

March 2005

**DB2.** Information Management Software



# **Technical Comparison of DB2 HADR and Oracle Data Guard**

*IBM Software Group  
Toronto Laboratory*

---

## Contents

---

1. Introduction
2. What is HADR and how does it work
3. Comparison with Oracle Data Guard
4. Conclusion
5. References

## Introduction

In a previous article entitled "[Why should you care about the cost of your High Availability solution?](http://ftp.software.ibm.com/software/data/highlights/ha.pdf)", ([ftp://ftp.software.ibm.com/software/data/highlights/ha.pdf](http://ftp.software.ibm.com/software/data/highlights/ha.pdf)) IBM DB2 Universal Database with its High Availability Disaster Recovery (HADR) feature was compared with Oracle Real Application Clusters (RAC). In that analysis, DB2 was shown to deliver equal or better availability for considerably lower costs, with a total solution that is easier to manage, and delivers better performance. A quick glance at the HADR technology may lead one to believe that a comparison with Oracle Data Guard is more in line with an apples-to-apples comparison. If this were actually true and Data Guard was able to deliver the same availability as HADR (and therefore equal or better availability compared to Oracle RAC), why would Oracle be trying to sell RAC to customers looking for higher levels of availability? After all, RAC is a 50% uplift to the base price of Oracle Enterprise Edition. If Data Guard could meet customer's availability needs at no additional charge, why would Oracle be proposing RAC?

The fundamental issue here is that, although HADR and Data Guard "look" like similar technologies, the solutions they deliver are not equivalent. Taking a superficial look at the architectures, one may conclude that they deliver the same capabilities. Looking at the comparison on page 4 quickly shows you the advantages of HADR. This paper will distinguish between the two architectures and describe in detail how HADR delivers superior availability and manageability compared with Data Guard and how you can use HADR in scenarios where Oracle would be proposing RAC.

## What is HADR and how does it work?

High Availability Disaster Recovery is a data replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby.

A failure can be caused by a hardware, storage, network, or software failure. With HADR, the standby database can take over the workload in a matter of seconds from any of these failures including disk failure on the primary. Furthermore, you can have clients automatically redirected to the standby database (now the primary database) without the need for changes to your application with the new automatic client reroute facility of DB2 v8.2.

Data changes made on the primary server are sent to the standby database directly from the log buffer of the primary server. Thus the two servers stay completely in sync with each other. There are three modes of operation for HADR; synchronous, near synchronous and asynchronous. In all of these modes, if the standby server or network fails, the primary is unaffected.

## Highlights

### Synchronous Mode

In synchronous mode, DB2 ensures that the log records being written to disk on the primary server are also written to disk on the standby server before an application receives a successful return code to its commit statements. In this mode, there is a guarantee that no committed transactions will ever be lost as both servers stay completely in sync.

### Near Synchronous Mode

In near synchronous mode, DB2 ensures that the log records being written to disk on the primary server are in memory at the standby server (but perhaps not on disk at the standby) prior to notifying an application that its commit statement was successful. In this mode, there will never be any transactions lost unless both the primary and standby fail simultaneously. If both servers fail simultaneously **and** you restart the standby as the new primary, there is potential that the last log buffers may not have been applied.

### Asynchronous Mode

In asynchronous mode, DB2 will write the log buffer to disk on the primary server and ensure the log buffer has been passed down to the TCP/IP socket to be sent over to the standby. In this case, it would be possible to lose a committed transaction if the primary failed and the packets containing the log buffer did not make it to the standby server prior to a takeover.

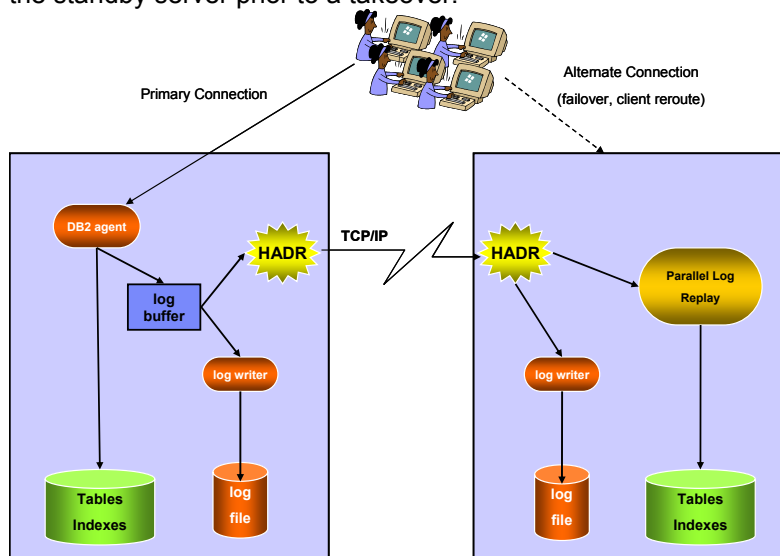


Figure 1

**“The DB2 automatic client reroute feature will automatically reconnect the application to the standby server so that the application can continue to function.”**

### Automatic Client Reroute

When a client is connected to a database and that server fails, the DB2 automatic client reroute feature will automatically reconnect the application to the standby server so that the application continues to function. Any in-flight transaction is rolled back and the application can then continue from where it left off. Automatic client reroute is configured on the database server, making mass deployments much simpler than competitor's products which require client side configuration for application failover.

---

## Highlights

---

When a client application connects to the database, the standby server information is pulled back to that client where it can then be used (at any point in the future even after the client disconnects) to automatically reestablish a connection to a standby server if the primary becomes unavailable.

### Failover and automation of failover

With HADR, the failover to a standby server is extremely simple. There is just a single command (or single click in the HADR monitor graphical interface) called TAKEOVER. The takeover command has two modes of operation. The first is a simple takeover in which the primary and standby switch roles. This allows for a graceful switch over in which the primary and standby coordinate with each other to change roles (primary becomes standby and standby becomes primary). This method is useful for rolling upgrades where you apply a fix or upgrade to the standby server, switch roles and then apply the fix or upgrade to the primary server.

In the event of a failure on the primary server you perform a TAKEOVER BY FORCE in which the standby server assumes the role of primary server without coordinating with the old primary. When this command is executed, the new primary will replay any logs it still has in memory, undo any in-flight transactions and open the database for new transactions. Note that because the standby was only processing insert/update/delete activity, most of the recently updated data pages will still be in memory on the standby and therefore the undo phase is exceptionally fast (no I/O is likely required in order to undo in-flight transactions). In fact, the undo phase of recovery completed in just 3 seconds for a 600 user SAP test described in the previous paper comparing HADR to Oracle RAC. <sup>(1)</sup>

**“the undo phase of recovery completed in just 3 seconds for the 600 user SAP test.”**

Note that there is no need to shut down and restart the standby instance during the takeover process. This, among other advantages will be discussed in detail in the following comparison with Data Guard.

## Comparison with Oracle Data Guard

Oracle Data Guard is a feature of 10g Enterprise Edition that allows for the creation of standby databases that can be kept transactionally consistent with a primary database. To achieve this, Oracle ships log buffers (or log files in some configurations) from the primary server to the standby server where the log records are replayed on the standby database. A Data Guard standby comes in two “flavours”, Logical Standby and Physical Standby. In logical standby mode, log records are converted to SQL statements and replayed on the standby database. This more closely resembles DB2’s SQL Replication and Q Replication capabilities and as such will not be discussed in this paper. In physical standby mode, log records are applied using redo logic which applies the records much in the same fashion as would occur when rolling forward a database through log files.

## Highlights

**“There are many differences that appear when you look just below the surface that make HADR a superior solution to Data Guard.”**

In this mode, both the primary and standby databases are exact physical copies of each other and the application of log buffers is similar to that of HADR (on the surface).

However, there are many differences that appear when you look just below the surface that make HADR a superior solution to Data Guard for high availability scenarios.

Below is a summary of the differences, each of which will be described in detail.

	DB2 V8.2 HADR	Oracle 10g Data Guard Physical Standby
Protection from software failure	Yes	Yes
Protection from server failure	Yes	Yes
Protection from storage failure	Yes	Yes
Protection from site failure	Yes	Yes
Support for rolling upgrades/fixes	Yes	No
Standby remains “hot” during failover	Yes	No
Sub-minute failover	Yes	No
Geographically separated	Yes	Yes
Support for multiple standbys	No	Yes
Available on Express and Standard One Editions	Feature	Not Available
Available on Workgroup and Standard Editions	Feature	Not Available
Simple to configure and monitor	Yes	No
Supports read on standby	No	Limited
Simple log management	Yes	No
Primary can be easily reintegrated after failover	Yes	No
License required on standby	1 CPU	All CPUs

### Protection from software, server, storage or site failure

Both HADR and Data Guard protect from failures such as software failure, primary server failure, storage failure or even from primary site failure. In both cases the configuration can include a second complete copy of the database in a remote location to protect from any or all of these forms of failure. It should be noted that Oracle

## Highlights

**“Real Application clusters only protects from server and software failure on a node in the cluster and has no protection for storage or site failure.”**

**“HADR allows for fixpacks (patch sets) and OS level fixes to be applied in a rolling fashion.”**

**“With Oracle Data Guard, a physical standby database does not support rolling upgrades.”**

**“With Data Guard, in order to convert a standby into a primary the standby database must be shut down and started up again.”**

Real Application Clusters only protects from server and software failure on a node in the cluster and has no protection for storage or site failure. To cover more failure scenarios, Oracle would likely recommend a combination of both Oracle RAC and Data Guard, which can significantly increase the cost of the solution (paying for all CPUs at both sites plus the 50% uplift for the RAC option).

### Support for rolling upgrades/fixes

HADR allows for fixpacks (patch sets) and OS level fixes to be applied in a rolling fashion. For example, the following steps are can be deployed to maintain maximum availability while patches are applied.

1. stop HADR on the standby
2. apply the DB2 fix or OS fix on the standby
3. start HADR on the standby – database will automatically resynchronize
4. perform a switch-roles takeover
5. stop HADR on the new standby (old primary)
6. apply the DB2 fix or OS fix to this server
7. start HADR on the new standby (old primary) – database will automatically resynchronize

At this point, since HADR is intended to be a peer to peer HA solution, you can leave the roles as they are above. Alternatively you can perform a takeover to switch roles again to get back to the original primary/standby configuration.

With Oracle Data Guard, a physical standby database does not support rolling upgrades. Both primary and standby servers must be using the same patch set.<sup>(2)</sup>

### Standby remains “hot” during failover

The following sequence of events occurs on the standby during an HADR takeover.

1. last log buffer is replayed (if not already done)
2. undo of in-flight transactions occurs – note that the buffer pool on the standby is likely full of all the recent updates so there is likely little to no random data page I/O during undo recovery.
3. new transactions are allowed to access the database

Note that the buffer pool and all other memory structures remain allocated.

With Data Guard, in order to convert a standby into a primary (during either failover or when switching roles) the standby database must be shut down and started up again<sup>(3)</sup>. This results in buffer caches, catalog caches and package caches (library caches) being torn down and recreated. Therefore a significant “brown out” period would follow a Data Guard failover. According to a presentation given by Angelo Pruscino, Principal Architect, Oracle Corporation, there is an issue with warming the buffer cache on a cold failover that can take “5+ minutes” to resolve.

**“Warm Buffer Cache** - On the new system, you start with a new system cache and it must be populated with blocks from disk. This is the less efficient because Disk I/Os are slower than Cache I/Os. It is a function of the cache size and the cache hit rate.”<sup>(4)</sup>

---

## Highlights

---

### Sub-minute failover

As demonstrated in the previous HADR paper, failover of a database supporting 600 concurrent SAP users was achieved in only 11 seconds. Clearly HADR is capable of supporting sub-minute failover.

One of the issues with Data Guard is that you must stop and restart the instance during failover which negatively impacts availability. The following quote is from a case study on Oracle's own internal processing systems which use Data Guard. It is significant to note that this system does not deliver sub-minute failover.<sup>(5)</sup>

**"Protects Against Data Loss and Downtime:** An exact replica of the Global Single Instance is maintained at a remote location. The transition of the standby server to the primary role can be completed in 15 minutes."

In a separate presentation, Marshall Presser, Principal Technologist from Oracle Corporation states that in a cold failover, the time required to "Restart Oracle" is "up to 5 minutes"<sup>(6)</sup>

### Geographically separated

Both HADR and Data Guard use TCP/IP to send log buffers from the primary server to the standby site. As such both allow for the servers to be separated by a large distance. In addition, both products offer asynchronous buffer transmission so that very large distances do not adversely affect the performance of the primary server.

### Support for multiple standby servers

With the first release of HADR, DB2 supports one primary and one standby database. The key objective is to provide the highest levels of availability possible. An issue with multiple standby servers is that the impact on the primary server becomes too great to efficiently support synchronous mode. Therefore in order to support multiple standby servers, the use of asynchronous mode is more appropriate. IBM's solution for asynchronous multi-site standby servers is Q Replication, with which you can have multiple targets for a given source database. The trade-off to consider when looking at asynchronous modes to multiple standby servers is the potential transaction loss in comparison to HADR in synchronous or near synchronous mode.

### Available on Express and Standard One Editions

For customers that are looking to deploy DB2 on servers with 2 or fewer CPUs, IBM offers DB2 Express edition

<http://www.ibm.com/software/data/db2/udb/db2express/>.

Customers looking to deploy on smaller servers do not necessarily have less interest in high availability. To satisfy this need, IBM offers the DB2 HADR Option for Express and Workgroup Editions. Oracle offers an entry level server offering called Standard Edition One which is available on 2 CPU servers.

**"IBM offers the DB2 HADR Option for Express and Workgroup Editions"**

---

## Highlights

---

**“Data Guard, along with many other features, are not available on Standard Edition One.”**

However, Data Guard, along with many other features, is not available on Standard Edition One. In order to use Data Guard, you must move up to the Enterprise Edition of Oracle.

### Available on Workgroup and Standard Editions

Similarly, DB2 Workgroup Server Edition (WSE) can be run on servers up to 4 CPUs. The DB2 HADR Option can also be purchased on top of WSE. Oracle Standard Edition also supports up to 4 CPUs, but as in the previous case, does not support Data Guard along with many other features.

### Simple to configure and monitor

HADR can be set up in one of two ways, either using the command line interface or by using the Control Center HADR Wizard. Using the GUI requires no additional software. The Control Center will connect directly to each server and perform the following tasks automatically for you (if using the command line interface, you would perform the same tasks manually):

1. Back up the primary database and restore that image on the standby
2. Set the following configuration parameters on each database
  - HADR\_LOCAL\_HOST – hostname of local machine
  - HADR\_LOCAL\_SVC – port for HADR to listen on
  - HADR\_REMOTE\_HOST – hostname of remote server
  - HADR\_REMOTE\_SVC – port for HADR to talk to
  - HADR\_REMOTE\_INST – instance name of remote database server
  - HADR\_SYNCMODE (optional) – use sync, nearsync or async communication modes
3. Start HADR on the standby
  - START HADR ON DB dbname AS STANDBY
4. Start HADR on the primary database
  - START HADR ON DB dbname AS PRIMARY

That's it! Now DB2 will take care of the rest. Any log files that are required on the standby will be sent automatically to the standby and the two databases will automatically synchronize themselves. If at any point in the future, the network between the two servers is down, the primary will stop sending records to the standby and when the standby re-establishes communications with the primary database, it will automatically resynchronize itself.

**“The manual configuration [with Data Guard] is significantly more complex”**

Oracle also allows Data Guard to be configured manually or through a graphical interface. The manual configuration is significantly more complex (as you will see below). In order to use the GUI tools, you must be using 10g Grid Control. However, Grid Control requires the installation of Oracle 10g Application Server as well as the creation of a new database to store the management and performance data. It is possible to store this information in an existing database. In either case, this management database also needs to be highly available, especially if you want to use it to drive a failover for Data Guard. Note that there is also the additional requirement to manually download the Java Cryptographic



---

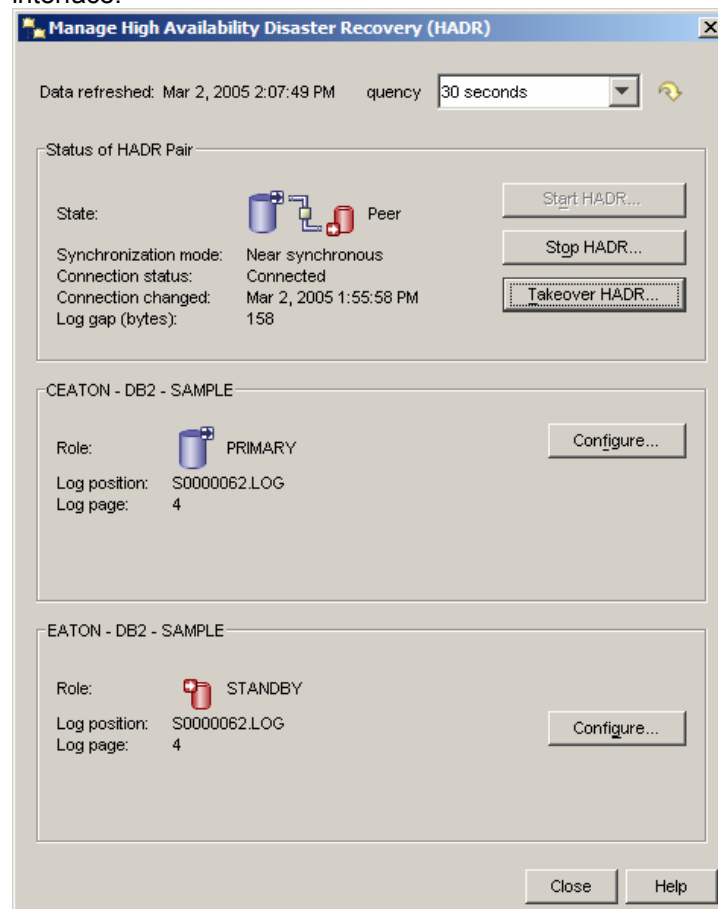
## Highlights

---

Extensions from Sun's website in order to install the application server or management agent on some UNIX servers. This management alone is significantly more than is required with DB2. Setting up Data Guard manually requires the following steps:

1. Enable forced logging so operations like load will log changes to the log buffer (DB2 allows load to run without logging while still loading changes into the standby)
2. Create a password file and ensure sys password is the same on all systems
3. Configure 17 initialization parameters per the Data Guard Administration guide
4. Create a backup of primary and restore it on standby
5. Create a separate "standby" control file and move it manually to the standby
6. Create a parameter file for the standby server then update the same 17 parameters for the standby
7. Start log apply on the standby with
  - ALTER DATABASE RECOVER MANAGED STANDBY DATABASE DISCONNECT FROM SESSION;
8. Create standby redo logs (more on this in the next section)

In terms of monitoring HADR, there are two methods. The graphical interface gives all the information you need at a glance and allows you to perform actions like starting or stopping HADR and performing a takeover. Here is a screenshot of the graphical interface.



## Highlights

This interface lets you see the state of the HADR pair (“Peer” state here indicates that the two servers are in sync). This display also tells you what log page each server is processing and a running average of the difference between the two servers in terms of bytes of log.

The second method to monitor HADR is to use the command line. A simple GET SNAPSHOT FOR DATABASE command will return the following HADR-related information:

HADR Status

```

Role                = Primary
State               = Peer
Synchronization mode = Nearsync
Connection status   = Connected , 03/02/2005
01:55:58.255208
Heartbeats missed   = 0
Local host          = CEATON
Local service       = DB2_HADR_1
Remote host         = EATON
Remote service      = DB2_HADR_2
Remote instance     = DB2
timeout(seconds)    = 120
Primary log position(file, page, LSN) =
S0000062.LOG, 4, 0000000014824226
Standby log position(file, page, LSN) =
S0000062.LOG, 4, 0000000014824188
Log gap running average(bytes) = 158
    
```

**“If one server goes down or if the log gap is increasing, DB2 will proactively notify the administrator of the issue and propose a resolution action.”**

This includes the same information displayed in the graphical interface. In addition to these interfaces, the DB2 Health Monitor automatically monitors an HADR pair. If one server goes down or if the log gap is increasing, DB2 will proactively notify the administrator of the issue and propose a resolution action.

### Support for Read on Standby

Oracle Data Guard allows for the log replay to be suspended on the standby server so that the database can be opened in read-only mode. This, however, elongates the failover times as the standby server cannot be both in read-only mode and be replaying logs at the same time. In some reports, delaying the log apply can add 15 minutes to the failover times. If read on standby is a higher priority then DB2 Q Replication would be a better alternative. Q Replication allows for reads and write on the remote databases. Combined with automatic client reroute, this solution provides “instant” failover as there is no need to recover in flight transactions after a failover.

### Simplified log management

DB2 HADR does not require any special log management in order to function. In fact, you simply specify where you want to archive your primary database log files and the rest is taken care of. The reason is that HADR was initially designed as a **log buffer** transport mechanism while Data Guard was initially designed as a **log file** transport mechanism that has since been adapted to also

## Highlights

**“This greatly simplifies the setup and management of database logging in favour of DB2.”**

**“[With Oracle] there is not only double archiving going on (for the same log files) but there is also twice as many log archive destinations that need to be configured.”**

allow for log buffer shipping. This greatly simplifies the setup and management of database logging in favour of DB2. As shown in Figure 1, there is an HADR process that takes the log buffer and passes it over to the HADR process on the standby machine. The log writer process is still responsible for writing log buffers to disk, and the archive process is still only responsible for archiving and retrieving log files locally. On the standby system the HADR process writes the log buffers to the log file and applies them to the database. While the standby database is in standby mode, log files are only required for restart recovery. There is no log archiving that occurs on the standby server and the log files are automatically configured in terms of size and number of files. As well, log files that are no longer required on the standby are simply reused.

Oracle Data Guard was initially designed to archive log files to the remote server and replay those logs on that server. As such, the default mode of operation is to configure a log archive destination to the remote server and allow the log archive process to send that file over at the time of a log switch. This mechanism has been enhanced to allow the log writer to directly send log buffers over to the standby but it still uses the original LOG\_ARCHIVE\_DEST configuration parameter. When log buffers are received on the standby, they are written to standby redo log files. When these redo log files are full, they are applied to the standby database. This elongates the failover time as there are more log records to be replayed. Oracle 10gR1 has added a feature called real time apply that will apply the log buffers directly to the database as they arrive as well as into standby redo log files. However, these standby redo log files are then archived on the standby server as well. So there is not only double archiving going on (for the same log files) but there are also twice as many log archive destinations that need to be configured. Here is a configuration example for a “simple” Data Guard physical standby database showing only the relevant logging parameters.

On the Primary Server

```
LOG_ARCHIVE_CONFIG='DG_CONFIG=(chicago,boston) '
DB_FILE_NAME_CONVERT=
'/arch1/chicago/', '/arch1/boston/', '/arch2/chicago/', '/arch2/boston/'
LOG_FILE_NAME_CONVERT=
'/arch1/chicago/', '/arch1/boston/', '/arch2/chicago/', '/arch2/boston/'
LOG_ARCHIVE_FORMAT=log%t_%s_%r.arc
LOG_ARCHIVE_DEST_1=
'LOCATION=/arch1/boston/
VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
DB_UNIQUE_NAME=boston'
LOG_ARCHIVE_DEST_2=
'SERVICE=chicago LGWR
VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE)
DB_UNIQUE_NAME=chicago'
LOG_ARCHIVE_DEST_STATE_1=ENABLE
LOG_ARCHIVE_DEST_STATE_2=ENABLE
FAL_SERVER=chicago
FAL_CLIENT=boston
```

## Highlights

**“With DB2 it is possible to reintegrate the old primary into the cluster as a standby.”**

**“In Oracle Data Guard, a failover requires the original primary to be rebuilt”**

**“Oracle licensing requires that all CPUs on a standby server be fully licensed using the same metric as the primary server.”**

A similar set of configuration parameters (with all the names of the services reversed) is required on the standby server.

### **Primary can be easily reintegrated after failover**

In the event of a failure on the primary server, the standby server can be forced into the primary role. With DB2, the command is called TAKEOVER BY FORCE, in which the standby does not need to communicate with the primary prior to taking over the role as the primary database. With DB2 it is possible to reintegrate the old primary into the cluster as a standby. When in synchronous (SYNC) mode, DB2 ensures that the logs on both servers are identical so reintegration only requires an HADR start command on the old primary in order for it to become the new standby. In the case of NEARSYNC, the only possible loss of transaction is if the primary and standby fail simultaneously. If this is not the case then a simple HADR start on the old primary will reintegrate that server as a new standby. In the case of ASYNC, there is the possibility that the failover to the standby occurred before log records made it to the database on that server. However, it is still recommended that the HADR start command be issued on the old primary after that server comes back up. DB2 will automatically check the log streams on both sites to determine if there were any transactions lost. If no transactions are missing, the old primary will automatically be reintegrated as a standby. If there are missing transactions, Recovery Expert can be used to list the missing transactions and the new standby can be rebuilt from a backup of the new primary.

In Oracle Data Guard, a failover requires the original primary to be rebuilt which adds additional work and elongates the time required to revert back to the original primary server. Here is a quote from the 10gR1 Data Guard Concepts and Administration manual.<sup>(7)</sup>

“During failovers involving a physical standby database:

- In all cases, after a failover, the original primary database can no longer participate in the Data Guard configuration.”

“To reuse the old primary database in the new configuration, you must re-create it as a standby database using a backup copy of the new primary database”

If flashback database was configured on the original primary in advance of the failure (which is not the default behavior) then Oracle provides a more complex procedure to reintegrate the failed primary.

### **License required on standby**

In DB2 an idle standby only requires a single CPU on that server to be licensed. This is true for all editions of DB2 (Express, Workgroup or Enterprise). Oracle licensing requires that all CPUs on a standby server be fully licensed using the same metric as the primary server.

From Oracle’s Database Licensing document <sup>(8)</sup>

**“Standby:** One or many copies of the primary database are maintained on separate server(s) at all times. These systems are configured for disaster recovery purposes. If the primary database

fails, the standby database is activated to act as the new primary database. **In this environment, the primary and the standby databases must be fully licensed.** Additionally, the same metric must be used when licensing the databases in a standby environment."

This difference can result in significant savings for a DB2 HADR configuration when compared to Oracle Data Guard.

### Conclusions


DB2 High Availability Disaster Recovery (HADR) was designed with two goals in mind. The first goal is ultra fast failover in the event the primary database fails for any reason. The second goal is simplicity of setup, management and monitoring. As demonstrated in this paper, not only is HADR simple to use, it delivers failover rates that exceed Data Guard's ability and in fact delivers failover rates that meet and exceed Oracle Real Application Clusters at a significantly lower cost.

It is important when evaluating availability solutions that you not simply look at the concepts but rather at how the underlying technologies deliver the availability and what the cost is to your business. HADR is designed to deliver both a highly available database and lower your overall total cost of ownership.

See a demonstration of HADR today at [http://demos.dfw.ibm.com/on\\_demand/Demo/IBM\\_Demo\\_DB2\\_HADR-Jan05.html](http://demos.dfw.ibm.com/on_demand/Demo/IBM_Demo_DB2_HADR-Jan05.html) or ask your IBM representative for a proof of technology.

### References

1. The SAP R/3 Sales and Distribution workload was not an SAP Benchmark or an SAP endorsed test. Rather the SAP workload was used to simulate a real application scenario.
2. Section 2.3.2 describes the software requirements for Oracle Data Guard. [http://download-west.oracle.com/docs/cd/B14117\\_01/server.101/b10823/standby.htm#50961](http://download-west.oracle.com/docs/cd/B14117_01/server.101/b10823/standby.htm#50961)
3. Step 6 Oracle Data Guard Concepts and Administration page 7-14 [http://download-west.oracle.com/docs/cd/B14117\\_01/server.101/b10823.pdf](http://download-west.oracle.com/docs/cd/B14117_01/server.101/b10823.pdf)
4. [http://download.oracle.com/owparis\\_2003/40256.ppt](http://download.oracle.com/owparis_2003/40256.ppt) page 8
5. Case Study: Oracle E-Business Suite with Data Guard Across a Wide Area Network <http://www.oracle.com/technology/deploy/availability/pdf/OracleGlobalITProfile.pdf>
6. <http://www.bwbug.org/docs/BWBUG-May2004.ppt> page 9..
7. Oracle Data Guard Concepts and Administration page 7-14 and 7-19 [http://download-west.oracle.com/docs/cd/B14117\\_01/server.101/b10823.pdf](http://download-west.oracle.com/docs/cd/B14117_01/server.101/b10823.pdf)
8. Oracle Database Licensing <http://www.oracle.com/corporate/pricing/databaselicencing.pdf>



© Copyright IBM Corporation 2004  
IBM Canada  
8200 Warden Avenue  
Markham, ON  
L6G 1C7  
Canada

Printed in United States of America  
03-2005  
All Rights Reserved.

IBM, DB2, DB2 Universal Database, OS/390, z/OS, S/390, and the ebusiness logo are trademarks of the International Business Machines Corporation in the United States, other countries or both.

UNIX and Unix-based trademarks and logos are trademarks or registered trademarks of The Open Group. Intel and Intel-based trademarks and logos are trademarks or registered trademarks of Intel Corp. Other company, product or service names may be the trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The information in this white paper is provided AS IS without warranty. Such information was obtained from publicly available sources, is current as of 03/01/2005, and is subject to change. Any performance data included in the paper was obtained in the specific operating environment and is provided as an illustration. Performance in other operating environments may vary. More specific information about the capabilities of products described should be obtained from the suppliers of those products.

.