Oracle Rdb7™ and Oracle CODASYL DBMS™

Guide to Hot Standby™ Databases

Release 7.0

Part No. A42860-1

ORACLE®

Oracle Rdb7 and Oracle CODASYL DBMS: Guide to Hot Standby Databases

Release 7.0

Part No. A42860-1

Contents

3 Replication Server Processing

4 Ongoing Database Management

5 Shutting Down Database Replication

6 Replication Commands

7 Monitoring Replication Activities

8 Handling Failure Conditions

Examples

Figures

Tables

Send Us Your Comments

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

You can send comments to us in the following ways:

- Electronic mail—nedc_doc@us.oracle.com
- FAX—603-881-0120
  Attn: Oracle Rdb and Oracle CODASYL DBMS Documentation
- Postal service
- Oracle Corporation
  Oracle Rdb and Oracle CODASYL DBMS Documentation
  110 Spit Brook Road, ZKO2-1/O19
  Nashua, NH 03062-2698
  USA.

If you like, you can use the following questionnaire to give us feedback. (Edit the online release notes file, extract a copy of this questionnaire, and send it to us.)

Name _____ Title _____

Company _____ Department _____

Mailing Address _____ Phone No. _____

_____

Book Title _____ Version No. _____

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?

- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available).

Preface

The Hot Standby optional software physically replicates a database at a local or geographically remote standby site. The standby database is fully accessible for read-only functions and for immediate processing capability in the event of a system or database failure at the master site.

This manual describes how to implement the Hot Standby software and how to set up and maintain a Hot Standby database for either an Oracle Rdb or an Oracle CODASYL DBMS database system.

Intended Audience

This manual is intended for anyone who is interested in replicating an Oracle Rdb database or an Oracle CODASYL DBMS database to a standby database site.

Structure

This manual contains eight chapters, two appendixes, a glossary, and an index:

Chapter 1	Introduces database availability concepts and provides an overview of replicated databases.
Chapter 2	Describes the steps you need to take to set up your master and standby sites for database replication.
Chapter 3	Explains how the Hot Standby software uses specialized servers to provide a replicated database environment.
Chapter 4	Describes some management tasks you might need to perform while replication operations are active.
Chapter 5	Walks you through the steps to stop replication.
Chapter 6	Provides the syntax and semantics for the Replicate commands and qualifiers.
Chapter 7	Describes how to evaluate your Hot Standby environment and its performance using several database monitoring tools.

Chapter 8	Explains how to detect failures and how to handle failure conditions.
Appendix A	Provides a table of the Hot Standby logical names and configuration parameters, including a brief description of each.
Appendix B	Contains the replication startup and shutdown checklists that you can copy for easy reference.

Related Documents

Refer to the following documents for supplemental information:

Oracle Rdb Documentation	Oracle CODASYL DBMS Documentation
Oracle Rdb7 Release Notes	*Oracle CODASYL DBMS Release Notes*
Oracle Rdb7 Installation and Configuration Guide	*Oracle CODASYL DBMS Installation Guide*
Oracle RMU Reference Manual	*Oracle CODASYL DBMS Database Administration Reference Manual*

Conventions

In this manual, OpenVMS refers to both the OpenVMS Alpha and the OpenVMS VAX operating systems.

Oracle Rdb refers to Oracle Rdb for OpenVMS and Oracle Rdb for Digital UNIX software.

Version 7.0 of Oracle CODASYL DBMS software and Oracle Rdb software may be referred to as V7.0.

The SQL interface to Oracle Rdb is referred to as SQL. This interface is the Oracle Rdb implementation of the SQL standard ANSI X3.135-1992, ISO 9075:1992, commonly referred to as the ANSI/ISO SQL standard or SQL92.

In examples, an implied carriage return occurs at the end of each line, unless otherwise noted. You must press the Return key at the end of a line of input.

The following conventions are also used in this manual:

Convention	Meaning
. . .	Vertical ellipsis points in an example mean that information not directly related to the example has been omitted.

. . .	Horizontal ellipsis points in statements or commands mean that parts of the statement or command not directly related to the example have been omitted.
boldface text	Boldface type in text indicates a term defined in the text, the glossary, or in both locations.
italic text	Italic text emphasizes important information and indicates complete titles of manuals and variables. Variables include information that varies in system messages (Internal error *number*), in command lines (/PRODUCER=*name*), and in command parameters in text (where *device-name* is the current disk name).
[ ]	Brackets enclose optional clauses from which you can choose one or none.
$	The dollar sign represents the DIGITAL Command Language prompt for OpenVMS systems and the Bourne shell prompt for Digital UNIX systems.

1
Database Availability

Environmental and operational faults are the major source of application downtime and unavailability. As computing systems and environments become increasingly complex, it is no longer sufficient that processors never fail and that disks are shadowed. The data must remain available at all times under all conditions, including disasters.

This chapter introduces key availability terms and concepts and describes how Oracle Rdb and Oracle CODASYL DBMS databases are designed to remain available and meet the computing needs of the business.

1.1 Introduction to Availability

Today's businesses depend on their computing resources for many key services. A computing system's **availability** is the measure of its ability to provide the desired service when required. The consequences of unavailability can range from inconveniencing users and dissatisfying customers, to impacting production and revenue, to even loss of life or limb.

Your business requirements affect the availability that you need from your computing system. Availability requirements for most businesses fall into the categories shown in Table 1-1.

Table 1-1 Business Requirements for Availability

Requirements	Description
Conventional	Business functions that can wait for a reasonable amount of time, with little or no effect, while a system or application is unavailable.
High Availability (24x365)	Business functions that require uninterrupted computing services, either during essential time periods or during most hours of the day throughout the year. Minimal downtime is acceptable.

Table 1-1 Business Requirements for Availability (Continued)

Requirements	Description
Fault Tolerance	Business functions that demand uninterrupted computing services and data access. Achieving the business's financial goals often depends directly on how well its computing systems and database support the firm's computerized business applications.
Disaster Tolerance	Business functions that demand fault tolerance as well as immunity from disasters such as earthquakes, fires, floods, hurricanes, power failures, vandalism, or acts of terrorism.

Oracle Rdb and Oracle CODASYL DBMS database management systems have been designed to provide high availability. Core functionality such as online backup and recovery operations, after-image journaling, online file and data restructuring, database consistency verification, statistics monitoring, and journal maintenance provide databases that support production systems with conventional through fault-tolerant availability requirements.

Using the optional Oracle Hot Standby software, you can implement an alternate operational database to extend the level of database availability to include disaster tolerance. By implementing the Hot Standby option, businesses with mission-critical requirements can prevent the catastrophic loss of data and data processing capability.

1.2 Reliable Data Replication

Although it is possible to manually move data (such as by using data manipulation language statements), *physical* methods of data replication are unreliable, error prone, and very expensive in terms of time and resources. Moreover, manual attempts to synchronize transactions, provide data sharing among different operating environments, and still maintain good performance are very difficult if not impossible.

The Hot Standby option solves these problems by:

- Automating the delivery of modified data to either a database on the local node or across a wide area network (WAN), and supporting both the OpenVMS and Digital UNIX operating systems
- Reliably delivering replicated data in a timely manner without affecting application performance
- Automatically and reliably synchronizing databases after failures
- Maintaining varying degrees of transactional database consistency and integrity

- Performing all replication operations transparently to users and applications

1.2.1 Introduction to Standby Databases

A **standby database** is a second running database that is created from and transactionally consistent with the primary or **master database**. Data modifications that are made to the master database are made simultaneously to the secondary database. The secondary database is sometimes referred to as a *hot* standby database because it is available immediately to pick up application processing if the primary database system fails.

Oracle Corporation offers an optional **Hot Standby** software solution that you can use to implement a standby database for mission-critical and disaster recovery functions. The Hot Standby software provides automatic, high-performance database replication for:

- Oracle Rdb databases running on OpenVMS or Digital UNIX operating systems
- Oracle CODASYL DBMS databases running on OpenVMS operating systems

The Hot Standby software prevents your Oracle Rdb database or Oracle CODASYL DBMS database from becoming a single point of failure by replicating the master database to a standby database. The Hot Standby software automatically performs coordinated database synchronization and verification with high performance and minimal impact on the master system resources.

You can locate the standby database either on a local node or cluster, or on a node or cluster that is geographically remote from the master database.

Figure 1-1 illustrates a company with a master database located in San Francisco and a remote standby database in Boston.

Figure 1-1 Site-to-Site Database Replication

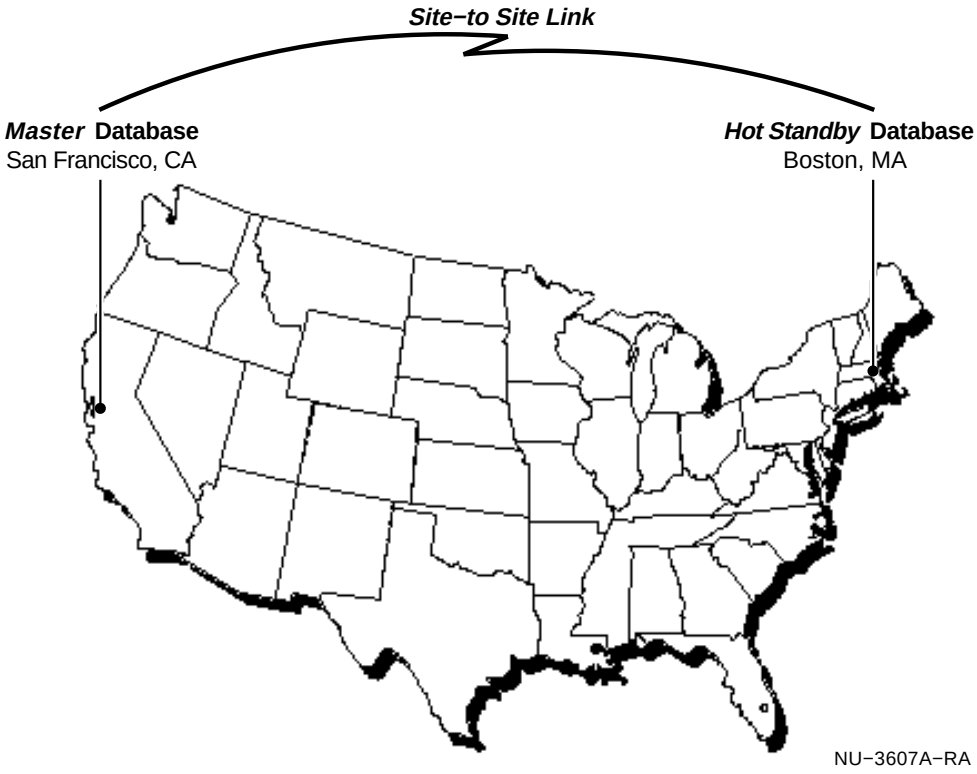

If the database in San Francisco becomes unavailable, the company can immediately fail over processing to the standby database in Boston. Thus, the Boston database becomes the master database. If the Boston database has sufficient processing capability, it can remain as the master database indefinitely or until you decide to fail back processing to the San Francisco database.

The Hot Standby software also provides these advantages:

- The master database resource requirements and performance impact are minimal, while all available resources on the standby system are fully utilized.
- The standby node configuration is distinct from the master database configuration. Thus, system failure of either the master or standby database is completely isolated; neither database is affected by the failure of the other. Plus, you can

make configuration changes to the master configuration without affecting the standby configuration.

Note

You cannot replicate Oracle7 databases using the Hot Standby software.

The Hot Standby software provides database replication by employing special server processes. Servers on the master database automatically ship database modifications, in the form of after-image journal records, to servers on the standby copy of the database. The after-image journal records (that have already been committed on the master database) are rolled forward to the standby database.

1.2.2 Replication Overview

The Hot Standby software is only one of several types of database replication architectures. Replication products come in numerous forms and combinations of techniques. The differences between replication architectures can be very subtle because the design is usually defined by varying levels of business requirements.

Business requirements for replicated databases can range from downloading data for easier localized data sharing to load balancing, to data warehousing, and to achieving disaster-tolerant levels of availability by maintaining a second (standby) database on a separate running computing system.

1.2.2.1 Distributing the Data

You can think of replication as a type of distributed database technology that attempts to keep the databases synchronized at all sites. Common to most replication architectures is the following process for distributing data:

1. Create identical copies of the same database at various locations throughout the organization
2. Make the databases transactionally consistent by applying the same modifications to each database

During the period of making the databases consistent, replication operations usually are not active.

Distributing the data is straightforward compared to managing data movement between the database sites.

1.2.2.2 Applying Data Modifications

Replication architectures vary with regard to how data modifications are shipped to all the databases. Data movement adds complexity because the updates must be applied to multiple locations, and, in some configurations, the updates may be streaming in from multiple locations. Data movement can be either unidirectional or bidirectional.

Unidirectional Data Movement

Data moves in only one direction, from a master (primary) database to one or more secondary databases. When a user changes or adds information to the master database, the local server automatically downloads the data to each secondary database.

This asymmetric method keeps the data centralized but access to data is distributed. Users can read data from all sites but perform write I/O to the master database only. Management is simplified with unidirectional data movement because there is almost no risk of data being overwritten by multiple servers.

Bidirectional Data Movement

Data moves back and forth between all servers on the network. Data updates made by any single server are replicated automatically to every other server on the network. With this method, data movement can flow in both a top-down and a bottom-up approach.

This symmetric method of processing distributes the data as well as distributing access to the data. Users can read and write data to and from all sites. Management becomes more complicated because the replication architecture must ensure that changes made by one server are not overwritten by another server.

1.2.2.3 Synchronizing the Databases

The replication architectures can update multiple copies of data using either synchronous (real-time) processing or asynchronous (deferred-update) processing. Some replication architectures allow you to use some combination of synchronous and asynchronous processing to dynamically customize replication operations. With this flexibility, you can achieve a better balance between high availability, high performance, and increased data integrity according to your business needs.

Synchronous Processing

Synchronous processing maintains copies of data at multiple sites on an event-driven or real-time basis. A change to any database is immediately propagated to other replication servers as a part of the same transaction. The databases are usually immediately identical and recoverable if a failure occurs.

Synchronous processing ensures data integrity and minimizes complexity but can be less available if the systems and networks involved are slow or unreliable. This technology is more suitable when you require read / write transactions at all of the database sites.

Asynchronous Processing

Asynchronous processing propagates changes to the other copies of the database based on a counter or on a time-driven basis. For example, updates are deferred until a predetermined time period elapses or a set number of transactions are processed.

Asynchronous processing maximizes availability and minimizes response time, but it requires careful planning and design to ensure application integrity. This technology is more suitable for applications that read data at many sites, and write data modifications at only one site.

1.2.3 Hot Standby Architecture

The Oracle Hot Standby option provides unidirectional data replication for the purpose of maintaining a standby database. To ensure reliable data delivery and availability, the software uses a combination of synchronous and asynchronous techniques.

1.2.3.1 Data Distribution and Movement

Replication is asynchronous in that servers on the master database first record database modifications in the master after-image journal before shipping the same modifications to the standby site. Applications can perform read / write and read-only transactions on the master database, and read-only transactions on the standby database.

This strategy provides high data integrity by ensuring that all database consistency and integrity verification occurs on the master database; data does not need to be reverified on the standby database because it is not replicated until it is already validated on the master database.

1.2.3.2 Database Synchronization

The Hot Standby software allows you to choose how closely you want committed transactions to be synchronized on the master and standby databases. You can explicitly choose from four levels that range from a very tightly consistent model of synchronization to a loosely consistent asynchronous model. (These models are discussed in Chapter 4.)

The level of database synchronization directly affects how the servers process transactions. For example:

- In the tightly consistent synchronization model, servers apply the changes to the master and standby databases within a single transaction. Servers lock the intended data at both locations until the successful commit of the records to both databases.
- In a loosely consistent synchronization model, the replication servers on the standby database delay processing the transaction. Replication servers on the standby database use the *store-and-forward* technique to delay applying changes to the standby database. With this method, transactions local to the master database are decoupled from their delivery to the standby database and distributed locks are not acquired.

The effect of different synchronization modes on the replication servers and on update performance is described in more detail in Section 4.1.

1.2.3.3 Server Model

Once you create a database on the standby node from a full backup copy of the master database and start the replication services, the Hot Standby software applies database modifications to the after-image journal on the master database, and then ships the same modifications over a network connection to the corresponding after-image journal for the standby database.

To do this, the Hot Standby software employs server processes that automate the after-image journal backup and rollforward operations. The server processes are designed to provide an automated, reliable, and high-performance mechanism for applying data modifications to both the master and standby databases. Moreover, the servers perform all replication operations transparently to applications and users on both master and standby database systems.

The Hot Standby software uses four server processes, two on the master database and two on the standby database, to implement client/server relationships that provide database replication services and network communications. Each server performs a distinct role in the replication process.

The master database servers include:

- Log catch-up server (LCS)

  Catches up the standby database with the after-image journal transactions recorded for the master database. The catch-up server does this by sending logs over the network to the replication servers on the standby database.
- AIJ log server (ALS)

Ships database modifications to servers on the standby database after applying the same modifications to the master database. Data shipments begin only after the master and standby databases are synchronized with each other.

Note:

The ALS process is not specific to replication operations; the ALS can be active for databases running with or without the Hot Standby option. When the Hot Standby option is enabled, the ALS process extends its capabilities to include shipping data modifications to both the master and the standby databases.

The standby database servers consist of:

- AIJSERVER

  A network object server that *receives* update records from the master database. This server does not perform any database access; rather, it is invoked by the network software to act as a communications gateway between the master and standby databases.

- Log rollforward server (LRS)

  A database server that *rolls forward* the database modification received from the master database and after-image journal.

There can be only one ALS, LRS, or LCS process on any single node in the cluster. However, there can be more than one AIJSERVER process on the standby database node where you invoke replication operations. (For example, there can be an AIJSERVER process communicating with the ALS process and another communicating with the LCS process). For simplicity, the figures in this manual depict only one AIJSERVER process.

Figure 1-2 shows the servers used for Hot Standby replication operations.

Note

The log catch-up server (LCS) operates on the master database to synchronize the master and standby databases. The LCS process is not shown in Figure 1-2 because once the databases are synchronized, the LCS process is in a local event flag (LEF) “wait” state while replication processing is performed by the ALS process. The LCS process is described in detail in Section 3.2.4.

Figure 1-2 Hot Standby Replication Server Processes[1]

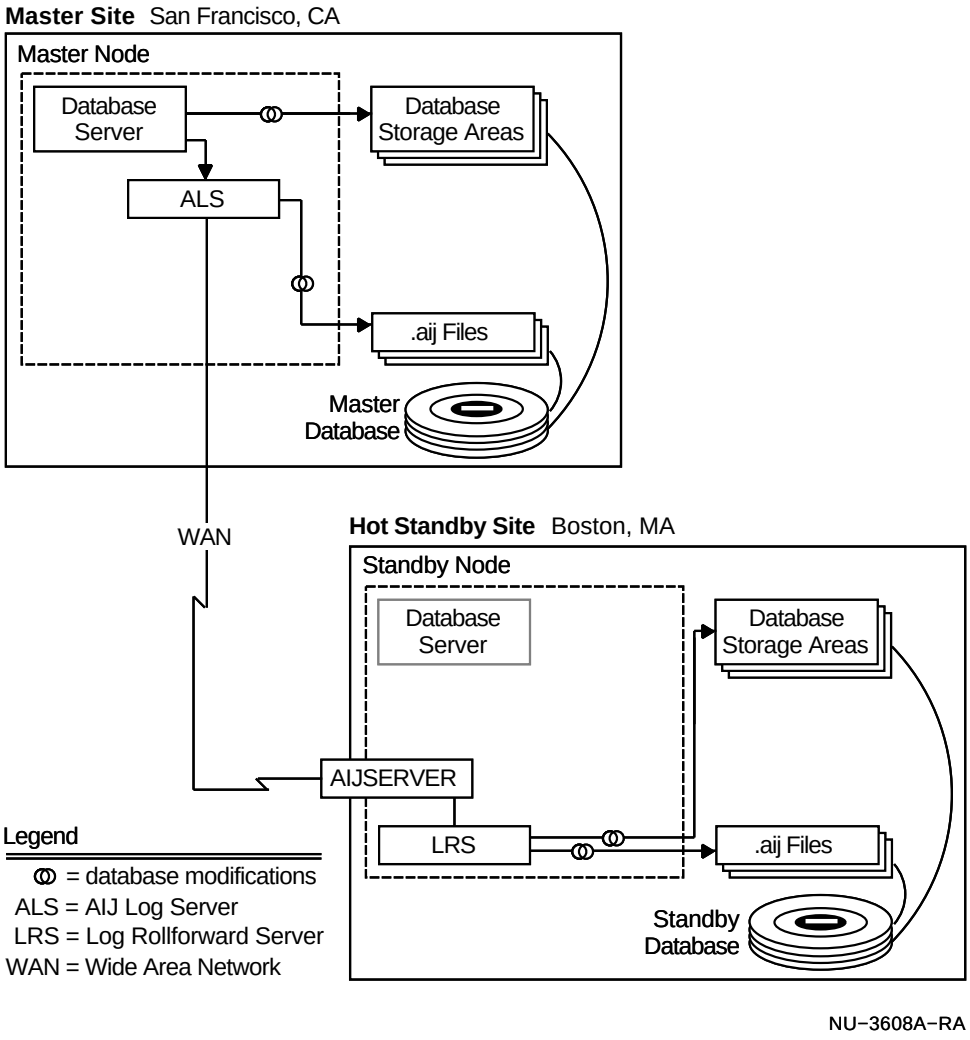

[1] The database server shown in figures in this manual represents the generic execution server process that is integral to the Oracle Rdb or Oracle CODASYL DBMS database product. The purpose of the database server is to execute application programming interface (API) requests for the communications server.

The replication server processes are heterogeneous in that they can transmit database updates between the OpenVMS and the Digital UNIX platforms. This manual often refers to the server processes as *replication servers* because they implement client/server relationships to provide database replication services.

See Chapter 3 for a complete discussion of the replication servers and functions.

1.2.4 Record-Level Replication

The Hot Standby software replicates data at the physical database record level, including application data, B-tree indexes, and hashed index data.

Replication occurs by transferring the maximum number of record-level modifications, as a group, in a single network communication. Although this method of communication requires tight synchronization between the master and standby databases, it efficiently utilizes the network bandwidth and maximizes performance.

1.2.5 Database Consistency and Synchronization

Because the updates to the master and standby databases are committed asynchronously, the databases might not always be identical. The transactions applied to the standby database can lag behind the master database.

You can specify the degree of database synchronization when you start replication operations. The amount of synchronization can range from a fully asynchronous "fire-and-forget" model with a large lag time, to a synchronous "transaction commit" level in which data modifications are made to both databases at the same time.

In general, a high degree of synchronization provides excellent recoverability from a master database failure, but might result in slower performance on the master database during replication operations.

When you start replication operations, you can either:

- Explicitly specify the desired level of synchronization
- Enable the Hot Standby replication governor function to automatically and dynamically adjust the database synchronization according to the current demands on the databases

Chapter 6 describes the replication commands for controlling database synchronization explicitly or using the replication governor.

1.2.6 Replication Commands

You start and stop replication operations by entering special commands on the master and standby database nodes. The replication commands are available using either Oracle RMU, the Oracle Rdb management utility, or DBO, the Oracle CODASYL DBMS Database Operator utility.

The Oracle RMU and DBO utilities provide syntax and semantics for the Replicate commands shown in the following table:

Command	Description
Replicate After_Journal Configure	Automatically sets up replication attributes for the master and standby databases
Replicate After_Journal Reopen_Output	Allows you to open a new replication output file
Replicate After_Journal Start	Initiates replication operations on the master node or the standby node
Replicate After_Journal Stop	Terminates replication operations

You can use the replication commands while the databases are open and active. Chapter 6 provides complete information regarding the use of the Replicate commands and qualifiers.

Note

There is no SQL syntax specific to the Hot Standby functionality.

1.2.7 What Is Not Supported

Because the primary purpose of the Hot Standby option is to keep the standby database consistent transactionally with the master database, the software does not provide general-purpose functionality for all facets of database replication. The following subsections describe functions that are not supported.

Multiple Standby Databases

The Hot Standby option permits the replication of data to only one standby database. Furthermore, the software does not allow a standby database to be replicated.

Offline, Unjournaled Database Operations

Some offline database operations either cannot be replicated or should not be used while replication operations are active. The following list shows the database operations you should avoid during replication because they are offline operations that require exclusive access to the master database and are not journaled:

- Add storage area slots
- Add after-image journal slots
- Add new index
- Replicate after-image journal syntax that is previous to Version 6.0
- Enable or disable after-image journaling
- Change fast commit parameters
- Change snapshot mode
- Enable or disable incremental backup operations
- Enable or disable global buffers
- Change buffer counts
- Enable or disable partitioned lock tree optimization
- Enable or disable adjustable lock granularity
- Change adjustable lock granularity factors
- Change the node count
- Change the user count
- Change the transaction sequence number (TSN) interval

Read/Write Transactions on the Standby Database

You cannot perform read / write transactions on the standby database. Therefore, during replication operations, the standby database is a read-only database. Any attempt to modify data in the standby database causes an exception or an error message.

Note

The standby database must be identical to the master. Thus, you cannot make changes to the standby database, such as adding new indexes, to better support report generation or other read-only queries.

Figure 1-3 shows how the servers handle read-only access to the standby database for ad hoc queries and to fetch database information for local business and report generation.

Figure 1-3 Performing Read Operations to the Standby Database

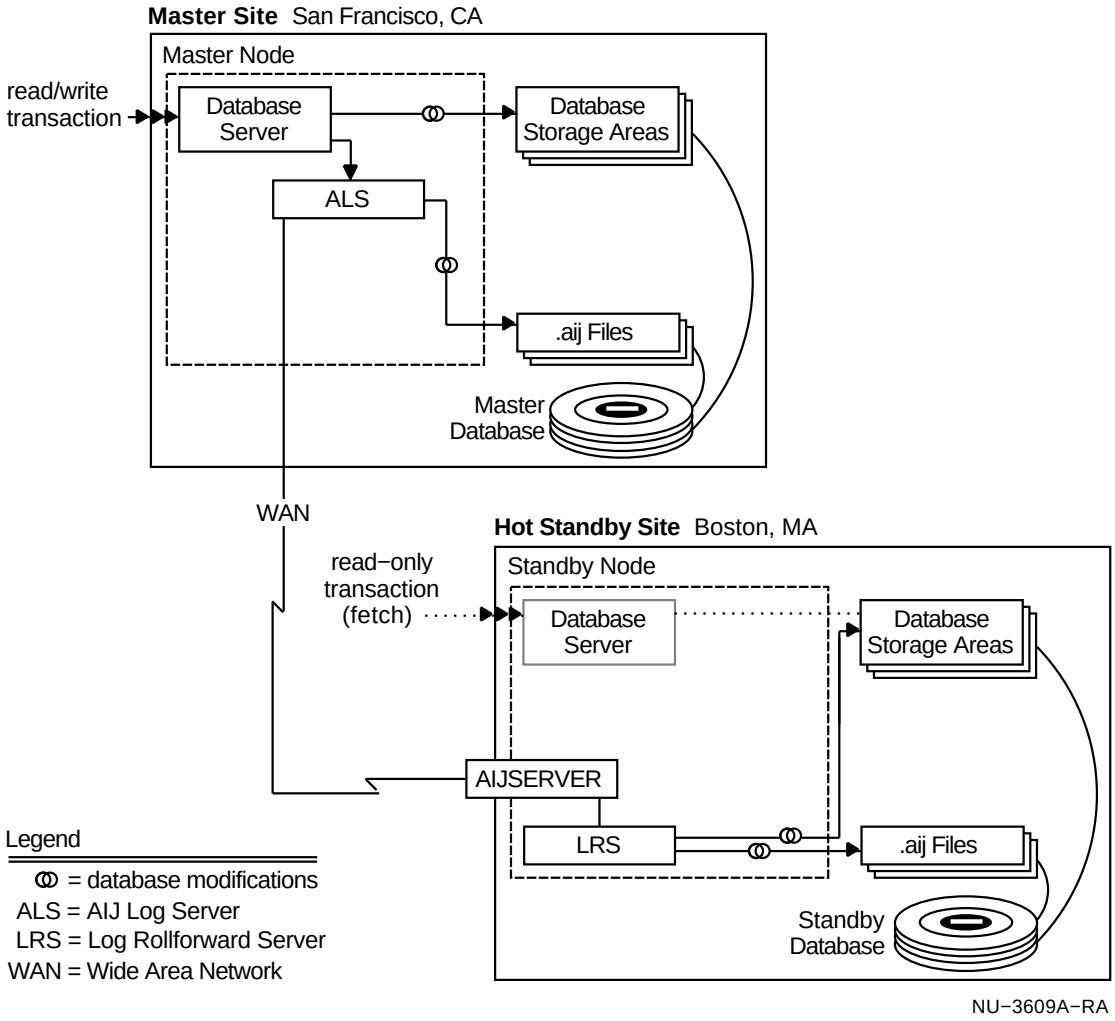

Automatic Failover

The Hot Standby option does not provide automatic failover of applications to the standby database. The standby database is a read-only database on the secondary node until a failure occurs. When this happens, you must manually fail over application processing to the standby database such that it becomes the new master database. The standby database takes over processing indefinitely or until you choose to *fail back* to the original database system. Similarly, there is no automatic failback mechanism; you must manually fail back to the original database system.

1.2.8 Failure Notification

The standby database configuration is separate and distinct from the master database configuration. System failure of either the master or standby database is completely isolated from the other database; neither database is affected by the failure of the other.

You must manually fail over applications to the standby database if a failure occurs. There is no automatic failover support. However, automatic failover detection is provided by the Hot Standby software, including notification for these types of failures:

- Network
- Cluster
- Node
- Disk
- Server
- Application
- Database monitor
- Other resources, such as process quotas, disk quotas, or disk space exhaustion

Failure detection and failover can be achieved with minimal interruption to database users and application processing.

Chapter 8 describes failure detection and handling in more detail.

1.3 Licensing and Installation

To obtain a Hot Standby license, you must contact your Oracle representative.

The Hot Standby software is included in the Oracle Rdb and Oracle CODASYL DBMS base kits. To enable this optional software, choose the Hot Standby option from the installation menu. The installation menu lists each of the licensed products, including products for which you do not have a license. You should choose the Hot Standby option only if you are licensed to install it.

1.4 Environment

The following sections describe the hardware configurations, software environments, operating systems, and network transports supported by the Hot Standby software.

1.4.1 Hardware

Database replication operations do not require specific hardware or configurations to operate. All functions can be performed on hardware components that support traditional Oracle Rdb and Oracle CODASYL DBMS databases. In fact, the Hot Standby option does not require that you configure the master and standby systems identically.

You can configure the master and standby databases to run on the full range of hardware configurations, from a single processor to multiple processors. Configuring the standby database on a node or cluster that is geographically remote from the master database results in the highest (disaster tolerant) availability.

You can configure a system or cluster in an unlimited number of ways. Uniprocessors, multiprocessors, clusters, parallel processors, and distributed computing all have trade-offs with regard to availability and performance. As you plan your configuration, evaluate the characteristics and costs of each type of system and decide what trade-offs you are willing to make to meet your business's availability requirements.

Note

For the best performance, Oracle Corporation recommends that you configure the standby database node on symmetric multiprocessing (SMP) hardware.

If you intend to use the standby database as a genuine “hot” standby system, you should maintain identical hardware configurations for the master and standby nodes. Physically identical configurations are important, especially if the standby database must be capable of serving as the master database for an indefinite amount of time. However, you can relax your configuration requirements if the only auxiliary function of the standby database is to generate reports and allow read-only database queries.

1.4.2 Software

The Hot Standby software is completely compatible with the standard Oracle Rdb and Oracle CODASYL DBMS database components. Applications and users perform database modifications using the same data definition language (DDL) and data manipulation language (DML) commands that are used for databases without the Hot Standby capability. All online database DDL and DML master database operations can be replicated on the standby database.

System managers and database administrators (DBAs) manage and monitor the master and standby databases and configurations using the same commands and utilities they use for databases without the Hot Standby capability.

Furthermore, you do not need to make any changes to application code. The Hot Standby software performs all database replication activities transparently to users and applications that are accessing the master database system.

1.4.3 Operating Systems

You can implement the Hot Standby software on systems running OpenVMS and Digital UNIX systems as follows:

- For Oracle CODASYL DBMS, you can implement the master and standby databases on systems running OpenVMS VAX, OpenVMS Alpha, or both.
- For Oracle Rdb, you can implement the master and standby databases on systems running OpenVMS VAX, OpenVMS Alpha, Digital UNIX, or any combination of these operating systems.

For example, you can implement the Hot Standby software with Oracle Rdb to replicate a master database running on OpenVMS Alpha to a standby database running on Digital UNIX. The only requirement is that the network connection is common to both operating systems.

1.5 Network and Local Communications

Synchronized database replication requires reliable and fast communications between highly available servers. You can configure the master and standby databases on the same (local) node or on geographically separate (remote) nodes. For the highest availability, Oracle Corporation recommends that you locate the databases on distinctly separate nodes connected by a high-speed network connection (for example, you might consider using the Asynchronous Transfer Mode (ATM), the Fiber Distributed Data Interface (FDDI), or the DS3 interconnects).

Note

Although the Hot Standby software can provide a high level of database availability when you replicate a database on the same local node, you can achieve disaster tolerance only when you locate the standby database at a geographically remote location.

When you locate the master and standby databases on the same node, the databases communicate using local interprocess connections. Databases located on separate nodes communicate using one of the network transport protocols shown in Table 1-2.

Table 1-2 Network Transport Protocols

Network	Implementation
DECnet for OpenVMS	Implements Phase IV of the Digital Network Architecture (DNA), and uses the Network Control Program (NCP) utility to configure and monitor the network.
DECnet / OSI	Implements Phase V of the Digital Network Architecture (DNA), and uses the Network Control Language (NCL) utility to configure and monitor the network.
TCP / IP	Implements Transmission Control Protocol / Internet Protocol networking, a common set of open networking standards that support open system interconnection.

1.5.1 Specifying the Default Network Transport

You can allow the Hot Standby software to use the network transport protocol that is the default for your system, or you can override the default network and specify the network of your choice by defining a logical name or configuration parameter.

Reference: Section 2.12 describes how to specify the default network transport protocol in more detail.

1.5.2 Configuring the DECnet Network Object

When you install your database software and choose the Hot Standby option, the installation procedure automatically configures the DECnet images necessary to use the Hot Standby capability.

You do not need to perform any special tasks to install or invoke the DECnet network object. The installation procedure automatically:

- Creates the network object for the AIJSERVER image in a restricted account. When you install your database software, the installation procedure prompts you to supply a valid user identifier for this account.
- Gives the AIJSERVER image (RDMAIJ.EXE or DBMAIJ.EXE located in SYS$COMMON:[SYSEXE]) netmbx and tmpmbx privileges.
- Invokes the database monitor, which:
  - Determines the type of network (DECnet for OpenVMS, DECnet/OSI, or TCP/IP) that is running between your database nodes.
  - Invokes the DBMAIJSERVER.COM or RDMAIJSERVER.COM procedure, which runs RDMAIJ.EXE or DBMAIJ.EXE to activate the appropriate DECnet network image.

For example, if you are using Oracle Rdb for OpenVMS, the RDMAIJSERVER.COM procedure can configure the AIJSERVER object for DECnet/OSI communications through the NCL command interface. The AIJSERVER object name is DBMAIJSRV for Oracle CODASYL DBMS databases, and RDMAIJSRV for Oracle Rdb databases.

If an error occurs (for example, if the database monitor is not started or if the Hot Standby software is not installed), a message displays to help you manually run the RDMAIJSERVER or DBMAIJSERVER command procedures. The message might require you to set up remote communications manually using the instructions in your database installation and configuration guide.

In addition, if you execute the monitor startup procedure interactively on other nodes after the installation, you do not have to invoke the DBMAIJSERVER or RDMAIJSERVER command procedure again to configure the network object.

Note

There are no variant names for the RDMAIJSERVER and DBMAIJSERVER procedure in multiple-version environments. The RDMAIJSERVER and DBMAIJSERVER names are generic for standard and multiversion environments, thus providing easy backward compatibility.

1.5.3 Configuring the TCP/IP Network Object

Although the installation procedure automatically creates and sets up the DBMAIJSERVER or RDMAIJSERVER procedures, the installation procedure does not set up the TCP/IP network object. Because many different TCP/IP services are available (for example, Digital UNIX systems use the UCX service), you must manually set up the TCP/IP network object.

To prepare for TCP/IP network communications, you must define the TCP/IP network transport protocol using the BIND_HOT_NETWORK_TRANSPORT logical name or configuration parameter and specify a UCX port number before you can enable TCP/IP network communications on Digital UNIX systems.

Reference: Sections 2.12 and 4.3 provide detailed information about the setup and management of TCP/IP network communications.

2
Starting Database Replication

This chapter provides a step-by-step procedure to implement a Hot Standby database environment. Although the Hot Standby software performs replication activities automatically and transparently, you must perform the tasks described in this chapter to enable the master and standby databases for replication.

All discussions are based on a configuration in which two identical computing sites and databases are connected by a network. Ideally, the standby site should be identical (or nearly identical) to the master site both in configuration and available resources. Oracle Corporation recommends identical configurations so you can fail over application processing to the standby site for an unlimited amount of time.

2.1 Database Dump Header

As you set up your Hot Standby environment, consider using the DBO or RMU Dump Header command to verify database attribute information. This chapter includes displays of database header information to show you the results of the task being performed. Example 2-1 shows the header information for a master database before it is prepared for replication.

Example 2-1 Dump Header Information Before Starting Replication Operations

```
$  RMU/DUMP/HEADER=HOT_STANDBY/OUT=master.lis mf_personnel
$  TYPE master.lis
*-------------------------------------------------------------------------------
* Oracle Rdb V7.0-00                                    24-MAY-1996 09:18:54.000
* Dump of Database header
*     Database: DISK1:[USER]MF_PERSONNEL.RDB;1
*-------------------------------------------------------------------------------
Database Parameters:
    Root filename is "DISK1:[USER]MF_PERSONNEL.RDB;1"
    Hot Standby...
      - WARNING: After-image journaling is disabled
      - WARNING: Fast commit is disabled
      - WARNING: Log server startup is MANUAL
```

```
- Informational: Operator notification is disabled
- Database is not currently being replicated
```

2.2 Replication Operation Checklist

Table 2-1 provides a checklist of tasks that you perform to start replication operations. Each step includes a reference to a section later in this chapter that provides additional information.

Table 2-1 Replication Startup Checklist

Step	Procedure	Reference
[1]	On the master database, set up the following after-image journal options: • Reserve five or more after-image journal slots • Add after-image journals and specify after-image journal names • Enable after-image journaling • Enable fast commit optimization • Enable the AIJ log server (ALS) as automatic • Enable system notification • Enable the AIJ backup server (ABS) as automatic • Set the database open mode to manual	Section 2.3
[2]	Open the master database manually (using the RMU Open command or the DBO Open command).	Section 2.4
[3]	Suspend the ABS process (automatic backup operations). You must suspend backup operations temporarily until you complete steps 4 through 12.	Section 2.5
[4]	Perform a full and complete backup operation on the master database. Oracle Corporation recommends that you perform an online, quiet-point backup operation (use the Quiet_Point qualifier).	Section 2.6
[5]	Save the after-image journal (.aij) configuration file to an output file (using the Output qualifier on the RMU or DBO Show After command).	Section 2.7
[6]	Create the standby database by using the Restore command to restore the most recent backup copy of the master database. The backup file that you restore ***must contain an exact copy*** of the master database. Also, be sure the options defined for the standby AIJ are ***identical*** to those for the AIJ on the master database. Oracle Corporation recommends using the Aij_Options qualifier on the Restore command.	Section 2.8
[7]	Open the standby database manually (using the RMU Open command or the DBO Open command).	Section 2.9

Table 2-1 Replication Startup Checklist (Continued)

Step	Procedure	Reference
[8]	Optionally, establish a default Hot Standby configuration quickly by using the Replicate After_Journal Configure command.	Section 2.10
[9]	Define the location of the after-image journal rollforward temporary work files by defining a logical name or configuration parameter.	Section 2.11
[10]	If more than one network transport protocol is used in your system configuration, specify the protocol to be used for Hot Standby communications.	Section 2.12
[11]	Start replication operations on the standby database.	Section 2.13
[12]	Start replication operations on the master database.	Section 2.14

Example 2-2 shows one way to implement the procedures outlined in Table 2-1 for an Oracle Rdb database running on an OpenVMS system.

Example 2-2 Sample Procedure to Start Replication Operations

```
$ SQL
ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
    RESERVE 5 JOURNALS
    ADD JOURNAL AIJ_1 FILENAME 'DISK1:[USER]aij_1.aij'
    ADD JOURNAL AIJ_2 FILENAME 'DISK1:[USER]aij_2.aij'
    ADD JOURNAL AIJ_3 FILENAME 'DISK1:[USER]aij_3.aij'
    ADD JOURNAL AIJ_4 FILENAME 'DISK1:[USER]aij_4.aij'
    ADD JOURNAL AIJ_5 FILENAME 'DISK1:[USER]aij_5.aij';

ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
        (FAST COMMIT IS ENABLED);
ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
        (LOG SERVER IS AUTOMATIC);
ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
        (NOTIFY IS ENABLED);
ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
        (BACKUP SERVER IS AUTOMATIC);
ALTER DATABASE FILE mf_personnel OPEN IS MANUAL;
EXIT;


$!
$ RMU/OPEN mf_personnel
$ RMU/SERVER BACKUP_JOURNAL/LOG SUSPEND mf_personnel

$ RMU/BACKUP/LOG/QUIET_POINT/ONLINE mf_personnel standby_personnel.rbf
$ RMU/SHOW AFTER_JOURNAL mf_personnel /OUTPUT=DISK1:[USER]aij_opt.dat

$! Edit the aij_opt.dat file to change the location (for example, node, disk,
$! and directory) of the backup files for the standby database.
```

```
$ RMU/RESTORE/NOCDD/LOG/NEW/DIR=DISK1:-
      /ROOT=DISK1:standby_personnel.rdb /AIJ_OPT=aij_opt.dat -
      standby_personnel.rbf

$!
$ RMU/OPEN standby_personnel
$!
$ DEFINE/SYS/LOG RDM$BIND_AIJ_WORK_FILE DISK1:[USER.HOTSTANDBY]
$!
$ RMU/REPLICATE AFTER_JOURNAL START standby_personnel -
    /MASTER_ROOT=ORANOD::DISK1:[USER]mf_personnel -
    /BUFFERS=1000 -
    /CHECKPOINT=100 -
    /GOVERNOR=ENABLED -
    /ONLINE -
    /LOG -
    /WAIT -
    /OUT=lrs_pid.out

$ RMU/REPLICATE AFTER_JOURNAL START mf_personnel -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]standby_personnel -
    /SYNCHRONIZATION=COLD -
    /QUIET -
    /CHECKPOINT=100 -
    /CONNECT_TIMEOUT=5 -
    /LOG -
    /WAIT -
    /OUT=lcs_pid.out

$!
$ RMU/REPLICATE AFTER_JOURNAL STOP mf_personnel
$ RMU/REPLICATE AFTER_JOURNAL STOP standby_personnel
```

Figure 2-1 provides a conceptual look at how the master and standby database are coordinated and synchronized during replication startup operations.

Figure 2-1 Replication Startup Operations Timeline

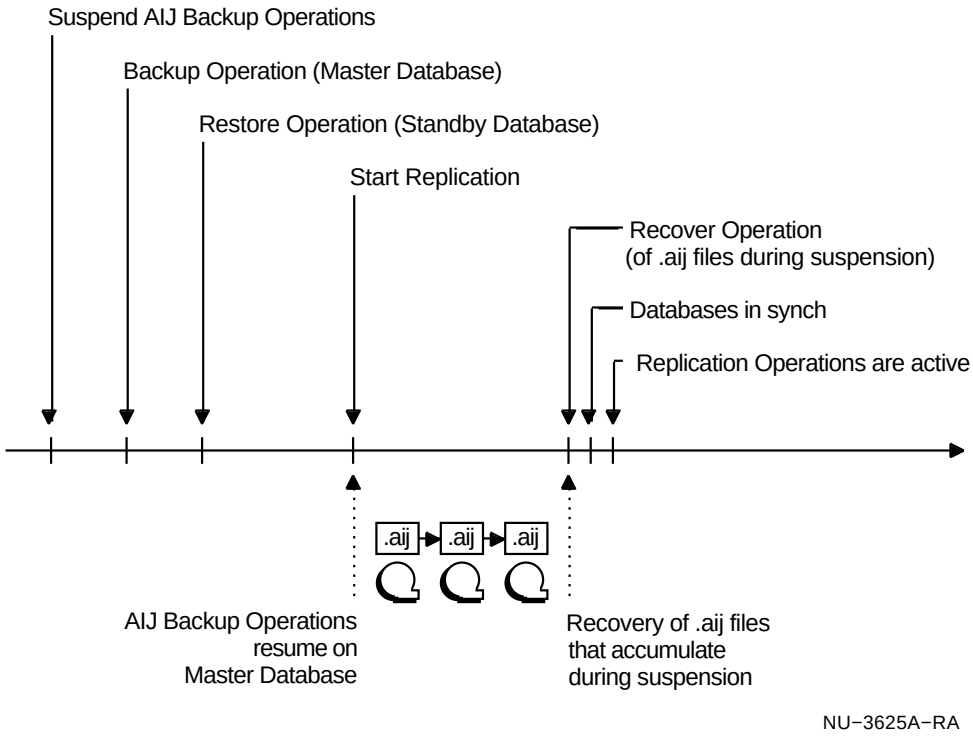

The following sections describe these steps in more detail. Each section includes a portion of the Oracle Rdb on OpenVMS syntax shown in Example 2-2, plus additional samples that show the commands to perform the same operations for Oracle Rdb on Digital UNIX and Oracle CODASYL DBMS.

2.3 Step 1: Set Up the After-Image Journal Options

Sections 2.3.1 through 2.3.8 describe how to set the after-image journal options to successfully set up the master database configuration. The options that you set now for the master database are eventually applied to the standby database when you create it with the Restore command in Step 6.

Reference: The following sections include examples. If you need more information about the SQL, RMU, or DBO syntax used, refer to the *Oracle Rdb7 SQL Reference Manual*, the *Oracle RMU Reference Manual*, or the *Oracle CODASYL DBMS Database Administration Reference Manual*, as appropriate.

2.3.1 Reserve After-Image Journal Slots

Reserving after-image journals slots is exactly the same for Hot Standby configurations as it is for nonreplicated database environments. However, you should give careful consideration to the number of slots you need, because the same number must be configured for the after-image journals on the master and the standby databases. The number of slots must be identical on both databases, or replication operations fail and return the fatal AIJSIGNATURE error message.

Reserve journals using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to reserve after-image journal slots:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/RESERVE=AFTER_JOURNAL=5 PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
   RESERVE 5 JOURNALS
```

Note: Oracle Corporation recommends that you reserve a minimum of five slots when you use the Hot Standby option. By providing several extra journals, you allow adequate time to preinitialize after-image journals on the standby database. Also, this allows you to add more journals during run time, and allows the database to create “emergency” after-image journals, if necessary.

2.3.2 Add After-Image Journals and Specify Journal Names

You must add journals and specify unique names for the after-image journal files. When you add journals, you can either specify a valid device and directory for the after-image journal (.aij) files or define a logical name or configuration parameter. Defining common symbols like logical names and configuration parameters can help you effectively manage the multiple files required by the master and standby directories. Storing shared information in a single, common place can also make failover easier.

Create and add after-image journals using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show the DBO syntax and SQL statements you use to add five after-image journals:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/AIJ_OPT=(CREATE,NAME=AIJ_1)/AFTER_JOURNAL=DISK1:[USER]AIJ_1.AIJ PARTS
$ DBO/MODIFY/AIJ_OPT=(CREATE,NAME=AIJ_2)/AFTER_JOURNAL=DISK1:[USER]AIJ_2.AIJ PARTS
$ DBO/MODIFY/AIJ_OPT=(CREATE,NAME=AIJ_3)/AFTER_JOURNAL=DISK1:[USER]AIJ_3.AIJ PARTS
$ DBO/MODIFY/AIJ_OPT=(CREATE,NAME=AIJ_4)/AFTER_JOURNAL=DISK1:[USER]AIJ_4.AIJ PARTS
$ DBO/MODIFY/AIJ_OPT=(CREATE,NAME=AIJ_5)/AFTER_JOURNAL=DISK1:[USER]AIJ_5.AIJ PARTS
```

Oracle Rdb on OpenVMS

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
    ADD JOURNAL AIJ_1 FILENAME 'DISK1:[USER]aij_1.aij'
    ADD JOURNAL AIJ_2 FILENAME 'DISK1:[USER]aij_2.aij'
    ADD JOURNAL AIJ_3 FILENAME 'DISK1:[USER]aij_3.aij'
    ADD JOURNAL AIJ_4 FILENAME 'DISK1:[USER]aij_4.aij'
    ADD JOURNAL AIJ_5 FILENAME 'DISK1:[USER]aij_5.aij';
```

2.3.3 Enable After-Image Journaling

Enabling after-image journaling is essential to using the Hot Standby option. Replication operations are designed to use after-image journaling to maintain an up-to-date copy of the master database at another location. The standby database is updated with the changes from the master database using a continuous rollforward operation, applying after-image journal records to the standby database.

Enable after-image journaling using the same syntax for Hot Standby configurations as for nonreplicated database environments. For Oracle CODASYL DBMS, you must explicitly enable journaling with the DBO Modify command. For Oracle Rdb databases, you have already enabled journaling as a part of the SQL statements (in Section 2.3.1) to reserve after-image journal slots and (Section 2.3.2) to add journals.

The following examples show how to enable after-image journaling:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/JOURNAL_OPTIONS=ENABLED PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
```

2.3.4 Enable Fast Commit Optimization

The fast commit optimization is important for Hot Standby configurations because it provides the ability to *redo* after-image journal logs to update the standby database. As the master database archives and ships its after-image journal logs to the standby database, the recovery operation uses the logs to roll forward data modifications to the standby database. The log rollforward server (LRS) process applies the updates and checkpoints the online database file information to update the time that the rollforward operation occurred.

Enable the fast commit option using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to enable the fast commit optimization:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/FAST_COMMIT=ENABLED PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
        (FAST COMMIT IS ENABLED);
```

2.3.5 Enable the AIJ Log Server As Automatic

You must make sure that the AIJ log server (ALS) startup mode is automatic so that the database monitor automatically starts the ALS process when you manually open the master and standby databases. The ALS process is active on both the master and standby databases while the standby database is being created and the databases are synchronized.

Enable the ALS process using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to enable the ALS process:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/SERVER=AFTER_JOURNAL=AUTOMATIC PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
            (LOG SERVER IS AUTOMATIC);
```

The Hot Standby software terminates the ALS process on the standby database as the LRS process takes over processing transactions against the standby database. (On OpenVMS systems, the message “AIJ Log Server terminated” displays on the standby operator console to notify you of the changeover.)

When replication operations are active, the ALS process applies data modifications to the after-image journal files for the master database and sends the same modifications to the LRS process on the standby database. However, the ALS process is reactivated if nodes (other than the master database) submit read-only transactions for processing on the standby database.

2.3.6 Enable System Notification (OpenVMS Only)

Oracle Corporation recommends using the after-image journal operator notification facility on OpenVMS systems. Because replication operations are performed using detached server processes, operator notification is the only proactive mechanism available for notification in the event of a server failure or replication termination. Each server process is also capable of maintaining an output file, similar to the monitor log file, which records information about the database replication operations.

Enable system notification using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to enable operator notification:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/JOURNAL_OPTIONS=(NOTIFY=OPER12) PARTS
```

Oracle Rdb on OpenVMS

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
            (NOTIFY IS ENABLED)
```

On Digital UNIX systems, use the RMU Show Users and Show Statistics commands to monitor the current database state.

2.3.7 Enable the AIJ Backup Server As Automatic

You should set up your Hot Standby system so that the AIJ backup server (ABS) process is enabled automatically. Automatically backing up the contents of after-image journal files is recommended for the Hot Standby option (and for any system that requires high availability).

Enable the ABS using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to enable the ABS:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/JOURNAL_OPTIONS=SPOOLER PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel JOURNAL IS ENABLED
            (BACKUP SERVER IS AUTOMATIC);
```

2.3.8 Set the Database Open Mode to Manual

You must manually open the master and standby databases prior to starting replication operations. By opening a database manually, you prevent users with insufficient privileges from opening and accessing the database while backup and restore operations proceed.

Set the database open mode to manual using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show how to set the database open mode to manual:

Oracle CODASYL DBMS

```
$ DBO/MODIFY/OPEN=MANUAL PARTS
```

Oracle Rdb on OpenVMS and Digital UNIX

```
$ SQL ALTER DATABASE FILE mf_personnel OPEN IS MANUAL;
```

2.4 Step 2: Open the Master Database

The master database must be open for access prior to starting replication operations. If the database is not already open and active on the node where you want to start replication operations, you must manually open the database using the DBO or RMU Open command.

Open the database using the same commands for Hot Standby configurations as for nonreplicated database environments. The following examples show how to open the master database:

Oracle CODASYL DBMS

```
$ DBO/OPEN PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -open mf_personnel
```

Oracle Rdb on OpenVMS

```
$ RMU/OPEN mf_personnel
```

2.5 Step 3: Suspend After-Image Journal Backup Operations

Even though you just enabled the ABS process (Section 2.3.7) for automatic after-image journal backups, you must temporarily suspend after-image journal backup operations. Suspending the backup operations prevents the ABS process from competing with the log catch-up server (LCS) process for access to the after-image journal.

Note: During replication startup, the LCS process synchronizes the databases by reading the master database after-image journal and forwarding the journal records to the standby database. Because this might take considerable time, an intervening backup operation could stall replication or prevent replication commands from operating properly.

To suspend backup operations, enter the RMU or DBO Server Backup_Journal Suspend command. The following examples show how to suspend the ABS process for Oracle Rdb and Oracle CODASYL DBMS databases.

Oracle CODASYL DBMS

```
$ DBO/SERVER BACKUP_JOURNAL/LOG SUSPEND PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -server backup_journal -log suspend mf_personnel
```

Oracle Rdb on OpenVMS

```
$ RMU/SERVER BACKUP_JOURNAL/LOG SUSPEND mf_personnel
```

Oracle Corporation recommends including the optional Log qualifier to display the status of the suspend operation. The log displays on your default output device.

Also, you can use the RMU or DBO Show Users command to show the current status of after-image journal backup operations. The following example shows this command for an Oracle Rdb database on an OpenVMS system:

```
$ RMU/SHOW USERS mf_personnel
Oracle Rdb V7.0-00 on node ORANOD 11-MAY-1996 08:22:53.98
  - monitor log filename is "DISK1:[USER.LOGS]RDMMON701_ORANOD.LOG;28"
  - 255 monitor buffers available [256 maximum]
database "DISK1:[USER.WORK.ALS]MF_PERSONNEL.RDB;1"
  - First opened 11-MAY-1996 08:22:26.65
  * database is opened by an operator
  * After-image backup operations temporarily suspended from this node
  - current after-image journal file is DISK1:[USER.WORK.ALS]AIJ2.AIJ;1
```

Suspending after-image journal backup operations is temporary and information about the suspension is not recorded in the database root file. For example, if the node on which you enter the Server Backup_Journal Suspend command fails, the ABS process automatically resumes after-image journal backup operations.

While the backup operations are suspended, you cannot use the manual after-image journal backup operation (using the Backup After_Journal command).

Backup operations are delayed until:

- Replication operations start

Note

For the ABS process to resume automatically, you must start replication operations on the same node where you entered the Server Backup_Journal Suspend command. If you start replication operations on another node, you must resume ABS backup operations manually using the Server Backup_Journal Resume command.

- Any of the following events occur:
  - You close the master database

- The master database node fails
- You enter the Server Backup_Journal Resume command

2.6 Step 4: Perform a Full Backup Operation

You must perform a full, online backup operation on the master database. You use the resulting backup file later, in Section 2.8, to create the standby database.

Oracle Corporation recommends that you perform online, quiet-point backup operations by including the Online and Quiet_Point qualifiers on the Backup command. You must perform an online backup operation if you opened the database manually as described in Section 2.4. (If you cannot perform an online backup operation, you must close the master database before you can perform an offline backup operation. Offline backup operations require exclusive access to the database.)

You perform backup operations using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show some sample backup commands:

Oracle CODASYL DBMS

```
$ DBO/BACKUP/MULTITHREAD/ONLINE/QUIET PARTS STANDBY_PARTS.DBB
```

Note: Oracle Corporation recommends you use the `/MULTITHREAD` qualifier.

Oracle Rdb on Digital UNIX

```
$ rmu -backup -quiet -online mf_personnel standby_personnel.rbf
```

Oracle Rdb on OpenVMS

```
$ RMU/BACKUP/LOG/QUIET/ONLINE mf_personnel standby_personnel.rbf
```

Before the database backup operation, a DBO or RMU Dump Header command displays the warning message shown in Example 2-3.

Example 2-3 Dump Header Information Before the Backup Operation

```
$ rmu -dump -header=hot_standby mf_personnel
    .
    .
    .

Database Parameters:
    Root filename is "/hotdisk/usr/mf_personnel.rdb"
    Hot Standby...
    - WARNING: Non-journalled database modifications have been made
    - Informational: Database backup has not been performed since AIJ modifications
    - Database is not currently being replicated
```

After the backup operation completes, the database root file header displays information similar to what is shown in Example 2-4.

Example 2-4 Dump Header Information After the Backup Operation

```
$ rmu -dump -header=hot_standby mf_personnel
    .
    .
    .

Database Parameters:
    Root filename is "hotdisk//usr/mf_personnel.rdb"
    Hot Standby...
    - Database is not currently being replicated
```

2.7 Step 5: Save the After-Image Journal File Configuration

You must save the after-image journal configuration information (in the AIJ options file) for the master database. Later, in Section 2.8, you restore the AIJ options file in conjunction with the master database backup file to create the standby database. Saving the after-image journal options to an output file is essential to ensure that the after-image journal configuration on the standby database is identical to the master database.

To save the after-image journal file, use the DBO or RMU Show After_Journal command with the Output qualifier on the master database. The following examples show some sample commands:

Oracle CODASYL DBMS

```
$ DBO/ SHOW AFTER_JOURNAL/OUTPUT=DISK1:[USER]AIJ_OPT.DAT PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -show after_journal -output=/hotdisk/usr/aij_opt.dat mf_personnel
```

Oracle Rdb on OpenVMS

```
$ RMU/SHOW AFTER_JOURNAL/OUTPUT=DISK1:[USER]aij_opt.dat mf_personnel
```

2.8 Step 6: Create the Standby Database

You create the standby database by restoring a backup copy of the master database and after-image journal options file using the RMU or DBO Restore command.

Your goal during this step is to create a standby database and environment that is identical to the master database. Differences between the databases or after-image journals between the master and standby sites could result in failure when you try to start replication operations.

To create the standby database:

1. Ensure that the after-image journal files on the standby database have the same physical characteristics as the after-image journal files on the master database. For example, configure the after-image journals for the master and standby databases with:
   - The same number of reserved after-image journal slots
   - The same number of active after-image journal files, in the same sequence (although it is acceptable for the journals to use different slot numbers)
   - The same physical block size allocation for each after-image journal file

   Note

   On OpenVMS systems, if the after-image journals configured for the standby database are allocated with a different physical block size (the file size both in blocks allocated and blocks used) from the after-image journals on the master database, replication operations might not be able to start. When this occurs, the Hot Standby software returns the AIJSIGNATURE error message. You can display and compare the physical block sizes of each after-image journal file using the DCL command Dir/Size=All.

2. Edit the AIJ options file to change the after-image journal file specification to reflect the correct location on the standby database configuration.
3. Use the DBO or RMU Restore command on the standby database node to:
   - Restore the backup file that you created in step 4 (Section 2.6) to restore an exact copy of the master database
   - Restore the AIJ options file that you created in step 5 (Section 2.7) to create an identical after-image journal configuration on the standby database
4. Recover (roll forward) all master database after-image journals that were backed up while you were performing the restore operation on the standby database.

The following examples show sample Restore commands, including the recommended qualifiers:

Oracle CODASYL DBMS

```
$ DBO/RESTORE/MULTITHREADED/NOCDD/NEW/DIR=DISK1:[USER] -
      /ROOT=DISK1:[USER]STANDBY_PARTS.ROO/LOG -
      /AIJ_OPT=AIJ_OPT.DAT STANDBY_PARTS.DBF
```

Note: Oracle Corporation recommends that you use the /MULTITHREADED qualifier.

Oracle Rdb on Digital UNIX

```
$ rmu -restore -new -dir=hotdisk \
> -root=/hotdisk/usr/standby_personnel.rdb -log \
> -aij_options=aij_opt.dat standby_personnel.rbf
```

Oracle Rdb on OpenVMS

```
$ RMU/RESTORE/NOCDD/NEW/DIR=DISK1:-
      /ROOT=DISK1:[USER]standby_personnel.rdb/LOG -
      /AIJ_OPT=aij_opt.dat standby_personnel.rbf
```

Note

The restored database cannot be used as the master database, the source database, or the backup file.

2.9 Step 7: Open the Standby Database

The standby database must be open for access before you start replication operations. If the database is not already open and active on the node where you intend to start replication operations, you must manually open the database using the DBO or RMU Open command.

Open a database using the same command for Hot Standby configurations as for nonreplicated database environments. The following examples show how to open the standby database:

Oracle CODASYL DBMS

```
$ DBO/OPEN STANDBY_PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -open standby_personnel
```

Oracle Rdb on OpenVMS

```
$ RMU/OPEN standby_personnel
```

2.10 Step 8: Configure Default Database Attributes

Prior to starting replication, you can predefine your Hot Standby database configuration using the Replicate After_Journal Configure command. This optional command allows you to:

- Establish default values for the qualifiers that you normally specify with the Replicate After_Journal Start command.

  When you are ready to start replication operations, you can enter the Replicate After_Journal Start command without specifying any qualifiers. Replication starts running using the previously defined qualifier values.

- Quickly start replication operations on both databases with a single command.

  You can enter the Replicate After_Journal Start command once on the master database to start replication on both the master and standby database nodes. There is no need to explicitly start replication operations on the standby database node.

  Reference: See the Replicate After_Journal Configure command reference section in Chapter 6 for information about the privileges required to start replication on the standby database from the master database node.

If you do not use the Replicate After_Journal Configure command to predefine your database configuration, the Hot Standby software starts replication operations using the default values defined in the master and standby database root files.

The following examples show the Replicate After_Journal Configure command:

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL CONFIGURE PARTS -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]STANDBY_PARTS -
    /SYNCHRONIZATION=COLD -
    /QUIET_POINT -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal configure mf_personnel \
> -standby_root=/hotdisk/usr/standby_personnel \
> -synchronization=cold \
> -quiet_point \
> -checkpoint=10 \
> -connect_timeout=1
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL CONFIGURE mf_personnel -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]standby_personnel -
    /SYNCHRONIZATION=COLD -
    /QUIET_POINT -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1
```

The first time you configure a database, you must include the Master_Root or Standby_Root qualifier to ensure that the standby database makes a network connection to the database you specify. The Master_Root or Standby_Root qualifier value is stored in the database root file. Therefore, it is optional on subsequent replication configuration commands.

Reference: See Chapter 6 for a complete description of the Replicate After_Journal Configure command. The default replication startup qualifier values are described in the reference section for the Replicate After_Journal Start command.

2.11 Step 9: Define the Location of the After-Image Work Files

Before you start replication operations on the standby database, you should define the location of the after-image rollforward work files using a logical name or configuration parameter from the following table:

Platform	Logical Name or Configuration Parameter
Oracle CODASYL DBMS	DBM$BIND_AIJ_WORK_FILE logical name
Oracle Rdb on Digital UNIX	RDB_BIND_AIJ_WORK_FILE configuration parameter
Oracle Rdb on OpenVMS	RDM$BIND_AIJ_WORK_FILE logical name

Oracle Corporation recommends that you use the BIND_AIJ_WORK_FILE logical names or configuration parameter to define the location. However, if you do not define the location of the temporary work files or if you specify an invalid location, the Hot Standby software locates the files in the home directory of the database monitor, as follows:

If . . .	Then . . .
You do not define the BIND_AIJ_WORK_FILE logical name or configuration parameter	The Hot Standby software uses the default directory for the database monitor as the location for the temporary after-image journal work files. As long as files can be created in the specified directory, there is no failure.
You define the BIND_AIJ_WORK_FILE logical name or configuration parameter, but the replication server is unable to create a work file because it cannot locate the work file directory	The Hot Standby software uses the default directory for the database monitor as the location for the temporary after-image journal work files, and returns an message stating there is an error in creating the work file or finding the directory.

Reference: Appendix A discusses how to define the logical names and configuration parameters that are specific to the Hot Standby software.

2.12 Step 10: Specify the Network Transport Protocol

The database software and the AIJSERVER network object server automatically facilitate communications between the master and the standby database using the default network transport protocol:

- On OpenVMS systems, DECnet for OpenVMS or DECnet/OSI is the default network
- On Digital UNIX systems, TCP/IP is the default network

For systems on which multiple network transport protocols are installed, you can define a logical name or configuration parameter to specify the network protocol (either DECnet or TCP/IP) that you want to use for communications between the master and standby database. The following table lists the logical names and configuration parameter:

Database and Platform	Logical Name or Configuration Parameter
Oracle Rdb for OpenVMS	RDM$BIND_HOT_NETWORK_TRANSPORT
Oracle CODASYL DBMS for OpenVMS	DBM$BIND_HOT_NETWORK_TRANSPORT
Oracle Rdb for Digital UNIX	RDB_BIND_HOT_NETWORK_TRANSPORT

Reference: On OpenVMS systems, you must define the logical names in the LNM$SYSTEM_TABLE logical name table. See Appendix A for more information about defining the logical names and configuration parameter.

The installation procedure creates and sets up the DBMAIJSERVER or RDMAIJSERVER account automatically for DECnet objects. (Section 1.5.2 describes how the network object is installed in detail.) You do not need to perform any additional tasks to install or invoke a DECnet network object.

The installation procedure does not set up the TCP/IP network object because so many different TCP/IP network transport protocols are available. For the TCP/IP network transport protocol, the installation procedure sets up the DBMAIJSERVER or RDMAIJSERVER file, but you must define a UCX port number, as shown in the following Oracle Rdb example:

```
$ CONFIG UCX AIJSERVER OBJECT
$ UCX SET SERVICE RDMAIJSRV
          /PORT=n
          /USER_NAME=RDMAIJSERVER
          /PROCESS_NAME=RDMAIJSERVER
          /FILE=SYS$SYSTEM:rdmaijserver_ucx.com
          /LIMIT=nn
```

In the example, set the variables *n* and *nn* as described in the following table:

Variable	Description
n	The variable *n* represents any available UCX port number. Use the UCX Show Service command to list all known services. Each service is assigned a port number (similar to a DECnet object number). Valid port numbers range from 1 to 65535. These numbers should be registered numbers. Because the DBMAIJSERVER (DBMAIJSRV) and RDMAIJSERVER (RDMAIJSRV) do not have registered numbers, you must pick an available (not currently in use) number.
nn	The variable *nn* represents the maximum number of image instances that can run simultaneously. In most cases, 10 (limit=10) is a reasonable number of simultaneous images. However, if you set the limit too low, additional connections to the network will fail.

To set the UCX port number for an Oracle CODASYL DBMS database, substitute "DBMAIJSRV" and "DBMAIJSERVER" for "RDMAIJSRV" and "RDMAIJSERVER" in the example.

After you set the UCX port number, you must enable the network service, as shown in the following example for an Oracle Rdb database:

```
$ UCX ENABLE SERVICE RDMAIJSRV
```

2.13 Step 11: Start Replication on the Standby Database

Note

This step is optional if you preconfigure your Hot Standby databases using the Replicate After_Journal Configure command (see Section 2.10). To configure your standby database attributes, continue reading this section. Otherwise, skip to Section 2.14.

You start replication on the standby database by entering the Replicate After_Journal Start command. Oracle Corporation recommends that you start replication operations on the standby database before the master database. Although it is not required that you initiate replication operations on the standby database first, it is advantageous to do so for several reasons:

- You can preset parameter values for when you subsequently start replication on the master database

- You can verify that the database restore operation on the standby database was successful
- Replication activity does not commence until the standby database creates the network connection and the databases are synchronized

Starting replication is an online operation that can occur while the standby database is open. However, replication operations cannot start on the standby database when these conditions exist:

- Database users are actively attached to the standby database
- Any storage area is inconsistent (for example, if you restore a storage area from a backup file but you have not rolled forward after-image journals to be consistent with respect to the rest of the database)
- Read / write transactions, including prestarted read / write transactions, are active on the standby database

  To disable prestarted transactions, define the appropriate logical name or configuration parameter from the following table:

Platform	Logical Name or Configuration Parameter
Oracle CODASYL DBMS	DBM$DISABLE_PRESTART_TXN logical name
Oracle Rdb on Digital UNIX	RDB_BIND_PRESTART_TXN configuration parameter
Oracle Rdb on OpenVMS	RDMS$BIND_PRESTART_TXN logical name

  Reference: Logical names and configuration parameter are described in the Oracle Rdb or Oracle CODASYL DBMS documentation.

- The DBMAIJSERVER or RDMAIJSERVER account has insufficient privileges to start replication on a remote node or to detect LRS process failure during a network connection

  Reference: Section 4.3 describes how to provide additional privileges to the DBMAIJSERVER or RDMAIJSERVER account.

When you enter the Replicate After_Journal Start command, you can specify a variety of command qualifiers to customize your standby database for the best performance and availability. However, the only qualifier you must specify (the first time you initiate replication operations) is the Master_Root qualifier; you can take the default values for all other qualifiers.

Note

Prior to starting replication, you can predefine the qualifier values using the Replicate After_Journal Configure command. This command allows you to establish default database configuration values without starting replication. Then, when you enter the Replicate After_Journal Start command (and do not include qualifiers), replication starts using the previously defined qualifier values. See Section 2.10 for more information.

To illustrate the richness of the Replicate After_Journal Start command, the following command examples include several qualifiers.

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL START STANDBY_PARTS -
    /MASTER_ROOT=ORANOD::DISK1:[USER]PARTS -
    /BUFFERS=1000 -
    /CHECKPOINT=100 -
    /GOVERNOR=ENABLED -
    /ONLINE -
    /LOG -
    /WAIT -
    /OUT=LRS_PID.OUT
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal start standby_personnel \
> -master_root=/ORANOD/hotdisk/usr/mf_personnel \
> -buffers=1000 \
> -checkpoint=100 \
> -governor=enabled \
> -online \
> -wait \
> -out=lrs_pid.out
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL START standby_personnel -
    /MASTER_ROOT=ORANOD::DISK1:[USER]mf_personnel -
    /BUFFERS=1000 -
    /CHECKPOINT=100 -
    /GOVERNOR=ENABLED -
```

```
    /ONLINE -
    /LOG -
    /WAIT -
    /OUT=lrs_pid.out
```

You must include the Master_Root qualifier the first time you enter the Replicate After_Journal Start command (unless you have preconfigured the Master_Root qualifier using the Replicate After_Journal Configure command). The Master_Root qualifier is necessary to ensure that the standby database makes a network connection to the master database you specify as the source of replication operations. The Master_Root qualifier is stored in the database root file. Therefore, it is optional on subsequent replication start commands.

Note

You must not include the Standby_Root qualifier. Any attempt to use the restored database as a master database fails.

Recommendation: By default, the asynchronous prefetch function is enabled on the standby database. For improved performance, Oracle Corporation recommends that you do not disable the detected asynchronous prefetch option.

Example 2-5 shows a portion of the Dump Header command output for the standby database. This command shows the database header after the Replicate After_Journal Start command was entered on the standby database.

Example 2-5 Dump Header Information for a Replicated Standby Database

```
$ RMU/DUMP/HEADER=HOT_STANDBY standby_personnel
    .
    .
    .
Database Parameters:
- Database has been configured as "replication standby"
    Master database is "ORANOD::DISK1:[USER]MF_PERSONNEL"
    Remote node name is "REMNOD"
    Replication last commenced on 15-MAY-1996 07:52:52.98
    Database replication is "offline"
    Server checkpoint interval is 100 messages
    Server timeout interval is 5 minutes
    Server buffer count is 1000
    Server 2PC transaction resolution is "commit"
```

2.14 Step 12: Start Replication on the Master Database

You can start replication while the master database is on line, open, and active. Online database access means that other database users can actively attach to the master database before, during, and after replication starts up. There is no need to close the master database or stop user access to the database when you initiate database replication. (Section 2.13 describes restrictions for starting replication on the standby database.)

You can start replication operations by entering the Replicate After_Journal Start command on the master database. At this time, automatic backup operations that you suspended in Section 2.5 are automatically resumed.

Note

If you used the Replicate After_Journal Configure command to preconfigure the master and standby database attributes (see Section 2.10), the Replicate After_Journal Start command starts replication operations on both the master and standby databases when you have sufficient privileges defined. Section 4.3 describes how to provide sufficient privileges to the DBMAIJSERVER or RDMAIJSERVER account.

When you specify the Replicate After_Journal Start command, you can include command qualifiers to customize your standby database for the best performance and availability. To illustrate the richness of the Replicate After_Journal Start command, the following command examples include a number of qualifiers:

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL START PARTS -
    /STANDBY_ROOT=REMNOD::STANDBY_PARTS -
    /SYNCHRONIZATION=COLD -
    /QUIET -
    /CHECKPOINT=100 -
    /CONNECT_TIMEOUT=5 -
    /LOG -
    /WAIT -
    /OUT=LCS_PID.OUT
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal start mf_personnel \
> -standby_root=/REMNOD/hotdisk/usr/standby_personnel \
> -synchronization=cold \
> -quiet \
> -checkpoint=100 \
> -connect_timeout=5 \
```

```
> -log \
> -wait \
> -out=lcs_pid.out
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL START mf_personnel -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]standby_personnel -
    /SYNCHRONIZATION=COLD -
    /QUIET -
    /CHECKPOINT=100 -
    /CONNECT_TIMEOUT=5 -
    /LOG -
    /WAIT -
    /OUT=lcs_pid.out
```

Note: You must include the Standby_Root qualifier the first time you enter the Replicate After_Journal Start command (unless you have preconfigured the Standby_Root qualifier using the Replicate After_Journal Configure command). This qualifier is necessary to ensure that the master database makes a network connection to the standby database you specify as the target of replication operations. The value of the Standby_Root qualifier is stored in the database root file. Therefore, it is optional on subsequent replication start commands.

Example 2-6 shows a portion of the output from a Dump Header command of the master database root file. The Dump Header command shows the database header after the Replicate After_Journal Start command was entered on the master database.

Example 2-6 Dump Header Information for a Replicated Master Database

```
$ RMU/DUMP/HEADER=HOT_STANDBY mf_personnel
    .
    .
    .
- Database has been configured as “replication master”
    Database is currently being replicated
    Master database is “MF_PERSONNEL.RDB;1”
    Database is local to this node (“ORANOD”)
    Replication commenced on 15-MAY-1996 09:36:05.16
    Database replication is “online”
    Server checkpoint interval is 100 messages
    Server timeout interval is 5 minutes
    Server buffer count is 256
    Server 2PC transaction resolution is “commit”
    Server database access is “read-write”
```

3
Replication Server Processing

This chapter describes how the server processes (introduced in Section 1.2.3) perform replication activities automatically and transparently to implement a Hot Standby database environment.

All of the discussions are based on two identical computing sites and databases connected by a network. Ideally, the standby site should be identical to the master site in configuration and resources. Oracle Corporation recommends identical configurations so that, if a failure occurs, you can switch application processing over to the standby site for an unlimited amount of time.

3.1 Server Processing During Replication Operations

When you enter the Replicate After_Journal Start command to start replication, the master database does not begin to actively ship database modifications to the standby database until these events occur:

- The standby database creates the network connection
- The master database attaches to the network connection
- The master and standby databases are synchronized with regard to committed transactions

The Hot Standby software meets these goals by employing special server processes. Servers on the master database automatically ship database modifications, in the form of after-image journal records, to servers on the standby database. The after-image journal records (that have already been committed on the master database) are rolled forward to the standby database.

Table 3-1 describes the Hot Standby server processes that perform replication operations.

Table 3-1 Server Functions

Server Name	Type of Server	Location	Function
Log catch-up server (LCS)	Database	Master database	Synchronizes the master and standby databases by automatically sending any after-image journal records (received since the last checkpoint interval) to the LRS process on the standby database. The LCS process is active only until the databases are synchronized. Then, the LCS process hibernates while replication processing activities for the master database are performed by the ALS process.
AIJ log server (ALS)	Database	Master database	Applies group commit buffers to the after-image journal for the master database and sends the same log data to the standby database. The ALS server does not begin shipping after-image journal records to the standby database until the standby database creates a network connection and the master and standby databases have been synchronized by the LCS process. Although the ALS process is optional for standard Oracle Rdb or Oracle CODASYL DBMS database operations, you must enable the ALS process when you use the Hot Standby option.
AIJSERVER	Network	Standby database	Receives after-image journal records from the master database and forwards them to the LRS process on the standby database. The AIJSERVER does not perform any database operations. Rather, it is invoked by the network software to act as a communications gateway between the master and standby databases.
Log rollforward server (LRS)	Database	Standby database	Receives after-image journal records (in group commit buffers) via the AIJSERVER from the master database and rolls the data modifications forward to the database and after-image journal on the standby database. Once invoked, the LRS process automatically goes into replay mode to initialize its transaction state information and roll forward any transactions that have been committed on the master database but have not been applied to the standby database.

The following sections provide a detailed description of how the servers perform replication operations.

3.2 Replication Servers

The following sections describe how the Hot Standby software implements servers to perform database replication operations. The discussions assume that you have prepared the master and standby database sites according to the checklist and recommendations in Chapter 2.

3.2.1 Setting Up Network Communications

When you invoke the Replicate After_Journal Start command to start replication operations, you indicate that you want to use network communications by including a node name when you start replication operations on the master and standby databases. For example:

```
$ RMU/REPLICATE AFTER_JOURNAL START mf_personnel -
     /STANDBY_ROOT=REMNOD::STANDBY_PERSONNEL
```

The Hot Standby servers parse the database name you supply to separate the node name (if any) from the database root file specification. Parsing the specification in this way identifies whether the database connection is local or remote, as described in the following table:

If the database specification . . .	Then the Hot Standby software uses . . .
Includes a remote node name	Remote (DECnet or TCP/IP) network communications
Includes the local node name	Local interprocess communications
Does not include a node name	Local interprocess communications

Databases located on the same node communicate directly using local interprocess connections. Databases located on separate nodes communicate using DECnet for OpenVMS, DECnet/OSI, or TCP/IP network transport protocols.

3.2.2 The ALS Process in a Nonreplicated Environment

Figure 3-1 shows typical activity and server processing for a database *before* you enable Hot Standby replication operations. The database configuration includes the ALS process that you must enable on the master (and the standby) database before starting replication activities.

Figure 3-1 Database Activity Before Replication Activities

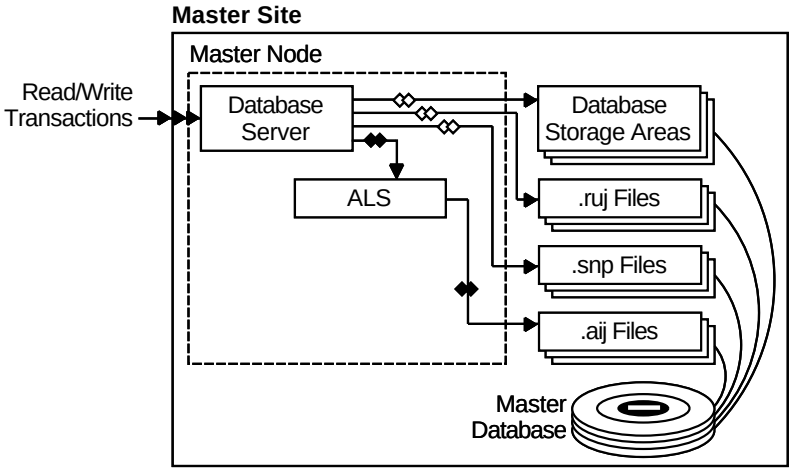

◆◆ = Group Commit Buffer
◇◇ = Database Updates

NU−3622A−RA

In the nonreplicated database environment shown in Figure 3-1, the ALS process is a dedicated server that applies log data (in a group commit buffer) to the after-image journal on behalf of all user processes. You can think of the ALS process as an after-image journal server that optimizes performance by eliminating the I/O bottlenecks caused when multiple user processes attempt simultaneous access.

Figure 3-1 is provided so you can compare the ALS in a nonreplicated environment to its function in a replicated environment. In a replicated database environment, the ALS process extends its functionality to send log data to the after-image journals for both the master and the standby databases.

3.2.3 Starting the LRS Process on the Standby Database

When replication operations start on the standby database, the database monitor invokes the special log rollforward server (LRS) process. The purpose of the LRS process is to initialize the transaction-state information, connect to the network, and roll forward committed transactions to the standby database.

Although you can initiate replication operations on either the master or the standby database, Oracle Corporation recommends that you start replication on the

standby database first. (See Section 2.13 for a complete discussion of this recommendation.)

3.2.3.1 Invocation

The Hot Standby software allows only one LRS process for a standby database. Figure 3-2 shows how the LRS process fits into the standby system configuration.

Figure 3-2 Invoking the LRS Process on the Standby Database

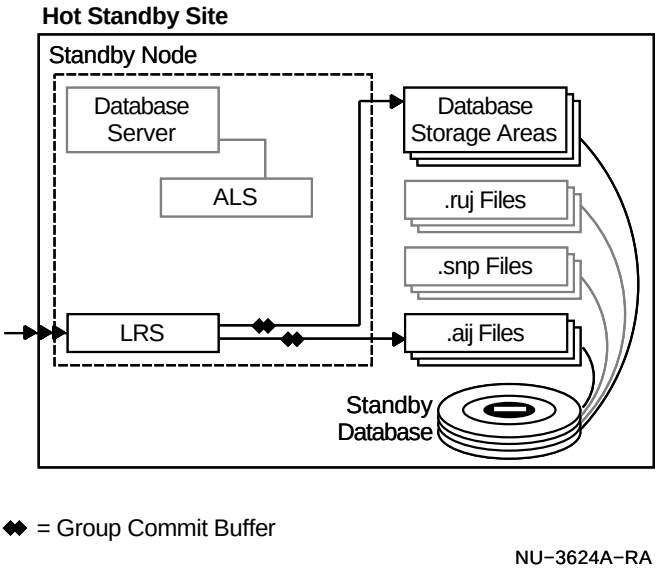

Once invoked, the LRS process automatically:

1. Initializes its transaction state information
2. Identifies the master database and makes a network connection (for eventual use by the AIJSERVER process)
3. Replays any journaled transactions recorded since the last checkpoint

Note

The LRS process does not start if the standby database:

- Has unjournaled modifications
- Has been modified by a process other than the LRS process

In Figure 3-2, note that the ALS process, the database server, the recovery-unit journal (.ruj) files and snapshot (.snp) files exist, but are not used when replication operations are active on the standby database.

The ALS process is active on both the master and standby databases while the standby database is being created and the databases are being synchronized. Then, the Hot Standby software terminates the ALS process on the standby database and the LRS process takes over processing transactions. (On OpenVMS systems, the message "AIJ Log Server terminated" displays on the standby operator console.)

When replication operations are active, the ALS process applies data modifications to the after-image journal files for the master database and sends the same modifications to the LRS process for processing on the standby database. In a cluster configuration, the ALS process is reactivated if nodes (other than the master database) submit read-only transactions for processing on the standby database.

3.2.3.2 Replication Startup Operations

Before replication activities start, the LRS process must:

1. Identify and attach to the master database from the remote standby location
2. Set up a network link between the servers on remote databases
3. Synchronize the standby database with the master database

Replication activity cannot begin until these steps complete successfully, and the standby database is transactionally up to date with the master.

Table 3-2 describes how the LRS process functions during the phases of replication startup.

Table 3-2 LRS Replication Startup Phases

Phase	LRS Function	Comments
Startup[1]	1. Identifies the master database and establishes a network connection. 2. Verifies that the after-image journal is enabled. Note that the LRS process does not verify the state of the after-image journals (for example, activity, fullness, transaction sequence numbers) until the replication restart phase. 3. Initializes the database transaction-state information.	To obtain exclusive access to the standby database, the LRS process acquires these locks: • Quiet-point lock—Ensures that all database verification operations and active transactions (such as the backup operation) finish before replication operations begin • After-image journal lock—Prevents processes other than the LRS process from making database modifications

Table 3-2 LRS Replication Startup Phases (Continued)

Phase	LRS Function	Comments
Network Connection	Establishes a network connection for eventual use by the AIJSERVER network object server process.	The AIJSERVER network object server is not active during the network connection phase; the AIJSERVER is invoked by the LCS process when replication is started on the master database (see Section 3.2.4).
Replication Restart[2]	Restarts replication operations by reading the after-image journals and rolling forward any committed transactions to the (restored) standby database.	Replication restart processing occurs when you first start replication operations, or after a system or database failure on a database that was being replicated.
Database Synchronization	Communicates with the LCS process across the network via the AIJSERVER network object server.	The LRS process actively receives after-image journal messages from the LCS process on the master database. The purpose is to make the standby database transactions consistent with regard to the master database.

[1] The LRS process stops running if storage areas or data pages in the standby database are corrupt or inconsistent, or if a process other than the LRS process attempts to update the after-image journal file.

[2] If you included the Wait qualifier on the Replicate After_Journal Start command, database operations do not proceed until the LRS process finishes this phase.

During the startup phase described in Table 3-2, local users and applications cannot access the standby database. As soon as database replication becomes active, local users and applications can access the standby database for read-only purposes. Read/write transaction activity on the standby database is not allowed after you restore the most recent backup copy of the master database and start database replication.

You can enter the Show Users command at any time to display the state of the standby database. See Example 7-2 for sample output from a Show Users command that was entered on the standby database.

3.2.4 Starting the LCS Process on the Master Database

When you start replication operations on the master database:

1. The database monitor invokes a special log catch-up server (LCS) process.

   The LCS process brings the standby database up to date with respect to the master database.

2. The LCS process waits for the AIJSERVER process to connect to the log rollforward server (LRS) process on the standby database.

The AIJSERVER process acts as a communications gateway on the standby database.

Figure 3-3 shows how the LCS and AIJSERVER processes fit into the Hot Standby configuration, and how the LCS sends group commit buffers via the AIJSERVER to the LRS process on the standby database.

Figure 3-3 Invoking the LCS Process on the Master Database

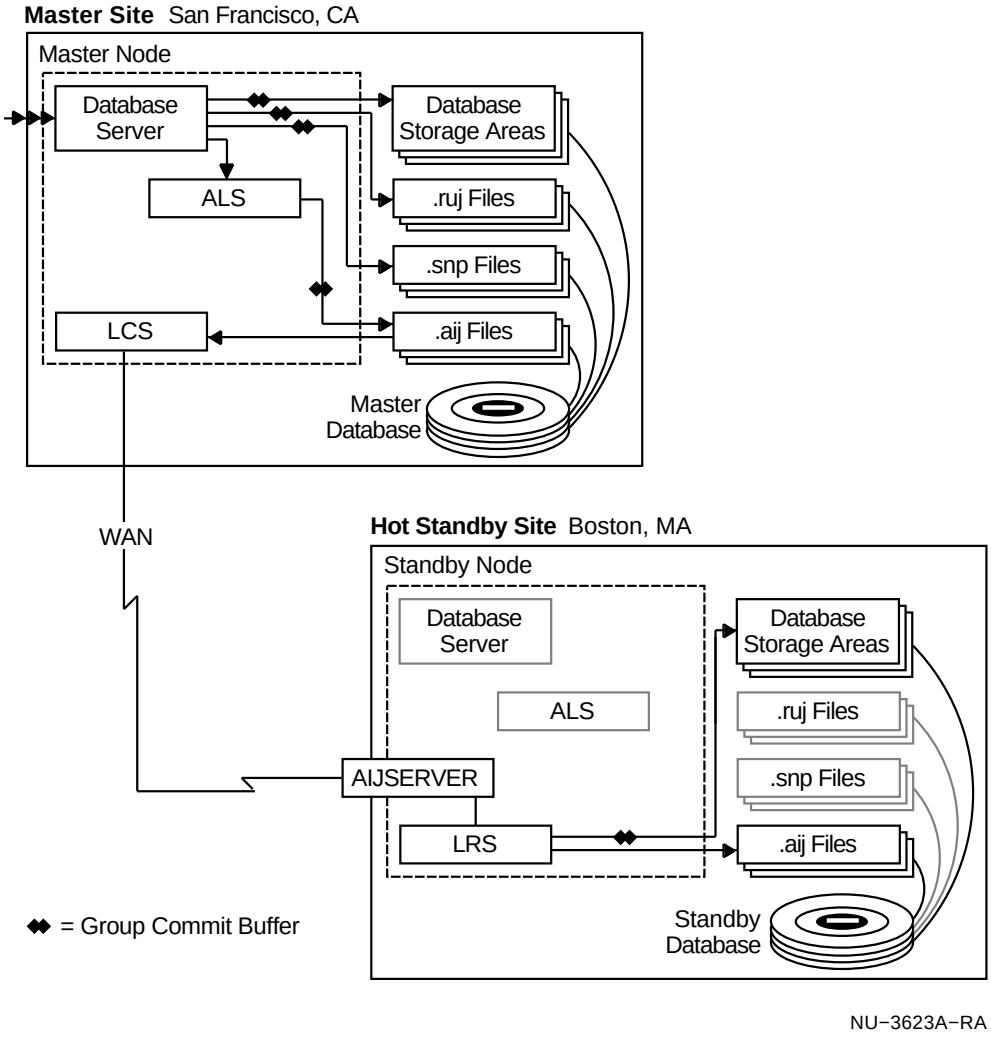

When you start database replication, the master database must be on line, open, and active. Note that:

- Read / write transaction activity can continue as you start database replication
- The ALS process and database server on the master database are still actively updating database storage areas, recovery-unit journals, snapshot files, and after-image journals

  The Hot Standby software terminates the ALS process on the standby database and the LRS process takes over processing transactions. (On OpenVMS systems, the message "AIJ Log Server terminated" displays on the standby operator console.)

  When replication operations are active, the ALS process applies data modifications to the after-image journal files for the master database and sends the same modifications to the LRS process for processing on the standby database. In a cluster configuration, the ALS process is reactivated if nodes (other than the master database) submit read-only transactions for processing on the standby database.

Invoking the LCS Process

The Hot Standby software allows only one LCS process to exist on the master database. The LCS process:

1. Waits until the LRS process (on the standby database) makes a network connection.
2. Attaches to the network and invokes the AIJSERVER network object server.
3. Automatically *restart*s the process of shipping after-image journal records that have not been sent since the last recorded checkpoint. The LCS process ships the records across the network, by means of the AIJSERVER process, to the LRS process on the standby database.
4. Waits while the LRS process on the standby database *replays* the after-image journal records to the standby database. If any of the transactions cannot complete successfully, then the standby database is *not recoverable* and replication operations shut down.

   While the LCS process is active, the ALS process and database servers on the master database can continue processing database modifications. Users and applications can continue to access the master database for read / write transactions while the LCS and LRS processes synchronize the databases. The LCS restart operation occurs parallel to, but independently of, ongoing database updates that are being performed by the ALS process.

AIJSERVER Network Object Server

During replication operations, the AIJSERVER process acts as a communications gateway on the standby database. When you enter the Replicate After_Journal Start command to start database replication, the log catch-up server (on the master database) waits for the AIJSERVER process to connect to the log rollforward server (LRS) process on the standby database. By default, the LCS process waits five minutes for the replication operations to start. However, you can specify a time limit using the Nowait and/or the Connect_Timeout qualifiers.

Oracle Corporation recommends that you use one or both qualifiers to limit the amount of time you allow for replication activity to begin. In effect, the Connect_Timeout command specifies the number of minutes that the log catch-up server (on the master database) waits for the AIJSERVER process to connect to the log rollforward server (LRS) process on the standby database.

The default setting, Connect_Timeout=5, gives the AIJSERVER a maximum of 5 minutes to connect to the LRS process on the standby database. If the AIJSERVER cannot make a connection within the specified time, the startup operation fails.

The value you specify for the Connect_Timeout qualifier affects the LCS process only. The ALS processes do not wait because the LCS process notifies the ALS process that the master and standby databases are properly synchronized and replication can commence. The Connect_Timeout interval for the ALS process is always 0 minutes.

Note

When you start database replication on the standby node, the ALS process on that node shuts down when the LRS process starts up. This is normal; the ALS process automatically restarts when database replication is shut down.

3.2.5 Server Processing When Replication Operations Are Active

When the master and standby databases are synchronized (with respect to the last checkpoint interval recorded in the after-image journal log), some of the server functions change. After the databases are synchronized:

- The LCS process on the master node no longer sends after-image journal records across the network. Instead, the LCS process periodically monitors the network, watching for gaps in the TSN sequence number. If a gap occurs, the LCS processing restarts the functions described in Section 3.2.4.

- The ALS process on the master node extends its capabilities to include forwarding all data modifications to the standby node so that the same group commit buffer is applied to the after-image journals for both the master and standby databases.

Note

You cannot terminate the ALS process on the master database while replication operations are active. If you attempt to use the Server After_Journal Stop command to stop the ALS process while replication operations are active, an error message is returned and replication processing continues.

- The ABS process on the master node is re-enabled so that after-image journals are backed up automatically.
- The AIJSERVER process on the standby node continues to receive after-image journal logs from the ALS process on the master node.
- The LRS process on the standby node continues to receive buffers from the AIJSERVER process and roll forward the after-image journal records to the standby database.

Figure 3-4 shows the active replication operations in which the ALS process on the master node sends transactions to both the master database and through the AIJSERVER process to the standby database.

Figure 3-4 Server Processing During Replication Operations

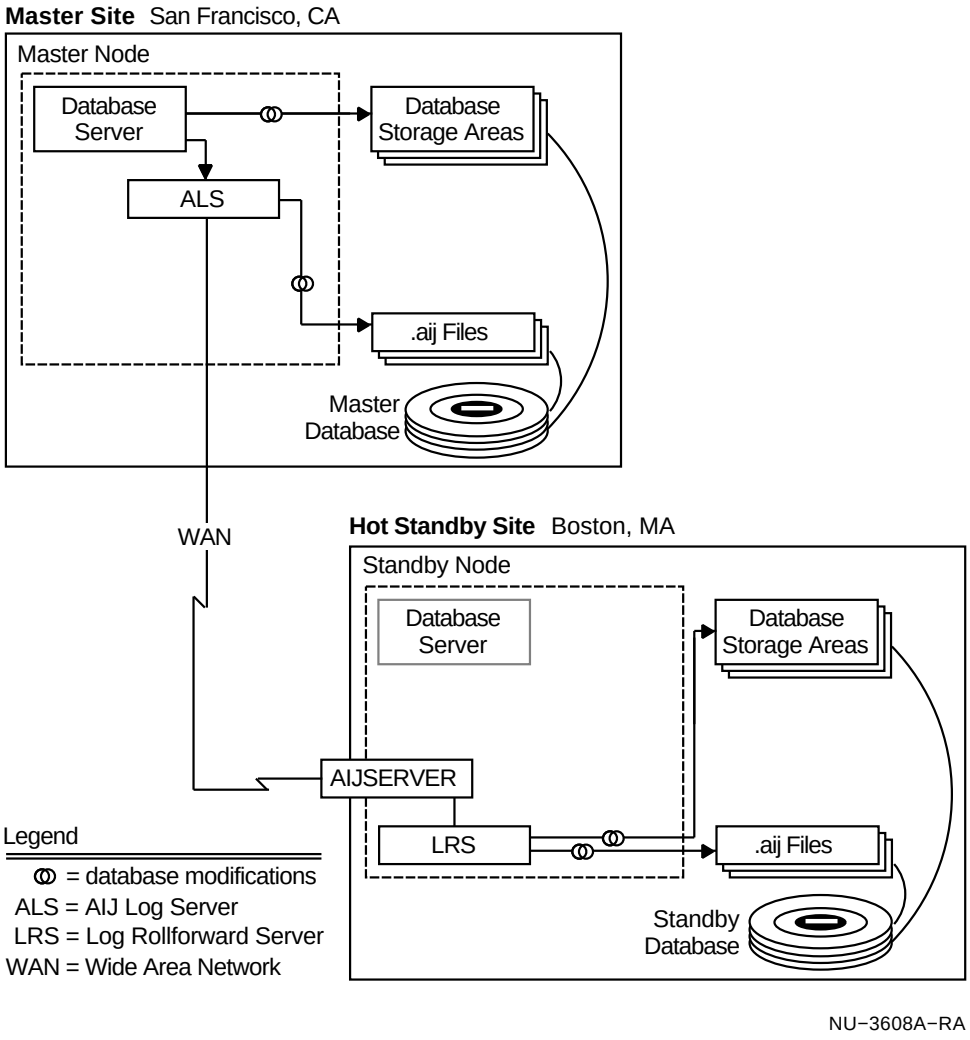

4
Ongoing Database Management

This chapter describes synchronizing the master and standby databases, and provides information about typical management tasks you might need to perform while replication operations are active.

4.1 Coordinating Database Synchronization

Once replication operations start, the main function of the Hot Standby servers is to keep the databases synchronized transactionally. All database consistency and integrity verification occurs on the master database. The Hot Standby software does not replicate data on the standby database unless it has already been validated on the master database. The Hot Standby software processes transactions in the following sequence:

1. Apply the database updates to the master database and after-image journal by means of the AIJ log server (ALS).
2. Ship the same database updates to the log rollforward server (LRS) on the standby database.
3. Commit the transactions on the master database.
4. Commit the transactions on the standby database.

It is possible for processing on the master database to outpace the rollforward operations on the standby database. For example, data partitioning, locks on the master database, caching characteristics, faster processors, more memory, or faster disks can make the performance of the master database exceed the ability of the standby database to roll the database changes forward.

Although the goal is to keep the databases identical at both database sites, the transactions applied to the standby database can sometimes lag behind the master database. This lag occurs because the standby database is open for read-only access and in a constant state of *recovery*. Thus, the standby database is updated with the changes from the master site in the form of *redo* operations after the transactions are already committed on the master database. Therefore, the standby database is in a

continuous rollforward mode of operation, applying after-image journal logs shipped across the network from the master database.

You can specify the degree of transaction lag when you start replication operations. The amount of synchronization can range from a fully asynchronous “fire-and-forget” model with a potentially large lag time, to a synchronous “transaction commit” level in which data modifications are made to both databases at the same time and in a coordinated manner.

In general, achieving a high degree of synchronization results in excellent recoverability from a master database failure, but it can occasionally affect the master database performance during periods of heavy I/O operations if the master database must wait for data modifications to commit on the standby database.

When you start replication operations, you can either:

- Explicitly specify the desired level of synchronization
- Enable the Hot Standby replication governor function to automatically and dynamically adjust the database synchronization according to the current demands on the databases

Chapter 6 describes the replication commands for controlling database synchronization explicitly or using the replication governor.

4.1.1 Performance and Availability Trade-Offs

You can control the amount of lag time between the databases, but first you need to determine the level of database consistency appropriate for your business. By weighing your requirements for availability against user demands for response time and performance, you can determine how tightly you want to synchronize the master and standby databases.

The Hot Standby servers support the following modes of synchronization:

- Commit

  This is the most tightly consistent and synchronous level of replication. The commit mode requires that the standby database be identical to the master database. To guarantee that the two databases are always consistent, the Hot Standby product employs this mode using a two-phase commit protocol. Performance can be affected because this mode can consume extensive system resources.

The following diagram shows the transmissions that occur for the commit mode of synchronization.

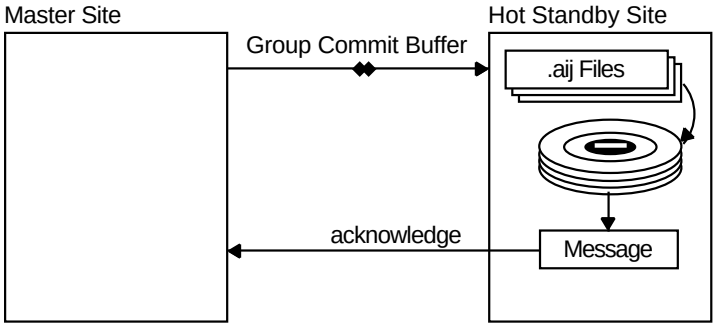

When the standby database receives the group commit buffer from the master database, the servers on the standby database:

- Write database modifications to the after-image journal on the standby system
- Apply the after-image journal modifications to the standby database
- Send a message back to the master database acknowledging the successful commit of the transaction

• Hot

The hot mode of synchronization brings the standby database extremely close to being transactionally identical to the master database. After-image journal records in transit are received and committed. Some restart processing may be required to synchronize the databases.

The following diagram shows the transmissions that occur for the hot mode of synchronization.

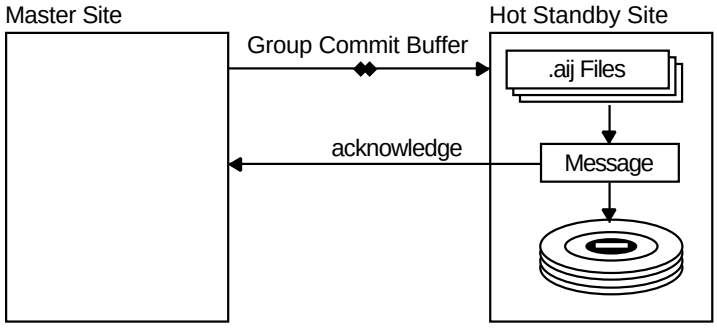

NU−3627A−RA

When the standby database receives the group commit buffer from the master database, the servers on the standby database:

- Write the database modifications to the after-image journal on the standby system
- Send a message back to the master database before applying the transaction to the standby database

• Warm

The warm mode of synchronization brings the standby database transactionally close to the master database, but the databases are not immediately identical.

The following diagram shows the transmissions that occur for the warm mode of synchronization.

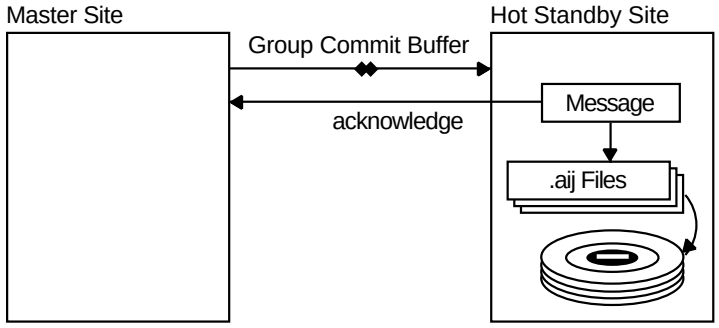

NU−3628A−RA

When the standby database receives the group commit buffer from the master database, the servers on the standby database:

- Send a message back to the master database before applying the transaction to either the after-image journal or the standby database.
- After-image journal records in transit might not be committed on the standby database.

In failover situations, it is possible for transactions to be committed on the master database, but the same transaction might be rolled back on the standby database. If this occurs, the Hot Standby software automatically resynchronizes the databases when you restart replication operations.

- Cold

This is the most loosely consistent and asynchronous mode of replication. Asynchronous processing updates the master database and then asynchronously propagates the changes to the standby database. During normal operations, the standby database lags behind the master database to some degree. With this mode, it is possible for the standby database to become identical to the master database only if updates to the master database stop or are stalled (such as through a database quiet point).

The following diagram shows the transmissions that occur for the cold mode of synchronization.

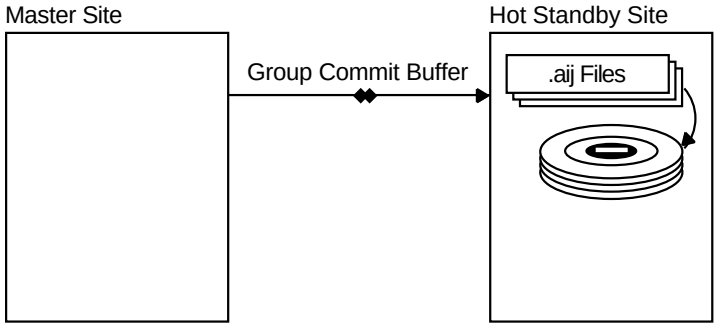

NU−3629A−RA

When the standby database receives the group commit buffer from the master database, the servers never return a message acknowledging the receipt of the information. The standby database is not immediately recoverable transactionally with respect to the master database. In failover situations, it is possible to have transactions rolled back on the standby database that were committed on the master database, but the Hot Standby software automatically synchronizes the database when you restart replication operations.

Additionally, the update performance of the servers on the master database can greatly exceed that of the standby database to roll forward changes, depending the synchronization mode that you choose.

4.1.2 Controlling Database Consistency

The Hot Standby option provides the following qualifiers for the Replicate After_Journal Start command to help you control the synchronization between the databases:

- Synchronization qualifier—Allows you to explicitly specify the synchronization mode as Commit, Hot, Warm, or Cold.
- Governor qualifier—Enables the replication governor process that automatically and dynamically adjusts the synchronization mode.

Because business environments might have different database consistency requirements within the same business day, Oracle Corporation recommends that you enable the replication governor to automatically determine the optimum synchronization mode for the current environment. The replication governor causes the serv-

ers to dynamically check and change the synchronization mode to balance performance against database transaction synchronization.

Also, you can define logical names or configuration parameters to specify the synchronization mode, or to enable or disable the replication governor. See Appendix A for information about the BIND_HOT_DATA_SYNC_MODE and the BIND_LRS_GOVERNOR_ENABLED logical names and configuration parameters.

4.1.3 Balancing Performance with Database Synchronization

When you set Synchronization=Cold and your master database is in a cluster configuration, the Hot Standby software potentially updates the master database faster on multiple nodes than the Hot Standby software is able to roll forward transactions to the standby database. However, this is balanced by the fact that the standby database:

- Does not write snapshot records
- Has fewer locking requirements
- Has no locking conflicts
- Maximizes buffer allocation

4.2 Database Management

The following sections provide information and recommendations regarding routine database management operations that you might perform while replication operations are active.

4.2.1 Enabling Read-Only Processing on the Standby Database

Once database replication operations are active, local users and applications can access the standby database for read-only purposes. To attach to the standby database, users must enter the following Database Query utility (DBQ) command or SQL statement:

Oracle CODASYL DBMS

```
$ DBQ READY CONCURRENT READ ONLY
```

Oracle Rdb on Digital UNIX and OpenVMS

```
$ SQL ATTACH 'FILENAME standby_personnel'
SQL> DECLARE TRANSACTION READ ONLY;
```

If you attempt to perform a read / write or batch-update transaction, the Hot Standby software returns an error message but replication operations continue processing transactions. For example:

```
SQL> SET TRANSACTION READ WRITE;
%RDB-E-BAD_TPB_CONTENT, invalid transaction parameters in the trans-
action parameter block (TPB)
-RDMS-F-ALSACTIVE, Database replication is active

SQL> SET TRANSACTION BATCH UPDATE;
%RDB-E-BAD_TPB_CONTENT, invalid transaction parameters in the trans-
action parameter block (TPB)
-RDMS-F-ALSACTIVE, Database replication is active
```

4.2.2 Moving a Database

You cannot move the standby database while replication is active. The following example shows the message that is returned when an RMU Move command is entered on an active standby database:

```
$ RMU/MOVE/ONLINE standby_personnel
%RMU-F-HOTNOONLINE, attempt to access standby database opened for
exclusive access
```

You can perform a Move operation on the master database while replication operations are active.

4.2.3 Closing a Database

When database replication is active, you cannot close the database on:

- The master node where the LCS process is running
- The standby node where the LRS process is running

However, you can close the database on any other participating node while replication is active. Closing nodes that do not have replication servers running does not affect the state of replication operations.

Note: You must stop replication before closing the database on either the master or standby nodes. See Chapter 5 for more information.

4.2.4 Performing Backup Operations

The database backup operations and after-image journal backup operations are not significantly different when replication operations are active. The following subsections provide recommendations regarding backup operations.

Database Backup Operations

Backing up the master database is not changed. You can perform all backup operations on the Hot Standby master database that you can perform in a standard database environment. However, because the standby database does not support

snapshot area updates, you cannot perform an online database backup operation on the standby database.

Note

Do not enable the Commit to Journal optimization on the standby database. This results in the backup operation marking the standby database as modified and halting all replication operations.

After-Image Journal Backup Operations

The AIJ backup server (ABS) is automatically re-enabled on the master database when you start replication, and backup operations proceed automatically as after-image journals fill up with data modifications. The ABS is enabled on the standby database to pre-initialize overwritten after-image journals.

Note the following about ongoing after-image journal backup operations:

- When replication operations are active, you cannot convert the after-image journaling technique from using a single extensible journal to using multiple fixed-size journals, or vice versa.
- If you have enabled the after-image journal overwrite capability, an after-image journal file that is waiting to be backed up or restored on the standby database will not be overwritten. However, the after-image journal file can be overwritten once its information is backed up or restored on the standby database node.
- You cannot suspend the ABS process on the standby database when replication operations are active. If the ABS process is suspended, the LRS process on the standby database tries to restart the ABS process. If the LRS process is unsuccessful, replication performance on the standby database degrades when a new after-image journal is selected.
- If an after-image journal switchover operation is suspended when replication operations are occurring, you must back up one or more of the modified after-image journals to add a new journal file.
- On OpenVMS systems, use a concealed system logical name when you specify the device for the default backup file name. Using concealed system logical names can help you avoid subsequent bugcheck errors due to invalid device names (such as when failover occurs making the standby database into a master database, and the ABS process is unable to translate the file specification of the after-image journals).

- Although it is possible to perform after-image journal backup operations on both the master and standby databases, you should perform backup operations only on the master database because:
  - You must perform backup operations on the master database
  - The after-image journals on the standby database are automatically overwritten by the LRS process

Manual After-Image Journal Switchover

Do not attempt to manually switch after-image journals on the standby database (using the Switch qualifier on the Set After_Journal command) when replication operations are active. The Hot Standby software returns the `ALSACTIVE` error message to indicate that this is an invalid operation on the standby database.

You can manually switch after-image journals on the master database at any time, including when replication operations are occurring.

4.3 Network Management

The following sections discuss network privileges, defining the minimum and maximum buffer sizes for transmitting after-image journal packets over the network, and managing network security. Refer to Sections 1.5 and 2.12 for additional information about network communications in a Hot Standby configuration.

4.3.1 DBMAIJSERVER and RDMAIJSERVER Privileges

When you install the Hot Standby software, the installation procedure automatically gives the AIJSERVER image netmbx and tmpmbx privileges. These privileges are sufficient for most Hot Standby configurations.

However, to perform the following functions, you must use the OpenVMS Authorize utility to provide additional privileges to the DBMAIJSERVER or RDMAIJSERVER accounts:

- On OpenVMS systems, if you have preconfigured your Hot Standby environment using the Replicate After_Journal Configure command and you plan to start replication operations remotely (for example, if you want to start replication on the standby database from the master database node), modify the DBMAIJSERVER or RDMAIJSERVER account to grant the SYSPRV privilege.
- On OpenVMS systems, to detect LRS process failures without waiting for the timeout period, modify the DBMAIJSERVER or RDMAIJSERVER account to grant GROUP and WORLD privileges.

If you do not grant these privileges, replication operations wait for a network connection to be made for the period of time specified by the Connect_Timeout qualifier or the BIND_HOT_NETWORK_TIMEOUT logical name. The default timeout period is 5 minutes.

Insufficient privileges can prevent the DECnet or TCP/IP network software from accessing the image.

4.3.2 Maximizing Network Transmissions

Data is exchanged over the network in a group commit buffer. Although the maximum after-image journal group commit input/output (I/O) size is 127 blocks, the Hot Standby software limits the after-image journal group commit I/O size to 96 blocks. Note that one block is used to package the buffer for transmission across the network. Figure 4-1 shows how the group commit buffer is packaged and applied to the after-image journals for both databases.

Figure 4-1 Packaging a Group Commit Buffer

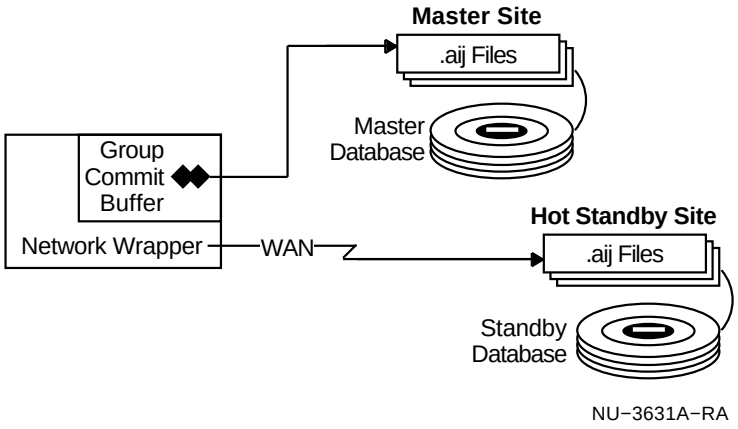

Because most after-image journal I/O operations are rarely as large as 96 blocks, this size is sufficient for most applications. When packaged into an after-image journal message buffer, the maximum message size fits within the 96-block limit.

You can adjust the minimum and maximum size of the group commit buffers and packet transmissions using the logical names or configuration parameters shown in the following table:

Database and Platform	Name
Oracle Rdb for OpenVMS	RDM$BIND_AIJ_IO_MAX RDM$BIND_AIJ_IO_MIN
Oracle Rdb for Digital UNIX	RDB_BIND_AIJ_IO_MAX RDB_BIND_AIJ_IO_MIN
Oracle CODASYL DBMS for OpenVMS	DBM$BIND_AIJ_IO_MAX DBM$BIND_AIJ_IO_MIN

The BIND_AIJ_IO_MAX logical name or configuration parameter allows you to override the maximum AIJ group commit I/O buffer size (the default buffer size is 127 blocks). The BIND_AIJ_IO_MIN logical name or configuration parameter allows you to override the minimum AIJ group commit I/O buffer size (the default buffer size is 8 blocks).

Reference: The *Oracle Rdb7 Guide to Database Performance and Tuning* provides more information about these logical names and configuration parameters for Oracle Rdb databases.

4.3.3 Securing Data on the Network

Network security should promote interoperability and uniform security approaches throughout the network. The following sections provide information about DECnet and TCP/IP security considerations.

DECnet Security

DECnet provides several security features that you can use to grant or restrict access to the standby database. To set limits on which users can gain network access to the standby database, Oracle Corporation recommends that you establish a proxy account that allows remote users to have access privileges on the standby database node without having an account on that node. Setting up proxies is preferable to other methods, such as using the SET PROTECTION command, because a proxy account does not require users to transmit login and password information over the network.

On OpenVMS systems using DECnet, the time allowed to start the AIJSERVER network object server is limited. The Hot Standby software requires the use of proxy accounts if a user's login information is to be used instead of the AIJSERVER net-

work login information. Therefore, when using DECnet proxies, add the following line as the first command in the LOGIN.COM procedure:

```
$ if “‘‘f$mode()‘”.eqs. “NETWORK” then goto EXIT
```

This command line allows replication operations to start. Note that the network section of the LOGIN.COM procedure must exit immediately.

Reference: Refer to the DECnet documentation for complete information about setting up proxy accounts.

TCP/IP Security

One method of controlling who can access operating system resources is by assigning identification codes to users. TCP/IP supports mechanisms for mapping OpenVMS user names to user identification (UID)/group identification (GID) pairs. Digital UNIX identifies users by user names and the UID/GID pairs, which are numbers that can identify a user on a system. The Digital UNIX system uses 32-bit UID/GID pairs. You might be able to use UID/GID pairs for user identification, especially if you are using the Hot Standby software across two different operating systems such as OpenVMS and Digital UNIX.

Another option is the Kerberos authentication system, which verifies the identity of users, thereby providing security in a network environment that otherwise might not be secure. You can use the Kerberos software to verify the authenticity of any user who requests access to a remote host.

Reference: Refer to your Digital UNIX documentation for more information about the Kerberos authentication system.

5
Shutting Down Database Replication

You must stop replication operations before you can stop the ALS process, move the database, or close the database. This chapter provides information about how to perform the orderly shut down of replication operations. It also describes how to perform an emergency database shutdown procedure while replication operations are active.

Table 5-1 provides a checklist of the shutdown tasks. Each step in the checklist includes a reference to a section later in this chapter that provides additional information.

Table 5-1 Replication Shutdown Checklist

Step	Procedure	Reference
[1]	Stop replication operations.	Section 5.1
[2]	Stop the ALS process on the master database.	Section 5.2
[3]	At this point, you can restart replication operations, move the database, or close the database:	
	• Move the database and restart replication operations	Section 5.3
	• Close the database using orderly or emergency shutdown procedures	Section 5.4

5.1 Stopping Replication Operations

You can stop replication operations by explicitly entering the Replicate After_Journal Stop command on either the standby or master database nodes. Stopping replication on either database terminates replication on both databases.

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL STOP PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal stop mf_personnel
```

Oracle Rdb on OpenVMS

```
$  RMU/REPLICATE AFTER_JOURNAL STOP mf_personnel
```

Reference: Chapter 6 describes the Replicate After_Journal Stop command in more detail.

5.2 Stopping the ALS Process

Once database replication has stopped, you can stop the ALS process on the master and standby database nodes using the RMU or DBO Server After_Journal Stop command. However, the ALS process must be running on the master database node whenever replication operations are occurring.

Note

You must stop database replication operations with the Replicate After_Journal Stop command before you can stop the ALS process. If you attempt to stop the ALS process on the master database while replication operations are active, the Hot Standby software returns an `ALSACTIVE` error message.

The following commands show the syntax you use to stop the ALS process:

Oracle CODASYL DBMS

```
$ DBO/SERVER AFTER_JOURNAL STOP /OUT=ALS_MSGS PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -server after_journal stop -out=als_msgs mf_personnel
```

Oracle Rdb on OpenVMS

```
$  RMU/SERVER AFTER_JOURNAL STOP /OUT=als_msgs mf_personnel
```

Reference: The Oracle Rdb and Oracle CODASYL DBMS documentation describe the Server After_Journal Stop command in more detail.

5.3 Restarting and Moving the Master Database

If you stop replication operations, you must wait for replication to stop completely before you attempt to restart replication operations or perform a database Move command to move the database to a different disk.

If you restart replication operations on the master database before replication has completely terminated on the standby database, you cannot move the master database. This is because a Move command causes the connection request on the master database to fail because it appears to be a different database from the master database that you started originally.

This situation occurs only when replication is restarted on the master before replication has fully terminated on the standby database (usually because of a failure on the master database).

Reference: The Oracle Rdb and Oracle CODASYL DBMS documentation describe moving a database in more detail. Chapter 6 describes the Replicate After_Journal Start command that you use to start (and restart) the replication operations.

5.4 Closing the Database

Stopping database replication operations can take several minutes, depending on the speed of the processors and the network traffic load. You should make sure that replication operations are stopped before you attempt to close either the master or standby database. Use the Oracle RMU or DBO Show Users command to determine whether or not replication operations are finished.

Note

You must stop replication operations before you can close any participating database.

You close the databases using the same syntax for Hot Standby configurations as for nonreplicated database environments. The following examples show some sample Close commands:

Oracle CODASYL DBMS

```
$ DBO/CLOSE PARTS
```

Oracle Rdb on Digital UNIX

```
$ rmu -close mf_personnel
```

Oracle Rdb on OpenVMS

```
$  RMU/CLOSE mf_personnel
```

On OpenVMS systems, you can close the master or standby database using the RMU Close command with the Nocluster qualifier on any node in the cluster *except* the node on which you started replication (where the LCS or LRS process is running).

Note

For databases running on OpenVMS systems, if you need to perform an emergency database shutdown procedure while replication operations are active, use the Abort=Delprc qualifier on the Close command. Do not attempt to close the database using the Abort=Forcex qualifier on the Close command. Using this qualifier returns an error message to the user.

Reference: The Oracle Rdb and Oracle CODASYL DBMS documentation describe closing a database in more detail.

6
Replication Commands

This chapter provides the syntax and semantics for the following Replicate commands and their parameters and qualifiers:

- Replicate After_Journal Configure
- Replicate After_Journal Reopen_Output
- Replicate After_Journal Start
- Replicate After_Journal Stop

These commands are available using either Oracle RMU (the Oracle Rdb database management utility) or DBO, the Oracle CODASYL DBMS Database Operator utility. The Oracle RMU and DBO command syntax described in this chapter follows the conventions provided in the following documentation:

Utility	Reference
DBO	*Oracle CODASYL DBMS Database Administration Reference Manual*
Oracle RMU	*Oracle RMU Reference Manual*

Privileges

You must have the following privileges to use the Replicate commands:

- To use the DBO utility on OpenVMS systems, you do not need any special OpenVMS privileges. All of the necessary DBO privileges are set up when the DBO image is installed. However, because all DBO commands are securable, your database administrator might have set up a command authorization list (CAL) previously for specific DBO commands. You can use the DBO Grant command to determine if the replication commands are included in the CAL.
- To use the Oracle RMU utility:
  - On Digital UNIX systems, you must have the RMU_OPEN privilege in the root file access control list (ACL) for the database or be superuser or the dbsmgr user

- On OpenVMS systems, you must have the RMU$OPEN privilege in the root file ACL for the database or the OpenVMS WORLD privilege

Server Names and Acronyms

The discussions in this chapter sometimes use acronyms to refer to the Hot Standby servers. The following table shows the server names and their acronyms, and the database where each server runs:

Server	Acronym	Database
AIJ log server	ALS	Master
Log catch-up server	LCS	Master
Log rollforward server	LRS	Standby
AIJSERVER	-	Standby

Default Command Qualifiers

The Hot Standby software supplies default values for most of the master and standby database attributes and maintains them in the database root file. Optionally, you can change one or more of the database attributes using qualifiers on the Replicate commands. When you specify a database attribute, the Hot Standby software updates the database root file so that the database root file always contains the most up-to-date qualifier values for the database.

You can specify database attributes using qualifiers on either of the following Replicate commands:

- Replicate After_Journal Configure—Preconfigures the master and standby database attributes without starting replication operations. This optional command allows you to preset database attributes that do not take effect until the next time you start replication operations using the Replicate After_Journal Start command.
- Replicate After_Journal Start—Configures database attributes at the same time you start replication for a database. If you preconfigured your database previously using the Replicate After_Journal Configure command, you can override the default settings by including one or more qualifiers on the Replicate After_Journal Start command.

Whenever you enter the Replicate After_Journal Start command, the Hot Standby software initiates database replication using the qualifier values specified on the Replicate After_Journal Start command line. If you do not specify qualifier values on the command line, the Hot Standby software uses values stored in the database root file or the default value for the qualifier.

Therefore, you do not need to respecify the qualifier values except to change a qualifier setting. For example, the following command examples show the Replicate After_Journal Start command the first time you enter it on the master database node:

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL START PARTS -
     /STANDBY_ROOT=REMNOD::DISK1:[USER]STANDBY_PARTS -
     /SYNCHRONIZATION=HOT
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal start mf_personnel \
> -standby_root=/REMNOD/hotdisk/usr/standby_personnel \
> -synchronization=hot
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL START mf_personnel -
     /STANDBY_ROOT=REMNOD::DISK1:[USER]standby_personnel -
     /SYNCHRONIZATION=HOT
```

The Hot Standby software saves the qualifier settings in the database root file (in this case, the database attributes are saved in the master database root file). The next time you start replication operations, you could enter the command line without the qualifier, as shown in the following Oracle CODASYL DBMS example:

```
$ DBO/REPLICATE AFTER_JOURNAL START PARTS
```

Examples 6-1 and 6-2 show examples of the header information from Oracle Rdb master and standby database root files.

Example 6-1 Header Information from the Master Database Root File

```
Hot Standby...
- Database is currently being replicated as “Master”
    Standby database is “_DISK1:[USER]STANDBY_PERSONNEL.RDB;1”
    Remote node name is "REMNOD”
    Replication commenced on 5-AUG-1996 08:13:30.57
    Synchronization obtained via quiet-point
    Server checkpoint interval is 100 messages
    Server connection-timeout interval is 5 minutes
    Replication synchronization is "hot"
```

Example 6-2 Header Information from the Standby Database Root File

```
Hot Standby...
- Database is currently being replicated as "Standby"
     Master database is "_DISK1:[USER]MF_PERSONNEL.RDB;1"
     Remote node name is "ORANOD"
     Replication commenced on  5-AUG-1996 08:13:23.91
     Database replication is "online"
```

```
Server checkpoint interval is 100 messages
Server gap-timeout interval is 5 minutes
Server buffer count is 256
Server 2PC transaction resolution is "commit"
```

6.1 Replicate After_Journal Configure Command

Allows you to preconfigure many of the master and standby database attributes (using qualifiers available with the Replicate After_Journal Start command) without starting replication operations.

You enter the Replicate After_Journal Configure command:

- On the master database to prespecify the Replicate After_Journal Start command qualifiers that are valid for the master database and store the qualifier settings in the master database root file
- On the standby database to prespecify the Replicate After_Journal Start command qualifiers that are valid for the standby database and store the qualifier settings in the standby database root file

Because the database attributes are stored in the respective database root files, the settings do not take effect until you start replication operations with the Replicate After_Journal Start command.

Format

Database and Platform	Command
Oracle CODASYL DBMS	DBO /Replicate After_Journal Configure *database-rootfile*
Oracle Rdb on OpenVMS	RMU /Replicate After_Journal Configure *database-rootfile*
Oracle Rdb on Digital UNIX	rmu -replicate after_journal configure *database-rootfile*

Description

The Replicate After_Journal Configure command is an optional command you can use to preconfigure the master and standby databases, one database at a time.

Note

You cannot preconfigure both the master and standby database attributes in a single Replicate After_Journal Configure command. Moreover, you cannot enter the Replicate After_Journal Configure command on the standby database to preconfigure master database attributes, or preconfigure standby database attributes from the master database.

You can specify one or more of the following qualifiers when you enter the Replicate After_Journal Configure command on the master database:

Master Database Qualifiers
Checkpoint
Connect_Timeout
{No]Log
[No]Quiet_Point
Standby_Root[1]
Synchronization

[1] You must specify the Standby_Root qualifier the first time you configure the master database.

The master database attributes that you specify are stored in the master database root file. (You cannot specify the Wait, NoWait, and Output qualifiers on the Replicate After_Journal Configure command. You can specify these qualifiers when you invoke the Replicate After_Journal Start command.)

You can specify one or more of the following qualifiers when you enter the Replicate After_Journal Configure command on the standby database:

Standby Database Qualifiers
Buffers
Checkpoint
Gap_Timeout
[No]Log
Master_Root[1]
[No]Online
Resolve

[1] You must specify the Master_Root qualifier the first time you configure the standby database.

The standby database attributes that you specify are stored in the standby database root file. (You cannot specify the Governor, Wait, NoWait, and Output qualifiers on the Replicate After_Journal Configure command. You can specify these qualifiers when you invoke the Replicate After_Journal Start command.)

Reference: See the Replicate After_Journal Start command for complete information about all qualifiers.

You should use the Replicate After_Journal Configure command if you want to:

- Preset qualifier values that you typically specify on the Replicate After_Journal Start command, but without starting replication operations.

The values you specify become the new default qualifier values that are stored in the database root file.

- Be able to quickly start replication operations by invoking a single Replicate After_Journal Start command on the master database.

  If you use the Replicate After_Journal Configure command to preconfigure the master and standby databases, you can start replication for both databases by entering one Replicate After_Journal Start command on the master database.

For example, if you have preconfigured both the master and standby databases and then invoke the Replicate After_Journal Start command on the master database node, the Hot Standby software:

1. Starts replication operations on the master database using default qualifier values from the master database root file
2. Creates the network connection to the standby database
3. Attaches the master and standby databases to the network
4. Starts replication operations on the standby database using default qualifier values in the standby database root file
5. Synchronizes committed transactions on the master and standby databases

Note

If you have not preconfigured database attributes using the Replicate After_Journal Configure command, the Hot Standby software uses either the system-supplied defaults or the values that you specified on a previous Replicate After_Journal Start command.

Command Parameter

database-rootfile

Specifies the name of the target database root file. For example, if you want to preconfigure the master database attributes, specify the master database root file. Similarly, you can specify the standby database root file to preconfigure the standby database.

Note

Do not include a node name when you specify the database-rootfile parameter. This parameter must specify a locally accessible database root file; the parameter cannot include a remote file specification.

Command Qualifiers

The command qualifiers are the same ones described for the Replicate After_Journal Start command. See Section 6.3 for qualifier information.

Usage Notes

- The first time you configure the standby database, you must include the Master_Root qualifier, and you must include the Standby_Root qualifier the first time you configure the master database.

  You must preconfigure the Master_Root or Standby_Root qualifiers because these qualifiers identify the “alternate” database for the database being configured. These qualifiers also identify whether a master or standby database is being configured (if the Replicate After_Journal Configure command includes the Master_Root qualifier, a standby database is being configured). The Master_Root and Standby_Root qualifiers are optional on subsequent replication configuration commands because the value is stored in the database root file.
- You can include a node name with the Master_Root or Standby_Root qualifiers.
- You cannot invoke the Replicate After_Journal Configure command when replication operations are active.
- You can override values you define with the Replicate After_Journal Configure command (and other the default values stored in the database root file) by specifying qualifiers on the Replicate After_Journal Start command.
- You cannot specify the Output qualifier on the Replicate After_Journal Configure command. Therefore, if you need to record Hot Standby server information to an output file when you start replication operations from the master database, specify an output file by:
  - Including the Output qualifier on the Replicate After_Journal Start command
  - Defining the BIND_ALS_OUTPUT_FILE, BIND_HOT_OUTPUT_FILE, BIND_LCS_OUTPUT_FILE, or BIND_LRS_OUTPUT_FILE logical name or configuration parameter as described in Appendix A

Note

If you plan to start replication operations remotely (for example, to start replication on the standby database from the master database node), you must have the following privileges:

- On OpenVMS systems, you must have GROUP, WORLD, and SYSPRV privileges
- On Digital UNIX systems, you must be superuser or the dbsmgr user

Examples

The following examples show how to use the Replicate After_Journal Configure command to configure replication attributes for the master database:

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL CONFIGURE PARTS -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]STANDBY_PARTS -
    /SYNCHRONIZATION=COLD -
    /QUIET_POINT -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal configure mf_personnel \
> -standby_root=/REMNOD/hotdisk/usr/standby_personnel \
> -synchronization=cold \
> -quiet_point \
> -checkpoint=10 \
> -connect_timeout=1
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL CONFIGURE mf_personnel -
    /STANDBY_ROOT=REMNOD:::DISK1:[USER]standby_personnel -
    /SYNCHRONIZATION=COLD -
    /QUIET_POINT -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1
```

6.2 Replicate After_Journal Reopen_Output Command

Closes the current informational file and reopens it as a new file. You can enter this command on either the master database node (to reopen the output file that records LCS information) or the standby database node (to reopen the output file that records LRS information).

Format

Database and Platform	Command
Oracle CODASYL DBMS	DBO /Replicate After_Journal Reopen_Output *database-rootfile*
Oracle Rdb on OpenVMS	RMU /Replicate After_Journal Reopen_Output *database-rootfile*
Oracle Rdb on Digital UNIX	rmu -replicate after_journal reopen_output *database-rootfile*

Description

The Hot Standby software dynamically and transparently switches from writing to the original output file to the new file. There is no need to stop or interrupt database replication operations during the transition to the new output file.

The Replicate After_Journal Reopen_Output command performs the following steps to reopen the output file:

1. Closes the current output file in which information about replication operations is recorded.
2. Reopens the output file by opening a new file using the original output file name. The new output file is written as follows:
   - On OpenVMS systems, the Hot Standby software opens a new output file using the originally specified file name and a new version number. Thus, you can view the original output file by specifying the older version number. If disk space is a problem, relocate the old output file to another disk.
   - On Digital UNIX systems, the new output file overwrites the original file contents. Therefore, if you need to save the original output file, you should copy or move it to another file before you enter the Replicate After_Journal Reopen_Output command.

You can enter the Replicate After_Journal Reopen_Output command on either the master or standby node as follows:

Enter the command . . .	To reopen the output file for the . . .
On the master database node	LCS server on the master database
On the standby database node	LRS server on the standby database

You must explicitly enable the ability to write replication startup information to an output file by including the Output qualifier when you start replication operations (see the Replicate After_Journal Start command for more information), or by specifying the BIND_ALS_OUTPUT_FILE, BIND_HOT_OUTPUT_FILE, BIND_LCS_OUTPUT_FILE, or BIND_LRS_OUTPUT_FILE logical name or configuration parameter. These logical names and configuration parameters are described in Appendix A.

The Replicate After_Journal Reopen_Output command is useful when:

- The output file becomes too large

  For example, as the output file grows over time, you might run out of disk space or notice that the database performance is slow. You can use the Replicate After_Journal Reopen_Output command to free up space on the disk. Once the new output file is open, you should relocate the old output file to a new location or delete the file.

  If the disk that contains the output file becomes full, the Hot Standby software stops writing information to the file (and on OpenVMS systems, a message is sent to the system operator). Note that replication operations continue, even when write I/O to the output file stops.

- You want to view the currently open output file

  By using the Replicate After_Journal Reopen_Output command, you can capture a snapshot of the output file and examine replication operations without interrupting processing. You can also view the contents of the current output file using the Type command at the OpenVMS system prompt, or the `more` command on Digital UNIX systems.

Note

You cannot use the Replicate After_Journal Reopen_Output command to change the size or location of the output file; the command is intended to create a new version of an existing output file.

- You want to open an output file for a server process that is actively performing replication operations

  Defining a logical name or configuration parameter is useful if you omitted the Output qualifier when you entered the Replicate After_Journal Start command to start replication. You can define a logical name or configuration parameter to specify an output file while replication operations are active. This can be done by defining the appropriate logical name or configuration parameter, and then invoking the Replicate After_Journal Reopen_Output command. This allows you to create an output file so the server can start writing to the file. The advantage to defining a logical name or configuration parameter is that you do not need to stop and restart the server.

  Reference: See the Output qualifier discussion under the Replicate After_Journal Start command.

Command Parameter

database-rootfile
Specifies the name of the master or standby database root file.

Usage Notes

- To write replication information to an output file, specify the Log and Output qualifiers on the Replicate After_Journal Start command.

  If you enter the Replicate After_Journal Reopen_Output command on a node where logging is not enabled, the Hot Standby software ignores the command; it does not return an error message if the Replicate After_Journal Reopen_Output command does not find an output file.

- The Replicate After_Journal Reopen_Output command is applicable only to the files that record activities for the LCS process or the LRS process. To reopen or view the output file that records information about the ALS process, use the RMU Server After_Journal Reopen_Output command or the DBO Server After_Journal Reopen_Output command, as appropriate.

  Reference: For more information about displaying ALS information, refer to:

  - The *Oracle RMU Reference Manual* for more information about the RMU Server After_Journal Reopen_Output command
  - The *Oracle CODASYL DBMS Database Administration Reference Manual* for more information about the DBO Server After_Journal Reopen_Output command

Examples

The following command examples show how to reopen an output file:

Oracle CODASYL DBMS

```
$ DBO /REPLICATE AFTER_JOURNAL REOPEN_OUTPUT PARTS.ROO
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal reopen_output mf_personnel.rdb
```

Oracle Rdb on OpenVMS

```
$ RMU /REPLICATE AFTER_JOURNAL REOPEN_OUTPUT mf_personnel.rdb
```

6.3 Replicate After_Journal Start Command

Initiates database replication operations.

Format

Database and Platform	Command
Oracle CODASYL DBMS on OpenVMS	DBO /Replicate After_Journal Start *database-rootfile*
Oracle Rdb on OpenVMS	RMU /Replicate After_Journal Start *database-rootfile*
Oracle Rdb on Digital UNIX	rmu -replicate after_journal start *database-rootfile*

OpenVMS Qualifiers (DBO and RMU)	Digital UNIX Qualifiers (RMU)	Defaults
/Buffers=*rollforward-buffer-count*	-Buffers=*rollforward-buffer-count*	Buffers=256
/Checkpoint=*checkpoint-interval*	-Checkpoint=*checkpoint-interval*	Checkpoint=100
/Connect_Timeout=*minutes*	-Connect_Timeout=*minutes*	Connect_Timeout=5
/Gap_Timeout=*minutes*	-Gap_Timeout=*minutes*	Gap_Timeout=5
/Governor=[Enabled or Disabled]	-Governor=[Enabled or Disabled]	Governor=Enabled
/[No]Log	-[No]Log	Nolog
/Master_Root=*master-rootfile*	-Master_Root=*master-rootfile*	None
/[No]Online	-[No]Online	Noonline
/Output=[*log-filename* or *log-filename*_PID]	-Output=*[log-filename* or *log-filename*_pid]	None
/[No]Quiet_Point	-[No]Quiet_Point	Noquiet_Point
/Resolve=[Commit or Abort]	-Resolve=[Commit or Abort]	Resolve=Commit
/Standby_Root=*standby-rootfile*	-Standby_Root=*standby-rootfile*	None
/Synchronization=[Commit or Hot or Warm or Cold]	-Synchronization=[Commit or Hot or Warm or Cold]	Synchronization=Cold
/[No]Wait	-[No]Wait	Wait

Description

To start database replication, you can enter the Replicate After_Journal Start command on both the standby node and the master node. Although you can initiate replication operations on either node, Oracle Corporation recommends that you start replication on the standby node before you start it on the master node. This is because replication activity does not begin until:

- The standby database creates the network connection
- The master database attaches to the network connection

- The master and standby databases are synchronized with regard to committed transactions

Note

If you used the Replicate After_Journal Configure command to preconfigure the master and standby database attributes (see Section 6.1), you can invoke a single Replicate After_Journal Start command to start replication operations on both the master and standby databases.

Starting Replication

You can start database replication while the master database, the standby database, or both databases are on line (open) and accessible for active use. There is no need to close either database to initiate database replication. The databases can process transactions during replication startup, as follows:

- The master database can process read / write transactions.

  Applications and users can continue to access data and make modifications to the master database whether or not replication activity has started. Waiting for the replication activity to begin does not inhibit access to, or interrupt modifications on, the master database.

- The standby database can process read-only transactions. There cannot be any active read / write transactions on the standby node.

  Starting replication is an online operation that can occur while the standby database is open. However, database users must not actively attach to the standby database prior to starting database replication if you perform offline backup operations.

In addition, replication operations cannot start when these conditions exist:

- Any read / write transactions, including prestarted read / write transactions, are active on the standby database (see Section 2.13 for information about disabling prestarted transactions)
- Any storage area is inconsistent (for example, if you restore a storage area from a backup file but you have not rolled forward after-image journals to be consistent with the rest of the database)

Note

On OpenVMS systems, if you have preconfigured your Hot Standby environment using the Replicate After_Journal Configure command and you plan to start replication operations remotely (for example, if you want to start replication on the standby database from the master database node), you must provide the SYSPRV privilege to the DBMAIJSERVER or RDMAIJSERVER account.

Qualifier Usage

Some of the qualifiers for the Replicate After_Journal Start command are applicable only when you start replication operations on the master database node, while others are applicable only to the standby database node. Table 6-1 categorizes the qualifiers according to usage.

Table 6-1 Qualifier Usage for the Replicate After_Journal Start Command

Master Node Qualifiers	Master and Standby Nodes	Standby Node Qualifiers
Connect_Timeout	Checkpoint	Buffers
[No]Quiet_Point	[No]Log	Gap_Timeout
Standby_Root	[No]Wait	Governor
Synchronization	Output	Master_Root
		[No]Online
		Resolve

The Hot Standby software does not allow you to use qualifiers that are not valid for the database where you enter the command. Therefore, when you enter the Replicate After_Journal Start command on the:

- Master node—you can specify any of the qualifiers listed in the first and second columns of Table 6-1
- Standby node—you can specify any of the qualifiers listed in the last two columns of Table 6-1

If you use an inapplicable qualifier (for example, if you use the Connect_Timeout qualifier when you start replication on the standby node), the Hot Standby software returns an error message.

Note

Whenever you specify a qualifier on the Replicate After_Journal Start command line, you must also include the Master_Root or Standby_Root qualifier, as appropriate, on the command line. For example, to change the value of the Synchronization qualifier on a master database node, you must specify both the Synchronization and Standby_Root qualifiers, as shown in the following example:

```
$ DBO /REPLICATE AFTER START PARTS /SYNCH=COLD -
_$ /STANDBY_ROOT=REMOTE::DISK1:[USER]PARTS
```

Command Parameters

database-rootfile

Indicates the root file specification for either the master or standby database where you want to start database replication.

Note

Do not include a node name when you specify the database-rootfile parameter. This parameter must specify a locally accessible database root file; the parameter cannot include a remote file specification.

The following table describes which database root file to specify depending on where you enter the command:

When you enter the command Replicate After_Journal Start . . .	Specify the database root file for the . . .
On the standby node	Standby database
On the master node	Master database

To ensure that the standby database accesses the correct master database as the source of replication operations, include the Master_Root qualifier on the command line. Similarly, to ensure that the master database accesses the correct standby database as the target of replication operations, include the Standby_Root qualifier on the command line.

Reference: See the Master_Root and Standby_Root qualifiers discussed later in this chapter.

Command Qualifiers

Buffers=*rollforward-buffer-count*

Specifies the number of database buffers available to roll after-image journals forward to the standby database.

Applicable to:	Standby database
Required or Optional:	Optional
Default Value:	256 buffers
Minimum Value:	2 buffers
Maximum Value:	524,288 buffers

During replication operations, the LRS process on the standby node receives after-image journal records from the master database and rolls them forward to the standby database. By default, the standby database has 256 buffers available to perform the rollforward operation.

You can use the optional Buffers qualifier to override the default number of database buffers.

Note

The LRS server ignores the Buffers qualifier if you defined global buffers on the standby database using either of the following methods:

- The SQL clause GLOBAL BUFFERS ARE ENABLED
- The Global_Buffers qualifier on the RMU Open command

The LRS server uses the number of buffers specified in the database, or specified when you open the database. The LRS process takes maximum advantage of the standby system by using the largest number of buffers possible to maximize replication performance and throughput.

For optimal performance, you should allocate a sufficient number of buffers so that the server process can roll the after-image journal records forward with a minimum number of I/O operations. To estimate an appropriate number of buffers, use the following equation:

```
(Number of Modified Buffers per Transaction * Number of Users) + 20%
```

For example, if the average number of modified buffers per transaction is 10 and there are 100 users on the database, then the server process needs 1000 buffers at one time. To ensure that you have an adequate number of buffers, add another 20 percent (200 buffers) for a total of 1200 buffers.

Note

The LRS process on the standby database does not use buffer values defined by the following:

- DBM$BIND_BUFFERS logical name
- RDB_BIND_BUFFERS configuration parameter
- RDM$BIND_BUFFERS logical name

When replication operations are active, you can use the RMU or DBO Show Users command to see the current number of database buffers allocated. If replication operations are not active or if you want to see the buffer value that was set on a previous Replicate After_Journal Start command (stored in the database root file), you can also use the Header and Dump_Select_Type=Hot_Standby qualifiers on the RMU or DBO Dump command. (See Chapter 7 for more information.)

Checkpoint=*checkpoint-interval*

Specifies, in terms of processed messages, how frequently the Hot Standby servers update information in the database root file. This qualifier can be set to different values on the master and standby databases.

Applicable to:	Master and standby database
Required or Optional:	Optional
Default Value:	100 messages
Minimum Value:	1 message
Maximum Value:	1024 messages

By default, the Hot Standby servers automatically perform checkpoint operations on both the master and standby databases after every 100 messages are processed. Checkpoints are essential to database availability because they:

- Enable the Hot Standby software to restart database replication operations more quickly in the event of a failure because frequent checkpoints limit the number of transactions that must be redone if a process or system fails.
- Cause all modified database cache buffers on the standby database to be flushed to the disk, making the buffers available for access by other users (when online database access is enabled)
- Improve the redo performance of the database recovery (DBR) process
- Allow after-image backup operations to back up older after-image journals on the master database

Note

In addition to performing checkpoint operations specified by the Checkpoint qualifier, the replication servers on the master database also checkpoint automatically after the following events:

- After two minutes of inactivity
- After a switchover to a new after-image journal (when you are using circular after-image journals)
- After an AIJ backup operation (when you are using extensible after-image journals)

On the standby database, the LRS process checkpoints after two minutes of inactivity if data has been processed since the last checkpoint.

These automatic checkpoints advance the oldest active checkpoint indicator to make older after-image journals available for backup operations. You cannot change or override these checkpoint intervals.

The default checkpoint interval usually is sufficient to effectively maintain synchronization between the master and standby database root files. However, you can override the default checkpoint interval by specifying the Checkpoint qualifier when you start replication on the master database, the standby database, or both.

For example, if you specify the qualifier Checkpoint=300 on the standby database, the LRS server process updates information in the standby database root file after every 300 messages are exchanged between the master and the standby database.

Table 6-2 describes how the frequency of the checkpoint operation can affect database synchronization.

Table 6-2 Setting the Frequency of Checkpoint Intervals

If you specify . . .	Then . . .
A small checkpoint interval	The Hot Standby software synchronizes the database root files more often, but uses less time to restart replication because fewer transactions need to be redone.
A large checkpoint interval	The Hot Standby software synchronizes the database root files less frequently, but requires more time to restart replication because more transactions must be redone.

In addition, the value you set for the checkpoint interval:

- Controls replication restart in the event of a failure on the master database. A side effect of this is that the ABS process cannot back up after-image journals that are needed to restart replication operations
- Effects how the after-image journals on the master database become available for backup

  Specifying a large value for the checkpoint interval can cause after-image journal backup operations to stall until the appropriate after-image journal file becomes available for a backup operation. This is because the after-image journal backup operation cannot back up any after-image journal file that is required for process recovery or replication restart.
- Effects the reinitialization of after-image journals on the standby database
- Effects the manner in which the LRS process on the standby database:
  - Releases page locks
  - Responds to page lock conflict messages from another attached database process

  Oracle Corporation recommends that you set a *reasonably* small checkpoint interval for the standby database. Specifying a checkpoint interval that is too large can prevent the LRS process from responding to requests for pages, and it is possible for other processes to become stalled.

For Oracle Rdb databases, you can monitor the effectiveness of the current setting of the Checkpoint qualifier by using the RMU Show Statistics command and examining the Checkpoint Information display.

Connect_Timeout=*minutes*

Specifies the maximum number of minutes that the LCS process on the master database waits for a network connection to the LRS process on the standby database.

Applicable to:	Master database
Required or Optional:	Optional
Default Value:	5 minutes
Minimum Value:	1 minute
Maximum Value:	4320 minutes (3 days)

When you start replication on the master database (before starting it on the standby database):

1. The Hot Standby software invokes the log catch-up server (LCS) process on the master database.
2. The LCS process invokes its corresponding network AIJSERVER process on the standby node.
3. The AIJSERVER process attempts to create a network connection to the LRS process on the standby node.

By default, the LCS process allows 5 minutes for the AIJSERVER to connect to the LRS process. You can override the default by specifying the Connect_Timeout qualifier when you start replication on the master database. (Note that if you specify the Connect_Timeout qualifier, you must specify a time value (in minutes).)

The Connect_Timeout qualifier is useful when you start replication operations on the master database before you start replication on the standby database. This is because the Connect_Timeout qualifier allows sufficient time for the network connection to be made before the LCS process begins sending after-image journal records across the network.

Note

While the LCS process on the master database waits for the replication activity to begin on the standby database, users and applications can continue to access and modify data in the master database.

Also, because the Connect_Timeout qualifier waits only for the network connection, you might consider using the Wait qualifier in addition to the Connect_Timeout qualifier. The Wait qualifier causes the Replicate After_Journal

Start command to wait for the server processes to be activated. See the Wait qualifier later in this chapter for additional information.

Gap_Timeout=*minutes*
Specifies the maximum number of minutes that the standby database (LRS process) should wait for a gap in the replication data sequence to be resolved.

Applicable to:	Standby database
Required or Optional:	Optional
Default Value:	5 minutes
Minimum Value:	1 minute
Maximum Value:	4320 minutes (3 days)

If a gap in the replication data sequence is not resolved in the period of time allowed, the LRS process:

1. Assumes that the node sending the message has failed
2. Terminates replication operations immediately

You must restart replication operations manually to resolve the situation.

Governor=Enabled
Governor=Disabled
Enables or disables the replication governor.

Applicable to:	Standby database
Required or Optional:	Optional
Default Value:	Governor=Enabled

The purpose of the replication governor is to coordinate database replication operations automatically between the master and the standby databases. With the replication governor enabled, you can effectively ensure that:

- The master and standby databases do not get too far out of synchronization with respect to each other
- The performance of the master database does not deviate greatly from that of the standby database
- The peak-time database requirements are handled automatically and dynamically by the Hot Standby software

The replication governor allows the ALS process on the master database and the LRS process on the standby database to automatically choose the synchronization mode that provides the best performance and ensures database replication synchronization.

To use the replication governor most effectively, ensure the Governor qualifier is Enabled and include the Synchronization=Cold qualifier when you start replication operations on the standby database. (Also, see the Synchronization qualifier discussed later in this section.)

Oracle Corporation recommends that you set the Synchronization qualifier to Cold mode. This setting is most effective because of the way the LRS process monitors its replication workload from the master database, as described in the following table:

If . . .	Then . . .
The replication workload increases at a rate that prevents the standby database from keeping up with the master database	The LRS process automatically upgrades to a *stronger* synchronization mode. For example, if the Synchronization qualifier was originally set to Cold mode, the LRS would change the synchronization mode to Warm (or higher, as required).
The replication workload shrinks	The LRS process automatically downgrades or *weakens* the synchronization mode. However, the synchronization mode is never weaker than the mode (Commit, Hot, Warm, Cold) that you specify with the Synchronization qualifier.

Because the synchronization mode changes dynamically, the LRS process transmits the current synchronization mode to the ALS process (on the master database) at every checkpoint interval (see the Checkpoint qualifier earlier in this chapter). For example, if the replication governor upgrades the synchronization mode from Cold to Warm, the LRS process transmits the information to the ALS process. Then, the ALS process uses the stronger mode on all subsequent messages to the standby database. (Note that the LRS process maintains a different synchronization mode for each master database node.)

Use the RMU or DBO Show Statistics command on the master database to monitor the dynamically changing synchronization mode required by the actual work load, and compare that to the mode you specified with the Synchronization qualifier.

Recommendation: Oracle Corporation recommends that you do not use the Governor=Disabled setting until the replication performance is well understood and constant. Severe performance deviations on the master database could stall or stop the database replication operations.

Reference: See Section 4.1 for more information about Hot Standby synchronization algorithms.

Log
Nolog

Indicates whether or not to log the status of, and information about, activities when you start replication operations.

Applicable to:	Master and standby database
Required or Optional:	Optional
Default Value:	Nolog

If you specify the Log qualifier, output showing the status of the replication startup is logged to:

- SYS$OUTPUT on OpenVMS systems
- stdout on Digital UNIX systems

Oracle Corporation recommends that you specify the Log qualifier.

Also, you can record status information to an output file by including the Output qualifier on the Replicate After_Journal Start command.

Reference: See the Output qualifier discussed later in this chapter.

Master_Root=*master-rootfile*

Identifies the name of the master database root file from which the replication servers on the standby database receive replication data.

Applicable to:	Standby database
Required or Optional:	Required the first time you enter the Replicate After_Journal Start command and any time you specify other Replication Startup qualifiers. Optional on all subsequent invocations.
Default Value:	None.

You must include the Master_Root qualifier the first time you enter the Replicate After_Journal Start command (unless you have preconfigured the Master_Root qualifier using the Replication After_Journal Configure command). This ensures that the standby database uses the master database you specify as the source of the replication operations. If you omit the Master_Root qualifier on subsequent Replicate After_Journal Start commands, the Hot Standby software retrieves the master

database name from the header information in the database root file (see Examples 6-1 and 6-2).

Whenever you specify the Master_Root qualifier, you must do the following to ensure the command executes successfully:

- Specify the name of the master database root file.

Note

Do not specify the name of the standby database on the Master_Root qualifier. Any attempt to use a *restored* database as a master database causes replication startup operations to fail.

- Include a node name and directory path for remote network communications. **Note**: You can define a logical name to identify the master node.
- Be able to access the master database.

Note

When the master database node is configured in a VMScluster system, the node name you specify with the Master_Root qualifier can be any participating node from which the master database can be accessed. Cluster aliases are acceptable when you use the Master_Root qualifier.

The master and standby databases communicate using network communications (for remote database access) or interprocess communications (for local database access) according to how you specify the master database name. The following table describes how the Hot Standby software chooses the method of communication:

If . . .	Then . . .
You include a node name when you specify the master database root file	The Hot Standby software uses remote network communications to receive the after-image journal log changes, unless the specified node is the current node
You do not include a node name when you specify the master database root file	The Hot Standby software uses local interprocess communications to receive the after-image journal log changes

The Hot Standby software compares and verifies the master database (that you specify with the Master_Root qualifier) against the standby database (that you spec-

ify with the Standby_Root qualifier when you start replication operations on the master database). This verification ensures that both databases are identical transactionally.

Online
Noonline
Allows or disallows users and applications to be on line (actively attached) to the standby database.

Applicable to:	Standby database
Required or Optional:	Optional
Default Value:	Noonline

Online database access means that database users and applications can be actively attached (and perform read-only transactions) to the standby database before, during, and after replication operations.

The default setting (Noonline) disallows applications and users from attaching to the standby database during replication operations. However, if the standby database is open on another node (thus, an ALS process is active on that node), the LRS process cannot start replication on the standby database and the error message STBYDBINUSE is returned.

Note

If record caching is enabled on the standby database, the Hot Standby software assumes the Online setting. Specifying the Noonline qualifier on the Replicate After_Journal Start command has no effect. Because record caching requires the record cache server to be an online server, you cannot override the Online setting.

Because the Replicate After_Journal Start command fails if you enter it on a standby node where read / write transactions are in progress (including prestarted read / write transactions), Oracle Corporation recommends that you choose the Noonline (default) setting. (See Section 2.13 for information about disabling prestarted transactions.)

The Online and Noonline qualifiers do not affect access to the master database.

Output=[*log-filename*.out or *log-filename*_pid.out]
Identifies the name of the file where you want the Hot Standby software to create an operational output file (log) for the LCS or LRS process:

- Specify the Output qualifier on the master database to create an output file and collect information about the LCS process.
- Specify the Output qualifier on the standby database to create an output file and collect information about the LRS process.
- Optionally, include "_PID" or "_pid" when you specify the output file name. This causes the software to create a unique file name because it includes the process identification (PID) number.

Applicable to:	Master and standby databases
Required or Optional:	Optional
Default Value:	None. If you do not specify the Output qualifier, the Hot Standby software does not record LCS or LRS process activities to an output file.

The Output qualifier overrides definitions you make with the BIND_LCS_OUTPUT_FILE or BIND_LRS_OUTPUT_FILE logical name or configuration parameter. If you enable replication operations for multiple databases, there will be multiple operational output files.

The purpose of the operational log is to record the transmittal and receipt of network messages, and to provide administrative and diagnostic information.

Note the following when you specify the Output qualifier:

- You must specify an output file name. When you include "_PID" in the output file specification, the command creates a unique file name that includes the process identification (PID). For example:

  ```
  $ DBO/REPLICATE AFTER_JOURNAL START /OUTPUT=DISK1:[USER]LRS_PID.OUT
  ```

  This command line creates a unique file name, for example, DISK1:[USER]LRS_25C02914.OUT.
- Do not include a node name designation when you specify the output file name.
- The default location is the database root file directory. You can optionally include a directory name when you specify a file name.
- The directory containing the output files must be readable and writable by all processes.
- The default file type is .out

- You can display the name of output files that you specify with the Output qualifier using the RMU or DBO Show Users command (shown in Example 7-1). Output file names are not displayed in Show Users output for files specified with a logical name or configuration parameter.

Note

All bugcheck dumps are written to a corresponding bugcheck dump file. Bugcheck dumps are not written to the Output operational log.

Although it is optional, Oracle Corporation recommends that you use the Output qualifier, or a logical name or configuration parameter, to collect information about the LCS and LRS processes during replication.

You can also collect information about the ABS, ALS, DBR, and AIJSERVER processes by defining a logical name or configuration parameter. The following table lists the logical names and configuration parameters you can define to collect server process information to an output file:

Logical Name or Configuration Parameter	Specifies an output file for the . . .
BIND_ABS_LOG_FILE	ABS process
BIND_ALS_OUTPUT_FILE	ALS process[1]
BIND_DBR_LOG_FILE	DBR process
BIND_HOT_OUTPUT_FILE	AIJSERVER process
BIND_LCS_OUTPUT_FILE	LCS process
BIND_LRS_OUTPUT_FILE	LRS process

[1] You can also collect information about the ALS process to an output file by including the Output qualifier on the RMU Server After_Journal command or the DBO Server After_Journal command, as appropriate. For more information about displaying ALS information, refer to the *Oracle RMU Reference Manual* or the *Oracle CODASYL DBMS Database Administration Reference Manual*.

Reference: See Appendix A for information about defining these system-level logical names and the configuration parameters.

Defining a logical name or configuration parameter is also useful if you omitted the Output qualifier when you entered the Replicate After_Journal Start command to start replication. You can define a logical name or configuration parameter to specify an output file while replication operations are active. This can be done by defining the appropriate logical name or configuration parameter, and then invoking the Replicate After_Journal Reopen_Output command. This allows you to create an

output file so the server can start writing to the file without you having to stop and start the server.

Quiet_Point
Noquiet_Point
Determines whether or not the log catch-up server (LCS) process acquires a transactional quiet point during the database synchronization phase of the replication restart procedure.

Applicable to:	Master database
Required or Optional:	Optional
Default Value:	Noquiet_Point

Oracle Corporation recommends using the Quiet_Point qualifier because it makes it easier to restart replication operations.

Resolve=*keyword*
Specifies whether unresolved distributed transactions should be aborted, committed, or ignored.

Applicable to:	Standby database
Required or Optional:	Optional
Default Value:	Resolve=Commit

Table 6-3 describes the keywords you can set with this qualifier.

Table 6-3 Keywords for the Resolve Qualifier

Keyword	Description
Abort	Aborts all unresolved transactions.
Commit	Commits all unresolved transactions. This is the default keyword.

The Hot Standby software does not record the resolution of the prepared transaction in the standby database after-image journal. This is because there may not be enough available free space in the journal to write the commit or rollback records.

Standby_Root=*standby-rootfile*

Identifies the name of the standby database root file to which the replication servers on the master database send replication data.

Applicable to:	Master database
Required or Optional:	Required the first time you enter the Replicate After_Journal Start command and any time you specify other Replication Startup qualifiers. Optional on all other invocations.
Default Value:	None

You must include the Standby_Root qualifier the first time you enter the Replicate After_Journal Start command (unless you have preconfigured the Standby_Root qualifier using the Replication After_Journal Configure command). This ensures that the master database communicates with the standby database you specify as the recipient of replication operations. If you omit the Standby_Root qualifier on subsequent Replicate After_Journal Start commands, the Hot Standby software retrieves the standby database name from the header information in the database root file (see Examples 6-1 and 6-2).

Whenever you specify the Standby_Root qualifier, you must do the following to ensure the command executes successfully:

- Specify the name of the standby database root file.
- Include a node name and directory path for remote network communications. **(**You can define a logical name to identify the master node.)

Note

When the standby database is configured in a VMScluster system, the node name you specify with the Standby_Root qualifier cannot be a cluster alias.

- Be able to access the standby database.
- Ensure that the standby database is opened for access prior to starting replication operations on the master database.

  You must open the standby database manually unless you preconfigured the standby database. If you preconfigured the database, you can start replication on both the master and standby databases by entering a single Replicate After_Journal Start command on the master database. The master database automatically opens the standby database, if necessary.

The master and standby databases communicate using network communications (for remote database access) or interprocess communications (for local database access) according to how you specify the database name. The following table describes how the Hot Standby software chooses the method of communication:

If . . .	Then . . .
You specify a node name (for access to a standby database on a remote node)	The Hot Standby software uses remote network communications to ship the after-image journal log changes, unless the specified node is the current node
You do not specify a node name	The Hot Standby software uses the following communications to ship the after-image journal log changes: • Local interprocess communications on the local node • Remote network communications on all other nodes and across the cluster

The Hot Standby software compares and verifies the master database (that you specify with the Master_Root qualifier) against the standby database (that you specify with the Standby_Root qualifier). The purpose of this verification is to ensure that both databases are identical transactionally.

If replication operations are not started on the standby database when you invoke the Replicate After_Journal Start command on the master database, the Hot Standby software attempts to start replication on the standby database using the default replication attributes configured in the database root file before starting replication on the master database.

Synchronization=*keyword*

Specifies the degree to which you want to synchronize committed transactions on the standby database with committed transactions on the master database.

Applicable to:	Master database
Required or Optional:	Optional
Default Value:	Synchronization=Cold

When you enable replication operations, server processes on the master database write transactions to the after-image journal for the master database and send them across the network to the after-image journal for the standby database. The standby database acknowledges the receipt of the transactional message and handles after-image journaling depending on the mode you have set with the Synchronization

qualifier. Table 6-4 describes the keywords you use to set the synchronization mode.

Table 6-4 Keywords for the Synchronization Qualifier

Keyword	Equivalence of Committed Transactions	Performance Impact on Master Database	Standby Database Recoverability
Commit	When the standby database receives the AIJ information from the master database, the servers on the standby database: 1. Write it to the after-image journal on the standby system 2. Apply the AIJ to the standby database 3. Send a message back to the master database acknowledging the successful commit of the transaction	Highest	The standby database is transactionally identical and recoverable with respect to the master database.
Hot	When the standby database receives the AIJ information from the master database, the servers on the standby database: 1. Write it to the AIJ on the standby system 2. Send a message back to the master database before applying the transaction to the standby database	High	The standby database is extremely close to being transactionally identical to the master database. After-image journal records in transit are received and committed. Some restart processing may be required to synchronize the databases.
Warm	When the standby database receives the AIJ information from the master database, the servers on the standby database: • Send a message back to the master database before applying the transaction to either the AIJ or the standby database • Might not commit after-image journal records to the database	Medium	The standby database is transactionally close to the master database, but the databases are not identical. There may be transactions rolled back on the standby database that have been committed on the master database.
Cold (default)	When the standby database receives the AIJ information from the master database: • The servers never return a message acknowledging the receipt of the AIJ information • In failover situations, it is possible that transactions rolled back on the standby database were committed on the master database	Low	The standby database is not immediately recoverable transactionally with respect to the master database. After-image journal records in transit could be lost.

For each level of database synchronization, you make a trade-off between how closely the standby and master databases match each other in regard to committed transactions against performance.

For example, the Synchronization=Cold level provides the fastest performance for the master database, but the lowest level of master and standby database synchronization. However, in some business environments, this trade-off might be acceptable. In such an environment, the speed of master database performance outweighs the risk of losing recent transactions in the event of failover; system throughput has greater financial importance and impact than the value of individual aij records (transactions).

Recommendation: For high-performance applications, Oracle Corporation recommends that you do not specify both the Synchronization=Cold and the Governor=Disabled qualifiers when you start replication on the standby system. This is because the master database can possibly outperform the standby database during updates. The replication governor should be enabled to prevent the master and standby databases from getting too far out of synchronization.

Reference: See Section 4.1 for more information about the modes of synchronization.

Note

You can define logical names or configuration parameters to specify the synchronization mode, or to enable or disable the replication governor. See Appendix A for information about the BIND_HOT_DATA_SYNC_MODE and the BIND_LRS_GOVERNOR_ENABLED logical names and configuration parameters.

Wait
Nowait

Indicates whether or not the Replicate command should wait for activation of the replication server processes before returning control to the user.

Applicable to:	Master and standby databases
Required or Optional:	Optional
Default Value:	Wait

The Wait qualifier has the following effects:

- On the master database node—replication waits for activation of the server processes on the master database node
- On the standby database node—replication waits for activation of the server processes on the standby database node

The following table describes the [No]Wait qualifier:

Qualifier	Description
Wait (default)	The Replicate command does not return to the user until the respective server process has successfully initiated the database replication operation. Replication waits indefinitely for the activation of the server process, even though activation might take substantial time. However, the server process might not actually start the replication operation.
Nowait	Control should be returned to the user as soon as the LCS or LRS server process has been invoked by the database monitor.

You can use the Connect_Timeout qualifier with the Wait qualifier to limit the amount of time replication waits for the server process to become active.

Note

You must wait for commands that include the Nowait qualifier to complete before you enter another command. This is because if the first command fails before the subsequent command executes, the second command might receive the HOTCMDPEND error. For example:

```
$ RMU/REPLICATE AFTER_JOURNAL START/NOWAIT mf_personnel
$ RMU/REPLICATE AFTER_JOURNAL STOP/WAIT mf_personnel
```

If the first command to start replication fails, the startup error might be returned to the waiting Replicate After_Journal Stop command.

Examples

The following examples show Replicate After_Journal Start commands that initiate database replication. The default qualifier values are read from the database root file header information.

Oracle CODASYL DBMS

```
$ DBO/REPLICATE AFTER_JOURNAL START PARTS -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]STANDBY_PARTS -
    /SYNCHRONIZATION=COLD -
    /QUIET -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1 -
    /LOG -
    /WAIT -
    /OUT=REMNOD::DISK1:[USER]LCS_PID.OUT
```

Oracle Rdb on Digital UNIX

```
$ rmu -replicate after_journal start mf_personnel \
> -standby_root=/REMNOD/hotdisk/usr/standby_personnel \
> -synchronization=cold \
> -quiet \
> -checkpoint=10 \
> -connect_timeout=1 \
> -log \
> -wait \
> -out=/REMNOD/hotdisk/usr/lcs_pid.out
```

Oracle Rdb on OpenVMS

```
$ RMU/REPLICATE AFTER_JOURNAL START mf_personnel -
    /STANDBY_ROOT=REMNOD::DISK1:[USER]standby_personnel -
    /SYNCHRONIZATION=COLD -
    /QUIET -
    /CHECKPOINT=10 -
    /CONNECT_TIMEOUT=1 -
    /LOG -
    /WAIT -
    /OUT=REMNOD::DISK1:[USER]lcs_pid.out
```

6.4 Replicate After_Journal Stop Command

Terminates database replication operations.

Format

Database and Platform	Command Format
Oracle CODASYL DBMS on OpenVMS	DBO /Replicate After_Journal Stop *database-rootfile*
Oracle Rdb on OpenVMS	RMU /Replicate After_Journal Stop *database-rootfile*
Oracle Rdb on Digital UNIX	rmu -replicate after_journal stop *database-rootfile*

OpenVMS Qualifiers (DBO and RMU)	Digital UNIX Qualifiers (RMU)	Defaults
/[No]Abort	-[No]Abort	Noabort
/[No]Log	-[No]Log	Nolog
/[No]Wait	-[No]Wait	Wait

Description

You can enter the command on either the master node or the standby node. Replication operations are stopped as described in the following table:

When . . .	Then . . .
You enter the command on the master database	Replication is terminated immediately. Active transactions are handled differently depending on whether you specify the Abort qualifier or take the default (Noabort).
You enter the command on the standby database	Replication is terminated after any pending after-image journal records are completely rolled forward. Any active transactions on the standby database are rolled back.

You can stop database replication while the master database, the standby database, or both databases are on line (open) and accessible for active use. There is no need to close either database to stop database replication.

If the database is not manually opened on the node where you entered the Replicate After_Journal Start command, you must enter the Replicate After_Journal Stop command on the node where the corresponding replication server is running, or first open the database manually.

When replication operations stop, the Hot Standby software automatically restarts the AIJ log server (ALS) processes on the standby node.

Command Parameter

database-rootfile
Specifies the database root file for which you want to stop replication operations.

Command Qualifiers

Abort
Noabort
Indicates whether pending after-image journal information is rolled forward on the standby database before database replication operations are shut down. The following table describes the qualifiers:

Qualifier	Description
Abort	Database replication shuts down as quickly as possible. Any after-image journal information waiting to be rolled forward on the standby database is discarded, and all active transactions on the standby database are rolled back.
Noabort (default)	Database replication shuts down after all after-image journal information waiting to be rolled forward on the standby database is completed. Note that this type of shutdown could still result in active transactions being rolled back on the standby database.

Log
Nolog
Enables or disables logging the results of the Replicate After_Journal Stop operation.

Applicable to:	Master and standby databases
Required or Optional:	Optional
Default Value:	Nolog

If you specify the Log qualifier, the log file output is written to:

- SYS$OUTPUT on OpenVMS
- `stdout` on Digital UNIX systems

Wait
Nowait
Indicates whether or not the Replicate command should wait for activation of the replication server processes before returning control to the user.

Applicable to:	Master and standby databases
Required or Optional:	Optional
Default Value:	Wait

The Wait qualifier has the following effects:

- On the master database node—replication waits for deactivation of the server processes on the master database node
- On the standby database node—replication waits for deactivation of the server processes on the standby database node

The following table describes the Wait and Nowait qualifiers:

Qualifier	Description
Wait (default)	The Replicate command does not return to the user until the respective server process has successfully stopped the database replication operation. Replication waits indefinitely for the termination of the server process, even though termination might take substantial time. However, the server process might not actually stop the replication operation.
Nowait	Control should be returned to the user as soon as the LCS or LRS server process has been stopped by the database monitor.

Note

You must wait for commands that include the Nowait qualifier to complete before you enter another command. This is because if the first command fails before the subsequent command executes, the second command might receive the HOTCMDPEND error. For example:

```
$ RMU/REPLICATE AFTER_JOURNAL START/NOWAIT mf_personnel
$ RMU/REPLICATE AFTER_JOURNAL STOP/WAIT mf_personnel
```

If the first command to start replication fails, the startup error might be returned to the waiting Replicate After_Journal Stop command.

Replicate After_Journal Stop Command

7
Monitoring Replication Activities

This chapter describes how to monitor database replication activities and describes some methodologies to enhance database replication performance.

7.1 Displaying Hot Standby Statistics

You can use the database Show and Dump commands to monitor and evaluate all database activity. In many cases, you can display the database state and replication performance information graphically or numerically through these methods:

- Show Users command to monitor the current database status
- Dump Header command to dump the contents of the database files
- Show Statistics command for online performance statistics
- Show Statistics Database Dashboard facility for viewing and dynamically changing certain database attributes
- Output files produced by the replication commands

7.2 Show Users Command

The Oracle RMU and Oracle CODASYL DBMS Show Users command is a powerful tool you can use to monitor the current status of the master and standby databases. For example, you can use the Show Users command to find out which replication servers are active or if after-image journal backup operations have been suspended.

Example 7-1 shows the output from an RMU Show Users command entered on the master database.

Example 7-1 Show Users Display of the Master Database

```
$ RMU /SHOW USERS

database DISK1:[USER]MF_PERSONNEL.RDB;1 [1]
    - First opened 24-MAY-1996 08:41:29.71
```

```
    * database is opened by an operator
    - current after-image journal file DISK1:[USER.MASTER]AIJ_1.AIJ;1
    - Hot standby "Log Catch-Up Server" is active [2]
    - Hot standby "Log Shipping Server" is active [3]
    - AIJ Log Server is active
    - 3 active database users
    - 24404893:1 - RDMLS - non-utility, RDBVMS - active user
        - image DSA1:[SYS4.SYSCOMMON.][SYSEXE]RDMALS70.EXE;811
        - output filename "DISK1:[USER.STANDBY]ALS_25C02918.OUT;" [4]
    - 24400F47:1 - HOT4 - utility, ORUSER - active user
        - image DISK1:[RDMS_X07000.VAX.][RMU]RMU70.EXE;1
    - 24404E96:1 - RDM_LCS - non-utility, RDBVMS - active user
        - image DSA1:[SYS4.SYSCOMMON.][SYSEXE]RDMLCS70.EXE;755
        - standby database           "
"REMNOD::DISK1:[USER]STANDBY_PERSONNEL" [5]
          - Replication is "active" [6]
          - output filename "DISK1:[ORAUSER]LCS.LOG;"
          - synchronization mode "cold"
          - checkpoint interval 10 messages
          - wait interval 1 minute
          - quiet-point synchronization
```

The following list describes the callouts shown in Example 7-1:

[1] The name of the master database.

[2] The LCS process is currently bringing the standby database up to date with the master database.

[3] The ALS process on the master database is active. Notice that the "Log Shipping Server" is another name for the ALS process when replication operations are active.

[4] The name of the ALS output file (specified with the BIND_ALS_OUTPUT_FILE logical name).

[5] The name of the standby database.

[6] Indicates that replication operations are active. The lines following this one show the replication parameter settings for the master database.

Example 7-2 shows the output from an RMU Show Users command entered on the standby database.

Example 7-2 Show Users Display of the Standby Database

```
$ RMU /SHOW USERS

database _DISK1:[USER]STANDBY_PERSONNEL.RDB;1 [1]
    - First opened 24-MAY-1996 08:41:07.12
    * database is opened by an operator
```

```
- current after-image journal file is
  DISK1:[ORAUSER.STANDBY]AIJ_1.AIJ;1
- Hot standby "Log Roll-Forward Server" is active [2]
- 2 active database users [3]
- 24400F48:1 - HOT3 - utility, ORAUSER - active user
    - image DISK1:[RDMS_X07000.VAX.][RMU]RMU70.EXE;1
- 24404E97:1 - RDM_LRS70 - non-utility, RDBVMS - active user
    - image DSA1:[SYS4.SYSCOMMON.][SYSEXE]RDMLRS70.EXE;551
      - master database "_DISK1:[USER]MF_PERSONNEL.RDB;1" [4]
      - Replication is "active"
      - checkpoint interval 100 messages
      - timeout interval 5 minutes
      - access mode "read only"
      - buffer count 256
      - 2PC resolve state "commit"
      - Replication governor is enabled
```

[1] The name of the standby database.

[2] The LRS process is active.

[3] There are 2 users active on the standby database. The lines following this one show the users.

[4] The name of the master database.

7.3 Dump Header Output

The Oracle RMU and DBO Dump command displays or writes the contents of the database, storage areas, and snapshot files including database root information.

Database header information is especially useful when you want to display database parameter settings. For example, the header information can indicate when database replication is not possible because of inappropriate database attribute settings.

Because the output from the Dump command is lengthy, you can include the Header qualifier to display only the information that you need. For example, to display only the Hot Standby information, you can specify the Hot_Standby keyword on the Header qualifier.

Examples 7-3 and 7-4 show Hot Standby information in the database headers for the master and standby databases, respectively.

Example 7-3 Dump Header Display for the Master Database

```
$ RMU/DUMP/HEADER=HOT_STANDBY HOTDISK:[ORAUSER.OE_MASTER]mf_personnel
*----------------------------------------------------------------
* Oracle Rdb V7.0-00                      24-MAY-1996 08:52:38.64
```

```
*
* Dump of Database header
*     Database: DISK1:[USER]MF_PERSONNEL.RDB;1
*----------------------------------------------------------------
Database Parameters:
    Root filename is “DISK1:[USER]MF_PERSONNEL.RDB;1”
    Hot Standby...
      - Database is currently being replicated as “Master”
          Standby database is “DISK1:[USER]STANDBY_PERSONNEL.RDB;1“
          Database is local to this node (“ORANOD”)
          Replication commenced on 24-MAY-1996 08:41:53.91
          Synchronization obtained via quiet-point
          Server checkpoint interval is 10 messages
          Server connection-timeout interval is 1 minute
          Replication synchronization is “cold”
```

Example 7-4 Dump Header Display for the Standby Database

```
$ RMU/DUMP/HEADER=HOT_STANDBY DISK1:[USER]standby_personnel.rdb;1
*----------------------------------------------------------------
* Oracle Rdb V7.0-00                        24-MAY-1996 08:51:10.97
*
* Dump of Database header
*     Database: DISK1:[USER]STANDBY_PERSONNEL.RDB;1
*----------------------------------------------------------------
Database Parameters:
    Root filename is “DISK:[ORAUSER.WORK.OE.OE_STANDBY]OE_RDB.RDB;1”
    Hot Standby...
      - Database is currently being replicated as “Standby”
          Master database is “DISK1:[USER]MF_PERSONNEL.RDB;1”
          Database is local to this node (“REMNOD”)
          Replication commenced on 24-MAY-1996 08:41:52.81
          Database replication is “online”
          Server checkpoint interval is 100 messages
          Server gap-timeout interval is 5 minutes
          Server buffer count is 256
          Server 2PC transaction resolution is “commit”
```

7.4 Performance Monitor Online Statistics

Oracle Rdb and Oracle CODASYL DBMS provide Performance Monitor facilities to help you display and diagnose your database statistics, and tune your database performance. To view the current database statistics and performance information, enter the Oracle RMU or DBO Show Statistics command. The following example shows how to use this command for an Oracle CODASYL DBMS database:

```
$ DBO /SHOW STATISTICS PARTS
```

The Performance Monitor displays a summary statistics screen. You can enter the letter “M” to display a menu of options. To view information specific to your Hot Standby environment, select the Hot Standby Information submenu or the Database Dashboard submenu from the Oracle RMU or DBO Show Statistics main menu.

7.4.1 Viewing Hot Standby and Synchronization Statistics

The following example shows the menu options you can select. You can highlight the Hot Standby Information submenu by typing the letter “J” or using the arrow keys to move to the Hot Standby selection.

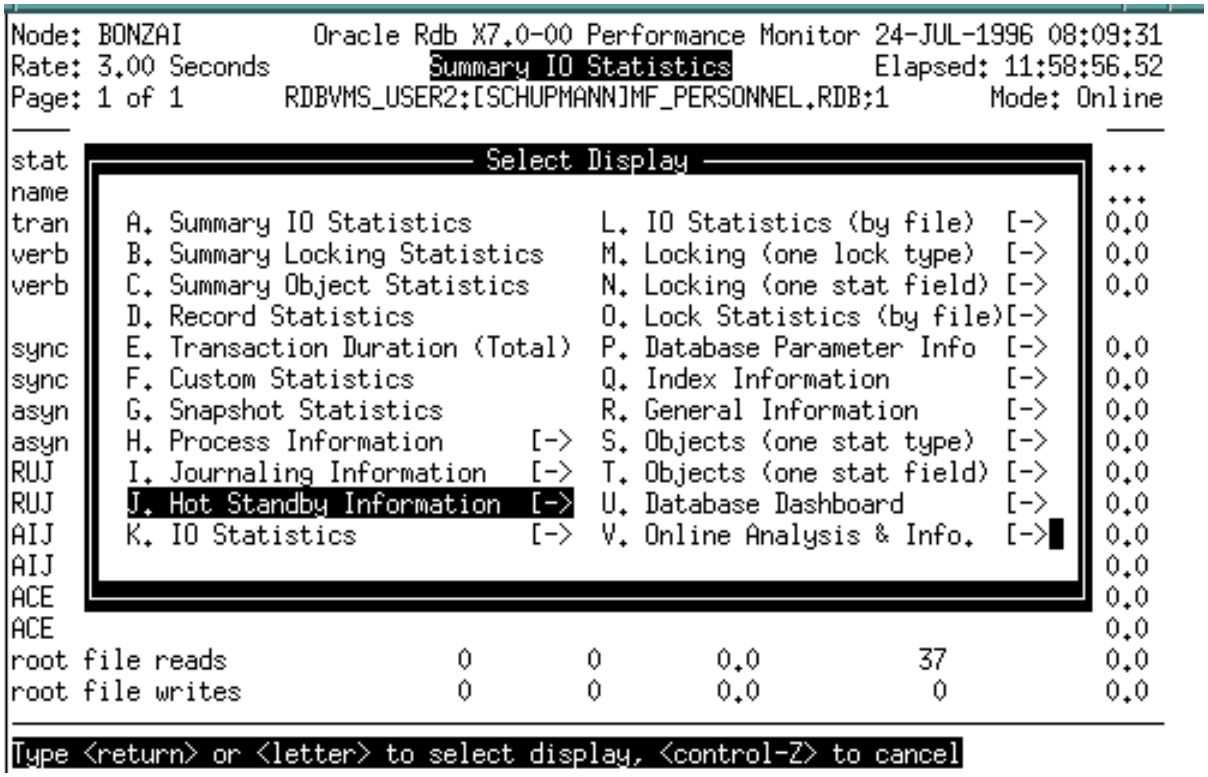

When you select the Hot Standby Information option and press the Return key, the following menu displays from which you can choose to view either Hot Standby statistics or synchronization mode statistics:

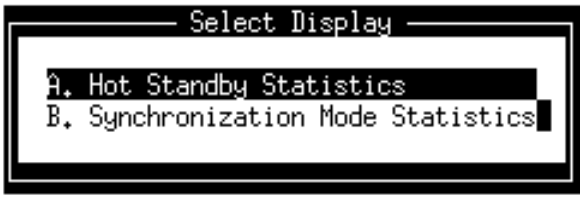

Example 7-5 shows an example of the Hot Standby Statistics display.

Example 7-5 Displaying Hot Standby Statistics

```
Node: ORANOD         Oracle Rdb V7.0-00 Performance Monitor  3-APR-1996 08:25:21
Rate: 1.00 Second            Hot Standby Statistics         Elapsed: 00:53:24.21
Page: 1 of 1            DISK1:[USER]MF_PERSONNEL.RDB;1             Mode: Online
------------------------------------------------------------------------------
State: Active     UserSync: Cold Current.Msg: 2539    Cl Mstr.AIJ: 22:4487
LagTime: 00:01:53 AutoSync: Warm Stalled.Msg:          1 Stby.AIJ: 22:3447
Stby.DB: DISK01$:[ORAUSER.WORK.OE.OE_STANDBY]OE_RDB.RDB;1

statistic.........      rate.per.second.............. total....... average......
name..............      max..... cur..... avg....... count....... per.trans....

AIJ network send              9       0       0.8         2822           3.5
AIJ network recv              7       0       0.4         1345           1.6
    data                      8       0       0.7         2494           3.1
    control                   2       0       0.1          328           0.4
    checkpoints               1       0       0.0          255           0.3
Stall time x100            1801      18      54.7       175432         220.1
blocks shipped              809      91      36.2       116248         145.8
      received                7       0       0.4         1345           1.6
Network Reconnect             0       0       0.0            0           0.0
Free Network Xmit            10       0       0.9         2508           2.4
Stalled MSN found             0       0       0.0            0           0.0
------------------------------------------------------------------------------
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Unreset Write X_plot
```

The last line of the display shows the options you can enter to control the display. For example, you can type the letter “H” to obtain online help information about any of the fields displayed. You can also type the letter “G” to obtain the Hot Standby Statistics in a graphical display.

Example 7-6 shows an example of the Synchronization Mode Statistics display.

Example 7-6 Displaying Synchronization Mode Statistics

```
Node: ORANOD          Oracle Rdb V7.0-00 Performance Monitor  3-APR-1996 08:25:29
Rate: 1.00 Second       Synchronization Mode Statistics     Elapsed: 00:53:31.40
Page: 1 of 1             DISK1:[USER]MF_PERSONNEL.RDB;1             Mode: Online
------------------------------------------------------------------------------
statistic.........      rate.per.second.............. total....... average......
name..............      max..... cur..... avg....... count....... per.trans....

transactions                  3        0        0.2          797           1.0

Cold sync send                8        0        0.4         1467           1.8
Warm sync send                7        3        0.2          961           1.2
Hot sync send                 3        0        0.0           70           0.0
Commit sync send              0        0        0.0            0           0.0

Cold stall x100               6        0        0.1          456           0.5
Warm stall x100            1088      463       33.8       108765         136.4
Hot stall x100             1010        0        2.4         7809           9.7
Commit stall x100             0        0        0.0            0           0.0
------------------------------------------------------------------------------
Exit Graph Help Menu Options Pause Reset Set_rate Time_plot Unreset Write X_plot
```

Note

In a cluster environment, the database replication information displayed by a Show Statistics command on one node might differ from the display on another node. For example, a Show Statistics display on your current node might show the standby database processing after-image journal information that is unknown by the current node.

Differences occur because delays in port-to-port communication between cluster nodes can result in slower delivery of update messages to some nodes. The lag time in message deliveries is normal and differences in Show Statistics information on any particular node should be negligible.

7.4.2 Viewing and Driving the Hot Standby Dashboard

The Performance Monitor provides a Database Dashboard facility that displays the actual database parameter and attributes settings being used by the processes attached to the database. To view information specific to your Hot Standby environment, select the Database Dashboard submenu from the Oracle RMU or DBO Show Statistics main menu.

The following example shows the menu options you can select. To choose the Database Dashboard option, highlight it by typing the letter “U” or using the arrow keys to move to that selection.

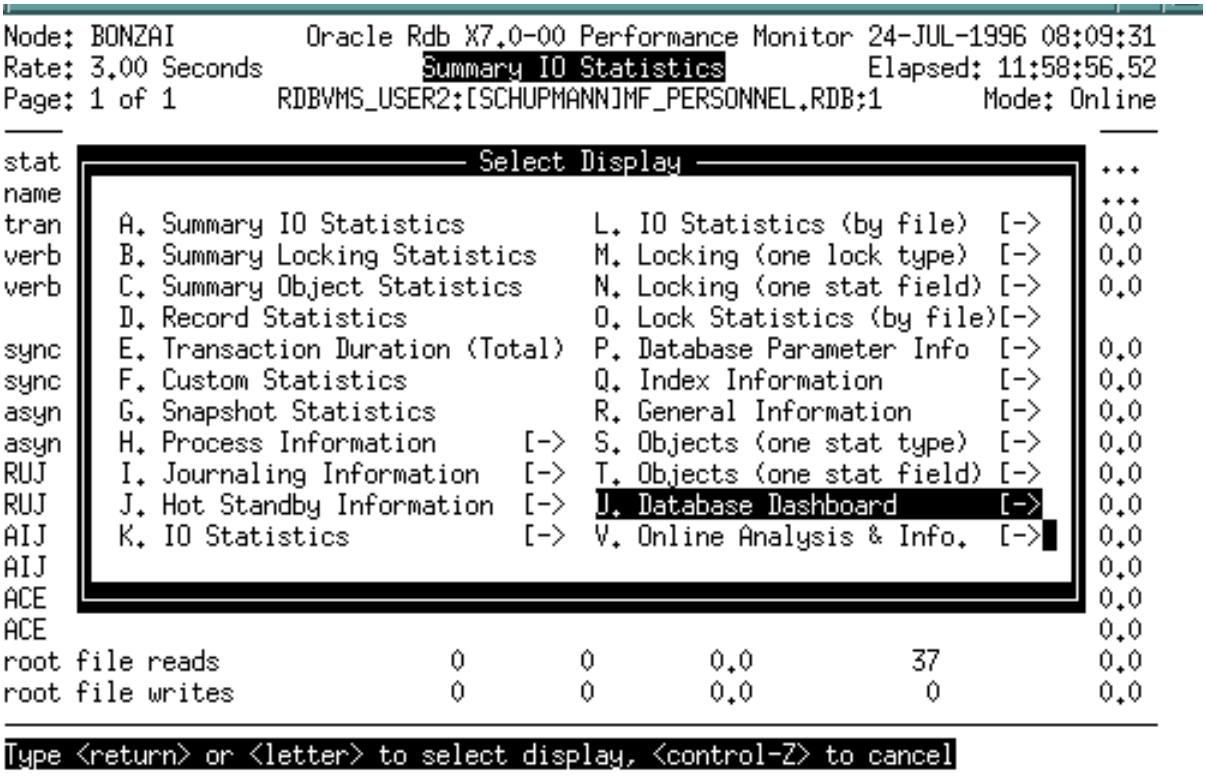

When you are viewing any of the Database Dashboard displays, you can examine logical name and configuration parameter values and other database attributes settings during run time. You can “drive” the database faster or slower by optionally making changes without having to restart database processes. The changes are in effect on a single node and only during run time.

When you select the Database Dashboard option and press the Return key, the following menu displays from which you can choose from a number of Dashboard options:

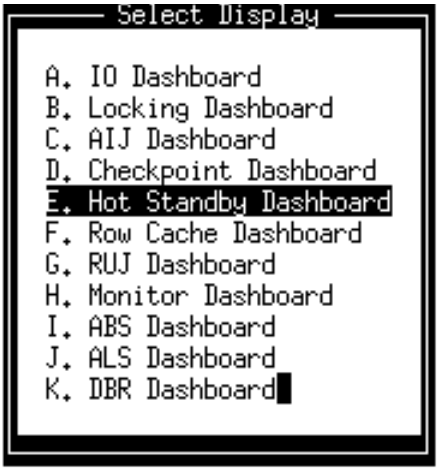

7.4.3 Monitoring Lag Time with the Hot Standby Dashboard Display

You can use the Hot Standby Dashboard display to monitor the amount of time the standby database lags behind the master database. Monitoring the synchronization between the databases can help you decide whether to try to recover the master database after a failure or fail over to the standby database.

The following Hot Standby events are tracked by the Hot Standby Dashboard:

- Network Timeout
- Connect Timeout
- Data Sync Mode
- Server Checkpoint
- Gap Timeout
- Governor Enabled

Example 7-7 shows an example of the Hot Standby Dashboard display.

Example 7-7 Displaying the Hot Standby Dashboard

```
Node: ORANOD         Oracle Rdb V7.0-00 Performance Monitor 23-JUL-1996 19:53:29
Rate: 9.00 Seconds           Hot Standby Dashboard          Elapsed: 00:37:47.21
Page: 1 of 1          DISK1:[USER]MF_PERSONNEL.RDB;1                Mode: Online
--------------------------------------------------------------------------------
Database.......... Current... Previous.. Lowest.... Highest... Original.. Chng
Attribute.Name.... Value..... Value..... Value..... Value..... Value..... Cnt.
```

```
Network Timeout [1]          120        120        120        120        120    0
Connect Timeout [2]            5          5          5          5          5    0
Data Sync Mode [3]             0          0          0          0          0    0

Server Checkpoint [4]        100        100        100        100        100    0

Gap Timeout [5]                5          5          5          5          5    0
Governor Enabled [6]           1          1          1          1          1    0
Suspend ABS [7]                0          0          0          0          0    0
--------------------------------------------------------------------------------
Config Exit Help Menu Options Set_rate Write !
```

In Example 7-7, you can use the Set_rate menu option at the bottom of the display to dynamically modify certain database parameter settings to be higher or lower and immediate see the impact of the change. The following list describes the fields in the Hot Standby Dashboard display and how you can change database attribute settings:

[1] Displays the amount of time, in seconds, after which the Hot Standby network times out. The default is 120 seconds. You can override the default with the BIND_HOT_NETWORK_TIMEOUT logical name or configuration parameter.

[2] Displays the amount of time, in minutes, to wait for the connection to be made between the master and the standby database. The default is 5 minutes. You can override the default with the BIND_LCS_CONNECT_TIMEOUT logical name or configuration parameter.

[3] Displays the current database synchronization mode:

- 0 = Cold
- 1 = Warm
- 2 = Hot
- 3 = Commit

You can override the default value with the BIND_HOT_DATA_SYNC_MODE logical name or configuration parameter.

[4] Displays the number of messages per server checkpoint interval. The default is 100 messages. You can override the default with the BIND_HOT_CHECKPOINT logical name or configuration parameter.

[5] Displays the amount of time, in minutes, to wait for stalled MSN (gap) resolution. The default is 5 minutes. You can override the default value with the BIND_LRS_GAP_TIMEOUT logical name or configuration parameter.

[6] Indicates if the replication governor is enabled or disabled.

- 0 = disabled
- 1 = enabled

You can override the default value with the BIND_LRS_GOVERNOR_ENABLED logical name or configuration parameter.

[7] Indicates if ABS backup operations are suspended.

Remember that the purpose of updating attributes is to test and measure the effects of changes on the database so that you can later make persistent changes to the appropriate database attributes.

Reference: Appendix A describes these logical names and configuration parameters in more detail.

7.4.4 Displaying Database Parameter Settings with the Monitor Dashboard

The Monitor Dashboard display shows the actual database parameter settings being used by the database. Just like the Hot Standby Dashboard display, most of the database parameter settings on the Monitor Dashboard display can be dynamically modified temporarily and nonpersistently as a means of experimenting with different values to identify their immediate effect on the overall operation of the database.

Example 7-8 shows an example of the Monitor Dashboard display.

Example 7-8 Displaying the Monitor Dashboard

```
Node: ORANOD         Oracle Rdb V7.0-00 Performance Monitor 24-JUL-1996 10:11:32
Rate: 1.00 Second              Monitor Dashboard           Elapsed: 14:00:57.89
Page: 1 of 1           DISK1:[USER]MF_PERSONNEL.RDB;1              Mode: Online
--------------------------------------------------------------------------------
Database.......... Current... Previous.. Lowest.... Highest... Original.. Chng
Attribute.Name.... Value..... Value..... Value..... Value..... Value..... Cnt.

Max DBR Count              50         50         50         50         50    0
ABS Priority               15         15         15         15         15    0
ALS Priority                8          8          8          8          8    0
DBR Priority               15         15         15         15         15    0
LCS Priority               15         15         15         15         15    0
LRS Priority               15         15         15         15         15    0
RCS Priority                4          4          4          4          4    0
--------------------------------------------------------------------------------
Config Exit Help Menu Options Set_rate Write !
```

Reference: You can obtain help on the Monitor Dashboard display and its fields using the Help menu option at the bottom of the display.

You can modify the database attributes during run time using the logical names and configuration parameters shown in Table 7-1.

Table 7-1 Setting Database Priorities with the Monitor Dashboard

Monitor Dashboard Field	Logical Names
ABS Priority	DBM$BIND_ABS_PRIORITY RDM$BIND_ABS_PRIORITY
ALS Priority	DBM$BIND_ALS_PRIORITY RDM$BIND_ALS_PRIORITY
DBR Priority	DBM$BIND_DBR_PRIORITY RDM$BIND_DBR_PRIORITY
LCS Priority	DBM$BIND_LCS_PRIORITY RDM$BIND_LCS_PRIORITY
LRS Priority	DBM$BIND_LRS_PRIORITY RDM$BIND_LRS_PRIORITY

Reference: Appendix A describes these logical names and configuration parameters in more detail.

7.5 Monitoring the Replication Server Processes

You can have the Hot Standby software record replication activities being performed by the ALS, LRS, and LCS server processes to an operational output file. Then, you can monitor the replication activities of these replication server processes by examining the appropriate log file.

You can make the Hot Standby software record this information by:

- Including the Output qualifier when you start replication operations with the Replicate After_Journal Start command. (Chapter 6 describes the Replication commands in detail.)

- Defining the logical names or configuration parameters shown in Table 7-2. (Appendix A provides more information about logical names and configuration parameters.)

Table 7-2 Logging Server Replication Activities to an Output File

To View the Replication Activities for the . . .	Oracle CODASYL DBMS	Oracle Rdb on OpenVMS	Oracle Rdb on Digital UNIX
ALS	DBM$BIND_ALS_OUTPUT_FILE	RDM$BIND_ALS_OUTPUT_FILE	RDB_BIND_ALS_OUTPUT_FILE
LCS	DBM$BIND_LCS_OUTPUT_FILE	RDM$BIND_LCS_OUTPUT_FILE	RDB_BIND_LCS_OUTPUT_FILE
LRS	DBM$BIND_LRS_OUTPUT_FILE	RDM$BIND_LRS_OUTPUT_FILE	RDB_BIND_LRS_OUTPUT_FILE

Example 7-9 provides an example of an ALS output file.

Example 7-9 Sample ALS Output File

```
--------------------------------------------------------------------------------

15-AUG-1996 09:52:06.20 - Oracle Rdb V7.0-00 database utility started

--------------------------------------------------------------------------------

This is a VAX 6000-630 running VMS V6.2

Current time is 15-AUG-1996 09:52:06.25

15-AUG-1996 06:35:35.97: Linked ALS ...
14-AUG-1996 20:20:55.80: Compiled ALS ...
14-AUG-1996 20:20:50.08: Compiled KOD$LIBRARY ...
    .
    .
    .
================================================================================
          Defined Product Logicals Information
================================================================================

    .
    .
    .
15-AUG-1996 09:52:06.50 - AIJ Log Server (ALS) startup
15-AUG-1996 09:52:07.69 - Updating dashboard information
15-AUG-1996 09:52:07.69 - Fast commit = 01
15-AUG-1996 09:52:07.69 - Commit-to-Journal = 00
15-AUG-1996 09:52:07.69 - Optimized Page Transfer = 00
15-AUG-1996 09:52:07.69 - Opening "_DISK1:[USER.MASTER]AIJ_1.AIJ;1"
15-AUG-1996 09:52:14.82 - Received LCS_STARTUP (1:0)
15-AUG-1996 09:52:25.61 - Received LSS_ACTIVE (1:0)
```

```
15-AUG-1996 09:52:25.61 - Changing LSS state from “Active” to “Net Bind”
15-AUG-1996 09:52:25.61 - Remote node: “REMNOD”
15-AUG-1996 09:52:25.61 - Standby database: “DISK1:[USER]STANDBY_PERSONNEL”
15-AUG-1996 09:52:25.61 - Server name: “554424313131245F0C000000380675”
15-AUG-1996 09:52:28.09 - Network protocol is “DECnet”
15-AUG-1996 09:52:28.09 - Sending LSS_ACTIVE (1:1)
15-AUG-1996 09:52:28.10 - Changing LSS_REF_COUNT from 0 to 1
15-AUG-1996 09:52:28.10 - AIJ Log Ship Server (ALS) activated on master database
15-AUG-1996 09:52:28.10 - sending “Attach_Req” message to server
15-AUG-1996 09:52:28.10 - Changing LSS state from “Net Bind” to “Connecting”
15-AUG-1996 09:52:28.10 - Sending “Attach_Req” (MSN 0) LRS reply expected
15-AUG-1996 09:53:10.54 - sending “Ckpt_Req” message to server
15-AUG-1996 09:53:10.55 - Checkpoint taken for inactivity
15-AUG-1996 09:53:10.55 - TROOT #3 AIJ VNO:VBN=6:1
15-AUG-1996 09:53:10.55 - Sending “Ckpt_Req” (MSN 0) LRS reply expected
15-AUG-1996 09:53:10.55 - Reading 49152 message
15-AUG-1996 09:53:10.56 - setting BUF_LEN=512
15-AUG-1996 09:53:10.56 - Received “Ckpt_Ack” reply (MSN 0)
15-AUG-1996 09:53:10.56 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 09:53:10.56 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 09:53:10.56 - Checkpointed at AIJ location 6:2
15-AUG-1996 09:53:10.57 - AIJ switch-over detected on another node
15-AUG-1996 09:53:10.57 - sending “Ckpt_Req” message to server
15-AUG-1996 09:53:10.57 - Checkpoint taken for inactivity
15-AUG-1996 09:53:10.57 - Sending “Ckpt_Req” (MSN 0) LRS reply expected
15-AUG-1996 09:53:10.57 - Reading 49152 message
15-AUG-1996 09:53:10.58 - setting BUF_LEN=512
15-AUG-1996 09:53:10.59 - Received “Ckpt_Ack” reply (MSN 0)
15-AUG-1996 09:53:10.59 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 09:53:10.59 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 09:54:57.13 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 09:54:57.13 - Checkpointed at AIJ location 7:2
15-AUG-1996 09:54:57.13 - Starting inactivity timer
15-AUG-1996 09:55:57.13 - sending “Ckpt_Req” message to server
15-AUG-1996 09:55:57.13 - Checkpoint taken for inactivity
15-AUG-1996 09:55:57.13 - Sending “Ckpt_Req” (MSN 0) LRS reply expected
15-AUG-1996 09:55:57.13 - Reading 49152 message
15-AUG-1996 09:55:57.42 - setting BUF_LEN=512
15-AUG-1996 09:55:57.42 - Received “Ckpt_Ack” reply (MSN 0)
15-AUG-1996 09:55:57.42 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 09:55:57.42 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 09:55:57.42 - Checkpointed at AIJ location 7:2
15-AUG-1996 09:55:57.42 - Starting inactivity timer
15-AUG-1996 09:56:57.42 - sending “Ckpt_Req” message to server
15-AUG-1996 09:56:57.42 - Checkpoint taken for inactivity
15-AUG-1996 09:56:57.42 - Sending “Ckpt_Req” (MSN 0) LRS reply expected
15-AUG-1996 09:56:57.42 - Reading 49152 message
15-AUG-1996 09:57:01.31 - setting BUF_LEN=512
15-AUG-1996 09:57:01.31 - Received “Ckpt_Ack” reply (MSN 0)
```

```
15-AUG-1996 10:01:25.36 - setting BUF_LEN=512
15-AUG-1996 10:01:25.36 - New AIJSERVER invocation detected
15-AUG-1996 10:01:25.36 - Received "Attach_Ack" reply (MSN 0)
15-AUG-1996 10:01:25.36 - AIJSERVER PID is 2661D013
15-AUG-1996 10:01:25.36 - Changing LSS state from "Connecting" to "Active"
15-AUG-1996 10:01:25.36 - sending "Ckpt_Req" message to server
15-AUG-1996 10:01:25.36 - Checkpoint taken for inactivity
15-AUG-1996 10:01:25.37 - TROOT #3 AIJ VNO:VBN=8:1
15-AUG-1996 10:01:25.37 - Sending "Ckpt_Req" (MSN 0) LRS reply expected
15-AUG-1996 10:01:25.37 - Reading 49152 message
15-AUG-1996 10:01:25.38 - setting BUF_LEN=512
15-AUG-1996 10:01:25.38 - Received "Ckpt_Ack" reply (MSN 0)
15-AUG-1996 10:01:25.38 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 10:01:25.38 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 10:01:25.38 - Checkpointed at AIJ location 8:2
15-AUG-1996 10:01:25.38 - AIJ switch-over detected on another node
15-AUG-1996 10:01:25.38 - sending "Ckpt_Req" message to server
15-AUG-1996 10:01:25.38 - Checkpoint taken for inactivity
15-AUG-1996 10:01:25.38 - Sending "Ckpt_Req" (MSN 0) LRS reply expected
15-AUG-1996 10:03:00.59 - Received "Ckpt_Ack" reply (MSN 0)
15-AUG-1996 10:03:00.59 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 10:03:00.59 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 10:03:00.59 - Checkpointed at AIJ location 9:2
15-AUG-1996 10:03:25.77 - sending "Ckpt_Req" message to server
15-AUG-1996 10:03:25.77 - Checkpoint taken for inactivity
15-AUG-1996 10:03:25.77 - Sending "Ckpt_Req" (MSN 0) LRS reply expected
15-AUG-1996 10:03:25.77 - Reading 49152 message
15-AUG-1996 10:03:25.78 - setting BUF_LEN=512
15-AUG-1996 10:03:25.78 - Received "Ckpt_Ack" reply (MSN 0)
15-AUG-1996 10:03:25.78 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 10:03:25.78 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 10:03:25.78 - Checkpointed at AIJ location 9:2
15-AUG-1996 10:03:25.78 - Starting inactivity timer
15-AUG-1996 10:04:25.78 - sending "Ckpt_Req" message to server
15-AUG-1996 10:04:25.78 - Checkpoint taken for inactivity
15-AUG-1996 10:04:25.78 - Sending "Ckpt_Req" (MSN 0) LRS reply expected
15-AUG-1996 10:04:25.78 - Reading 49152 message
15-AUG-1996 10:04:25.80 - setting BUF_LEN=512
15-AUG-1996 10:04:25.80 - Received "Ckpt_Ack" reply (MSN 0)
15-AUG-1996 10:04:25.80 - LAG_BLK_CNT=0, AVG_BLK_STALL_MS=0
15-AUG-1996 10:04:25.80 - LAG_TAD=   0 00:00:00.00
15-AUG-1996 10:04:25.80 - Checkpointed at AIJ location 9:2
15-AUG-1996 10:04:25.80 - Starting inactivity timer
15-AUG-1996 10:04:53.66 - Received LCS_STARTUP (0:1)
15-AUG-1996 10:04:53.66 - Received LSS_SHUTDOWN (1:0)
15-AUG-1996 10:04:53.66 - sending "Shutdown_Req" message to server
15-AUG-1996 10:04:53.67 - Sending LSS_ACTIVE (1:0)
15-AUG-1996 10:04:53.67 - Changing LSS_REF_COUNT from 1 to 0
15-AUG-1996 10:04:53.67 - AIJ Log Ship Server (ALS) deactivated for master database
```

```
15-AUG-1996 10:04:53.69 - Changing LSS state from “Shutdown” to “Inactive
```

Example 7-10 provides an example of an LCS output file.

Example 7-10 Sample LCS Output File

```
--------------------------------------------------------------------------------

16-AUG-1996 10:10:51.59 - Oracle Rdb V7.0-00 database utility started

--------------------------------------------------------------------------------
This is a VAX 6000-630 running VMS  V6.2

Current time is 16-AUG-1996 10:10:51.63

16-AUG-1996 09:19:00.81: Linked LCS ...
16-AUG-1996 08:46:21.06: Compiled LCS ...
16-AUG-1996 08:46:14.14: Compiled KOD$LIBRARY ...
    .
    .
    .
================================================================================

          Defined Product Logicals Information

================================================================================
    .
    .
    .
16-AUG-1996 10:10:51.86 - AIJ Log Catch-Up Server (LCS) activated
16-AUG-1996 10:10:51.86 - Remote node: “REMNOD”
16-AUG-1996 10:10:51.86 - Master DB: “_DISK1:[USER]MF_PERSONNEL.RDB;
16-AUG-1996 10:10:51.86 - StandbyDB: “_DISK1:[USER]STANDBY_PERSONNEL.RDB;1”
16-AUG-1996 10:10:51.86 - Server name: “”
16-AUG-1996 10:10:51.91 - Debug information initialized
16-AUG-1996 10:10:51.93 - Updating dashboard information
16-AUG-1996 10:10:51.94 - Changing LSS state from “Inactive” to “DB Bind”
16-AUG-1996 10:10:51.94 - Fast commit = 01
16-AUG-1996 10:10:51.94 - Commit-to-Journal = 00
16-AUG-1996 10:10:51.94 - Optimized Page Transfer = 00
16-AUG-1996 10:10:51.94 - Changing LSS state from “DB Bind” to “Net Bind”
16-AUG-1996 10:10:51.97 - Opening standby database _DISK1:[USER]STANDBY_PERSONNEL
16-AUG-1996 10:10:51.97 - Starting standby database replication
16-AUG-1996 10:10:55.05 - sending “Bind_Req” message to server
16-AUG-1996 10:10:55.05 - setting network timer to    0 00:02:00.00
16-AUG-1996 10:10:55.08 - “Bind_Req” message complete
16-AUG-1996 10:10:55.09 - Changing LSS state from “Net Bind” to “Connecting”
16-AUG-1996 10:10:55.09 - Sending LCS_ACTIVE (0:1)
16-AUG-1996 10:10:55.09 - Sending LRS_SHUTDOWN (0:0)
16-AUG-1996 10:10:55.09 - Automatically suspended AIJ backup operations
```

```
16-AUG-1996 10:10:55.12 - sending “Connect_Req” message to server
16-AUG-1996 10:10:55.12 - AIJDSC[0]: VNO=0 AIJID=0 ALLOC=512
16-AUG-1996 10:10:55.12 - AIJDSC[1]: VNO=1 AIJID=1 ALLOC=512
16-AUG-1996 10:10:55.12 - AIJDSC count=2
16-AUG-1996 10:10:55.12 - AIJFB_CNT=9
16-AUG-1996 10:10:55.14 - ALLOCATION=512, PEOF=512
16-AUG-1996 10:10:55.17 - ALLOCATION=512, PEOF=512
16-AUG-1996 10:10:55.23 - AIJ_SIGNATURE=00000624 FILID_SIGNATURE=000E001C
16-AUG-1996 10:10:55.23 - Sending “Connect_Req” (MSN 0) LRS reply expected
16-AUG-1996 10:10:55.23 - setting network timer to    0 00:02:00.00
16-AUG-1996 10:10:55.95 - Received “Connect_Ack” reply
16-AUG-1996 10:10:55.95 - LRS replied
16-AUG-1996 10:10:55.95 - Master database replication restarts at 1:35
16-AUG-1996 10:10:55.95 - AIJSERVER PID is 00000000
16-AUG-1996 10:10:55.95 - Resuming suspended AIJ backup operations
16-AUG-1996 10:10:55.96 - Changing LSS state from “Connecting” to “DB Synch.”
16-AUG-1996 10:10:56.06 - After-image journal 1 switch-over in progress (to 2)
16-AUG-1996 10:10:56.07 - No more after-image journals available
16-AUG-1996 10:10:56.10 - Scanning AIJ sequence 1:35-36
16-AUG-1996 10:10:56.10 - Reading AIJ sequence 1:35-36
16-AUG-1996 10:10:56.13 - sending “Sync_Data” message to server
16-AUG-1996 10:10:56.14 - Sending “Sync_Data” (MSN 6) no reply expected
16-AUG-1996 10:10:56.14 - setting network timer to    0 00:00:05.00
16-AUG-1996 10:10:56.15 - sending “Sync_Req” message to server
16-AUG-1996 10:10:56.15 - SYNC_REQ: VNO=1, VBN=35
16-AUG-1996 10:10:56.15 - Sending “Sync_Req” (MSN 0) LRS reply expected
16-AUG-1996 10:10:56.15 - setting network timer to    0 00:00:05.00
16-AUG-1996 10:10:56.16 - Received “Sync_Ack” reply
16-AUG-1996 10:10:56.16 - LRS replied
16-AUG-1996 10:10:56.16 - sending “Info_Req” message to server
16-AUG-1996 10:10:56.16 - Sending “Info_Req” (MSN 0) LRS reply expected
16-AUG-1996 10:10:56.16 - setting network timer to    0 00:00:05.00
16-AUG-1996 10:10:56.24 - Received “Info_Ack” reply
16-AUG-1996 10:10:56.24 - LRS replied
16-AUG-1996 10:10:56.24 - Verified replicated journal commit TSN 0:384
16-AUG-1996 10:10:56.24 - NUM_ACTIVE_AIJ=2
16-AUG-1996 10:10:56.24 - AIJFB_CNT=9
16-AUG-1996 10:10:56.26 - ALLOCATION=512, PEOF=512
16-AUG-1996 10:10:56.28 - ALLOCATION=512, PEOF=512
16-AUG-1996 10:10:56.28 - AIJ_SIGNATURE=00000624
16-AUG-1996 10:10:56.28 - LAG_BLK_CNT=1, AVG_BLK_STALL_MS=0
16-AUG-1996 10:10:56.28 - LAG_TAD=   0 00:00:00.00
16-AUG-1996 10:10:56.28 - Changing LSS state from “DB Synch.” to “Active”
16-AUG-1996 10:10:56.28 - Sending LSS_ACTIVE (0:1)
16-AUG-1996 10:10:56.28 - Sending LSS_REF_COUNT (0:0)
16-AUG-1996 10:12:12.95 - Received LCS_STARTUP (0:1)
16-AUG-1996 10:12:12.96 - Received LSS_SHUTDOWN (1:0)
16-AUG-1996 10:12:12.96 - sending “Shutdown_Req” message to server
16-AUG-1996 10:12:12.96 - Sending “Shutdown_Req” (MSN 0) no reply expected
```

```
16-AUG-1996 10:12:12.96 - setting network timer to    0 00:00:05.00
16-AUG-1996 10:12:12.96 - Sending LSS_SHUTDOWN (1:1)
16-AUG-1996 10:12:12.97 - Sending LCS_ACTIVE (1:0)
16-AUG-1996 10:12:12.97 - Disconnecting from tincan
```

Example 7-11 provides an example of an LRS output file.

Example 7-11 Sample LRS Output File

```
--------------------------------------------------------------------------------
27-JUL-1996 16:09:39.31 - Rdb7 V7.0-00 database utility started
--------------------------------------------------------------------------------
This is a VAXstation 3600 Series running VMS  V6.2

Current time is 27-JUL-1996 16:09:39.60

27-JUL-1996 15:25:31.60: Linked LRS ...
27-JUL-1996 06:41:23.11: Compiled LRS ...
27-JUL-1996 06:41:14.95: Compiled KOD$LIBRARY ...
    .
    .  (System and configuration information has been omitted from this example.)
    .
--------------------------------------------------------------------------------
27-JUL-1996 16:09:40.59 - AIJ Log Roll-Forward Server (LRS) activated
27-JUL-1996 16:10:14.00 - Received "Bind_Req" from 66000A9B
27-JUL-1996 16:10:14.13 - Sending "Bind_Ack" to 66000A9B
27-JUL-1996 16:10:14.33 - Received "Connect_Req" from 66000A9B
27-JUL-1996 16:10:15.05 - Received "Sync_Data" from 66000A9B
27-JUL-1996 16:10:17.75 - Received "Sync_Data" from 66000A9B
27-JUL-1996 16:10:17.76 - Received "Sync_Data" from 66000A9B
27-JUL-1996 16:10:47.15 - Received "Sync_Req" from 66000A9B
27-JUL-1996 16:11:01.64 - Received re-submitted "Sync_Req" from 66000A9B
```

8
Handling Failure Conditions

A **failure** is the inability of a computing component to perform its function correctly, due to one or more internal faults whose effects cannot be contained. When an extended failure prevents continuous transaction processing, you must determine how to quickly recover with minimal data loss and downtime. For example, can you quickly restart the master database or should you *fail over* to the standby database?

Failover is the ability to reconfigure a computing system to utilize an alternate active component when a similar component fails. The Hot Standby software keeps both the master and standby databases identical and synchronized so that, at any point in time, you can fail over processing to the standby database.

8.1 Detecting Failures

How do you know when something has gone wrong? It is critically important that you use one or more of the following methods to continuously review the state of replication operations:

- On OpenVMS systems, enable the operator notification capability (described in Section 2.3.6).

  The Hot Standby software performs database replication using detached server processes, so operator notification is the best mechanism available to notify the database administrator in the event of server failure or replication termination.

- On OpenVMS and Digital UNIX systems, use monitoring tools such as the Show Users command and the Show Statistics utility (described in Chapter 7) to view the current database state and performance information.

  As a DBA or system manager, you need to use tools to monitor and report on error conditions so that component faults do not interfere with requirements for 24x365 availability. Also, for Digital UNIX systems or other systems that do not provide an operator notification facility, Oracle Corporation recommends displaying system and database information to detect failures.

- On OpenVMS and Digital UNIX systems, enable an operational output file that records server activities when replication operations are active (described in Section 6.3).

  The Hot Standby software is capable of maintaining an operational output file, similar to the monitor log file, for each replication server process. The output file contains important up-to-date information about the database replication operations. You can record operational output by specifying the Output qualifier on the Replicate After_Journal Start command, or by using a logical name or configuration parameter (for example, RDM$BIND_LRS_OUTPUT_FILE for LRS process output or DBM$BIND_HOT_OUTPUT_FILE for AIJSERVER process output).

  Reference: See Appendix A for complete information about logical names and configuration parameters you can use to record operational output.

- On OpenVMS and Digital UNIX systems, diagnose problems that occur when you start up or shut down replication operations by setting the value of the BIND_KODA_TRACE logical name or configuration parameter to ACYZM (for example, RDM$BIND_KODA_TRACE=ACYZM). Define the logical name or configuration parameter on both the master and standby database nodes.

  Setting this logical name or configuration parameter allows the replication servers to output diagnostic information to the corresponding output file. Then, following a failure condition, you can examine the output file to find comprehensive details regarding the cause of the failure. Because the BIND_KODA_TRACE logical name or configuration parameter produces long output files, you should disable it except when you need to perform troubleshooting.

8.2 Events That Trigger Failures

In general, the master database site fails as a result of a hardware or software failure, a catastrophic disaster, or when there is a need to separate the current master and standby databases. (For example, you might want to replicate the master database to another standby database site.)

Replication failures on the standby database occur typically because the database is modified by a read / write transaction while replication operations are in progress. Updating the standby database stops replication operations, and, once modified, the standby database cannot be resynchronized with the master database.

Section 8.3 provides methodologies for handling and recovering from all types of failures.

Section 8.4 describes specific failures that can affect the Hot Standby configuration and provides information about recovering from each failure.

8.3 Recovering from Failures

What do you do when something goes wrong? If a failure occurs, you must quickly identify the problem and decide the best recovery strategy. The following list describes recovery options:

- Shut down database replication temporarily and restart replication operations later using the same (original) master database and standby database configuration.

  When you restart replication operations, the standby database attempts to resynchronize itself with the master database while the master database continues transaction processing.

- Fail over to the standby database and use it as the new master database. Update operations can continue on the standby database despite master database system failures.

  In many cases, you can resolve failures at the master database site within a reasonable amount of time. However, failures that require complete restore and recovery operations can be resolved more quickly by failing over to the standby database.

There are always trade-offs and decisions to be made when you must recover from a failure. Because recovery options do not occur instantaneously or automatically, it is likely that you must base your decision on which database you can restart and access most quickly. The following sections discuss recovery procedures in more detail.

8.3.1 Shutting Down and Restarting Replication Operations

Recovery from some failures might require that you shut down database replication operations until you can perform database backup and restore operations, and restart replication operations. This section describes how you can shut down database replication operations in either an orderly manner or rapidly for emergency conditions:

- Orderly shutdown—terminates replication operations after rolling forward any unprocessed after-image journal information
- Rapid shutdown—immediately terminates replication on both the master and standby databases, discarding any active transactions on the standby database

You can perform an orderly shutdown by entering the Replicate After_Journal Stop command on either the master or standby database. The following table describes the actions the Hot Standby software takes when you enter the Replicate After_Journal Stop command to perform an orderly shutdown:

Database where you enter the command:	Replication on the master database . . .	Replication on the standby database . . .
Master database	Terminates immediately	Terminates after rolling forward any unprocessed after-image journal information. When the after-image journals are rolled forward, active transactions on the standby database are aborted.
Standby database	Terminates after notifying any LCS or ALS processes on the master database	Terminates after: • The master database server processes are notified that a replication shutdown is going to take place • Rolling forward any processed after-image journal information After all rollforward operations complete, active transactions on the standby database are aborted.

For rapid shutdown, you include the Abort qualifier when you enter the Replicate After_Journal Stop command. The following table describes the actions the Hot Standby software takes to perform a rapid shutdown:

Database where you enter the command:	Replication on the master database . . .	Replication on the standby database . . .
Master database	Terminates immediately	Terminates immediately. Unprocessed after-image journal information on the standby database is discarded, and active transactions on the standby database are aborted.
Standby database	Terminates immediately	Terminates immediately. Unprocessed after-image journal information on the standby database is discarded, and active transactions on the standby database are aborted.

For availability, you should restart replication operations as soon as possible. You can restart replication operations using the same master and standby databases. The Hot Standby software attempts to resynchronize the standby database with the master database, assuming that the necessary after-image journals are available on the master database.

If the after-image journals are not available, perform either of the following operations to restart replication operations:

- Perform a recovery operation on the standby database to manually roll forward the master database after-image journal backup files onto the standby database.
- Perform a backup operation on the master database and restore the backup file on the standby database node.

See Chapter 2 for complete information about starting replication operations.

8.3.2 Failing Over to the Standby Database

In some cases, the decision to fail over to the standby database is obvious. For example, if the master database fails and is not immediately recoverable, you should fail over all processing to the standby database. Failures such as loss of the current rollforward log, loss of the system or rollback tablespace, or loss of the entire master node and database are all valid reasons to fail over to the standby database.

Database failover is a manual operation in which the standby database becomes a new master database that is no longer synchronized with the original master database.

The following sections provide suggestions for failover in a Hot Standby environment. Once you successfully fail over to the standby database, you should restart replication operations as soon as possible to reestablish the disaster tolerance and fault resilience provided by the Hot Standby software.

8.3.2.1 Temporary Versus Permanent Failover

When you fail over to the standby database, you must decide if the standby database is going to serve as the master database temporarily or permanently:

- A temporary failover means that you want to eventually fail back to the original master database (for example, after the failure situation is resolved).

  The standby database takes over processing until you choose to *fail back* to the original database system. Failback is achieved by performing a failover process in reverse. There is no automatic failback mechanism; failback requires that you follow failover procedures (described in Table 8-1) to return processing back to the original database system.

- A permanent failover means that you want to continue processing on the standby database indefinitely.

  Standby databases are based on the theory that two identical database sites are connected by a network. The standby site should be identical in configuration and resources to the master database site. That way, if a failover is required, the standby database site can continue processing transparently and indefinitely.

8.3.2.2 Failover Procedure

Table 8-1 provides a checklist of tasks that you can perform to fail over processing to the standby database. Each step includes comments that provide additional information about the step.

Table 8-1 Failover Checklist

Step	Procedure	Comments
[1]	Before you fail over to the standby database, transfer and roll forward all available after-image journals from the master database to the standby database.	Applying the after-image journals to the standby database makes it current to the same point in time and transactional history as the master database before the failure occurred. If you cannot apply the after-image journals to the standby database, you must fail over processing without recovering the transactions from the master database.
[2]	Define logical names (for OpenVMS systems) or configuration parameters (for Digital UNIX systems) with which you can specify the master and standby databases.	On clustered systems, define the logical names and configuration parameters on each node. For example, the names you define could be in the following format: *database-name*_ALS_MASTER_SERVER *database-name*_ALS_STANDBY_SERVER
[3]	On each node, write two procedures for each database: one procedure starts replication on the master node and one starts replication on the standby node. Each procedure should translate the logical names or configuration parameters you created in step 2.	When you start the boot process, you can set the logical names or configuration parameters to start up the proper replication configurations (master or standby) on the local node. As you proceed through the boot process, you can easily change the values depending on how the same logical names or configuration parameters are defined on the remote node.
[4]	Check to see if replication operations are already running on the standby database.	Use the Show Users command as described in Chapter 7.

Table 8-1 Failover Checklist (Continued)

Step	Procedure	Comments
[5]	Switch over all transaction processing so that applications access the standby database.	The Hot Standby configuration allows you to enable the standby database to function as the primary copy until the failure is resolved.
[6]	Back up the "new" master database.	Because the database is no longer considered a standby database, it should be backed up immediately. (Also, if you are ready to restart replication activities, you can use the backup file to create a new standby database.)
[7]	Restart replication operations on the standby database following the steps outlined in Table 2-1.	At this point, the standby database is the new master database.

Because all database failover activity must occur with a minimum of downtime and manual effort, you might want to automate some of the work described in Table 8-1 to provide quick recovery, flexibility, and resilience.

8.4 Failure Scenarios and Recovery

The following sections discuss the various database replication failure scenarios and describe possible recovery solutions to each failure. The recovery scenarios include automatic recovery by some component of the Hot Standby software, or manual recovery by the database administrator (DBA).

8.4.1 Network Failures

This section discusses the various types of network failures that affect the Hot Standby configuration and describes how to recover from each failure. Because the Hot Standby software depends on networking, failure scenarios are visible as a network problem. Table 8-2 describes possible network failures.

Table 8-2 Network Failures and Recovery Scenarios

Failure	Description
Network Disconnect	One of the most prevalent network problems is when the network connection is "dropped," either temporary or permanently. When the network connection is dropped, any process trying to receive a message receives an "EOF" error. Any process trying to send a message receives a "SEVERED" error.
False Network Disconnect	Failure of the AIJSERVER process appears to both the master and standby databases as a false network disconnect problem. Section 8.4.5 discusses the AIJSERVER process failure problem.
Network Hang	The second most common network problem is when a message send or receive request never returns control to the calling process. This condition, sometimes referred to as a network *hang*, is caused typically by network hardware problems.
	You can include the Gap_Timeout qualifier on the Replicate After_Journal Start command to identify potential network hang situations.
Network Timeout Failures	In the event of a network request timeout, the Hot Standby software assumes the network connection is unavailable. Because the Hot Standby software relies on the network connection to transmit after-image journal information, your only course of action is to terminate database replication operations. Oracle Corporation recommends that you use the Replicate After_Journal Stop command.
Network Performance	Slow network performance, even if it is for the duration of a single transaction being sent across the network, can dramatically affect the relative synchronization of the master and standby databases. This is especially true if you specified the Synchronization=Cold qualifier. Oracle Corporation recommends that you use the Governor=Enabled qualifier on the Replicate After_Journal Start command.
	The replication governor reduces the database differential by dynamically adjusting the master database synchronization mode to allow the standby database to maintain timely synchronization.
Standby Database Performance	Slow standby database performance, even if it is for the duration of a single transaction being sent across the network, can affect the performance of the standby database platform and of the rollforward operation.
	The goal of the replication governor is to reduce the database differential by dynamically adjusting the master database synchronization mode to a level where the replicated database is able to maintain timely synchronization.

8.4.2 Cluster Failures

This section discusses the various types of cluster failures that affect the Hot Standby configuration and how to recover from each failure. Table 8-3 describes possible cluster failures.

Table 8-3 Cluster Failures and Recovery Scenarios

Failure	Description
Local Master Database Cluster	When both the master and standby database reside on nodes in the same cluster, the viability of using the Hot Standby software for a true "hot standby" database is severely compromised. Cluster failure provides no capability to fail over the database. You must recover both the master and standby databases when you reopen the databases. It is likely that the recovery operation for the standby database will proceed extremely quickly because there are no update processes other than the LRS process, which does not have a recovery-unit journal file. Therefore, you might prefer to fail over to the standby database while the master database is being recovered.
Remote Master Database Cluster	The failure of the master database is exactly the situation for which you should implement the Hot Standby feature. Typically, total cluster failure occurs for one of two reasons: catastrophic disaster (natural or otherwise), or cascading node failure. Note that catastrophic disaster seldom occurs, while cascading node failure is fairly prevalent. • In a catastrophic disaster, it is unlikely that you can recover the master database in the near term. Therefore, failing over to the standby database is your only recovery option. • In a cascading node failure (which leads to total cluster failure), you must make a difficult decision. Rebooting the cluster often corrects the situation that caused the cluster failure. However, the time required to recover the master database could be substantial, depending on the update activity in progress at the time of the failure. The standby database detects the master database cluster failure via one of the network failure conditions, typically network disconnect. Database replication is terminated when there is a network disconnect.
Remote Standby Database Cluster	Failure of the standby database cluster is not as severe a situation as failure of the master database cluster, because ongoing database activity is not directly affected. However, it does leave the master database vulnerable to failure without chance of failover. The proper course of action is to restart the database replication operation when the standby database cluster is rebooted. It is also possible to "transfer" the database replication operation to a third cluster, or (as a last resort) to the same cluster as the master database.

8.4.3 Node Failures

This section discusses the various types of node failures that affect the Hot Standby configuration and how to recover from each failure. Table 8-4 describes possible node failures.

Table 8-4 Node Failures and Recovery Recommendations

Failure	Description
Master Database Node	Failure of a single node in the master database cluster should not cause database replication to terminate. The database recovery process (DBR) can recover the failed LCS or ALS process on another node. For the ALS process, the DBR process does not perform much additional work. Because the ALS process does not store the last allocated message sequence number in a nonvolatile storage buffer, the DBR process must determine the last allocated message sequence number by examining the after-image journals on the master database. The DBR process does this by starting from the ALS process' last checkpoint location that is stored in the RTUPB data structure. To do this, the DBR process finds the last successfully allocated message sequence number, by acquiring the message lock value block in concurrent retrieval mode. The last successfully allocated message sequence number identifies the end of the search of the after-image journals whose beginning location is identified in the RTUPB data structure. Also contained in the message lock value block is the process ID of the last ALS process to successfully update the message sequence number. If the lock value block process ID (PID) matches the PID of the failed ALS process located in the RTUPB data structure, then you must invalidate this message sequence number. Otherwise, the DBR process reads the master database after-image journals, starting at the failed ALS process' checkpoint location. Any gap in sequence numbers might be caused by the failed ALS process, and must be invalidated. In this case, you cannot invalidate the last allocated message sequence number because that might cause a valid after-image journal message buffer to be improperly discarded by the LRS process if it is to receive the invalidate message prior to the arrival of the real message packet.
Standby Database Node	Because database replication operates on a single node of the standby database cluster, failure of any other node in the standby database cluster cannot adversely affect database replication operations. However, a node failure on the node where you invoked the LRS process is similar to standby database cluster failure. Database replication is immediately terminated, primarily as a result of the network disconnect that occurs.
Database Monitor	You manage database monitor failure identically to the way you manage master and standby database node failures.

8.4.4 Hardware Failures

This section discusses the various types of hardware failures that affect the Hot Standby configuration and how to recover from each failure. Table 8-5 describes possible hardware failures.

Table 8-5 Hardware Failures and Recovery Scenarios

Failure	Description
Master Database Disk	The failure of one or more storage area devices on the master database should not cause database replication failure, because no more changes can be made to the failed storage areas. Furthermore, the ability to perform online "by-area" or "by-page" database restore and recover operations on the master database is supported even during active database replication. There is no need to resynchronize the failed areas with the standby database, because, in effect, the restore and recover operations performs this function.
Standby Database Disk	The failure of one or more storage area devices on the standby database causes the Hot Standby software to immediate terminate database replication. Although the standby database is not functioning properly, updates to the master database continue uninterrupted. Once the storage area devices are back online, the affected areas on the standby database are automatically resynchronized with the master database, if the necessary after-image journals are still available on the master database. You might be required (or it might be faster) to perform a manual "by-area" database restore and recover operation to the affected areas of the standby database, using the backed up master database after-image journals.

8.4.5 Process Failures

This section discusses the various types of process failures that affect the Hot Standby configuration and how to recover from each failure. Table 8-6 describes possible process failures.

Table 8-6 Process Failures and Recovery Scenarios

Failure	Description
DBR Process	In all process failure scenarios, the failed process is recovered by the database recovery process (DBR). If the DBR process fails, then the database is immediately shut down and cannot be accessed until the DBR process successfully completes the recovery operation.
Application Process	The Hot Standby software has no impact on application user process failures. User process failures are handled as they are handled for systems without the Hot Standby option.
	However, the master database has some special processing requirements for user process failure. Because of the database freeze that is in effect while the user process is being recovered, the ALS process is unable to process the transaction resolution determined by the DBR process. The DBR process commits a prepared two-phase commit transaction if it is directed to do so by the DECdtm software, or the DBR process rolls back the transaction if database modifications have been made by the failed user process. The DBR process transmits the appropriate after-image journal record to the LRS process on the standby database.
ALS Process	The ALS process failure is described in the topic about master database node failure in Section 8.4.3.
LCS Process	The recovery of the LCS process requires special handling by the DBR process. Failure of the LCS process always terminates database replication on the master database; it does not affect the LRS process on the standby database.
LRS Process	The recovery of the LRS process requires special handling by the DBR process. Failure of the LRS process always terminates database replication on both the master and standby databases. The master database is notified due to network disconnect failure.
	If the DBR process fails, then the database is recovered when the database is reopened.
AIJSERVER Process	Because the AIJSERVER process is not a database process and runs in user mode on OpenVMS systems, explicit recovery operations are unnecessary. The failure of the AIJSERVER process appears to both the master and standby databases as a false network disconnect problem.
Monitor Process	The failure of the database monitor is identical to failure of the database node. Refer to Section 8.4.3 for complete discussions on failure recovery operations.

8.4.6 Resource Failures

This section discusses the various types of resource failures that affect the Hot Standby configuration and how to recover from each failure. Table 8-7 describes possible resource failures.

Table 8-7 Resource Failures and Recovery Scenarios

Failure	Description
Process Quota Exhaustion	On OpenVMS systems, the Hot Standby server processes invoked by the database monitor have exceptionally large quotas and rarely fail due to process quota exhaustion. The effect of a Hot Standby server process failure on replication operations is described in Section 8.4.5.
Disk Quota Exhaustion	On OpenVMS systems, exhausting the process disk quota causes the affected process to either terminate or mark the specific database page as corrupt. Do not use disk quotas on any database device.
Disk Space Exhaustion	Failure of any storage area or after-image journal extension due to inadequate disk space causes immediate process termination. The use of extensible after-image journals requires that the after-image journal on the standby database be extended at run time. This is because new after-image journal messages are received beyond the current logical end of file. However, this does not affect the recoverability of the standby database.

8.4.7 Rollforward Failures

The replication of physical database operations, such as the addition of new after-image journals or storage areas, can cause database replication failures. To avoid problems during rollforward operations, ensure that the physical file specifications are unique and translatable on both the master and standby databases.

Other rollforward failures are often related to:

- Corruption of the after-image journal either:
  - On the master database, resulting in corruption of the standby database
  - Because of a failure on the standby database
- Unresolved distributed transactions that cause rollforward problems on the standby database.
- Receiving after-image journal messages out of order on the standby database because of the asynchronous nature of the network transmissions. Failure of the standby node during an out-of-sequence message can prevent you from backing up the standby database after-image journals.

Table 8-8 describes possible rollforward failures.

Table 8-8 Rollforward Failures and Recovery Scenarios

Failure	Description
After-Image Journal Wraparound	After-image journal *wraparound* occurs on the standby database when a process attempts to write an entry to an after-image journal that has not finished rolling forward all of its current transactions. This problem usually occurs when the Synchronization qualifier mode is set to either Cold or Warm. The problem occurs only for new or reused after-image journals whose after-image journal sequence number is unknown to the LRS process. (A reused journal is an after-image journal that becomes available after a backup operation.) To resolve this problem, I/O operations must wait for the following information to be rolled forward: • For a new after-image journal, the standby database must roll forward the after-image journal creation record and also create a new after-image journal • For a reused after-image journal, the standby database must wait until all pending rollforward records have been processed
New After-Image Journal Creation	The creation of new after-image journals occurs as an online operation on the master database. To successfully duplicate the creation of new after-image journals on the standby database, the file name specification requires that you use concealed system logicals on OpenVMS systems and unexpanded environment variables on Digital UNIX systems.
New Storage Area Creation	The creation of new storage areas occurs as an online operation on the master database. To successfully duplicate the creation of the new storage areas on the standby database, the file name specification requires that you use concealed system logicals on OpenVMS systems and unexpanded environment variables on Digital UNIX systems.
Two-Phase Commit Transactions	For unresolved transactions (for example, if a failure occurs on the master database while it is waiting for a transaction to commit or roll back), the standby database tries to resolve the unfinished transaction. However, the standby database typically is unable to resolve two-phase commit transactions because they are normally resolved on the master database; the standby database is only one part of the distributed transaction. To help the LRS process on the standby database resolve transactions, specify the preferred transaction resolution (for example, the COMMIT statement) in the application program. However, because some transactions might require a commit resolution while others require a rollback resolution, this strategy might not work for all applications. For example, if you specify a COMMIT statement in your application, the application applies the commit resolution to *all* unresolved two-phase commit transactions. Use discretion if you decide to resolve all transactions from within the application program, because a single COMMIT or ROLLBACK statement might not be the correct action to take in all cases.

Table 8-8 Rollforward Failures and Recovery Scenarios (Continued)

Failure	Description
Message Sequence Number Gap Timeout	Under most circumstances, there are no gaps in the message sequence numbers of after-image journal messages sent from the master to the standby database. However, during node or cluster failure, it is possible that a gap might occur due to the loss of information stored in the affected node's global section. To guarantee that a gap in the message sequence number never affects rollforward operations, the Hot Standby software performs several proactive and reactive operations. The LCS process proactively queries the LRS process for the state of the standby database. If there are any gaps, the LRS (reactively) notifies the ALS processes of any gaps. The ALS processes then notifies the LCS process.
Standby Database After-Image Journals Hard Data Loss Resolution	A failure of the standby database after-image journal subsystem is likely to make the affected after-image journal inaccessible and possibly mark the after-image journal subsystem as having hard data loss. To resolve this problem, back up the database. However, you cannot back up the standby database while replication is active. Therefore, you must stop database replication prior to performing the database backup operation.
Loss of Database Synchronization	Loss of database synchronization occurs when the last commit TSN stored in each of the master database's after-image journal open record does not match the last commit TSN stored in the standby database's corresponding after-image journal open record. This failure usually indicates that either the replication operation somehow missed applying a master database modification to the standby database or a master database modification was not properly journaled.

8.4.8 Other Failures

Table 8-9 describes a few additional failures that can affect replication operations and how to recover from each of the failures.

Table 8-9 Other Failures and Recovery Scenarios

Failure	Description
Database Freeze	It is possible that a database freeze operation of an extremely long duration could exceed the network operation timeout interval and create the impression of a false network hang. Because the after-image journaling fast commit optimization means that transaction redo is possibly performed by the DBR process on the master database, a redo operation that takes an exceptionally long time, or an exceptionally long duration undo operation, is possible. Proper database tuning can prevent this failure.
TSN Not Found	A TSN mismatch occurs only during the restart synchronization phase of the database replication operation. A TSN mismatch occurs because of an unjournaled database modification on either the master or standby databases, which causes the "Last Commit TSN" on the master database to differ from the "Last Commit TSN" on the standby database.
Modified Database	Any unjournaled database modification makes the database replication operation invalid.

A

Logical Names and Configuration Parameters

This appendix describes the Hot Standby logical names and configuration parameters available to manage database replication operations.

A.1 Facility Prefixes

Except for the facility prefix, the logical names and configuration parameters are the same regardless of the database or operating system that you use.

Table A-1 shows the prefixes that apply to the logical name or configuration parameter for each database and operating system.

Table A-1 Prefixes for Logical Names and Configuration Parameters

Database and Platform	Prefix	Example
Oracle CODASYL DBMS on OpenVMS	DBM$	DBM$BIND_AIJ_WORK_FILE
Oracle Rdb on OpenVMS	RDM$	RDM$BIND_AIJ_WORK_FILE
Oracle Rdb on Digital UNIX	RDB_	RDB_BIND_AIJ_WORK_FILE

The logical names and configuration parameters described in this appendix are completely compatible with other operating system logical names and configuration parameters that you define for standard database operations.

A.2 Logical Name Tables and Configuration Files

For Oracle Rdb or Oracle CODASYL DBMS databases running on OpenVMS systems, define the logical names in the LNM$SYSTEM_TABLE logical name table.

For Oracle Rdb databases on Digital UNIX systems, define the configuration parameters in the `rdb.conf` configuration file.

Also, you can define many of the logical names and configuration parameters using the Database Dashboard facility available with the RMU or DBO Show Statistics command.

Note

Oracle Corporation recommends that you define all logical names or configuration parameters before opening the master and standby databases. The logical names or configuration parameters are translated when you open each database, and the values are stored in the Database Dashboard facility (described in Chapter 7).

Defining logical names and configuration parameters before opening the database is especially useful in a multiple-database configuration; you can customize each database configuration by dynamically assigning, opening, and then deassigning the logical names or configuration parameters before you open another database.

Reference: See Chapter 7 for more information about the using the Database Dashboard facility.

Table A-2 provides a listing of the logical names and configuration parameters for Hot Standby configurations, and includes a brief description of each.

Table A-2 Logical Names and Configuration Parameters

Name	Defines
DBM$BIND_ABS_LOG_FILE RDM$BIND_ABS_LOG_FILE RDB_BIND_ABS_LOG_FILE	The location and name of the informational output file for the after-image journal backup server (ABS). The default file extension is .log.
DBM$BIND_ABS_PRIORITY RDM$BIND_ABS_PRIORITY	The base priority at which the ABS process is invoked by the database monitor. Define this logical name on OpenVMS systems only. If you do not define this logical name, the current priority of the database monitor is used.
DBM$BIND_AIJ_WORK_FILE RDM$BIND_AIJ_WORK_FILE RDB_BIND_AIJ_WORK_FILE	The directory location of the temporary after-image journal rollforward work files. Define this logical name or configuration parameter for the standby database only. If you do not include a file extension, it defaults to .tmp. See Section 2.11 for more information.

Table A-2 Logical Names and Configuration Parameters (Continued)

Name	Defines
DBM$BIND_ALS_OUTPUT_FILE RDM$BIND_ALS_OUTPUT_FILE RDB_BIND_ALS_OUTPUT_FILE	The location and name of the informational output file for the AIJ log server (ALS). You can include "_PID" in the output file specification to create a unique file name that includes the process identification (PID). For example: `$ DEFINE DBM$BIND_ALS_OUTPUT_FILE ALS_PID.OUT` This creates a unique file name that includes the PID, for example, `DISK1:[USER]ALS_25C02914.OUT.` The default file extension is .out. Alternatively, you can define the location of the output file using the Output qualifier on the Replicate After_Journal Start command. The Output qualifier overrides definitions you make with the logical name or configuration parameter.
DBM$BIND_ALS_PRIORITY RDM$BIND_ALS_PRIORITY	The base priority at which the AIJ log server (ALS) process is invoked by the database monitor. This logical name is applicable to OpenVMS systems only. If you do not define this logical name, the Hot Standby software uses the current priority of the database monitor.
DBM$BIND_DBR_LOG_FILE RDM$BIND_DBR_LOG_FILE RDB_BIND_DBR_LOG_FILE	The location and name of the database recovery process (DBR) log file. The default file extension is .log.
DBM$BIND_DBR_PRIORITY RDM$BIND_DBR_PRIORITY	The base priority at which the database recovery process (DBR) is invoked by the database monitor. This logical name is applicable to OpenVMS systems only. If you do not define this name, the Hot Standby software uses the current priority of the database monitor.
DBM$BIND_DBR_WORK_FILE RDM$BIND_DBR_WORK_FILE RDB_BIND_DBR_WORK_FILE	The directory location of the temporary DBR redo work file. If you do not define this logical name or configuration parameter, the default location is the recovery-unit journal directory. If you do not include a file extension, it defaults to .tmp.
DBM$BIND_HOT_CHECKPOINT RDM$BIND_HOT_CHECKPOINT RDB_BIND_HOT_CHECKPOINT	How frequently, in terms of processed messages, the Hot Standby software performs a checkpoint operation to update information about the database in the root file. Note the following: • Define this logical name or configuration parameter on the master database node, the standby database node, or both nodes • Specify a value from 1 to 1024 messages • The default checkpoint interval is 100 messages Alternatively, you can specify the Checkpoint qualifier on the Replicate After_Journal Start command.

Table A-2 Logical Names and Configuration Parameters (Continued)

Name	Defines
DBM$BIND_HOT_DATA_SYNC_MODE RDM$BIND_HOT_DATA_SYNC_MODE RDB_BIND_HOT_DATA_SYNC_MODE	The degree to which you want to synchronize committed transactions on the standby database with committed transactions on the master database. Define this logical name or configuration parameter on the master database system only. Valid settings include the following values: Value — Synchronization Mode 3 — Commit 2 — Hot 1 — Warm 0 — Cold (default) Alternatively, you can specify the Synchronization qualifier on the Replicate After_Journal Start command.
DBM$BIND_HOT_DECNET_OBJECT RDM$BIND_HOT_DECNET_OBJECT RDB_BIND_HOT_DECNET_OBJECT	The name of the DECnet network object for the AIJSERVER process on the remote standby database. Define the DECnet object names as follows: • Specify DBMAIJSRV for Oracle CODASYL DBMS databases. Alternatively, you can specify the DECnet object number value 52 for DECnet for OpenVMS. • Specify RDMAIJSRV for Oracle Rdb databases. Alternatively, you can specify the DECnet object number 0 for DECnet for OpenVMS (Phase IV) or DECnet object number 35 for DECnet/OSI.
DBM$BIND_HOT_NETWORK_TIMEOUT RDM$BIND_HOT_NETWORK_TIMEOUT RDB_BIND_HOT_NETWORK_TIMEOUT	The maximum number of seconds that replication operations wait for a network transmission between the master and standby databases. You can specify a value from 30 to 1800 seconds (30 minutes) on the master and standby databases. The default timeout interval is 120 seconds. If a network transmission does not occur in the time allowed, replication operations shut down. **Note**: You can specify a value of 0 seconds to disable the network timeout capability. This results in the replication servers waiting indefinitely for a network transmission to occur. However, Oracle Corporation recommends disabling network timeouts only in a well-controlled database environment.
DBM$BIND_HOT_NETWORK_TRANSPORT RDM$BIND_HOT_NETWORK_TRANSPORT RDB_BIND_HOT_NETWORK_TRANSPORT	The network transport protocol you want to use to ship after-image journal modifications to the remote database.Valid values are DECNET and TCPIP. Do not use the " / " (for example, TCP/IP) when you define this logical name or configuration parameter.

Table A-2 Logical Names and Configuration Parameters (Continued)

Name	Defines
DBM$BIND_HOT_OUTPUT_FILE RDM$BIND_HOT_OUTPUT_FILE RDB_BIND_HOT_OUTPUT_FILE	The location and name of the output file that records the replication activities of the AIJSERVER process. You can include "_PID" in the output file specification to create a unique file name that includes the process identification (PID). For example: `$ DEFINE DBM$BIND_HOT_OUTPUT_FILE HOT_PID.OUT` This creates a unique file name that includes the PID, for example, `DISK1:[USER]HOT_25C02915.OUT`. The default file extension is .out.
DBM$BIND_HOT_SNAP_QUIET_POINT RDM$BIND_HOT_SNAP_QUIET_POINT RDB_BIND_HOT_SNAP_QUIET_POINT	Enables or disables quiet-point locks for snapshot transactions on the master database. The valid settings are as follows: Value — Description 0 — Disables quiet-point locks for snapshot transactions. 1 — Enables quiet-point locks for snapshot transactions. This is the default setting. **Note**: Quiet-point locks must be enabled (the default setting) on the standby database whenever replication operations are occurring.
DBM$BIND_LCS_CONNECT_TIMEOUT RDM$BIND_LCS_CONNECT_TIMEOUT RDB_BIND_LCS_CONNECT_TIMEOUT	The number of minutes that the LCS process on the master database waits for a network connection to be made so replication operations can start. The default value is 5 minutes. You can set this logical name or configuration parameter from 1 to 4320 minutes on the master database only. Alternatively, you can specify the Connect_Timeout qualifier on the Replicate After_Journal Start command.
DBM$BIND_LCS_OUTPUT_FILE RDM$BIND_LCS_OUTPUT_FILE RDB_BIND_LCS_OUTPUT_FILE	The location and name of the log catch-up server (LCS) information output files. You can include "_PID" in the output file specification to create a unique file name that includes the process identification (PID). For example: `$ DEFINE DBM$BIND_LCS_OUTPUT_FILE LCS_PID.OUT` This creates a unique file name that includes the PID, for example, `DISK1:[USER]LCS_25C02916.OUT`. The default file extension is .out. Alternatively, you can define the location of the output file using the Output qualifier on the Replicate After_Journal Start command. The Output qualifier overrides definitions you make with the logical name or configuration parameter.

Table A-2 Logical Names and Configuration Parameters (Continued)

Name	Defines
DBM$BIND_LCS_PRIORITY RDM$BIND_LCS_PRIORITY	The base priority at which the log catch-up server (LCS) process is invoked by the database monitor. This logical name is applicable to OpenVMS systems only. If you do not define this name, the Hot Standby software uses the current priority of the database monitor.
DBM$BIND_LRS_GAP_TIMEOUT RDM$BIND_LRS_GAP_TIMEOUT RDB_BIND_LRS_GAP_TIMEOUT	The maximum number of minutes (from 1 to 4320) that the standby database waits for a gap in the sequence of messages from the master database to be resolved. The default value is 5 minutes. Specify this logical name or configuration parameter on the standby database only. Alternatively, you can specify the Gap_Timeout qualifier on the Replicate After_Journal Start command.
DBM$BIND_LRS_GOVERNOR_ENABLED RDM$BIND_LRS_GOVERNOR_ENABLED RDB_BIND_LRS_GOVERNOR_ENABLED	Enables the replication governor that automatically and dynamically chooses the synchronization mode that balances performance and database replication synchronization. The valid settings are as follows: Value — Description 0 — Disables the replication governor. 1 — Enables the replication governor (default). Alternatively, you can specify the Governor=Enabled qualifier on the Replicate After_Journal Start command
DBM$BIND_LRS_OUTPUT_FILE RDM$BIND_LRS_OUTPUT_FILE RDB_BIND_LRS_OUTPUT_FILE	The location and name of the output file that contains information about the replication activities of the log rollforward server (LRS). You can include "_PID" in the output file specification to create a unique file name that includes the process identification (PID). For example: `$ DEFINE DBM$BIND_LRS_OUTPUT_FILE LRS_PID.OUT` This creates a unique file name that includes the PID, for example, `DISK1:[USER]LRS_25C02917.OUT.` The default file extension is .out. Typically, you define this name on the standby database only. However, you can define this name on the master database to specify the LRS output file on the standby database only if replication is not already active on the standby database. Alternatively, you can define the location of the output file using the Output qualifier on the Replicate After_Journal Start command. The Output qualifier overrides definitions you make with the logical name or configuration parameter.

Table A-2 Logical Names and Configuration Parameters (Continued)

Name	Defines
DBM$BIND_LRS_PRIORITY RDM$BIND_LRS_PRIORITY	The base priority at which the log rollforward server (LRS) process is invoked by the database monitor. This logical name is applicable to OpenVMS systems only. If you do not define this name, the Hot Standby software uses the current priority of the database monitor.

B
Startup and Shutdown Checklists

This appendix contains the replication startup and shutdown checklists that are described in Chapters 2 and 5.

B.1 Replication Startup Checklist

Table B-1 provides a checklist of tasks that you perform to start replication operations. Each step includes a reference to a section in Chapter 2 that provides additional information.

Table B-1 Replication Startup Checklist

Step	Procedure	Reference
[1]	On the master database, set up the following after-image journal options: • Reserve five or more after-image journal slots • Add after-image journals and specify after-image journal names • Enable after-image journaling • Enable fast commit optimization • Enable the AIJ log server (ALS) as automatic • Enable system notification • Enable the AIJ backup server (ABS) as automatic • Set the database open mode to manual	Section 2.3
[2]	Open the master database manually (using the RMU Open command or the DBO Open command).	Section 2.4
[3]	Suspend the ABS process (automatic backup operations). You must suspend backup operations temporarily until you complete steps 4 through 12.	Section 2.5
[4]	Perform a full and complete backup operation on the master database. Oracle Corporation recommends that you perform an online, quiet-point backup operation (use the Quiet_Point qualifier).	Section 2.6
[5]	Save the after-image journal (.aij) configuration file to an output file (using the Output qualifier on the RMU or DBO Show After command).	Section 2.7

Table B-1 Replication Startup Checklist (Continued)

Step	Procedure	Reference
[6]	Create the standby database by using the Restore command to restore the most recent backup copy of the master database. The backup file that you restore ***must contain an exact copy*** of the master database. Also, be sure the options defined for the standby AIJ are ***identical*** to those for the AIJ on the master database. Oracle Corporation recommends using the Aij_Options qualifier on the Restore command.	Section 2.8
[7]	Open the standby database manually (using the RMU Open command or the DBO Open command).	Section 2.9
[8]	Optionally, establish a default Hot Standby configuration quickly by using the Replicate After_Journal Configure command.	Section 2.10
[9]	Define the location of the after-image journal rollforward temporary work files by defining a logical name or configuration parameter.	Section 2.11
[10]	If more than one network transport protocol is used in your system configuration, specify the protocol to be used for Hot Standby communications.	Section 2.12
[11]	Start replication operations on the standby database.	Section 2.13
[12]	Start replication operations on the master database.	Section 2.14

B.2 Shutdown Checklist

Table B-1 provides a checklist of the replication and database shutdown tasks. Each step in the checklist includes a reference to a section in Chapter 5 that provides additional information.

Table B-2 Shutdown Checklist

Step	Procedure	Reference
[1]	Stop replication operations.	Section 5.1
[2]	Stop the ALS process on the master database.	Section 5.2
[3]	At this point, you can restart replication operations, move the database, or close the database:	
	• Move the database and restart replication operations	Section 5.3
	• Close the database using orderly or emergency shutdown procedures	Section 5.4

Glossary

24x365

24 hours a day, 365 days a year.

AIJ log server (ALS)

The AIJ log server is a database server that flushes log data to the after-image journal (AIJ) file for databases. There is one ALS process per database per node. Although enabling the ALS process is optional for standard database operations, you must invoke the ALS process if you plan to use the Hot Standby option.

AIJSERVER

An optional network-object server invoked by the network software. The AIJSERVER receives after-image journal records from the master database and forwards them across the network to the LRS process on the standby database. The AIJSERVER is not required if the master and standby databases reside on the same node or the same cluster.

asymmetric processing

See “unidirectional.”

asynchronous processing

Processing that propagates changes to the other copies of the database based on a counter or on a time-driven basis. For example, updates are deferred until a predetermined time period elapses or a set number of transactions are processed.

See also “synchronous processing.”

availability

The percentage or amount of scheduled time that a computing system provides application service.

bidirectional

A form of replication in which replication servers on both the master (originating) and standby sites can perform database updates. This technique, also referred to as symmetric processing, is in direct contrast to the unidirectional form of data movement.

See also “unidirectional.”

data warehouse

A decision-support architecture that physically separates the decision-support databases from the operational databases. The data warehouse obtains and integrates its data from any number of independent operational systems, and provides its data to a collection of departmental databases and desktop systems where business managers and other knowledge workers can use it for analysis and decision making.

disaster tolerance

The ability of a computing component to achieve fault tolerance as well as immunity from disasters such as earthquakes, fires, floods, hurricanes, power failures, vandalism, or acts of terrorism.

failover

The ability to reconfigure a computing system to utilize an alternate active component when a similar component fails.

failure

The inability of a computing component to perform its function correctly, due to one or more internal faults whose effects cannot be contained.

fault tolerance

The ability of a computing system to withstand faults and errors while continuing to provide the required services.

log catch-up server (LCS)

A database server that synchronizes the master and standby databases. The LCS process synchronizes the database by automatically sending any after-image journal records (received since the last checkpoint interval) to the LRS process on the standby database. As soon as the databases are synchronized, the ALS process takes over processing on the master database, and the LCS process terminates.

log rollforward server (LRS)

A database server that receives after-image journal records from the master database, and applies them to the standby database after-image journal and to the database itself. Once invoked, the LRS process automatically goes into replay mode to initialize its transaction state information and roll forward any transactions that have been committed on the master database, but have not been applied to the standby database.

master database

A database on the master node that contains the source of the data modifications. This database is the active database. Also referred to as the primary database.

mission critical

Business functions that require a fault-tolerant or disaster-tolerant level of availability. Functions of this nature are referred to as *mission critical*, because achieving the business's financial goals often depends directly on how well its computing systems and database support the firm's computerized business applications.

primary database

See "master database."

redundant

Duplicate or extra computing components that safeguard the integrity of a computing system.

reliability

The ability of a computing system to operate without failing. Reliability is measured by mean time between failures (MTBF).

replay

A mode of operation during replication startup or failure recovery when the LRS process brings the standby database up to date with the master database. The LRS replays database modifications by initializing its state information and rolling forward after-image journal records to the standby database that have not yet been applied.

replicated database

See "standby database."

replication governor

A Hot Standby process that coordinates database replication operations on the master and standby databases. The replication governor automatically ensures that the databases do not get too far out of synchronization with respect to each other, and that the performance of the master database does not deviate greatly from that of the standby database.

restart

A mode of operation during replication startup or failure recovery when the LCS process resends after-image journal records to the LRS process for processing on the standby database. The records include those that have not been sent since the last recorded checkpoint.

Note: Restart operations occur in parallel with, and independently of, ongoing database updates that are being performed by the ALS process.

rollforward

The process of using an after-image journal to restore a database to a known state. This process replaces updates to the database that were lost because a system, program, or disk failure required the installation of backup media.

scalability

A measure of how well the software or hardware product is able to adapt to future business needs.

secondary database

See “standby database.”

standby database

A read-only database on the standby node. The standby database is physically identical to the master database and can immediately perform application processing if the master computing system or database fail. As a result of catastrophic failures, data modification activities fail over to the standby database such that it becomes the new master database.

standby node

A second running computing system that is ready to pick up application processing if the primary computing system fails. That is, the secondary system takes over the processing at the point where the original computing system stopped and the secondary system continues the processing.

symmetric multiprocessing (SMP)

A multiprocessing system configuration in which all processors have equal access to operating system code residing in shared memory and can perform all, or almost all, system tasks.

symmetric processing

See “bidirectional.”

synchronous processing

Processing that maintains copies of data at multiple sites on an event-driven or real-time basis. A change to any one database is immediately propagated to other replication servers as a part of the same transaction. The databases are usually immediately identical and recoverable if a failure occurs.

See also “asynchronous processing.”

unidirectional

A form of replication in which only the servers on the master (originating) database can perform database updates. Servers on the standby database may or may not be able to perform limited access, such as read-only access for the purpose of writing reports. This technique, also referred to as asymmetric processing, is in direct contrast to the bidirectional form of data movement.

See also “bidirectional.”

Index

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W