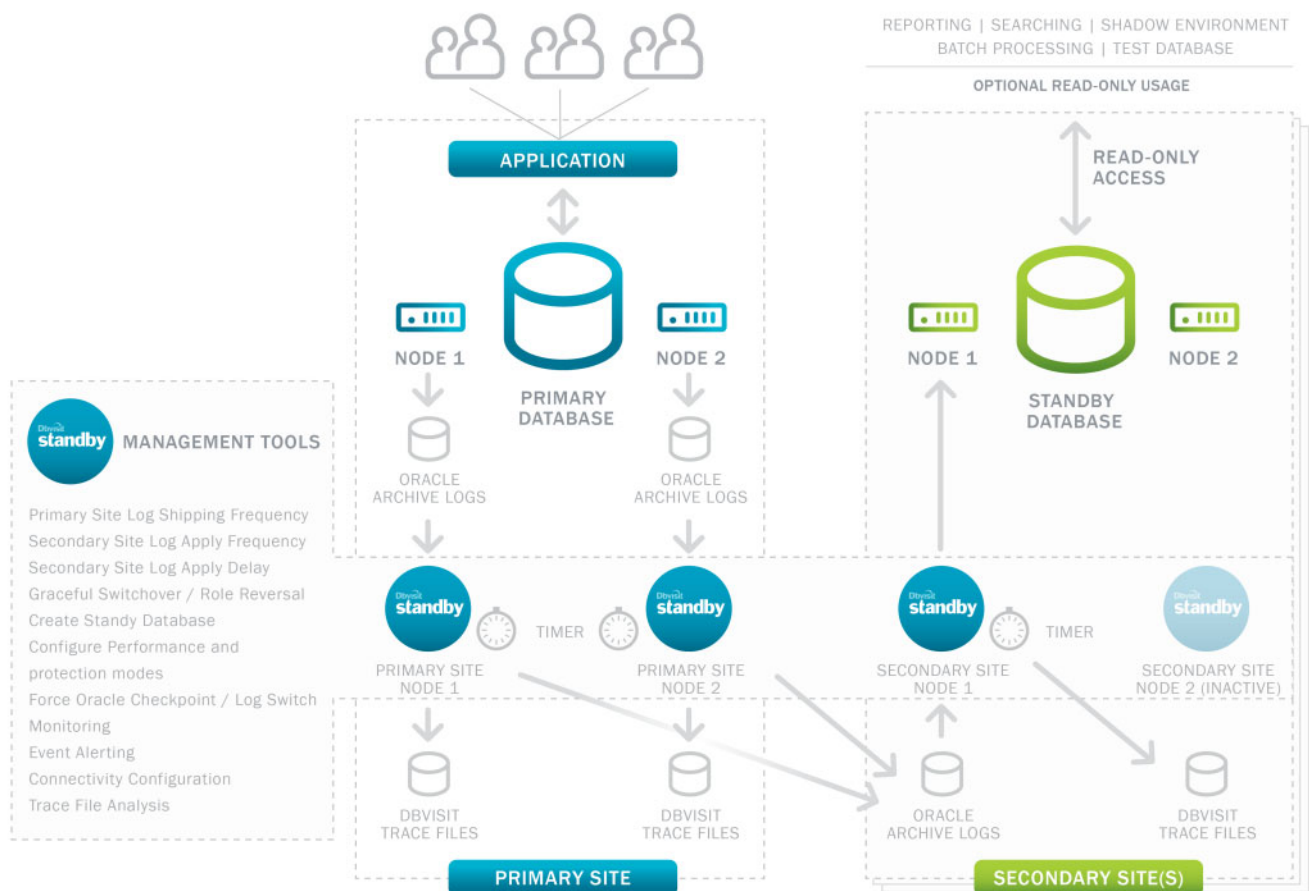
# White Paper

Dbvisit Standby provides specific support for Oracle RAC implementations, including support for RAC at both primary and secondary sites, or RAC installed only at the primary site. Configuration of Standby for RAC implementations is very similar to that used for standard single Oracle instances, and the RAC database does not need to be restarted, or any Oracle parameters changed.

With a RAC implementation on the primary site, if a RAC node fails and is unavailable, Standby on another node will automatically take over the log shipment from the failed node.

## RAC to RAC

The following diagram illustrates the architecture of the Dbvisit Standby solution when implemented in a RAC environment on both the primary and secondary sites:

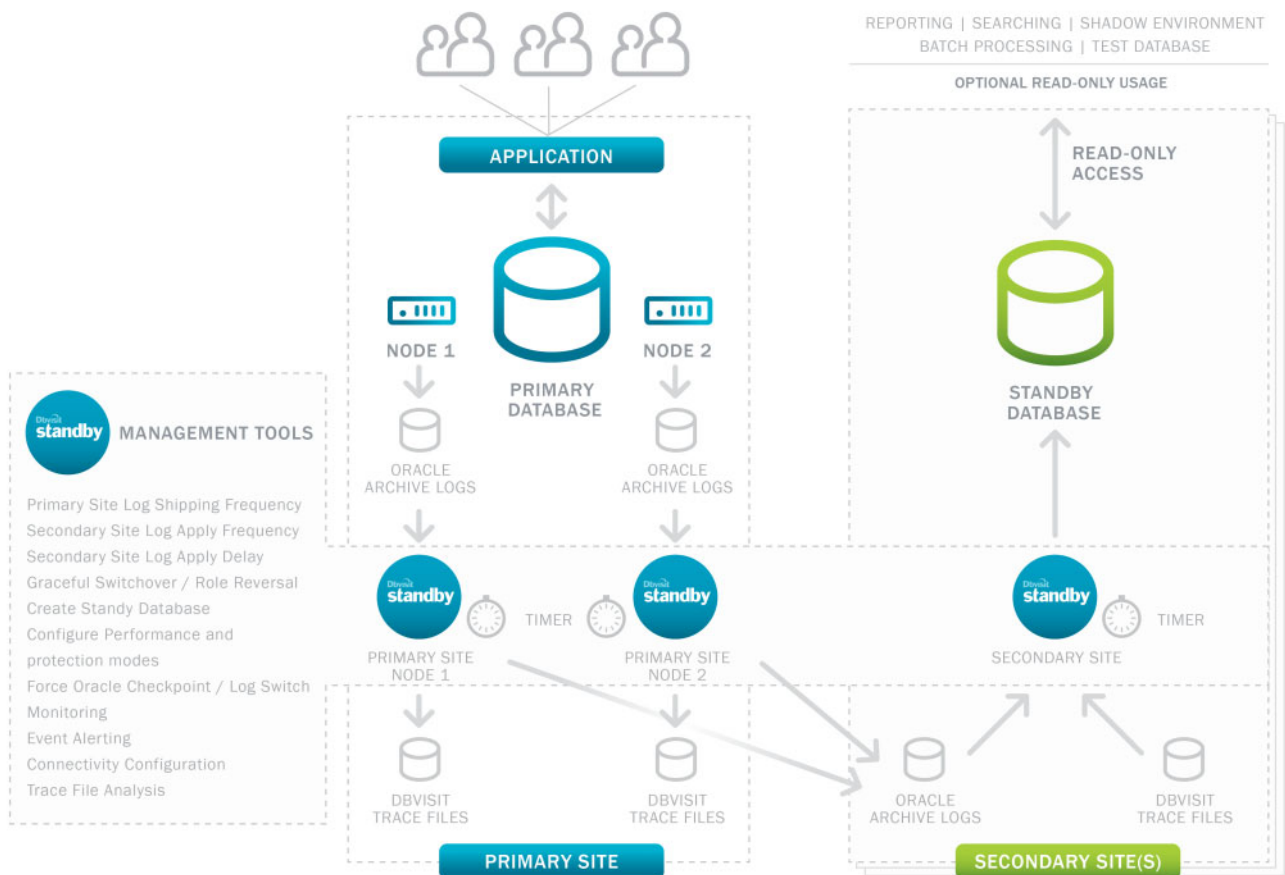


For implementations with RAC installed at both the primary and secondary sites, Standby operates in the following manner:

- 1) Dbvisit Standby, running on each of the RAC nodes at the primary site, independently transfers each node's archive logs to a shared location on the secondary RAC cluster.
- 2) Dbvisit Standby is scheduled to run on a single node on the secondary site's RAC cluster, with that single node applying the archives from all primary nodes to the secondary database.
- 3) In the event that the active node on the standby site becomes unavailable, Dbvisit Standby can be activated on an alternative node, applying logs sourced from the shared location.

# RAC to Single Instance

The following diagram illustrates the architecture of the Dbvisit Standby solution when a RAC environment is implemented on the primary site, and the secondary site consists of a single Oracle instance only:



For implementations with RAC installed at the primary site and a single instance at the standby site, Standby operates in the following manner:

- 1) *Dbvisit Standby, running on each of the RAC nodes at the primary site, independently transfers each node's archive logs to the single standby server.*
- 2) *Dbvisit Standby is scheduled to run on the standby database server, with the single node applying the archives from all primary nodes to the secondary database.*

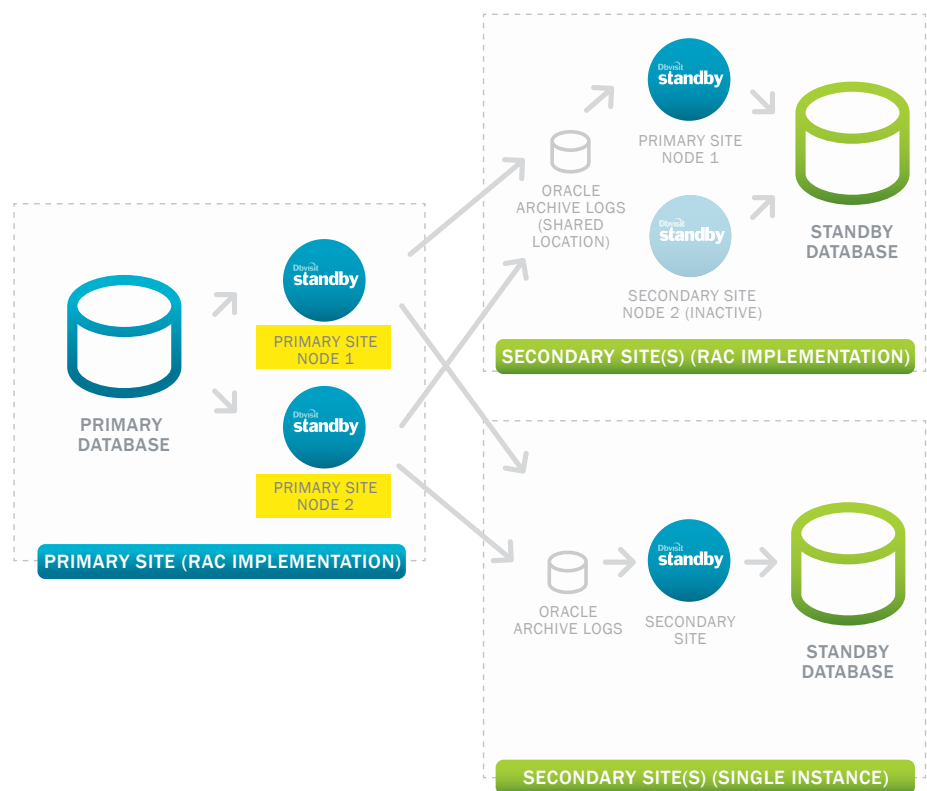
## Log Extract

*RPO is generally a definition of what an organisation determines is an acceptable loss in a disaster situation.*

Key to the replication of data is the ability to extract changes from the primary production database and make these available for delivery to the standby databases. Dbvisit Standby is built on top of the tried and proven Oracle logging mechanism. Dbvisit Standby is configured to periodically pick up the Oracle archive logs and prepare these for transmission to the standby databases.

In a RAC implementation, Standby runs independently on each database server node and transmits the Oracle Archive Logs from that server to the standby server. If the standby server is also running a RAC configuration then each primary node transfers files to a single shared location on the secondary server. If the standby server is running a single Oracle instance then the primary nodes all send their log files to the single server at the secondary site.

The following diagram illustrates the alternative extract processes for sending to RAC and single-instance Oracle implementations from a RAC implementation, with the extraction processes highlighted in yellow. It also illustrates the mechanism used by Standby on the secondary servers to apply these logs (discussed below).

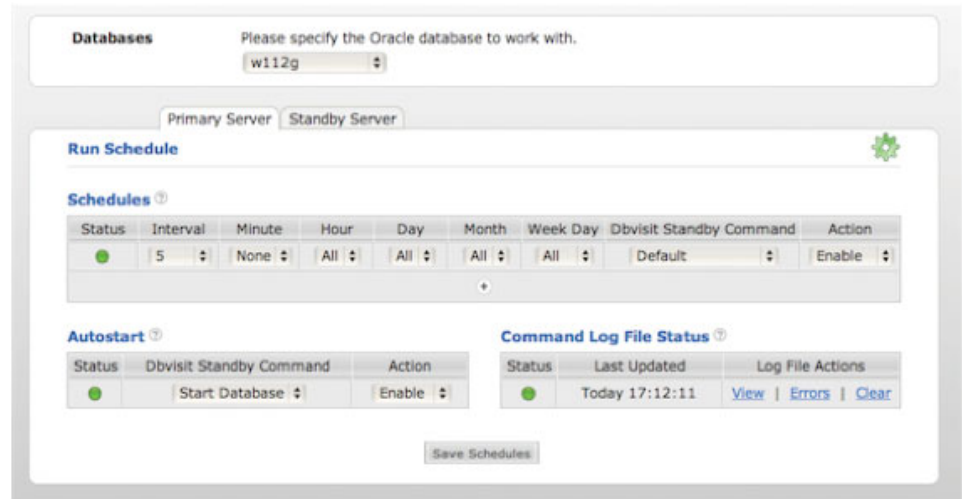


Organizations set the extraction frequency from the primary site based on their Recovery Point Objective (RPO), which describes the acceptable amount of data loss measured in time. The RPO is expressed in time and is generally a definition of what an organization determines is an “acceptable loss” in a disaster situation.

# White Paper

*Logs are compressed and transported securely, reducing the bandwidth required and transfer times.*

The following screenshot illustrates the transfer schedules on the server managed by Dbvisit Standby:



This replication scheduling can be achieved using the Dbvisit scheduler or via operating system tools such as cron or the Microsoft Windows task scheduler.

## Transport

The primary (production) server initiates the transfer of extracted logs at the same frequency as the extraction process. Dbvisit Standby uses a secure transport channel to send the logs to the standby server where they are available to be applied by a separate and independent process. The logs are compressed prior to being transferred, thus reducing the bandwidth required, and transfer times.

In the situation where communication with the standby server is lost for a period of time, the primary server continues to accumulate logs. Dbvisit Standby monitors the network connection with the standby server, transmitting these logs to the standby server once communication is restored.

## Log Apply

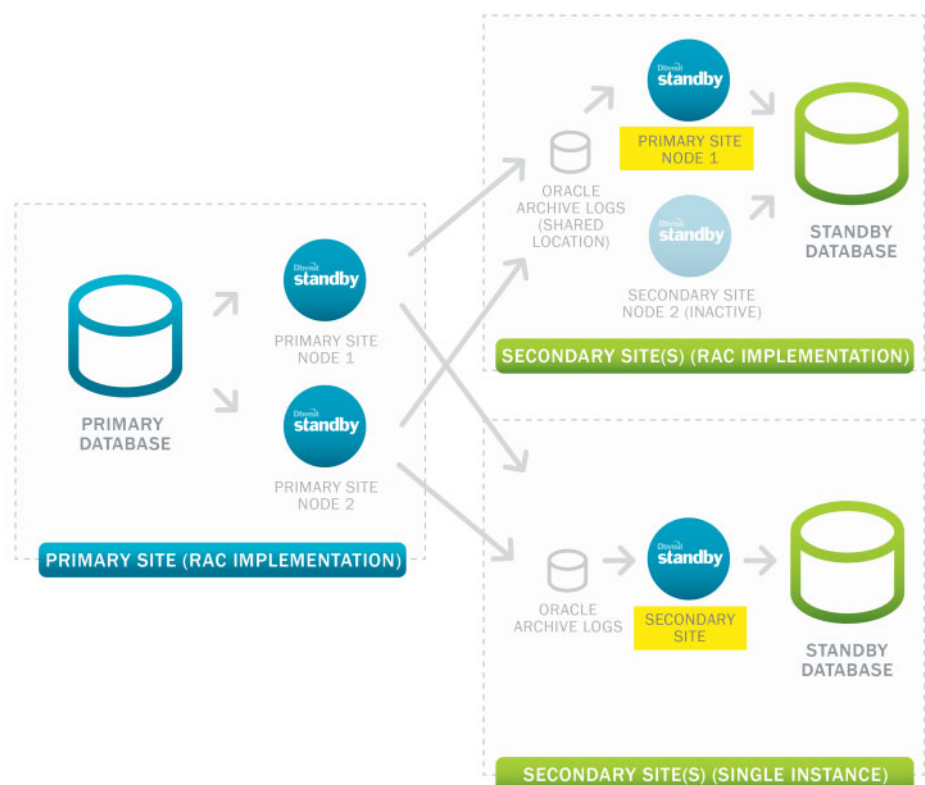
Dbvisit Standby running on the secondary server is configured to periodically check for new update logs delivered by the primary server. The standby server's process runs completely independently of the primary server's process, allowing both operating in the absence of the other, and thereby identifying and alerting for situations where the other has failed.

At a defined frequency the standby server applies all update logs received since the last time it ran.

*In the event of a failure, one of the remaining nodes in a standby cluster can still retrieve and apply the logs.*

For a standby environment operating a RAC implementation, a single node within the cluster is scheduled to collect the logs received from the primary server and apply these to the standby database. The primary server sends all logs to a single shared location so that, in the event of the failure of one of the nodes within the standby cluster, one of the remaining nodes can still retrieve and apply the logs.

The following diagram illustrates the alternative log application processes on the standby sites for RAC and single-instance Oracle implementations, with the active log application processes highlighted in yellow.



## Exception Management

Key functionality of any disaster recovery solution is the operation required to switch processing from the primary database to the standby database in the event of a failure. In order to ensure full disaster recovery and continuity of service to users, as defined by the agreed service levels for the application, all components must be subject to redundancy, and the database is only one component of this. In this paper however, we deal solely with the database.

### Failover

Dbvisit Standby provides alerting to assist a monitoring system in identifying a failure of the primary database. This includes failure by the standby server to receive updates from the primary server for a configured period of time.



# White Paper

*The database is only one component that must be subject to redundancy to ensure full DR and continuity of service.*

The `MAX_TIMES_TRIED` parameter is used to set to the number of empty update cycles to perform prior to alerting (e.g. A two minute frequency and `MAX_TIMES_TRIED` set to five, will have the effect of raising an alert when no updates have been received for ten minutes).

Other monitoring and alerting within the application stack would also be used to identify failure of the primary server.

Once failure of the primary site is identified, the following command is used on the standby server to force it to take over as primary database:

```
dbv_orasStartStop activate oracle_database [Yes].
```

Activation cannot be reversed, and a standby database will have to be built before the replication service can be restarted.

## **Graceful Switchover**

Dbvisit Standby supports a graceful switchover function that is used to transition the primary database to a standby database and the standby database to a primary database. There is no loss of data during the transition and the standby database does not have to be rebuilt. Graceful switchover may be used to enable planned DR testing as well as maintenance of servers, operating systems, database management systems and applications etc.

A role swap, or graceful switchover, is initiated using the command:

```
dbv_orasStartStop switchover oracle_database [unique_key]
```

# More than Technology

*Disaster recovery encompasses process, policies, procedures and equipment.*

As a part of overall Business Continuity, Disaster Recovery (DR) is aimed at ensuring users continue to obtain service in the event of a disaster that makes the primary (production) site unavailable. Disaster recovery encompasses process, policies, procedures and equipment.

A full strategy for business continuity and disaster recovery must include the following considerations:

## *(i) Processes*

- a. Rules for declaring a disaster, including critical factors, thresholds and pre-emptive invocation*
- b. Processes for invoking DR*
- c. Rules for operation whilst under DR (change freeze, cancellation of staff leave, halting the use of test and training systems etc.)*
- d. Steps to recover from reduced operation (disaster mode) back to normal operation*
- e. Staff, supplier and customer communication plans*

## *(ii) Hardware Design*

- a. Design and sizing for the standby site such that it will provide the required levels of performance and resilience in a cost-effective manner.*
- b. Use of standby site for other purposes (reporting, testing etc.) during normal operation.*
- c. Design of each site such that they offer local redundancy and thereby reduce the likelihood of invoking full disaster recovery mode as the result of failures within a site.*

## *(iii) Site Location*

- a. Proximity to natural and man-made hazards and avoidance of high-risk regions*
- b. Geographical separation between primary and standby sites.*

*(iv) Replication approach including frequency, in-built delays, monitored events, alerting processes etc.*

## Disaster Management

The following diagram illustrates the steps around a disaster event from normal operation until normal operation is resumed:



*Key in overall DR is the ability to identify that a disaster has taken place.*

In order to achieve DR, components across the entire site must be replicated, kept updated, and available for activation in the event of a failure, not just the database. The database is a key component of this however, and it is the database that is addressed by Dbvisit Standby.

### Identifying Failure

*With a disaster identified it is critical that an organisation has the systems and processes in place to respond inline with its agreed service objectives.*

Key in overall disaster recovery is the ability to identify that a disaster has taken place. This identification can be as simple as awareness of a major physical event that renders the primary site inoperable. It can however require more detailed monitoring and alerting when exceptions are indicated.

Dbvisit Standby provides monitoring of its replication operations, and the availability of the primary and secondary sites to participate in the replication. Alerting of exceptions outside of the configured acceptable bounds provide advance warning to operators of a failure or disaster.

Typically this alerting would be fed into an organization-wide network monitoring solution such as IBM Tivoli or Zabbix. These systems are able to collect observations from multiple sources and analyze them to report and alert operators based on pre-configured rules.

Having observed a disaster event that renders the primary site inoperable, the disaster recovery process must be initiated.

## **Responding to Failure**

*Once replicated systems are operable and stable the task turns to one of recovery to the normal state.*

With a disaster identified it is critical that an organization has the systems and processes in place to respond, and regain service in line with its agreed service levels. These will be measured in terms of the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). To meet the RTO the standby site must be operational as the new primary site within the agreed timeframe. In order to meet the RPO, the restarted systems must be in service at the standby site with the agreed levels of data loss (measured in time).

Dbvisit Standby aids in the achievement of both RTO and RPO.

- *Standby helps achieve RTO through monitoring, alerting and fast failover. Failover can be automated through the implementation of a simple custom script, or can be left as a manual process to retain a level of control over when failover is invoked.*
- *RPO is delivered through configuration of appropriate replication frequencies, and the provision of tools to allow the rapid application of pending logs on the standby server (particularly in the event of a configured delay in the log file application).*

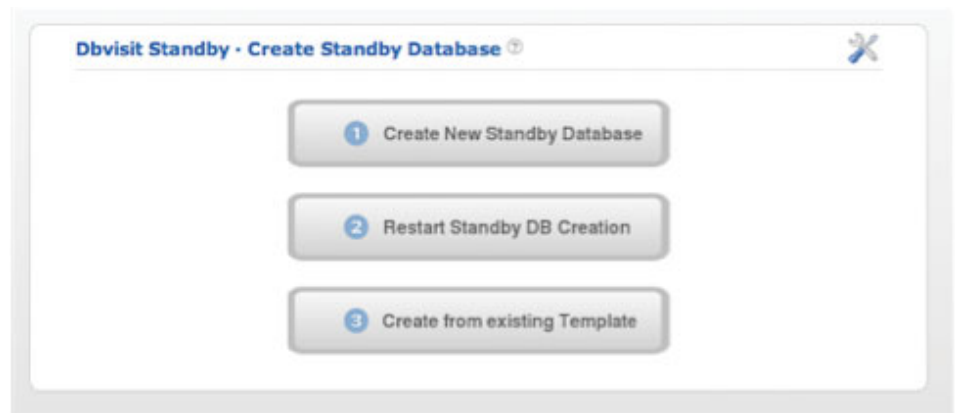
## **Recovery**

Once disaster recovery has been invoked and the replicated (standby) systems are operable and stable, the task turns to one of recovery to the normal state where all environments (primary and standby) are again available. In the event of a true disaster striking the primary site, this will typically involve considerable effort and take some considerable period of time.

Dbvisit Standby aids in the recovery steps through tools to enable the recreation of new standby databases from the operational production database. This process is highly automated and can be quickly used to establish standby databases at any number of sites.

The following screenshot illustrates the creation of a new standby database from the primary server:

# White Paper



Once all preparations are complete, and it is deemed appropriate to divert production operation back to the primary site, Dbvist Standby's Graceful Switchover facility allows an operator to do this. Graceful switchover is used to switch roles, transitioning the primary database to a standby database and the standby database to a primary database. This occurs with no loss of data, and the standby database does not have to be rebuilt.

## ROI

*A solution such as Dbvisit Standby can increase the returns available by allowing secondary sites to serve additional roles during normal operations.*

The primary objective of a disaster recovery solution is to ensure continued service in the event of a disaster. The investment, or cost, for the DR solution is relatively easy to calculate as it is based on measurable items such as hardware and systems, as well as the ongoing support costs for the standby sites. The return, or value, however is largely measured in terms of costs that will be saved during an unplanned outage, including direct staff costs to recover from the failure and lost company revenue.

ROI can be improved through a combination of increased returns and reduced costs.

### **Increased Returns**

A solution such as Dbvisit Standby can increase the returns available from a DR solution by allowing secondary sites to serve additional roles that add value during normal operation. For example, Standby supports the ability for a single primary site to push updates to multiple standby sites.

This allows organizations to maintain multiple synchronized standby databases, opening the opportunity to utilize secondary sites not only for DR purposes, but also for other read-only activities such as test and shadow environments, reporting systems and searching. These additional environments, kept continuously synchronized, can add significant value to an organization, including:

- (i) Improved solution quality through production-scale test environments*
- (ii) Improved solution scalability through segmentation of non-core activities (such as reporting) to standby servers*
- (iii) Improved production performance through the offloading of processing to these standby databases*
- (iv) Extended life of production hardware through reduced load on production servers*

### **Reduced Costs**

In terms of the investment, or cost, of a disaster recovery solution it is important to consider the total cost of ownership, including the establishment, operation and execution in a disaster:

# White Paper

*A high risk DR solution has little value if an organisation cannot be certain it will work in a disaster.*

- (i) Hardware costs (e.g. servers, network infrastructure)*
- (ii) System software costs (e.g. operating systems)*
- (iii) License costs (e.g. replication software, database management system)*
- (iv) Ongoing operational costs during normal operation (e.g. staff and equipment)*
- (v) Cost of system downtime during failover to DR site (based on RTO)*
- (vi) Cost of lost data in the event of a disaster (based on RPO)*
- (vii) Cost of recovery from a disaster*

A proven solution such as Dbvisit Standby reduces costs in a number of these areas:

- (i) Licensing costs for the Dbvisit Standby product are lower than alternative products.*
- (ii) Dbvisit Standby supports Standard Edition, unlike products such as Oracle Data Guard that require an upgrade to Enterprise Edition, saving Oracle license fees.*
- (iii) Dbvisit Standby is a proven product that reduces the operational costs compared with a home built system through reliability, automated operation and regular updates to support new Oracle releases.*
- (iv) Dbvisit Standby can assist in delivering a recovery point objective of as little as 60 seconds, thereby minimizing the cost of lost data*
- (v) Dbvisit Standby provides tools to reduce the effort and increase the reliability of recovering from a disaster and reinstating the normal operational state. These tools reduce the cost of recovery from a disaster.*

## **Risk Management**

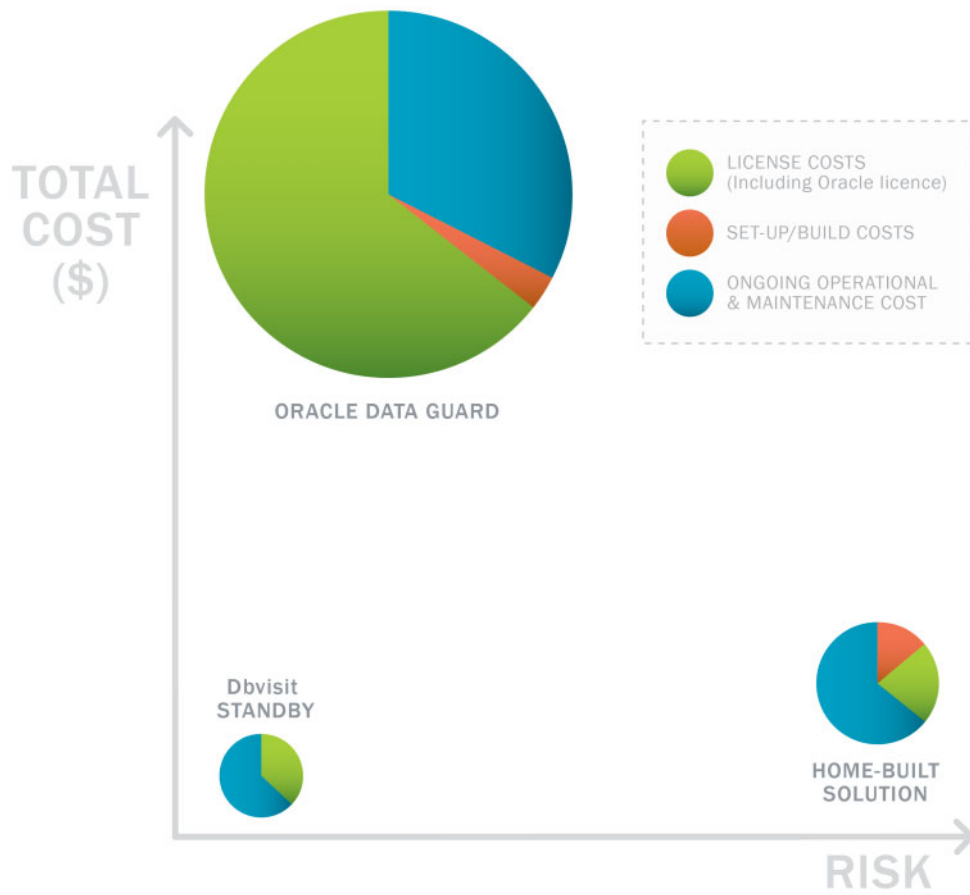
By its nature, disaster recovery is all about managing risk. Reducing risk typically increases cost, leading organizations to balance the cost of their DR solution against the benefit that it provides.

Some solutions may come at a lower cost, but if these are unproven then their value must be questioned. For this reason it is critical when evaluating disaster recovery options to consider the risk profile for each, both in absolute terms and relative to the cost. For example, a high risk DR solution has little value if an organization cannot be certain it will work in a disaster.

*It is important to consider the total cost of ownership, including the establishment, operation and execution in a disaster.*

As described above, Dbvisit Standby delivers a proven and low risk solution at a relatively low total cost, including licencing, implementation and ongoing operation.

The following diagram demonstrates this concept, with axis for cost and risk, and the size of the individual circles illustrating relative cost of the solutions.



## Conclusion

*A true DR solution requires physical replication of the primary database, not just logical replication.*

This paper has discussed the Oracle RAC environment and identified the fact that while it provides high availability and scalability, RAC does not provide Disaster Recovery (DR). To provide DR for any Oracle implementation, including RAC, requires a remote site with automated and regular synchronization from master database to the standby site (or sites). A true DR solution requires physical replication of the primary database, not just logical replication.

We then discussed three alternatives for adding DR to a RAC implementation including Oracle Data Guard, a home built solution, or a third party product. We demonstrated the high total cost for Oracle Data Guard, including the requirement to deploy Oracle Enterprise Edition. We discussed how home built systems are unproven, risky, dependant on the continuity of in-house resources, and are likely to incur higher than expected overall costs, including development, operation and maintenance. Finally, we discussed the approach provided by a third party tool, and used Dbvisit Standby as an example.

We showed how a cost effective DR implementation, with Oracle RAC, is achievable using Dbvisit with Standard Edition. We showed how this solution delivers opportunities for reduced cost and for increased value, thereby significantly improving the return on investment that can be demonstrated for a DR solution.



## Further information

For further information on this  
White Paper please contact us:  
via e-mail: [info@dbvisit.com](mailto:info@dbvisit.com)  
by phone: +64 9 950 3301

Dbvisit Software Limited  
PO Box 48180, Blockhouse Bay  
Auckland 0644, New Zealand