October 2013

Technical Comparison Oracle Database 12c vs. IBM DB2 10.5: Focus on High Availability

Executive Overview 1
Introduction 2
Unplanned and Planned Downtime 3
Overview: Oracle's High Availability Solutions 3
Minimizing Unplanned Downtime 5
Minimizing Planned Downtime 6
HA Comparison: Oracle and DB2 7
Oracle vs. DB2 – Addressing Unplanned Downtime 12
Addressing System Failures 12
Addressing Data Failures 18
Addressing Disaster Recovery 24
Addressing Human Errors 41
Oracle VS. DB2 – Addressing Planned Downtime 44
Addressing System Maintenance 44
Addressing Data Maintenance 48
Oracle High Availability Best Practices 52
Oracle High Availability Customers 52
Conclusion 53
Reference 54

Executive Overview

Today's businesses depend heavily on their databases. Should applications and data become unavailable, the entire business may halt. Revenue and customers may be lost and penalties may be incurred. Bad press can have a lasting effect on both customers and stock prices. Certainly, providing continuous data availability is essential for today's businesses.

Oracle Database 12c comes with an integrated set of High Availability (HA) capabilities that help organizations ensure business continuity by minimizing the various kinds of downtime that can affect their businesses. These capabilities take care of most scenarios that might lead to data unavailability, such as system failures, data failures, disasters, human errors, system maintenance operations and data maintenance operations. As the rest of this document shows, IBM DB2 10.5 database (for Linux, Unix and Windows) provides rudimentary functionality for both high availability and data protection, and is multiple releases behind Oracle in terms of the breadth and depth of HA functionality.

Oracle is the database that runs mission critical, highly available enterprise applications for well-known global companies such as NTT DOCOMO, CERN, VocaLink, Turkcell, National Australia Bank, HDFC Bank, Coca Cola, Thomson Reuters, MetLife, Starbucks, Sabre, Monsanto and Telefonica. When it comes to the ability to provide reliable, highly available and continuous service to customers, Oracle is the database of choice over competing solutions such as DB2.

.

Introduction

Any organization evaluating a database solution for enterprise data must also evaluate the High Availability (HA) capabilities of the database. Data is one of the most critical business assets of an organization. If this data is not available and/or not protected, companies may stand to lose millions of dollars in business downtime as well as negative publicity. Building a highly available data infrastructure is critical to the success of all organizations in today's fast moving economy.

This document provides an in-depth comparative assessment of the HA capabilities available with Oracle Database 12c and IBM DB2 10.5. The intended audience of this document are IT managers, architects and executives who are evaluating these two databases for their businesses, and are interested in knowing to what extent the HA capabilities of these databases can protect their data and maintain business continuity.

Unplanned and Planned Downtime

One challenge in designing a high availability IT infrastructure is examining and addressing all possible causes of downtime. Downtime can be classified into two primary categories: unplanned and planned. IT organizations should consider potential causes of both unplanned and planned downtime while designing a fault tolerant and resilient IT infrastructure.

Unplanned downtime primarily results from system failures or data failures (e.g. because of human errors, disasters, data corruptions). While such failures may be infrequent, the magnitude of their adverse impact on business operations is significant, leading to high costs of downtime. Planned downtime, on the other hand, is caused by scheduled maintenance activities (e.g. data changes, system upgrades), which are part of the daily operations of any data center. The challenge is to complete the maintenance activity as transparently as possible causing minimal to no disruptions to business operations.

IT managers interested in a high availability solution to meet the demands of their businesses must assess these solutions based on some key metrics, such as:

- The comprehensiveness of their HA capabilities to address various causes of downtime
- Ease-of-use to manage and adapt these solutions to changing business requirements
- Ability to utilize redundant components for effective business use and maximum return on investment.

Any solution that is comprised of a disjointed set of technologies that address HA issues in an isolated manner but not in an integrated fashion, and/or a solution comprised of a lot of redundant, but basically idle components, will not meet the demanding HA requirements of today’s enterprises. With this perspective, the remaining sections perform an HA-based analysis of the DB2 and Oracle databases.

Overview: Oracle’s High Availability Solutions

Oracle Database 12c comes with an integrated set of high availability capabilities (ref. Fig. 1) that help organizations minimize the various kinds of downtime that can affect their businesses. The next few sections provide an overview of these capabilities. For further details on each of these capabilities, please refer to [1] and [2].

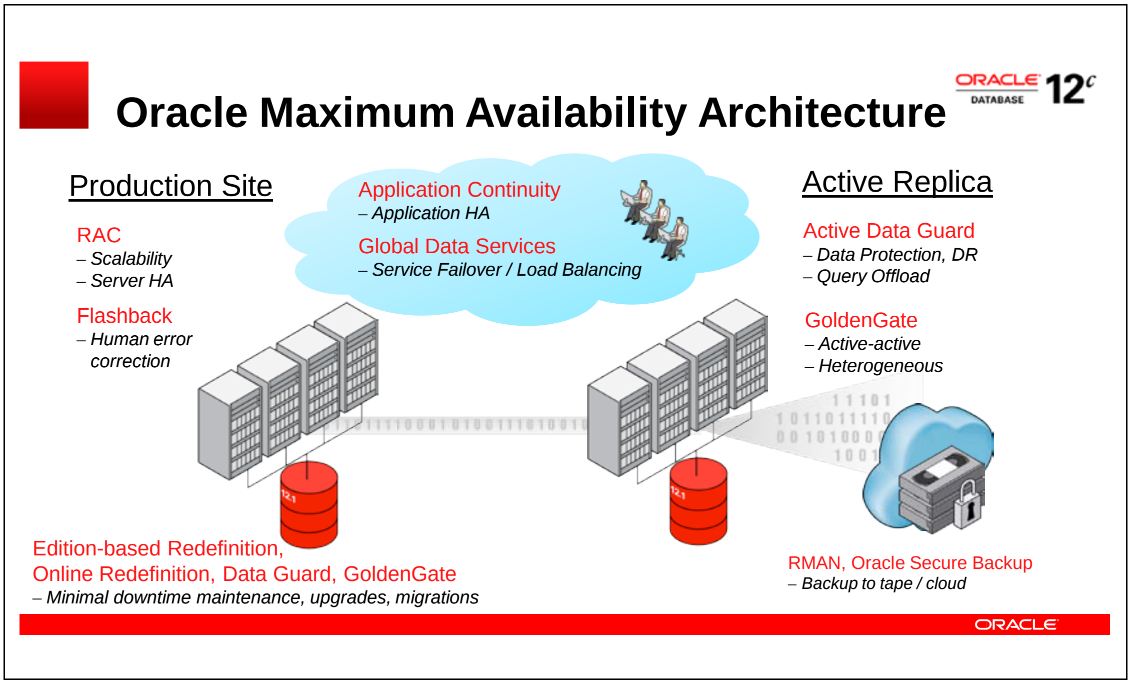

Figure 1: Integrated High Availability Features of Oracle Database 12c

Note

Oracle Multitenant, a new option for Oracle Database 12c, delivers groundbreaking technology for database consolidation and cloud computing. The Multitenant architecture drives down IT costs by enabling a true 'manage-as-one' architecture for consolidation and virtualization of the database tier. The Multitenant architecture also makes extreme high availability a fundamental requirement when database consolidation is applied to business-critical applications. By definition, database consolidation is an exercise of 'putting all eggs in one basket.' The more successful you are at driving down cost through consolidation, the more eggs are in a single basket, and the greater is the operational and financial impact to the business should an outage occur.

New high availability (HA) capabilities in Oracle Database 12c are designed to provide the extreme level of availability required for consolidating databases onto Private Clouds. This includes support for multitenant architecture across all Oracle HA features, new levels of redundancy, transparent failover of in-flight transactions, zero-data loss disaster protection at any geographic distance. The Oracle Multitenant architecture represents the next-generation in database technology, and long-standing and time-proven Oracle HA design principles are ready from day one to provide the extreme availability required by consolidated environments.

Minimizing Unplanned Downtime

To protect from server failures, Oracle offers Real Application Clusters (RAC), which allows multiple servers to access a single Oracle database in a clustered environment. A benefit of this approach is scalability and high availability without requiring application code changes.

To protect against data failures of various kinds – e.g. those that result from storage failures, human errors, corruptions and site failures, Oracle database offers a suite of features. One of them is Automatic Storage Management (ASM), which offers integrated volume manager capabilities for Oracle. ASM provides native mirroring of database files for extra protection. To protect from human errors, Oracle database offers the Flashback suite of features (e.g. Flashback Database, Flashback Table, etc.) with which it is very easy to rewind the state of the database to a known safe point in time, and undo the effects of human errors without requiring long downtimes.

For protection of data from various media failures, Oracle database offers Recovery Manager (RMAN), which is a comprehensive backup, restore and recovery solution for the Oracle database. With RMAN, backups of the Oracle database can be taken online, without requiring expensive downtime. Furthermore, Oracle Database offers the Fast Recovery Area, which is a unified disk-based storage location for all recovery-related files and activities in an Oracle database. The automation and integration between RMAN and Fast Recovery Area provide an enhanced disk-based backup and recovery solution integrated with the Oracle Database. In addition, Oracle offers a tape and cloud backup solution as well – namely, Oracle Secure Backup (OSB), and Oracle Secure Backup Cloud Module. OSB is integrated with RMAN, providing performance optimizations and faster tape backups not available with other solutions. The cloud backup module allows administrators to use the familiar RMAN interface to backup Oracle database data changes over the cloud to Amazon S3 storage.

To protect from site or storage failures that could result from localized or regional disasters – such as fire, earthquakes, hurricanes, malicious acts, etc., Oracle offers Data Guard. Data Guard also protects against server failure for configurations where Oracle RAC has not yet been deployed. In a Data Guard configuration, multiple standby databases are connected to, and kept synchronized with, the production or primary database over a network. In the event of an unforeseen disaster at the primary data center, the processing can be easily switched to one of the standby databases, with no data loss, if desired. Data Guard standby databases can routinely be utilized for activities such as reporting, backups and quality assurance testing, without compromising data protection. Through the Active Data Guard option, this reporting could be done in a real-time manner using physical standby databases, and those databases could also be utilized for fast incremental backups. This enables active utilization of all Data Guard standby databases, and an instant ROI on the Data Guard investment.

Beyond Data Guard's disaster recovery (DR) solution, Oracle offers Oracle GoldenGate, focused on information sharing and data integration. GoldenGate allows data changes to be distributed from one or more source databases – which could be Oracle or non-Oracle, to one or more target databases, which also could be Oracle or non-Oracle. GoldenGate offers flexible capabilities such as data subsetting, data transformations, multi-master replication, active-active configurations with conflict detection, zero downtime upgrades, etc.

Oracle Database 12c introduces two new innovative features that extend the benefits of HA to the application layer. One is called Application Continuity that masks outages from end users and applications by transparently replaying the in-flight database sessions following recoverable outages within the infrastructure. The other new feature, Global Data Services (GDS) is both an availability and scalability solution that enables workload routing, connect-time and run-time load balancing and application service orchestration across replicated databases (using replication technologies such as Active Data Guard and GoldenGate).

Minimizing Planned Downtime

Planned downtime, which includes activities such as routine operations, periodic maintenance, and new deployments, can be just as disruptive to operations, especially in global enterprises that support users in multiple time zones. As with minimizing unplanned downtime, Oracle database offers a suite of capabilities that help eliminate or minimize planned downtime.

With the Online Table Redefinition capability, Oracle database supports many data maintenance operations without disrupting database operations or users updating or accessing data. For example, database tables can be redefined – changing table types, adding, dropping or renaming columns, changing storage parameters, etc. – all without interruption to end-users who are viewing or updating the underlying data. With the Rolling Upgrades capability, using Data Guard or GoldenGate, upgrades of database patchsets or major releases can be done in a rolling manner, minimizing application downtime. Patches can be installed on running Oracle instances in a fully online manner using the Online Patching capability. With the Edition-based Redefinition capability, application upgrades can be facilitated with minimal downtime by enabling the upgrades of the database components of the application while they are in use.

The Oracle Database dynamically accommodates changes to hardware configurations such as adding and removing processors from an SMP server, adding and removing nodes in a RAC cluster, dynamically growing and shrinking its shared memory allocation, adding and removing database disks online without disrupting database activities using ASM, etc. Finally, Data Guard or GoldenGate can be used for minimizing downtime during large-scale migrations such as data center moves, SAN migration, technology refresh, etc.

HA Comparison: Oracle and DB2

IBM DB2 v10.5 [3] cannot match the depth and breadth of Oracle's high availability capabilities. This paper takes each high availability challenge and compares how Oracle Database and IBM DB2 10.5 address the challenge, demonstrating how DB2 still continues to lag Oracle significantly in this regard.

Note

The previous version of this document [4] compared DB2 Version 9.7 with Oracle Database 11g Release 2. Since DB2 Version 9.7, IBM has released the following database versions and – as can be seen below – these new versions have only limited incremental enhancements in the area of availability.

DB2 VERSION	*HA ENHANCEMENTS*
9.8	*Introduction of DB2 pureScale in a limited fashion – non-integrated with other DB2 HA features such as HADR, supports only AIX & Linux. Ref. [5]. In contrast, Oracle has offered the RAC capability since Oracle9i (2001).*
10.1	*HADR now has multiple standby support (but only up to 3 standbys), archived log file compression, delayed replay, spool incoming data to disk on standby to overcome a performance limitation. Replication: schema-level replication. pureScale: better install support, RDMA over Ethernet support for AIX, table partitioning support. Ref. [6]. In general, these are all incremental enhancements addressing some of the existing product deficiencies.*
10.5	*HADR now supported with pureScale, but restrictions remain (see later); pureScale: better support for online operations, restore between DB2 pureScale Feature and DB2 Enterprise Server Edition (although with an offline backup operation), better reorganization capabilities, apply fix pack updates in a rolling manner. Ref. [7]. Net net – no significant new HA feature.*

In the rest of this document, Oracle refers to Oracle Database 12c Enterprise Edition, and unless otherwise stated, DB2 refers to IBM DB2 Version 10.5 Advanced Enterprise Server Edition for Linux, Unix and Windows (LUW). Descriptions of Oracle Database 12c features are available at the Oracle Database 12c documentation site [8]. Unless otherwise noted, descriptions of DB2 10.5 are based on the IBM DB2 Version 10.5 Online Documentation for Linux, Unix and Windows [9].

For easy reference, the following table provides a list of the major differentiators between Oracle and DB2 for every category of downtime.

TABLE 1: KEY HIGH AVAILABILITY DIFFERENTIATORS – ORACLE VS. DB2

ADDRESSING SYSTEM FAILURES	ORACLE	DB2
Recovery advisories	Yes	No
Data availability during rollback	Yes	No
Query performance unaffected by data skew in partitions	Yes	No
Integrated clustering technology for all major OS and server platforms	Yes	No
Unified clustering for both OLTP and Data Warehouse applications	Yes	No
Integrated application replay to mask recoverable outages	Yes	No

ADDRESSING DATA FAILURES	ORACLE	DB2
Built-in database failure detection, analysis, and repair	Yes	No
Built-in incrementally updated backup strategy	Yes	No
Unused block compression	Yes	No
Flexible backup compression levels	Yes	No
Automatic restore failover to next available backup during recovery	Yes	No
Restore preview	Yes	No
Trial recovery	Yes	No
Integrated backup encryption	Yes	No
Database cloning over the network	Yes	No
Block media recovery	Yes	No
Resumable backup	Yes	No
Incremental backups of LOBs	Yes	No
Cross-platform backup and restore	Yes	No

Integrated fine-grained table recovery	Yes	No
Fine-grained backup & recovery performance tuning capabilities	Yes	No
Integrated Mirroring	Yes	No
Proactive disk health checks with automatic corruption repair	Yes	No

ADDRESSING DISASTER RECOVERY – DATA PROTECTION AND AVAILABILITY	ORACLE	DB2
Standby apply progress has no impact on primary database performance or data protection	Yes	No
Silent corruptions due to hardware/software faults are detected at both primary and standby	Yes	No
Corrupt blocks are automatically repaired online, transparent to users and applications	Yes	No
Fast recovery from human error and logical corruptions	Yes	No
Integrated automatic database failover with guarantees of zero data loss and no split-brain	Yes	No
Controlled automatic failover for ASYNC to comply with user configurable data loss SLA's	Yes	No
Connect-time failover of applications in all cases – controlled by database role	Yes	No
Fast reinstatement of primary after failover without needing a full restore	Yes	No
Flexible support of multiple standbys	Yes	No
Real-time cascaded standbys	Yes	No
Database rolling upgrades across major versions and subversions	Yes	No
Support for some mixed primary/standby configurations	Yes	No
Integrated redo transport compression for efficient network utilization	Yes	No
Fast, non-disruptive recovery from network and standby database outages	Yes	No
Replicate stored procedures to standby database	Yes	No
Support database partitioning	Yes	No
Complete support of active-active database clusters	Yes	No
Long-distance, zero-data loss standby	Yes	No
Proactive readiness health checks on standbys	Yes	No

ADDRESSING DISASTER RECOVERY ROI - STANDBY DATABASE UTILIZATION	ORACLE	DB2
Full read consistency for queries on active standby	Yes	No
No datatype limitations for active standby	Yes	No
Continuous user access to an active standby database during apply of DDL changes	Yes	No
Continuous user access to an active standby database during resynchronization	Yes	No
Automatic monitoring and enforcement of query SLAs	Yes	No
Enhanced reporting and global temporary table support on active standby	Yes	No
Complete support between active standby and active-active database clusters	Yes	No
Automatic memory management on active standby	Yes	No
Backup and recovery operations on active standby	Yes	No
Dual-purpose standby database for Dev/Test	Yes	No
Integrated Load Balancing and Workload Management across replicated configurations	Yes	No

ADDRESSING HUMAN ERRORS	ORACLE	DB2
Built-in capability to recover dropped objects	Yes	No
Built-in capability to mine logs and audit changes using a SQL interface	Yes	No
Built-in capability to unwind granular transactions	Yes	No
Flexible tablespace point-in-time recovery	Yes	No
Built-in ability to unwind a table to a point in time in the past	Yes	No
Built-in ability to unwind the database to a prior point in time without restoring a backup	Yes	No

ADDRESSING SYSTEM MAINTENANCE	ORACLE	DB2
Simple online addition of cluster nodes that requires no data redistribution	Yes	No
Automatic rebalance with online adding or dropping of disks	Yes	No
Online patching	Yes	No
Rolling database upgrades for full patch-sets and major releases	Yes	No
Extensive support to adjust memory online	Yes	No
Helpful advisories on memory management	Yes	No
Most configuration parameters may be modified online	Yes	No

ADDRESSING DATA MAINTENANCE	ORACLE	DB2
Flexible and comprehensive capabilities to reorganize and redefine tables online	Yes	No
Online move of data files across storage	Yes	No
Online move of partitions	Yes	No
Online application upgrades with an Edition view of the underlying data	Yes	No
Fast online add column, with default value	Yes	No
Online add/modify constraint, add column, index create/rebuild do not require exclusive lock	Yes	No
Invisible indexes	Yes	No
DDL operations wait for user-specified time, if underlying resource is busy	Yes	No

Oracle vs. DB2 – Addressing Unplanned Downtime

Addressing unplanned downtime can be discussed in the context of addressing system failures and addressing data failures.

Addressing System Failures

System failures are the result of hardware failures, power failures, and operating system or server crashes. The amount of disruption these failures cause depends upon the number of affected users, and how quickly service is restored. The challenges with system failures lie in ensuring fast recovery, or better still, a higher level of fault tolerance.

As shown in the following table, Oracle provides a set of features that clearly differentiate Oracle from DB2 in terms of how effectively it addresses system failures.

TABLE 2: ADDRESSING SYSTEM FAILURES – ORACLE VS. DB2

ADDRESSING SYSTEM FAILURES	ORACLE	DB2
Recovery advisories	Yes	No
Data availability during rollback	Yes	No
Query performance unaffected by data skew in partitions	Yes	No
Integrated clustering technology for all major OS and server platforms	Yes	No
Unified clustering for both OLTP and Data Warehouse applications	Yes	No
Integrated application replay to mask recoverable outages	Yes	No

The following sections provide further details on these differentiators.

Fast-Start Fault Recovery

The Oracle Fast-Start Fault Recovery scheme is designed to minimize downtimes related to system failures. It has two components – Fast-Start Checkpointing that optimizes roll forward recovery by continually and incrementally advancing the checkpoint position, and Fast-Start Rollback that eliminates the delays associated with the rollback phase of recovery.

Fast-Start Checkpointing – Predict Average Recovery Time

To control the time to recover from system failures, Oracle allows Mean Time To Recover (MTTR) to be directly specified via a dynamic parameter, FAST_START_MTTR_TARGET. Oracle continuously estimates the recovery time and automatically adjusts the checkpointing rate to meet the target recovery time [10].

Beyond this, Oracle takes it one step further by providing real-time feedback on the cost of the target MTTR through the `v$instance_recovery` dynamic view, as well as a GUI-based advisory through Oracle Enterprise Manager. Oracle also provides an advisory through the `v$mttr_target_advice` view that simulates the cost of a range of recovery scenarios. The simulation runs in real-time based on the current production workload. Based on the output of the advisory the administrator can choose the best tradeoff between very fast recovery time and extra I/O overhead. This takes the guesswork and risk out of configuring for fast recovery.

DB2 provides no means to effectively predict or control recovery time. In DB2, there are two parameters: `page_age_trgt_mcr` or the *page age target member crash recovery configuration parameter* [11] and `page_age_trgt_gcr` or the *page age target group crash recovery configuration parameter* [12]. These two parameters are applicable to the Enterprise Server Edition and pureScale Edition respectively, and neither can be configured online.

Fast-Start Rollback – Shorten Worst-Case Recovery Time

Oracle's crash recovery time is immune to long transactions, because Oracle allows users to access the database before instance recovery rollback operation is complete through a unique on-demand rollback technology. With Oracle, once the roll-forward processing completes, the database opens for user access. Oracle does not wait until all transactions have been rolled back. Instead, transactions are rolled back in the background while new user transactions access the data. If one of these new user transactions encounters data that was locked by a dead transaction, the user transaction instantly rolls back the change to the data made by the dead transaction and continues.

In contrast, DB2 only allows *member crash recovery* in a pureScale configuration to be online as far as other members go, while group crash recovery or regular crash recovery in a non-pureScale environment requires the database to be offline:

> *Group crash recovery is also usually initiated automatically (as a part of a group restart) and the database is inaccessible while it is in progress, as with DB2 crash recovery operations outside of a DB2 pureScale environment [13].*

Fault Tolerance with Real World Clustering

The cornerstone of Oracle's high availability solutions that protects from system failures is Oracle Real Application Clusters (RAC). Oracle RAC is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-nothing and shared-disk approaches, to provide a highly scalable and available database solution for all business applications.

In contrast to Oracle's unified clustering solution, DB2's clustering solutions are somewhat disparate and non-integrated. DB2 offers two distinct kinds of cluster-like solutions:

- DB2 Database Partitioning Feature (DPF) aimed at Data Warehouse apps
- DB2 pureScale aimed at OLTP apps, and available with DB2 version 9.8

DB2 Database Partitioning Feature (DPF)

DB2 provides a shared-nothing clustered database available through Database Partitioning [14], [15]. Each node in such a DB2 cluster houses one or more database partitions (or database partition groups).

> ***Note***
>
> *DB2 database partitioning used to be formally available through a capability called Database Partitioning Feature (DPF). With the advent of pureScale, IBM has stopped making much noise on DPF, and currently it is believed to be sold as part of IBM InfoSphere Warehouse [16].*

In the event of an unexpected node failure, both RAC and DB2/Database Partitioning can transparently recover the database. However, recovery times are faster with RAC due to the capabilities described above, plus the following:

- DB2 relies on the database manager to restart the partition on a surviving node. This requires the DB2 processes to be started, shared memory to be initialized, and database files to be opened.
- It is not uncommon for large systems today to have large buffer cache. Warming up a buffer cache this large can take a long time. With Oracle Real Application Clusters, failover occurs to an instance that has a warm cache. With DB2/Database Partitioning, failover always involves starting a new instance from scratch with a cold cache. That is why, after the database has been recovered, applications can be expected to obtain their original response times faster in RAC, because the data and the packages needed by the application may have already been cached in the surviving nodes.

Besides the way failures are treated, DB2 Database Partitioning is not performance-optimal for OLTP apps. This is because of the way that DB2 maintains distribution keys, which is a column (or group of columns) that is used to determine the database partition in which a particular row of data is stored [17]. A distribution key is defined on a table using the `CREATE TABLE` statement (default is the first column of the primary key, or – in the absence of a primary key, the first non-long field column defined on that table). Once distribution keys are defined, they are used to determine the location of each row of a table (for queries or updates) in the following way:

- A hashing algorithm is applied to the value of the *distribution key*, and generates a number between 0 and 32767.
- When a database partition group is created, a corresponding *distribution map* (that it is associated with) is also created. Distribution map is an internally generated array containing up to 32768 entries for multiple-partition database partition groups, and used by the database manager to find the data that it needs.
- Each of the partition numbers in the partition group (e.g. 0, 1, 2, 3 for a 4-member partition group) is sequentially repeated in a round-robin fashion to fill the distribution map, e.g. {0, 1, 2, 3, 0, 1, 2, 3, 0, 1, 2, 3, …}.
- When the database manager wants to access a table, the distribution key for the corresponding row is hashed to a number between 0 and 32767, which in turn is used as an index into the distribution map to select the database partition where that row is stored.

This obviously impacts query/update performance especially for OLTP applications accessing a DB2/partitioned environment (the reason that the DB2 documentation stresses that "*Choosing a good distribution key is important*" [17]) – e.g. when the data access involves joins across multiple partitions. Furthermore, an inappropriate partitioning key can cause uneven data distribution, leading to uneven application performance. Finally, since every query/update has to be applied with the partitioning hash algorithm, the cost of a query/update increases, with the cost being proportional to the size of the partitioning key.

To alleviate this, IBM recommends using a separate tool called the DB2 Design Advisor [18], which can be used to identify the objects needed to improve the performance of the DB2 workload in the presence of database partitions, although it comes with its own set of limitations and restrictions (e.g. "*The Design Advisor provides advise about database partitioning only for DB2 Enterprise Server Edition.*").

On top of this, IBM recognizes that redistributing data across partitions is necessary when partitions are added / removed or when data is not balanced across partitions [19]. To mitigate this, DB2 offers a `REDISTRIBUTE DATABASE PARTITION GROUP` command-line utility that has to be manually run when necessary. Data redistribution has its own set of restrictions [20] – e.g. no queries / updates allowed on the tables undergoing data redistribution.

Since RAC is based on a shared-disk subsystem that does not require partitioning for scalability and high availability, Oracle does not suffer from any of these design and performance drawbacks.

DB2 pureScale

In autumn 2009, IBM announced a new version of DB2 v9 known as DB2 9.8 or DB2 pureScale. This is IBM's first release on a UNIX platform of a shared disk solution. Considering that Oracle RAC was released as part of Oracle 9*i* Release 1 (June, 2001), this shows that DB2 has been playing catch-up and following Oracle for several years and this is one more validation that they feel Oracle has provided the right direction.

The pureScale capability is aimed only at OLTP systems, i.e. it is not recommended for Data Warehouse apps. The platform support [21] for pureScale is still limited to IBM System z, IBM Power Systems and IBM System x. DB2 pureScale is also available on x86-based System x servers running SUSE Linux Enterprise Server and Red Hat Enterprise Linux. No other OS support such as Solaris, HP-UX or Windows, is available.

The main technology that DB2 pureScale brings to the table [22] is centralized lock management and group buffer pool. These components reside on their own server or paired servers for redundancy. This is a high systems overhead, e.g. a 4 node cluster would require a 50% uplift in hardware costs alone for the pureScale components. The size of the global buffer pool is limited by the size of memory on a single server.

pureScale still lacks integration with many other DB2 features, as the rest of this document will demonstrate.

Oracle Real Application Clusters (RAC)

In contrast to the two completely different clustering architectures for DB2, Oracle provides a unified shared-disk clustering solution through Oracle Real Application Clusters (RAC). RAC supports the transparent deployment of a single database across a cluster of active servers, providing fault tolerance from hardware failures or planned outages. RAC supports mainstream business applications of all kinds – these include packaged products such as Oracle E-Business Suite, PeopleSoft, Siebel, SAP, as well as custom applications. RAC provides very high availability for these applications by removing the single point of failure with a single server. In a RAC configuration, all nodes are active and serve production workload. If a node in the cluster fails, the Oracle Database continues running on the remaining nodes. Individual nodes can also be shut down for maintenance while application users continue to work. RAC is integrated with mid-tier clients such as Oracle JDBC, Oracle Data Provider for .NET and Oracle Call Interface (OCI) to enable automatic and coordinated connection-pool and application failover to surviving nodes, in the event of individual node failures.

A RAC configuration can be built from standardized, commodity-priced processing, storage, and network components. RAC also enables a flexible way to scale applications, using a simple scale-out model. Using sophisticated load-balancing algorithms (e.g. runtime connection pool load-balancing integrated with server-side load-balancing advisories), user sessions can be routed to the least loaded node in the cluster. On top of that, RAC supports mixed workload environments, enabling the same database to be shared by various kinds of OLTP and Data Warehousing applications. Using a "Services" concept, RAC provides a simple solution to the challenges of managing different application workloads on the same database – DBAs have the power to control which processing resources are allocated to which Service during both normal operations and in response to failures. Resource allocations to Services can be made easily and dynamically enabling flexible enterprise grid environments.

Oracle Automatic Storage Management (ASM) and Oracle Clusterware complement RAC, providing an integrated storage management and cluster software solution for enterprise grids. Unlike pureScale, RAC does not have any OS or hardware restrictions, which has led to RAC being deployed at more than ten thousand customer sites around the world.

Oracle RAC extends protection to applications with Application Continuity, a new feature in Oracle Database 12c that protects application from instance and session failures by re-playing affected "in-flight" transactions on another database instance in the cluster. It is complex for application development to mask outages of the database session and as a result errors and timeouts are often exposed to the end-users leading to user frustration, lost productivity, lost opportunities and complex software coding. Application Continuity performs its recovery beneath the application so that the outage appears to the application as a slightly delayed execution.

DB2 pureScale doesn't have similar functionality regarding application continuity. If a member fails, DB2 pureScale can provide crash recovery to roll-back uncommitted redo so hopefully the application can restart the in-flight transaction(s). For further details on the technical merits of Oracle RAC compared to the clustering solutions from IBM, refer to [23].

Addressing Data Failures

It is vital to design a solution to protect against, and recover from, data and media failure. A system or network fault may prevent users from accessing data, but media failures without proper backups can lead to lost data that cannot be recovered.

As shown in the following table, Oracle offers a wide-ranging set of capabilities to address data failures, which differentiates itself from DB2.

TABLE 3: ADDRESSING DATA FAILURES – ORACLE VS. DB2

ADDRESSING DATA FAILURES	ORACLE	DB2
Built-in database failure detection, analysis, and repair	Yes	No
Built-in incrementally updated backup strategy	Yes	No
Unused block compression	Yes	No
Flexible backup compression levels	Yes	No
Automatic restore failover to next available backup during recovery	Yes	No
Restore preview	Yes	No
Trial recovery	Yes	No
Integrated backup encryption	Yes	No
Database cloning over the network	Yes	No
Block media recovery	Yes	No
Resumable backup	Yes	No
Incremental backups of LOBs	Yes	No
Cross-platform backup and restore	Yes	No
Integrated fine-grained table recovery	Yes	No
Fine-grained backup & recovery performance tuning capabilities	Yes	No
Integrated Mirroring	Yes	No
Proactive disk health checks with automatic corruption repair	Yes	No

Much of the brains behind Oracle's data failure protection lie with Oracle Recovery Manager (RMAN), which is a database-integrated technology to facilitate efficient backup and recovery of Oracle databases. RMAN optimizes performance and space consumption during backup with file multiplexing and compression. It also ensures that backups are corruption-free by validating production data block integrity during backup operations and validating integrity of backups when they are restored. RMAN leverages the Fast Recovery Area and integrates with Oracle Secure Backup [24] and other tape backup management products for a centralized disk-to-disk-to-tape backup strategy.

The following sections provide further details on the Oracle differentiators mentioned in the above table.

Built-in Database Failure Detection, Analysis, and Repair

When faced with data failures, a DBA first invests time to diagnose the issues and plan an appropriate recovery strategy. Depending on the nature of the failure, this investigation and planning time can often comprise a large percentage of the total recovery time. The Data Recovery Advisor (DRA) dramatically reduces this time by automatically detecting failures in real-time (e.g. block corruptions, missing files), reporting failure analysis results, and generating a feasible recovery strategy (e.g. RMAN recovery script) that can be run as-is or customized for running at a later time. In addition, Data Integrity Checks allow proactive monitoring of database integrity, thereby catching and repairing data issues before users even come across them.

In large environments where DBAs manage hundreds of databases and thousands of data files, the DRA can dramatically simplify recovery diagnosis and management tasks. By self-diagnosing Oracle failures, the DRA also reduces the chances that a user could develop an improper recovery strategy or commit errors, while under the pressure of recovering a critical production system.

DB2 offers partial comparable functionality with the Recovery Expert [25], a separately licensed product. Recovery Expert can propose optimal procedures for a specific recovery action, but does not detect failures in real-time, nor link failure diagnosis with the proper recovery strategy.

Built-in Incrementally Updated Backup Strategy

With integrated fast incrementally updated backups, RMAN rolls forward an image copy by applying incremental backups. The image copy is updated with block changes up through the SCN at which the latest incremental backup was taken. Incrementally updated backups eliminate the need and overhead of performing a full database backup every day.

IBM recommends doing this with DB2 Merge Backup, which offers similar functionality as the RMAN Incrementally Updated Backup Strategy. However, DB Merge Backup is a separately installed and separately licensed product and only offers a command-line utility [26].

Unused Block Compression

During full backups, RMAN can reduce backup sizes considerably by eliminating blocks that are not currently being used. For example, if a 1 TB table is dropped and purged, the next full backup would not backup those 1 TB worth of blocks. DB2 does not offer this block elimination capability. This is a very smart way to accomplish database-integrated source-side deduplication.

Flexible Backup Compression Levels

RMAN has offered efficient native compression of backup data since Oracle Database 10g. With Oracle Database 11g Release 2, RMAN offers more backup compression levels to flexibly suit a particular environment's compression ratio and backup performance needs. Users can choose from the new HIGH, MEDIUM, and LOW levels, based on the desired degree of compression. DB2 offers only one backup compression setting (i.e. library) by default and if additional settings are desired, the user must acquire third party compression libraries [27].

Automatic Restore Failover during Recovery

During a restore, when RMAN finds corruption in a backup, or finds that a backup cannot be accessed, RMAN tries to restore the file from all known backups before returning an error. This is done automatically whenever RMAN restores file(s) with the RESTORE or RECOVER command, eliminating the need to search for valid backups and re-trying the operation when a restore failure occurs. DB2 lacks this capability.

Restore Preview

Before a database restore operation, a DBA can request to view the list of backup files needed to complete the operation. The RMAN restore preview capability ensures that all required backups are available or to identify situations in which the DBA may want to direct RMAN to use or avoid specific backups. DB2 does not offer this capability.

Trial Recovery

A test recovery can be very useful to first ensure that all required archived logs are present, corruption-free, and can be successfully applied to the restored data files, without having to perform actual media recovery. Oracle trial recovery provides just that and does not modify the restored data files. DB2 does not provide this capability.

Integrated Backup Encryption

Protecting backups of highly sensitive and confidential information is vital to the stability of many companies. Backups should only be able to be opened and read by their creators. RMAN provides the ability to encrypt backups as they are created, in a database-integrated manner, with 128, 256, or 512-bit versions of the Advanced Encryption Standard (AES). DB2 does not provide native backup encryption.

For encryption, IBM recommends using DB2 Database Encryption Expert – a standalone solution, which enables encryption of offline database backups as well as encrypts online database files [28]. This product is separate from the core DB2 Database, does not work / is not supported with DB2 pureScale, and this lack of database-level integration is perhaps not surprising considering that this product is based on a partner technology [29].

Database Cloning over the Network

A common DBA task is to clone production databases for test, QA, reporting, and disaster recovery purposes. A backup can be used to restore and create the clone database, but this requires the backup to be copied or made accessible to the clone server. Starting with Oracle Database 11g, RMAN offers a method (Active Duplicate) to clone an online production database over the network directly to a clone server, without requiring a backup. In Oracle Database 12c, the workload is moved to the destination server via auxiliary channels, relieving resource bottlenecks on the source database server. This capability can use RMAN compression and multi-section capabilities to further increase performance. Unused block compression happens automatically.

DB2 does not provide these comprehensive cloning capabilities.

Block Media Recovery

With Oracle's block-level media recovery feature, if only a single block is damaged then only that block is recovered. The rest of the file, and thus the table containing the block, remain online and accessible, increasing data availability. DB2 cannot recover data in single-block units, thus requiring the entire file to be taken offline, restored, and recovered.

Furthermore, in Oracle Database 11g Release 2, when an Active Data Guard standby database has been configured, a corruption that occurs on the primary database will be automatically detected and transparently repaired in-line with a good block from the standby database. A user or application query accessing the corrupt block will only experience a short pause while the block is repaired and then the query results are returned. DB2 does not have online, automatic block repair capabilities.

Resumable Backups

Another time saving feature Oracle provides through RMAN is resumable backup operations. With Oracle, if these operations fail, they can be restarted from the point of failure. Because DB2 has no such capability, problems during backup means time lost while the entire operation starts from the beginning. To further compound the problem, in DB2 "*a table space backup operation and a table space restore operation cannot be run at the same time, even if different table spaces are involved*." [30]

Incremental Backups of Large Objects

Large objects (LOBs) often store images, sound files, etc., that never change. Incremental backup is critical for these. While Oracle can perform incremental backups of LOBs, DB2 is unable to do so:

> *"If a table space contains long field or large object data and an incremental backup is taken, all of the long field or large object data will be copied into the backup image if any of the pages in that table space have been modified since the previous backup."* [31]

Cross-platform Backup and Restore

New RMAN Cross-platform functionality in Oracle Database 12c enables backup and restore across different platforms, for the most efficient tablespace and database migration which can improve application availability. In earlier releases, moving a database across platforms required either import/export or cross-platform transportable tablespaces procedures, thereby, affecting application availability. On the source platform, BACKUP creates backup sets of user tablespaces, including Data Pump metadata dump file, in read-only mode. RESTORE on the destination platform automatically performs data file Endian conversion and plugs-in tablespaces. To minimize read-only impact, Oracle recommends taking incremental backups, that are then converted and applied to restored data files. Only the final incremental must be taken while tablespaces are in read-only mode, with separate Data Pump metadata export and import.

For Linux/Unix, DB2 can provide cross-platform backup and restore capability on when the source and destination platforms are of the same Endian [32].

Integrated Fine-grained Table Recovery

In Oracle Database 12c, RMAN can recover individual database tables from backup, via a simple `RECOVER TABLE` command. This recovers one or more tables (the most recent or an older version) from an RMAN backup. Tables can be recovered in-place or to a different tablespace. Optionally, RMAN can create a Data Pump dump file of the table(s). This efficient functionality replaces error-prone manual processes and reduces recovery time. It extends the range of recovery in areas where Flashback is not applicable, for example when a dropped table has been purged out of the Recycle Bin, or when the desired point to recover is outside the window given by the `UNDO_RETENTION` parameter.

DB2 can restore only at the database or table space level [33]. The separately licensed DB2 Recovery Expert [25] does offer capability to recover at the table level.

Fine-grained Backup & Recovery Performance Tuning Capabilities

Recovery Manager provides optimal performance tuning capabilities to meet the broad spectrum of database size, datafile make-up and RTO. The number of RMAN channels, data streams, determines the degree of parallelism for backups and restoration. Channels may be automatically or manually allocated with each channel corresponding to one server session. Backup performance may be further optimized by levering RMAN multiplexing which defines the number of files read concurrently by each channel; default is 8. Taking into account that some files may be very large, RMAN provides the capability to further parallelize the backup by spreading the processing of large files across multiple channels referred to as a multi-section backup.

DB2 does not offer this level of performance tuning [34]. DB2 parallelism defines how many table spaces may be read concurrently which is in sharp contrast to RMAN's capability to parallelize a single datafile.

ASM – Integrated Data Mirroring

Oracle Automatic Storage Management (ASM), which is an integrated volume management and file system available with the Oracle Database, provides a native mirroring mechanism based on the concept of *disk failure groups*, which can be used to protect against storage failures. An ASM failure group is a set of disks sharing a common resource (disk controller or an entire disk array) whose failure can be tolerated. With ASM mirroring, when database extents are allocated, a primary copy and a secondary copy are created, with the disk for the secondary copy chosen to be in a different failure group than the primary copy. This ensures that the data is available and transparently protected against the failure of any component in the storage subsystem. Not only that, if there are read errors associated with reading a corrupted block, ASM will transparently read a good copy of the extent, and copy it to the disk that had the read error.

DB2 does not provide any such integrated mirroring mechanism for additional data protection.

Addressing Disaster Recovery

Oracle Data Guard

Oracle Data Guard provides the management, monitoring, and automation software infrastructure to create and maintain one or more standby databases to protect Oracle data from failures, disasters, errors, and data corruptions. In the event of a planned or unplanned outage at the primary site, Data Guard ensures that a standby database can be easily switched to the production database role without data loss, client connections are automatically redirected, and the new production database commences with serving enterprise data needs. The advanced capabilities of Data Guard in Oracle Database 11g – Oracle Active Data Guard and Data Guard Snapshot Standby, and in Oracle Database 12c – Data Guard Far Sync, deliver the highest levels of availability, data protection, operational transparency, and return on investment (ROI) in standby software, servers, and storage.

IBM HADR

HADR (High Availability Disaster Recovery) was first released in DB2 version 8.2 and is based on a similar feature, called High Availability Data Replication (HDR for short) from IBM's Informix Dynamic Server acquisition [35].

DB2 HADR replicates data changes from a source database (primary), to a target database (standby). In the event of a partial or complete site failure, the standby database can take over for the primary database.

So at a high level, Oracle Data Guard and IBM HADR appear similar. However, the devil is in the details.

Data Guard: Comparative Strengths

A closer look at the two technologies uncovers a number of areas where Data Guard functionality is superior to DB2 HADR. Let’s begin by reviewing key requirements for enterprise data protection and availability to provide context for a detailed comparison.

Data protection and data availability

The first requirement of enterprise data protection is to guarantee that the standby database is a reliable and complete replica of the primary database. This means that should a primary and standby database ever be compared at the same relative point in time – e.g. a point in time when they have each applied up through the same insert, update, delete or DDL transaction – a comparison would always confirm that the databases are exact replicas of each other. The infrastructure of the replication mechanism must guarantee zero risk of inconsistency between primary and standby databases – to the degree that there is never a need to perform a comparison to validate this fact. This characteristic enables the standby database to be a surrogate for the primary in every respect, to the point that a backup of either database can be used to restore the other to its original state. This enables customers to meet any regulatory statute or business requirement for data protection by guaranteeing that data can never diverge.

Both Data Guard Redo Apply (physical standby) and IBM HADR score high marks on this requirement. Each uses native database recovery mechanisms to insure that the primary and standby databases are exact physical replicas of each other. No data divergence between the two databases as of the same relative point in time of processing is possible. No regularly scheduled "compare" of the primary and standby database is required.

Data Guard users, however, benefit from advanced capabilities that provide superior data protection and data availability compared to IBM HADR. These advantages are discussed in later sections.

Isolation

The second requirement is made more challenging by virtue of the first. While the standby database must be an exact replica of the primary, the two databases must also be loosely coupled in order to isolate the standby from any fault that might impact the primary database.

Data Guard and IBM HADR score high marks in this regard. Both solutions use database-aware processes to independently validate all change data before it is applied to the standby database. Standby databases also use a different code path than the primary database when applying changes to the standby database. This prevents software and hardware faults that can lead to primary database corruption or failure from impacting the standby database. This approach is far superior to the alternative of storage array remote-mirroring that is limited to replicating physical blocks.

As is the case with data protection, however, Data Guard users benefit from advanced capabilities that provide superior isolation between primary and standby databases than is possible using IBM HADR. These advantages are also discussed in later sections.

Do no harm

The third requirement of data protection solutions is identical to the oath taken by medical doctors – do no harm. This means primary database performance and availability must not be impacted when the data protection solution encounters challenges typical of real life environments, including:

- Wide area networks that can deliver unpredictable levels of service.
- Sudden peaks in workload than can exceed network capacity.
- Standby servers, storage subsystems, and networks that fail.
- Human errors that can interrupt service and/or corrupt data.

Data Guard provides numerous advantages over IBM HADR that protect the primary database from events that can impact remote data transmission or the progress of the apply process at the standby database. Data Guard's integration with other Oracle HA features provide additional levels of HA protection by enabling fast recovery from human error or logical corruptions not possible with IBM HADR. Many of these advantages are attributes inherent to Data Guard architecture; no management intervention or implementation-specific configuration is required. Details are discussed in the sections below.

Return on investment – utilization of standby servers, storage, and software, while in standby role

The fourth requirement reflects the reality of today's economic environment. As the cost of downtime has continued to climb so has the need to provide comprehensive data protection and high availability WHILE maximizing return on investment (ROI) in standby servers, storage and software.

Oracle Active Data Guard was introduced with Oracle Database 11g Release 1. It provides customers the ability to have one or more physical standby databases in a Data Guard configuration open for read-only access while they continue to apply changes received from the primary database. Active Data Guard transparently provides all of the same capabilities as any Oracle Database that is open read-only - there are no special restrictions or additional operational complexity associated with an Active Data Guard standby database. Active Data Guard in Oracle Database 12c has added several advanced quality-of-service features that further enhance the value of an Active Data Guard standby database.

IBM HADR introduced a new Active Standby capability with Fix Pack 1 of DB2 9.7. Unlike Active Data Guard, IBM HADR capability in this area has a number of substantial limitations. These are discussed in subsequent sections.

Summary Comparison

The following table summarizes Data Guard strengths compared to HADR:

TABLE 4: ADDRESSING DISASTER RECOVERY – ORACLE VS. DB2

ADDRESSING DISASTER RECOVERY – DATA PROTECTION AND AVAILABILITY	ORACLE	DB2
Standby apply progress has no impact on primary database performance or data protection	Yes	No
Silent corruptions due to hardware/software faults are detected at both primary and standby	Yes	No
Corrupt blocks are automatically repaired online, transparent to users and applications	Yes	No
Fast recovery from human error and logical corruptions	Yes	No

Integrated automatic database failover with guarantees of zero data loss and no split-brain	Yes	No
Controlled automatic failover for ASYNC to comply with user configurable data loss SLA’s	Yes	No
Connect-time failover of applications in all cases – controlled by database role	Yes	No
Fast reinstatement of primary after failover without needing a full restore	Yes	No
Flexible support of multiple standbys	Yes	No
Real-time cascaded standbys	Yes	No
Database rolling upgrades across major versions and subversions	Yes	No
Support for some mixed primary/standby configurations	Yes	No
Integrated redo transport compression for efficient network utilization	Yes	No
Fast, non-disruptive recovery from network and standby database outages	Yes	No
Replicate stored procedures to standby database	Yes	No
Support database partitioning	Yes	No
Complete support of active-active database clusters	Yes	No
Long-distance, zero-data loss standby	Yes	No
Proactive readiness health checks on standbys	Yes	No

ADDRESSING DISASTER RECOVERY ROI - STANDBY DATABASE UTILIZATION	ORACLE	DB2
Full read consistency for queries on active standby	Yes	No
No datatype limitations for active standby	Yes	No
Continuous user access to an active standby database during apply of DDL changes	Yes	No
Continuous user access to an active standby database during resynchronization	Yes	No
Automatic monitoring and enforcement of query SLAs	Yes	No
Enhanced reporting and global temporary table support on active standby	Yes	No
Complete support between active standby and active-active database clusters	Yes	No
Automatic memory management on active standby	Yes	No

Backup and recovery operations on active standby	Yes	No
Dual-purpose standby database for Dev/Test	Yes	No
Integrated Load Balancing and Workload Management across replicated configurations	Yes	No

The following sections expand upon Data Guard's comparative strengths.

Disaster Recovery – Data Protection and Availability:

Data Guard architecture and capabilities provide superior data protection and data availability without impacting the primary database. This includes:

Standby apply has zero impact on primary performance or data protection

Fault isolation is a key Data Guard design criteria. Data Guard standby apply services function completely asynchronous to the primary database in all protection levels and transport modes (SYNC and ASYNC). If circumstances cause standby apply to fall behind or stop, there is zero impact on primary database performance, availability, or data protection. The Data Guard primary continues to ship data to the standby. The standby continues to protect transactions by archiving data to log files local to the standby database while the apply process catches up or the fault is repaired.

DB2 HADR does not implement this level of isolation or resiliency. If standby apply falls behind, the buffer used by an HADR standby to receive primary redo data becomes full. This blocks the standby from receiving any more data (increasing the risk of data loss), and immediately blocks any new primary transactions in HADR sync or near-sync modes [36]:

> *"By default, the log receive buffer size on a standby database is two times the value that you specify for the logbufsz configuration parameter on the primary database. This size might not be sufficient. For example, consider what might happen when the HADR synchronization mode is set to ASYNC and the primary and standby databases are in peer state. If the primary database is also experiencing a high transaction load, the log receive buffer on the standby database might fill to capacity, and the log shipping operation from the primary database might stall."*

HADR standbys blocking the primary database even in ASYNC mode is really a terrible design. To alleviate this, IBM suggests tuning two parameters (`DB2_HADR_BUF_SIZE` and `hadr_spool_limit`), making things even more complicated.

Detecting silent corruptions caused by lost writes

Data Guard provides industry-unique protection against lost writes. A lost write occurs when an I/O subsystem acknowledges to Oracle that a write is complete, while in fact the write did not occur in the persistent storage. On a subsequent block read, the I/O subsystem returns the stale version of the data block, which can be used to update other blocks of the database, thereby corrupting it. Oracle provides a `DB_LOST_WRITE_PROTECT` initialization parameter which when set, will record buffer cache block reads in the redo log and use this information to detect lost writes. When a Data Guard standby database applies this redo during managed recovery, it reads the corresponding blocks and compares the SCNs with the SCNs in the redo log to determine if there has been a lost write. The procedure recommended for repairing lost writes on a primary database is to failover to the physical standby and recreate the primary. If the lost write has occurred on the standby, simply recreate the standby database or the affected files.

DB2 HADR does not have the ability to detect lost writes and avoid data loss and production downtime that results from such an event.

Automatic repair of corrupt blocks at either primary or standby database

Active Data Guard 11g Release 2 enables the automatic repair of corrupt blocks. Block-level data loss usually results from intermittent, random I/O errors, as well as memory corruptions that are written to disk. When Oracle discovers a corruption it marks the block as media corrupt, writes it to disk, and typically returns an `ORA-1578` error to the application. No subsequent read of the block will be successful until the block is recovered manually. However, if the corruption occurs on a primary database that has an Active Data Guard standby, block media recovery is performed automatically, transparent to the application, using a good copy of the block from the standby database. Conversely, bad blocks on the standby database are automatically recovered using the good version from the primary database.

DB2 has no such advanced data protection capabilities.

Built-in mechanisms to undo corruptions related to human errors

Human error is one of the leading causes of downtime. Such errors lead to logical corruptions that can be widespread and result in significant downtime for point-in-time recovery operations. Data Guard can quickly recover from such corruptions using delayed apply – which delays the processing of redo on the standby by a configurable period of time. This provides administrators the opportunity to failover to the standby before the error is applied to the standby database, or to export still valid data from the standby and use it to surgically repair the primary database. A typical example of such an error is an incorrect batch job run on the primary database.

A compromise that accompanies using a delay is that the standby database must first apply the backlog of log data before it can assume the primary database role – increasing downtime should a failover be required. Oracle Flashback Database offers Data Guard users a second method to achieve the same level of protection by enabling the primary and standby databases to be quickly recovered to a previous point in time. Flashback Database eliminates any delay at failover time because the standby database is always up to date. Flashback Database is the recommended approach when also using an Active Data Guard standby for up-to-date queries and reporting.

DB2 HADR supports the capability to delay apply at the standby although it falls short of the integrated solution offering of Oracle Data Guard and Oracle Flashback Database.

Integrated automatic database failover with zero data loss

Data Guard Fast-Start Failover automatically detects primary database outages and executes failover to a previously chosen, synchronized standby database – no manual intervention or external integration with clusterware software is required. Once the original primary has been repaired, mounted, and is able to establish a connection with the new primary, Data Guard automatically converts the failed primary to a standby database and resynchronizes it without requiring a time-consuming restore from backup (both SYNC and ASYNC configurations).

Fast-Start Failover also guarantees that automatic failover can never result in "split brain" (a condition when 2 or more databases in a primary/standby configuration act as primaries simultaneously). It also supports advanced capabilities such as doing such automatic failovers even in asynchronous transport conditions (as long as RPO SLAs are not violated), or upon designated health check violations.

DB2 HADR does not automatically detect a failed primary and issue a takeover operation. This process is manual. The administrator must determine the primary database is unavailable, and then must issue the takeover command.

IBM recommends using a cluster manager as a workaround for implementing automatic failover for an HADR configuration. This requires a separate integration and scripting, and becomes even more complex for typical disaster recovery architectures where WAN deployment requires a geo-cluster implementation. Given that automatic failover is not an HADR feature, care must be taken to make sure that external integration and scripting prevents the occurrence of split-brain condition for every possible failure scenario.

Connect-time failover of applications in all cases

Data Guard combined with Oracle Transparent Application Failover (TAF) enables client-side connect time failover using an alternate connect descriptor. Oracle Data Guard 11g Release 2 implements role-specific database services – ensuring that services are automatically started and are appropriate for the then-current role of the database (primary, standby, or snapshot standby). Automatic client failover in a Data Guard configuration is not dependent upon the client having already connected to a primary database in order to determine the failover partner. This is important in situations where the primary database is unavailable to new client connections or in configurations having multiple standby databases, any one of which is able to assume the role of primary database at failover time.

DB2 HADR automatic client reroute (ACR) does not support connect-time failover. ACR functions by registering the standby database with the primary database with an "`update alternate server for database...`" command. After this command is issued, the client must successfully connect to the primary database to obtain the alternate server information. If the client is never able to connect to the primary database, because it is down to begin with, it will not be automatically rerouted [37].

Furthermore, ACR does not implement the concept of database role; this can lead to applications being directed to the wrong database [38].

Easy reinstatement of primary to a new standby database after a failover

Data Guard supports fast reinstatement of a failed primary as a new standby database without requiring a full restore from backup – for both SYNC and ASYNC failovers. This quickly returns the configuration to a protected state, practically eliminating any extra effort on the part of DBAs.

DB2 HADR requires a full restore from backup any time that the primary and standby databases are not completely synchronized at failover time [39]:

> *"If the primary database fails when the primary and standby databases are in peer state, it is possible that the original primary database cannot to rejoin the HADR pair as a standby database without being reinitialized using a full restore operation."*

This can be a daunting task given the common occurrence of multi-terabyte databases and WAN topologies where high latency or bandwidth limitations make a remote restore even more time consuming and expensive. Several undesirable outcomes result from this limitation:

- There is a tendency to failover to the standby database only as a last resort after significant downtime has elapsed.
- There is a high network cost – disrupting other applications that share the same network bandwidth, while a backup of the new primary is transmitted to the original primary site.
- Time, effort, frustration, and potential for human error during the restore process, often result in DBA's expending even more time, effort, and frustration.
- There is an extended period of time when the new primary database is in an unprotected state and thus is vulnerable to data loss and downtime should there be a second failure event.

Flexible support for multiple standbys

Data Guard supports multiple standby databases with very flexible configuration options within a single Data Guard configuration. For example, a primary database may have a local SYNC standby database and a second remote ASYNC standby (up to 30 directly connected standby databases are supported). Failover to any physical standby database results in the remaining standby database automatically recognizing the new primary database, thus providing continuous data protection throughout the failure event.

In addition to increased data protection and availability, this also makes it very easy to deploy multiple Data Guard standby databases to serve other uses e.g. offload ad-hoc queries, reports, scale read performance, backups, perform QA testing, execute database rolling upgrades and migrations.

HADR in DB2 v10.5 has started supporting multiple standbys too, but this support comes with a set of restrictions [40] which yet again prove how far HADR is lagging compared to Data Guard:

- HADR supports only up to three standby databases for each primary – one principal standby and two auxiliary standbys.
- DB2 pureScale configuration does not support multiple standbys.
- Only the principal standby supports all HADR synchronization modes; all auxiliary standbys have to be in SUPERASYNC mode.
- If the primary server fails, any one standby can take over, but if a standby received more logs from the original primary compared to the new primary, then this standby has to be completely reinitialized. This differs from the Data Guard case when one can use a simple Flashback on this standby and resume operations.

Real-time cascaded standby(s)

To reduce the load on the primary system, Data Guard supports cascading standbys. A cascading standby database is a standby database that receives it redo logs from another standby database, not from the original primary system. The standby database would send redo to the cascaded standby upon a log switch. Oracle Database 12c has taken cascaded standby databases to the next level providing real-time cascading capabilities in that redo logs are immediately forwarded from the standby to the cascaded standby instead of waiting for a log switch to occur.

DB2 10.1 introduced the capability to have multiple standbys although cascaded standby databases are still not supported. The DB2 principal and auxiliary standbys receive redo directly from the primary database.

Database rolling upgrades across major versions and subversions

Data Guard supports rolling database upgrades across major versions and subversions. A Data Guard standby database never needs to be recreated to execute an upgrade.

DB2 HADR does not support rolling upgrades between DB2 version or major subversion migrations, such as version 8.2 to 9.1, or 9.1 to 9.8 or from 10.1 to 10.5 [41]. Rolling upgrades are only supported from one IBM Fix Pack or Modification Level, to the next.

Support for mixed primary/standby configurations to reduce planned downtime

Data Guard has the flexibility to support a number of configurations where primary and standby systems may have different CPU architectures, operating systems (e.g. Windows and Linux), operating system binaries (32-bit/64-bit), and Oracle database binaries (32-bit/64-bit). This offers a method of reducing planned downtime and risk when executing certain platform migrations and technology refresh.

Data Guard can also be used to migrate to Automatic Storage Management (ASM), from single instance Oracle Databases to Oracle RAC, or to Oracle Exadata storage.

Finally, a Data Guard standby database can be configured with fewer resources (e.g. CPU, memory, I/O) than its primary database. This provides customers flexible deployment options. For example, one deployment model is to consolidate the hosting of multiple standby databases on a single server or Oracle RAC. This can reduce the investment required in standby systems if such a model meets business objectives.

DB2 HADR does not have such flexibility – there are no differences allowed between primary and standby in steady-state situations [42]. In many cases the reasons for this are architectural, in other cases it is due to concern that an under-resourced standby database can impact primary database performance and data protection. Data Guard does not have these limitations.

Integrated network transport compression

Data Guard 11g Release 2 with the Oracle Advanced Compression Option enables automatic compression of all redo transmission between primary and standby databases. Transport compression enables more efficient utilization of available network bandwidth. This enables customers with limited bandwidth to achieve their recovery point objectives even if their redo volume exceeds available bandwidth. It also results in faster resynchronization of primary and standby databases following any outage that interrupts network transmission.

DB2 HADR does not provide an integrated capability for network compression

Fast, non-disruptive recovery from replication and standby database outages

In addition to transport compression, Data Guard is well architected for high-performance log gap resolution to quickly recover from any outage that impacts replication, such as network or standby database outages. Multiple parallel Data Guard background processes automatically transmit large volumes of archive log data that may be required to resynchronize a standby database following an extended outage. While this occurs in the background, Data Guard also transmits current log file data thereby preventing transmission from falling any further behind. High volume, parallel shipment is possible due to the high degree of isolation between primary and standby systems in a Data Guard configuration. If there is a large backlog of data to be transmitted, a Data Guard standby database is able to receive data faster than it can be processed without any negative impact on network transport, primary database performance or availability. The faster the data can get to the standby, the faster the data is protected and the sooner the database returns to a protected state.

IBM HADR does not have an equivalent capability.

Support database partitioning and replication of stored procedures

Data Guard Redo Apply transparently supports database partitioning and replicates all stored procedures. It has no datatype restrictions.

DB2 HADR does not support DB2's partitioning feature and does not replicate stored procedures – they must be manually recreated at the HADR standby [43].

Complete support for active-active database clusters

Data Guard is completely integrated with Oracle RAC. Any of the primary or the standby databases can be multi-node Oracle RAC instance. All Data Guard protection modes, transport modes, and apply modes are supported. Automated transmission of redo data and recovery are available for all configurations. Any physical standby database, whether Oracle RAC or single node, can also be an Active Data Guard Standby database and support read-only queries while it applies updates received from its primary database. Data Guard and Oracle RAC support all platforms that are supported by the Oracle Database.

DB2 HADR did not support DB2 pureScale till DB2 v10.1. Starting with DB2 v10.5, IBM has announced HADR support for pureScale, but even that has several restrictions [44]:

- The HADR SYNC or NEARSYNC synchronization mode cannot be used in a DB2 pureScale environment. This means zero data loss is not supported in a pureScale environment.
- In a pureScale configuration, one cannot have more than one HADR standby database. In other words, pureScale does not support multiple standbys – it can have only one principal standby.
- The reads on standby feature is not supported for a pureScale environment [45]. For a pureScale standby database, only one member (the replay member) is activated, which replays all of the logs. Other members are not activated. This implies that the combination of pureScale and HADR leads to a waste of system resources.
- The topology of the primary and the standby is very restrictive in a pureScale + HADR configuration. It must be synchronized – which means that the primary and standby have to have identical configurations. If one adds a member on the primary, that operation is replayed on the standby. However, if one drops a member on the primary, the standby must be reinitialized by using a backup or a split mirror from the primary's new topology. No such restriction exists for Oracle RAC + Data Guard.
- With HADR, there are two types of takeover: role switch and failover. Role switch can be performed only when the primary is available and it switches the role of primary and standby. Failover can be performed when the primary is not available. It is commonly used in primary failure cases to make the standby the new primary. Both types of takeover are supported in a DB2 pureScale environment, and both can be issued from

any of the standby database members and not just the current replay member. However, after the standby completes the transition to the primary role, the database is only started on the member that served as the replay member before the takeover. The database can be started on the other members by issuing an ACTIVATE DATABASE command or implicitly through a client connection. No such manual / downtime restriction exists for Oracle RAC + Data Guard – in this case, Data Guard Broker automates this process and brings all the RAC instances online after a role transition.

Long-distance, Zero-data Loss Standby

Oracle Database 12c introduces Data Guard Far Sync functionality that enables extending the zero-data loss standby protection to any distance – no longer limited by latency. This is achieved by the introduction, between a primary and a standby database in an Active Data Guard configuration, of a lightweight Oracle instance – the Far Sync instance. A Far Sync instance has a standby control file, standby redo logs, and archive redo logs – but no data files. It is deployed at a distance from the primary so that the primary can tolerate the network latency of synchronous transport to the Far Sync, which appears to the primary as a Data Guard destination. Far Sync receives redo synchronously from the primary database, and forwards redo asynchronously in real-time to its final destination, which is a full standby database.

A Far Sync configuration supports a one-step, zero data loss failover, via the same failover/switchover commands used for any Data Guard configuration. At failover, Far Sync transparently ensures that all the redo that had been synchronously transferred to the Far Sync instance is applied at the standby, thus delivering the zero-data-loss guarantee.

DB2 has no such capability. For DB2, the zero data loss feature is constrained by the usual limitations of network distance and latency.

Proactive Readiness Health Checks on Standbys

In Oracle Database 12c, Data Guard Broker includes new Role Change Readiness functionality, which automatically performs comprehensive health tests to assess readiness for a failover or switchover. This leads to a reduced, more predictable Recovery Time Objective (RTO).

DB2 HADR does not have this comprehensive capability. DB2 HADR provides some limited monitoring via the `db2pd` command and/or `MON_GET_HADR` table function [45].

Disaster Recovery ROI – Standby Database Utilization

Data Guard Redo Apply (physical standby) has advanced capabilities that enable customers to easily utilize their standby database servers, storage, and software, for productive purposes while in standby role – hence obtaining an effective return on their DR investment. One of the premier Data Guard standby database utilization capabilities is offered through Active Data Guard – available since Oracle Database 11g Release 1. Active Data Guard enables the Data Guard physical standby database to be used as an extremely high performance real-time-synchronized replica of the production database that is simultaneously available for read-only access. This enables read-only application modules (e.g. reporting applications) to be offloaded to the Active Data Guard physical standby that in turn improves overall application throughput. In addition, Active Data Guard also allows offloading fast incremental backups to the physical standby database, improving the production server performance even further. Furthermore, as discussed earlier, Active Data Guard also serves as a real-time repository of active data blocks that can be used to transparently repair block corruptions on the primary database without bringing down the production application.

In contrast, DB2 HADR only recently – in late 2009, through Fix Pack 1 for DB2 v9.7, has begun to offer a read-capability on their HADR standby. IBM calls this capability the HADR Reads on Standby Feature, and refers to the corresponding standby as Active Standby – but it has several restrictions [46] making it either unusable or unable to provide adequate quality of service for reporting applications. This DB2 feature also lacks the advanced capabilities of Active Data Guard such as supporting online backups or online transparent block corruption protection.

In addition to the capabilities of Active Data Guard, a Data Guard physical standby database can be utilized as a read-write test system while in standby role with zero compromise to data protection – this capability is called Data Guard Snapshot Standby. DB2 has nothing equivalent to this capability.

The following sections provide further details on Data Guard advantages relative to DB2 HADR in terms of standby database utilization.

Full read consistency for queries executing at an active standby database

An Active Data Guard standby database is accessible for read-only queries and reporting while it applies changes received from primary database transactions. Active Data Guard standby databases also implement the same full read consistency model used by a primary database. An Active Data Guard standby database is able to return accurate, up-to-date results, just as if the query had been executed at the primary database. This makes an Active Data Guard standby database a reliable method for offloading workload from the primary database, substantially improving primary database performance and increasing return on investment in standby systems.

IBM DB2 HADR, on the other hand, is only able to support an Uncommitted Read (UR) isolation level, which means it cannot provide read consistency for queries executing on an HADR Active Standby database [47]. Any application, statement, or sub-statement requesting read consistency from an HADR Active Standby database will receive an error. Queries that conform the more limited UR isolation level will receive unreliable results (such queries are able to access uncommitted data and have non-repeatable reads and phantom reads). This makes the HADR Reads on Standby feature absolutely unusable by most reporting applications.

No datatype limitations when querying an Active Data Guard standby database

Data Guard provides an unprecedented level of transparency in operation. Data Guard redo apply (physical standby) fully supports replication and read-only access for all datatypes, there are no restrictions.

In contrast, HADR Active Standby has various datatype restrictions [46]. With IBM HADR, XML and large object (LOB) datatypes must be inline to be successfully queried on the read-only standby. Long field (LF), a distinct type based on one of these data types, or structured type columns, cannot be queried on an HADR Active Standby database. Attempts to query these data types will receive an error. Similarly, access to Global Temporary Tables is also not supported on the HADR Active Standby.

Continuous read-only access to an active standby database during replay of DDL

DDL operations are transparently replicated to an Active Data Guard standby database without any impact to applications that may be querying the active standby while DDL log records or maintenance operations are applied.

In contrast, when an HADR Active Standby database is replaying DDL log records or maintenance operations, the standby enters the replay-only window [48]. When the standby is in the replay-only window, existing connections to the standby are terminated and new connections to the standby are blocked (error SQL1224N). New connections are allowed on the standby after the replay of all active DDL or maintenance operations has completed.

Continuous read-only access to an active standby database during synchronization

Read-only users have continuous access to a Data Guard Active Standby Database during all phases of replication and synchronization of the standby with its primary database – whether the source of the log data is local or remote to the standby database.

HADR does not allow connections to an Active Standby database while it is being synchronized by applying log files that are in its local log file path, a state IBM refers to as *local catchup state* [46]. Active Data Guard does not have this limitation.

Automatic monitoring and enforcement of query SLAs on an active standby database

Beginning with Oracle Database 11g Release 2, Active Data Guard enables configurable service level agreements (SLA) that can be implemented using the session parameter, `STANDBY_MAX_DATA_DELAY`. The value for this parameter specifies a limit for the amount of time (in seconds) allowed to elapse between when changes are committed on the primary and when they can be queried on an active standby database. The Active Data Guard Standby will return an ORA-3172 error code if the limit is exceeded. Applications can respond to this error similar to a disconnect, and redirect the query to another active standby database or to the primary database to achieve the required SLA. This relieves the administrator from monitoring standby apply progress or responding to events (such as interruption in network connectivity between primary and standby database) that can impact the ability of an active standby database to be current for reporting requirements.

DB2 HADR does not have any equivalent functionality for an active standby database.

Enhanced reporting and global temporary table support on read-only standby

Active Data Guard in Oracle Database 12c enables richer reporting functionality on the standby, by enhancing support for DML on global temporary tables, separating temporary table undo and more flexible sequence support which grants a unique range of sequences that can be accessed on each standby.

In contrast, DB2 does not support reads of created or declared temporary tables on the standby [46]. Creating a created temporary table (CGTT) on the primary will trigger a replay-only window on the standby making it unavailable for read-only access. Similarly, one cannot use the `NEXT VALUE` expression on an HADR Active Standby to generate the next value in a sequence.

Fully supported with active-active database clustering

Active Data Guard is fully integrated with Oracle RAC. This means that the Active Data Guard standby can be a multi-node RAC database, with one node designated as the redo apply node, while the other nodes designated as scale-out reporting / query instances, thereby leading to very effective systems utilization across the entire Data Guard configuration.

In contrast, the reads on standby feature is not at all supported in DB2 pureScale environments

Automatic memory management

Automatic memory management is supported by Data Guard standby databases.

IBM's self-tuning memory manager (STMM) is not supported on an HADR Active Standby [46]. Manual effort is required to tune the standby database (either to suit running a read-only workload or to perform well after a failover has occurred).

Ability to utilize standby databases for backups

A Data Guard physical standby database can be used to easily offload backups from the primary database. Backup operations (full or fast incremental and merge) are performed online, with zero impact to data protection or the availability of the standby database to support read-only transactions. Backups taken on the standby can be used to restore either the primary or standby databases.

In addition, Data Guard's integration with Oracle Recovery Manager (RMAN) has taken the support for comprehensive backup and recovery operations on either primary or standby to a whole new level. Some highlights include:

- RMAN DUPLICATE FROM ACTIVE DATABASE eliminates the need for intermediate storage to instantiate a remote standby database.
- Backup and restore operations are transparent across primary and standby role transitions without requiring manual synchronization of the equivalent of DB2 recovery history file (db2rhist.asc).
- Ability to perform table level restore on primary or standby.
- Quick synchronization of a standby database with the primary database using a new Oracle Database 12c simple RMAN command: RECOVER DATABASE .. FROM SERVICE.

In contrast, for DB2, backup operations are simply not supported on an HADR standby database [43]. IBM does document a workaround that requires additional procedures and compromises data protection. IBM's Storage Copy function must be used to first take a snapshot of the storage image from the standby database. This requires the administrator to deactivate the standby database, stop the database instance, unmount all file systems and deactivate (*varyoff*) volume groups in order to disallow any write activity during the FlashCopy.

Dual-purpose standby database as test system

A single command can convert a Data Guard 11g physical standby database to an open read-write database without any impact to primary performance or data protection. This functionality is called Data Guard Snapshot Standby, enabling direct read-write access to a complete point-in time snapshot of the primary database. A Data Guard Snapshot Standby is an ideal QA system for pre-production testing or for any other activity that requires temporary read-write access to a copy of the primary database.

During operation, a Snapshot Standby continues to receive and archive, but does not apply, redo data transmitted by its primary database. Should a failover be necessary or when QA testing is complete, it is converted back into a standby database via a single command that discards all local updates made during testing. The standby is automatically resynchronized with the primary database using the redo data archived locally while it was open read-write. There is no limit to the number of times this process can be repeated.

A Snapshot Standby can also be "reset" as many times as needed by using a guaranteed restore point, to enable multiple iterations of test runs against an identical database – making this an ideal complement to Oracle Real Application Testing.

DB2 HADR does not allow a standby database to transparently and with a single command serve dual-purposes of being open read-write for test and other uses that require read-write access to the standby database, while also providing data and disaster protection for current transactions executing at the primary database. A DB2 HADR standby is also unable to transition back and forth to a read-write database as many times as the user desires, while maintaining disaster protection.

Database-integrated Load Balancing of Application Connections

Many customers have offloaded Read-only and Read-mostly workloads to their Active Data Guard Standby replicas, and Oracle GoldenGate replication also enables distributing workloads over multiple databases, both within and across datacenters. In Oracle Database 12c, Global Data Services (GDS) extends the familiar notion of Database Services to span multiple database instances over replicated Active Data Guard / GoldenGate configurations, in near and far locations, and provides RAC-like failover, service management, and service load balancing to these replicated database configurations. GDS benefits include:

- Higher Availability by supporting service failover across local and global databases.
- Better Scalability by providing load balancing across multiple databases.
- Better Manageability via centralized administration of global resources.

GDS provides very sophisticated database-integrated inter- and intra-region load balancing across replicated databases. For example, it can distribute load across a reader farm composed of standby instances, and even direct some read traffic to the primary if conditions warrant it. Data Guard performs role transitions, but GDS is aware of them and directs load as appropriate. With Active Data Guard, GDS supports:

- Service failover and load balancing across replicated databases in local and remote data centers.
- Automatic role-based services upon Data Guard role transitions.
- Load balancing for reader farms.

With GoldenGate, GDS supports failover and load balancing for local and remote data centers. When Active Data Guard and Oracle GoldenGate allow offloading production workloads to the replication assets, GDS enables better replica utilization, yielding better performance, scalability and availability.

With database-integrated Global Data Services, Oracle customers do not need to install separate hardware load balancers and traffic managers for their database infrastructure. In contrast, DB2 has no such sophisticated capabilities.

Addressing Human Errors

A leading cause of data failure and application downtime is human error, be it accidental or malicious. Traditionally, while it might take minutes to damage a database, it would take hours to recover the original data. Moreover, human errors generally cannot be detected by the database, as erroneous changes are processed just like any other changes while the database remains operational. The real challenge is in identifying such errors and taking the fastest route for recovery.

As shown in the following table, Oracle provides clearly differentiating capabilities compared to DB2 in terms of how effectively it addresses human error situations. Oracle Flashback Technologies reduce recovery time from hours to minutes, unlike traditional solutions that may take hours as they require restoring a lot or all of the database state as it was prior to the incident. Flashback Technologies are a revolution in recovery technology, because they deliver on the principle that the time to recover a database must be independent of the size of the database.

TABLE 5: ADDRESSING HUMAN ERRORS – ORACLE VS. DB2

ADDRESSING HUMAN ERRORS	ORACLE	DB2
Built-in capability to recover dropped objects	Yes	No
Built-in capability to mine logs and audit changes using a SQL interface	Yes	No
Built-in capability to unwind granular transactions	Yes	No
Flexible tablespace point-in-time recovery	Yes	No
Built-in ability to unwind a table to a point in time in the past	Yes	No
Built-in ability to unwind the database to a prior point in time without restoring a backup	Yes	No

The following sections provide further details on the Oracle differentiators to address human errors.

Oracle Flashback Technologies

Oracle Flashback technologies provide point-in-time viewing and quick recovery at the row, transaction, table, and database level. With Flashback, the time to fix a logical error is no greater than the time it took to make the error. Additionally, it is extremely easy to use, e.g. a single SQL command can recover the database instead of performing complex media recovery. Flashback provides fine-grained surgical analysis and repair for localized damage, e.g. when the wrong customer order is deleted. It also allows for correction of more widespread damage yet does it quickly to avoid long downtime, e.g. when all of this month's customer orders have been deleted.

Oracle's Flashback technologies support recovery at all levels including the row, transaction, table, and the entire database [49].

IBM DB2 10.1 introduced Time Travel Query which leverages temporal logic and analysis. This is similar to Oracle Flashback Data Archive, using "`AS OF`" query based on valid (business) time and/or transaction (system) time [50].

Data Recovery

Oracle's Flashback Query allows an administrator or user to view the state of the data at a point in time in the past without requiring any structural changes to the database. This powerful feature can be used to view and reconstruct lost data that may have been deleted or changed by accident. Developers can use this feature to build self-service error correction into their applications, empowering end-users to undo and correct their errors without delay, rather than burdening administrators to perform this task. Flashback Query is extremely simple to manage, as the Oracle server automatically keeps the necessary information to reconstruct data for a configurable time in the past.

DB2 Time Travel Query – specifically System-period Temporal Tables [51], provides similar functionality as that provided by Flashback Query although different in its implementation. Time Travel Query captures table changes at the row level versus Flashback Query's capturing changes at the more granular block level. Flashback Query is available by default leveraging Oracle undo. On the other hand, Time Travel Query is user-configured for each table as desired. Each table using system-period temporal tables has its own history table which contains row changes and thereby needs additional storage.

To recover accidentally dropped objects, Oracle provides Flashback Drop, which accesses the Recycle Bin, a logical container for all dropped objects and their dependent objects. The Recycle Bin uses the free space in each tablespace, and purges objects on a first-in, first-out basis. When a table is dropped, it is not actually deleted, but simply "moved" to the Recycle Bin, renamed with a prefix of BIN$$. All associated table attributes are also renamed. The dropped objects are still accessible by their new name, and each user retains the same rights and privileges on them, as before. The Recycle Bin is always available by default, and no additional setup is required.

DB2 provides no such sophisticated database-integrated table recovery features – it only supports time consuming database-level or tablespace-level restore and recovery [52]. Otherwise customers have to use a separately licensable product – Recovery Expert [25].

Built-in Transaction Recovery

Oracle provides simple queries for viewing a history of changes to a set of rows and examining a transaction and all its effected changes. Flashback Versions Query allows the administrator to see a record of all changes, row version by row version, between two different times and associated metadata such as transaction id and operation type. With Flashback Transaction Query, an administrator can retrieve a listing of all operations (and corresponding undo) effected by a particular transaction. Both of these operations use the existing undo data, whose retention time can be easily configured.

Oracle LogMiner is a powerful audit tool that enables a DBA to find and correct unwanted changes, and can be used in conjunction with Flashback Versions Query. Its simple SQL interface allows searching by user, table, time, type of update, value in update, or any combination of these. It supports a multi-versioned dictionary that gives it the ability to track a database table even as the table goes through DDL-related structural changes. LogMiner provides SQL statements needed to undo the erroneous operation. Additionally, the EM interface graphically shows the change history. It is much easier and quicker than restoring a backup to perform a recovery.

A "bad" transaction's changes can also be easily backed out with Flashback Transaction, accessed via PL/SQL package or EM, along with backout of changes from any of its dependent transactions (i.e. transactions that later modified the same table rows as originally modified by the problem transaction). Once the backout completes, the data can be verified and the changes committed or rolled back. This operation relies on the available redo logs. DB2 does not have single-command backout of erroneous or malicious transactions.

The DB2 Recovery Expert, a separately licensed product, does support rolling back specific transactions and some log analysis [25]. However, DB2 Recovery Expert comes with its own set of restrictions (e.g. unsupported DDLs or datatypes) [53] which Oracle Flashback Technologies do not have. Of course, Oracle offers all Flashback capabilities built into the database itself and are not separately licensed.

Point-in-Time Recovery

When performing a point-in-time recovery, Oracle allows querying the database without terminating recovery. This is useful to determine whether errors affect critical data or non-critical structures (such as indexes). Oracle also allows trial recovery in which recovery continues but can be backed out if an error occurs. It can also be used to undo recovery if point-in-time recovery has gone on for too long. DB2 does not have this capability.

Oracle allows full tablespace point-in-time recovery with no limit on the operations that can be backed out. DB2 imposes a *minimum recovery time* [54], which is the earliest point-in-time for which point in time recovery can be performed for a tablespace. In addition, tablespace point in time recovery is not supported in an HADR configuration.

To quickly recover from human error at the table level, Oracle provides Flashback Table , which allows fast, online recovery for a table or set of tables, to a specified point-in-time with just a single SQL, such as: `FLASHBACK TABLE orders, order_items TIMESTAMP time`. Similar to Flashback Query, Flashback Table relies on the undo data to recover the tables. DB2 does not have this type of built-in fine-grained recovery. IBM requires the use DB2 Recovery Expert, a separately licensed product, for similar capabilities.

In the case of database-wide logical corruption, a point-in-time restore and recovery may be required to get the data back to a state before the corruption occurred. Oracle circumvents the time-consuming traditional restore and recovery procedure by providing Flashback Database, in which the entire database is rewound to a specific point-in-time. Because Flashback Database operates on just the changed blocks, the manually operated traditional restore and recover is not needed, and can complete within seconds or minutes. By issuing the Flashback Database command, blocks are retrieved from the flashback logs, which maintain a history of changed blocks, and then archived redo logs are used to recover to the specific point-in-time.

DB2 has no such built-in continuous data protection (CDP) capability and continues to rely on storage-level functionality for snapshot and CDP. IBM Recovery Expert could use undo to rollback the database although does not have the single command for ease of use like Flashback Database.

Oracle VS. DB2 – Addressing Planned Downtime

Addressing System Maintenance

As business needs change, system changes may also be required. For example, business growth often entails growth in data processing volume. This may translate into a requirement for additional processing power through hardware upgrades of disks, memory, CPUs, nodes in a cluster, or entire systems. Oracle is unique in the ability to change any system resource dynamically, as shown in the following table.

TABLE 6: ADDRESSING SYSTEM MAINTENANCE – ORACLE VS. DB2

ADDRESSING SYSTEM MAINTENANCE	ORACLE	DB2
Simple online addition of cluster nodes that requires no data redistribution	Yes	No
Automatic rebalance with online adding or dropping of disks	Yes	No
Online patching	Yes	No
Rolling database upgrades for full patch-sets and major releases	Yes	No
Extensive support to adjust memory online	Yes	No

Helpful advisories on memory management	Yes	No
Most configuration parameters may be modified online	Yes	No

The following sections provide further details on these capabilities provided by Oracle.

Adding a Cluster Node

While DB2 allows online addition / removal of members in a pureScale environment or in a Data Partitioning environment, the issue with Data partitioning in a shared-nothing environment is that it makes adding new servers to a cluster time consuming and costly, because redistribution of partitioned data according to the new partitioning map is required. Here's what a DBA or System Administrator has to do to add a node to a DB2 database with the Partitioning option:

- Add hardware
- Configure a new partition (set partition-specific parameters, etc.)
- Restart the database (i.e. shut down and restart all nodes)
- Re-distribute the data to spread it across a larger number of partitions

Re-distributing the data in a DB2 system with Data Partitioning involves DBA work and downtime for the database. There are three ways to redistribute this data, but all of them interrupt database operations:

- *Redistribute existing nodegroup* – The data in the nodegroup is inaccessible until the command completes. The time taken for the command to complete grows with the amount of data to be redistributed. Because this is an in-place redistribution, the operation is logged and prone to running out of log space.
- *Create new nodegroup* – Replicas of the old table are created in the new nodegroup. This requires sufficient space to store the data stored in the old nodegroup. Even with this space, the data in the old nodegroup will not be available for modification while it is being copied to the new nodegroup. Further, all dependencies such as indexes, triggers, constraints and privileges will need to be recreated.
- *Piecewise redistribution* – This is similar to the first option except that data in the hash buckets can be redistributed one at time. This spreads the re-distribution over a longer period of time controlled by the user, thereby limiting the window of unavailability at any given time. This is an immense management burden, and will require that the database be taken off-line for a non-trivial amount of time.

Consider, on the other hand, the management tasks needed when you a node has to be added to Oracle RAC:

- Add hardware
- Configure new instance (set instance-specific parameters, etc.)

That's it! No data re-partitioning, no offline maintenance, no database restart – just a seamless scale-out. RAC allows nodes to be added without interrupting database access.

Automatic Rebalance with Online Adding or Dropping of Disks

With ASM, it is possible to add disks to, or drop disks from the disk group that the Oracle database is actively using, without causing any downtime to the database. ASM automatically rebalances a disk group whenever disks are added or dropped, ensuring that database files are spread evenly across all disks in a disk group. This means that administrators do not need to search for hot-spots in a disk group and manually move data around to restore a balanced I/O load.

DB2 does not have any such integrated capability. In IBM pureScale environments, disks may indeed be added or deleted while the database is online although disk rebalancing is a manual step for the DBA and recommended to be performed when there is less load on the system as it is an intensive input / output operation [55].

Online Patching

With Oracle Database 11g, it is possible to install certain one-off database patches for some selected platforms, completely online, without requiring the database instance to be shut down, and without requiring RAC or Data Guard configurations. With online patching, which is integrated with OPatch, each process associated with the instance checks for patched code at a safe execution point, and then copies the code into its process space.

DB2 does not have any similar capability. DB2 at best can perform rolling application of fix packs, requiring an HADR configuration.

Comprehensive Rolling Database Upgrades

For Oracle, Data Guard may be used for rolling database upgrades across major versions and subversions. Data Guard can also be used to minimize downtime when migrating a production database to a new system, storage or architecture. Such examples include Data Guard support for different operating system versions, mixed Linux/Windows, mixed 32bit/64bit, and other select mixed primary/standby combinations.

DB2 HADR does not support rolling upgrades between DB2 version or major subversion migrations, such as version 8.2 to 9.1, or 9.1 to 9.7 or from 10.1 to 10.5. With DB2 HADR, rolling upgrades are only supported from one IBM Fix Pack or Modification Level to the next.

Configuring Memory Online

Oracle allows dynamic resizing of all memory structures without shutting down and restarting the database. IBM, on the other hand offers limited abilities in this regard, such as adding a new buffer pool, altering the size of an existing buffer pool and dropping a buffer pool without requirement to restart the database.

Oracle also helps administrators in ensuring optimal use of available memory. It includes a number of advisories to help administrators precisely determine the amount of memory required to maximize database performance, using patent-pending simulation algorithms to generate the accurate advisory data out-of-the-box with absolutely minimal overhead. Thanks to these advisories, Oracle DBAs no longer need to indulge in time-consuming bouts of trial-and-error to determine memory allocations, nor do they waste system memory due to over-allocation. DB2 does not offer such advisories, which implies DB2 administrators either rely on their experience or use empirical approaches to tune the database performance.

Configuring Parameters Online

In an HA environment, it may be necessary to change database configuration parameters to tune the database's operations for specific operating and systems environments. A significant number of Oracle's parameters can be dynamically modified without requiring any database restart, and thus without causing any downtime. In DB2 version 10.5, a large number of its Database Manager Configuration parameters and Database Configuration parameters still cannot be configured online [56], thus requiring the database to be stopped and restarted for the new parameter values to take effect. This causes application downtime, which can be very detrimental in an HA environment.

Addressing Data Maintenance

As business requirements and processes change, the underlying data has to be maintained and transformed to suit the new environment, and done in such a way that there are minimal or no disruptions to the business. Maintaining, re-defining and transforming the data that supports a business is a critical activity for any DBA – this may be required unexpectedly with new business conditions, or this may even be a regularly scheduled activity. As shown in the following table, Oracle, in contrast to DB2, provides a suite of capabilities in this regard, generally supporting the online maintenance of the production data in support of high availability.

TABLE 7: ADDRESSING DATA MAINTENANCE – ORACLE VS. DB2

ADDRESSING DATA MAINTENANCE	ORACLE	DB2
Flexible and comprehensive capabilities to reorganize and redefine tables online	Yes	No
Online move of data files across storage	Yes	No
Online move of partitions	Yes	No
Online application upgrades with an Edition view of the underlying data	Yes	No
Fast online add column, with default value	Yes	No
Online add/modify constraint, add column, index create/rebuild do not require exclusive lock	Yes	No
Invisible indexes	Yes	No
DDL operations wait for user-specified time, if underlying resource is busy	Yes	No

Online Operations

Oracle offers a wide range of online index and table reorganization operations, from the ALTER INDEX and ALTER TABLE commands to management of more complex reorganization tasks via Online Redefinition. In particular, Oracle's unique Online Redefinition capability allows one to:

- Modify the storage parameters of a table
- Move a table to a different tablespace
- Add, modify, or drop one or more columns in a table
- Add or drop partitioning support

- Change partition structure
- Change physical properties of a single table partition, including moving it to a different tablespace in the same schema
- Change physical properties of a materialized view log or an Oracle Streams Advanced Queueing queue table
- Add support for parallel queries
- Re-create a table to reduce fragmentation
- Change the organization of a normal table (heap organized) to an index-organized table, or do the reverse
- Convert a relational table into a table with object columns, or do the reverse
- Perform data transformations during online table redefinition, e.g. convert LONG data types to LOB, compute new column default values in redefined table based on original table.

While DB2 9.7 introduced a capability to move and redefine tables online, via ADMIN_MOVE_TABLE, this capability has several deficiencies and restrictions when compared with the Oracle Database capabilities [57]. For example, IBM warns of performance issues, lock timeouts and deadlocks while using this procedure.

Online Data File Move and Online Partition Move

With Oracle Database 12c, a data file can be moved online while it is open and being accessed with a simple `ALTER DATABASE MOVE DATAFILE` command. Being able to move a data file online means that many maintenance operations, such as moving data to another storage device or moving databases into Oracle Automatic Storage Management (ASM), can be performed while users are accessing the system. This ensures that continuity of service and service-level agreements (SLA) on uptime can be met.

Similarly, with Oracle Database 12c, the `ALTER TABLE … MOVE PARTITION … ONLINE` statement enables one to relocate data of a partition of a partitioned table into a new segment, or optionally into a different tablespace with additional quota, or modify any of the storage attributes (e.g. compression) of the partition – all in an online manner, which enables DML operations to run uninterrupted on the partition or subpartition that is being moved. Global indexes are also maintained during the move partition, so a manual index rebuild is no longer required.

DB2 has none of these simple, yet sophisticated and online capabilities.

Online Application Upgrades using Editions

Oracle Database also supports online upgrade of applications with uninterrupted availability, via the use of the Edition-Based Redefinition capability introduced with Oracle Database 11g Release 2. When the installation of the upgrade is complete, the pre-upgrade application and the post-upgrade application can be used at the same time. Therefore an existing session can continue to use the pre-upgrade application until its user decides to end it; and all new sessions can use the post-upgrade application. As soon as no pre-upgrade application sessions are left, it can be retired. The application as a whole enjoys hot rollover from the pre-upgrade version to the post-upgrade version. Editions-based Redefinition works as follows:

- Code changes are installed in the privacy of a new edition.
- Data changes are made safely, by writing only to new columns or new tables not seen by the old edition. An editioning view exposes a different projection of a table into each edition to allow each to see just its own columns.
- A crossedition trigger propagates data changes made by the old edition into the new edition's columns, or (in hot-rollover) vice-versa.

IBM's DB2 lacks similar support for upgrading applications online.

Fast Online Add Column

With Oracle, adding new columns with DEFAULT value and NOT NULL constraint does not require the default value to be stored in all existing records. Instead, default values of columns are simply maintained in the data dictionary. This not only enables a schema modification in sub-seconds and independent of the existing data volume, it also consumes virtually no space. DB2 does not offer this type of fast add column ability.

Online No-Lock Add/Modify Constraint, Add Column, Create/Rebuild Index

Oracle's online index and rebuild operations do not use exclusive locks at any time during the operation. This means that ongoing DML (i.e. update, insert, delete) operations on the table work transparently and do not wait for the index operations to finish, thereby minimizing the drops and spikes in system usage that can occur with locks/waits. DB2 acquires a super exclusive 'Z' lock for any alter table or create index operation, which not only restricts updates, but also reads to the table [58]:

> *This lock is acquired on a table under certain conditions, such as when the table is altered or dropped, an index on the table is created or dropped, or for some types of table reorganization. No other concurrent application can read or update the table.*

Invisible Indexes

An Oracle invisible index is an alternative to making an index unusable or dropping it. An invisible index is maintained for any DML operation, but is not used by the optimizer unless the index is explicitly specified with a hint.

Invisible indexes have great uses in application development and testing. Applications often have to be modified without being able to bring the complete application offline. Invisible indexes enable you to leverage temporary index structures for certain operations or modules of an application without affecting the overall application. Furthermore, invisible indexes can be used to test the removal of an index without dropping it right away, thus enabling a grace period for testing in production environments. DB2 has no such equivalent index capabilities.

DDL Wait

Oracle DDL (Data Definition Language) operations affect the logical structure of database objects and may require a lock on the underlying object, e.g. drop column, drop table. Oracle allows the user to specify a wait time for a DDL operation to complete, rather than just immediately failing if a needed lock cannot be acquired. The WAIT option allows developers to define grace periods so that DDL operations can eventually succeed, instead of raising an immediate error, and thus requiring additional application logic to handle such errors. DB2 offers no such grace period for database structure modification operations. While DB2 offers generic lock wait settings for concurrent applications at the database, table, row, partition levels, it does not provide a grace period specifically for database structure modification operations.

Oracle High Availability Best Practices

A well-integrated set of technologies is essential in meeting the stringent high availability demands of today's complex business operations. What is also important is a set of best-practices guidelines that accompany these technologies, so that the implementation can be done in the most efficient manner.

The Oracle Maximum Availability Architecture (MAA) [59], is Oracle's best practice blueprint focused on High Availability to achieve this goal. It provides the following benefits:

- MAA reduces the implementation costs for a highly available Oracle system by providing detailed configuration guidelines. The results of performance impact studies for different configurations are highlighted to ensure that the chosen highly available architecture can continue to perform and scale accordingly to business needs.
- MAA provides best practices and recovery steps to eliminate or minimize downtime that could occur because of scheduled and unscheduled outages such as human errors, system faults and crashes, maintenance, data failures, corruptions, and disasters.
- MAA gives the ability to control the length of time to recover from an outage and the amount of acceptable data loss under disaster conditions thus allowing mean time to recovery (MTTR) to be tailored to specific business requirements.

MAA is designed, architected and validated by Oracle's High Availability Systems Engineering group, which is comprised of individuals with deep-domain design and implementation expertise in high availability best practices, as well as various Oracle and related systems technologies.

Oracle High Availability Customers

Oracle's high availability capabilities have strong customer adoption and are a critical differentiator when the time comes for a prospective customer to choose a database technology that can support the 24x7 uptime requirements of today's businesses. A long list of customers who have implemented various Oracle high availability solutions, along with detailed implementation case studies, is available at [60].

Conclusion

Recognizing the high availability challenges every business faces, Oracle provides comprehensive, unique, powerful, and simple-to-use capabilities that protect businesses against all forms of unplanned downtime, including system faults, data corruption, disasters, and human errors. Oracle achieves this in an environment where the downtime that occurs during planned maintenance activities is also minimized.

Oracle offers a well-integrated high availability solution stack – comprised of components such as RAC, ASM, Data Guard, GoldenGate, Global Data Services, Application Continuity, RMAN, Flashback, etc. that work across multiple platforms in an integrated and efficient manner. This saves customers time, money and system/people resources – factors that are extremely critical in today's economy. Oracle has gone one step further by publishing best practice guidelines for configuring a High Availability solution through its Maximum Availability Architecture framework, and making it available for its customers. The long list of Oracle customers who have embraced its High Availability solutions is a testimonial to Oracle's unparalleled technical leadership and vision in this area.

In contrast to the Oracle Database, IBM DB2 Database offers just a basic set of high availability features and lacks the completeness and depth of High Availability functionality required by most businesses today. DB2 continues to lag several releases behind Oracle in this regard and is not an appropriate choice for today's business applications demanding high levels of uptime.

Reference

1. Oracle Database High Availability Overview 12c – http://docs.oracle.com/cd/E16655_01/server.121/e17601/toc.htm
2. Overview: Oracle Database High Availability – http://www.oracle.com/goto/availability
3. IBM developerWorks: "What's new in DB2 10.5 for Linux, UNIX, and Windows": http://www.ibm.com/developerworks/data/library/techarticle/dm-1304whatsnewdb2105/
4. Technical Comparison – Focus on High Availability - DB2 Version 9.7 vs. Oracle Database 11g Release2 – http://www.oracle.com/technetwork/database/features/availability/cwp-ha-oracle11gr2-db2-9-131679.pdf
5. IBM DB2 Documentation – What's New in Version 9.8 – http://pic.dhe.ibm.com/infocenter/db2luw/v9r8/topic/com.ibm.db2.luw.welcome.doc/doc/welcome.html
6. IBM DB2 Documentation – What's New in Version 10.1 – http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.wn.doc/doc/r0051514.html
7. IBM DB2 Documentation – What's New in Version 10.5 – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.wn.doc/doc/r0051514.html
8. Oracle Database 12c documentation library – http://www.oracle.com/pls/db121/portal.all_books
9. IBM DB2 Version 10.5 Information Center – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/index.jsp
10. Oracle Database Performance Tuning Guide - Tuning Instance Recovery Performance: Fast-Start Fault Recovery, http://docs.oracle.com/cd/E16655_01/server.121/e15857/pfgrf_instance_tune.htm#i1009414
11. IBM DB2 Documentation – Reference – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.config.doc/doc/r0061025.html
12. IBM DB2 Documentation – Reference – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.config.doc/doc/r0061026.html

13. IBM DB2 Documentation – Crash Recovery – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0005962.html
14. IBM DB2 Documentation – Database Partition Groups – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.partition.doc/doc/c0004888.html
15. IBM DB2 Documentation – Database partition and processor environments – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.partition.doc/doc/c0004569.html
16. IBM DB2 Wikipedia – http://en.wikipedia.org/wiki/IBM_DB2
17. IBM DB2 Documentation – Distribution keys – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.partition.doc/doc/c0004906.html
18. IBM DB2 Documentation – The Design Advisor – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0005144.html
19. IBM DB2 Documentation – Redistributing data across database partitions – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.partition.doc/doc/c0053335.html
20. IBM DB2 Documentation – Restrictions on data redistribution – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.partition.doc/doc/c0051446.html
21. IBM White Paper - IBM DB2 pureScale: Extreme capacity with continuous availability – ftp://public.dhe.ibm.com/common/ssi/ecm/en/imb14061usen/IMB14061USEN.PDF
22. IBM White Paper - IBM DB2 Transparent application scaling with IBM DB2 – http://public.dhe.ibm.com/common/ssi/ecm/en/imw14253usen/IMW14253USEN.PDF
23. Oracle White Paper - "Technical Comparison of Oracle Real Application Clusters 11g and IBM DB2 v9 for Linux, Unix, and Windows", http://www.oracle.com/technetwork/products/clustering/overview/wp-oracleibm-2009-130764.pdf
24. Oracle Secure Backup - http://www.oracle.com/goto/osb
25. DB2 Recovery Expert - http://www-03.ibm.com/software/products/us/en/db2recoexpeforlinuunixandwind/
26. DB2 Merge Backup Overview - http://pic.dhe.ibm.com/infocenter/idm/docv3/topic/com.ibm.datatools.mbk.doc/mbkucon_overview.html

27. IBM DB2 Documentation – Backup Compression – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0056483.html
28. IBM DB2 Documentation – IBM Database Encryption Expert for encryption of data at rest – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.sec.doc/doc/c0053466.html
29. Article – "IBM beefs up DB2 with encryption from Vormetric" – http://www.it-director.com/business/content.php?cid=10066
30. DB2 Documentation – Backing up data, http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/t0006192.html
31. DB2 Documentation – Incremental backup and recovery, http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0006069.html
32. DB2 Documentation – Backup and restore operations between different operating systems and hardware platforms, http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0005960.html
33. DB2 Documentation – Using restore – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/t0006242.html
34. DB2 Documentation – Optimizing backup performance – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0006205.html
35. IBM Developerworks article: "A look at the new functions in DB2 Universal Database V8.2", http://www.ibm.com/developerworks/data/library/techarticle/dm-0404zikopoulos/
36. DB2 Documentation – Database configuration for high availability disaster recovery (HADR) – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011761.html
37. DB2 Documentation – Configuring automatic client reroute and high availability disaster recovery (HADR) – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011558.html
38. DB2 Documentation –Enabling reads on standby – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/t0054260.html

39. DB2 Documentation – High availability disaster recovery (HADR) synchronization mode – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011724.html
40. DB2 Documentation – Restrictions for multiple standby databases – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0059997.html
41. DB2 Documentation – Performing rolling updates in a DB2 high availability disaster recovery (HADR) environment – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/t0011766.html
42. DB2 Documentation – System requirements for DB2 high availability disaster recovery (HADR) – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011759.html
43. DB2 Documentation – Restrictions for high availability disaster recovery (HADR) – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0011760.html
44. DB2 Documentation – High availability disaster recovery (HADR) in DB2 pureScale environments – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0061084.html
45. DB2 Documentation – High availability disaster recovery (HADR) monitoring for multiple standby databases – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0060226.html
46. DB2 Documentation – Reads on standby restrictions – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0054258.html
47. DB2 Documentation – Isolation level on the active standby database – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0056191.html
48. DB2 Documentation – Replay-only window on the active standby database – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0056190.html
49. Oracle Database Flashback Technologies, http://www.oracle.com/goto/flashback
50. DB2 Documentation – Time Travel Query using temporal tables – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.dbobj.doc/doc/c0058476.html

51. DB2 Documentation – System-period temporal tables – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.dbobj.doc/doc/c0058477.html
52. DB2 Documentation – Recovering a dropped table – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/t0006318.html
53. IBM Documentation – IBM DB2 Recovery Expert 4.1, Considerations and restrictions – http://pic.dhe.ibm.com/infocenter/db2ars/v1r1/topic/com.ibm.datatools.db2re.doc/contents/amacrest.html
54. DB2 Documentation – Rolling forward changes in a table space – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.ha.doc/doc/c0006312.html
55. DB2 Documentation – Rebalancing a file system – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.sd.doc/doc/t0056703.html
56. DB2 Documentation – Configuration parameters summary – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.config.doc/doc/r0005181.html
57. DB2 Documentation – ADMIN_MOVE_TABLE procedure - Move tables online – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.sql.rtn.doc/doc/r0055069.html
58. DB2 Documentation – Lock attributes – http://pic.dhe.ibm.com/infocenter/db2luw/v10r5/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0005270.html
59. Oracle Maximum Availability Architecture – http://www.oracle.com/goto/maa
60. Oracle High Availability Customer Case Studies – http://www.oracle.com/technetwork/database/features/availability/ha-casestudies-098033.html

ORACLE

Technical Comparison: Oracle Database 12c vs. IBM DB2 10.5: Focus on High Availability

October 2013

Author: Oracle HA Product Management

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Oracle is committed to developing practices and products that help protect the environment