

Oracle9i

Database Performance Tuning Guide and Reference

Release 2 (9.2)

October 2002

Part No. A96533-02

ORACLE®

Oracle9i Database Performance Tuning Guide and Reference, Release 2 (9.2)

Part No. A96533-02

Copyright © 2000, 2002 Oracle Corporation. All rights reserved.

Primary Author: Connie Dialeris Green

Graphic Designer: Valarie Moore

Contributors: James Barlow, Eric Belden, Qiang Cao, Sumanta Chatterjee, Benoit Dageville, Vinayagam Djegaradjane, Harvey Eneman, Bjorn Engsig, Cecilia Gervasio, Ray Glasstone, Leslie Gloyd, Lester Gutierrez, Karl Haas, Brian Hirano, Andrew Holdsworth, Mamdouh Ibrahim, Christopher Jones, Srinivas Kareenhalli, Stella Kister, Herve Lejeune, Yunrui Li, Juan Loaiza, George Lumpkin, Joe McDonald, Bill McKenna, Sujatha Muthulingam, Gary Ngai, Michael Orlowski, Kant C. Patel, Richard Powell, Shankar Raman, Vinay Srihari, Sankar Subramanian, Margaret Susairaj, Hal Takahara, Nitin Vengurlekar, Stephen Vivian, Simon Watt, Andrew Witkowski, Graham Wood, and Mohamed Zait.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent and other intellectual and industrial property laws. Reverse engineering, disassembly or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Oracle Corporation.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

Restricted Rights Notice Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the Programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and Oracle Corporation disclaims liability for any damages caused by such use of the Programs.

Oracle is a registered trademark, and Oracle Expert, Oracle Store, Oracle7, Oracle8, Oracle9i, PL/SQL, SQL*Net, SQL*Plus, and iSQL*Plus are trademarks or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

Contents

Send Us Your Comments	xix
Preface.....	xxi
Audience	xxii
Organization.....	xxiii
Related Documentation	xxvi
Conventions.....	xxvii
Documentation Accessibility	xxx
What's New in Oracle Performance?	xxxii
Part I Writing and Tuning SQL	
1 Introduction to the Optimizer	
Overview of SQL Processing.....	1-2
Overview of the Optimizer.....	1-3
Features that Require the CBO	1-4
Optimizer Operations	1-5
Choosing an Optimizer Approach and Goal.....	1-5
How the CBO Optimizes SQL Statements for Fast Response.....	1-9
Understanding the Cost-Based Optimizer.....	1-10
Components of the CBO.....	1-11
Understanding Execution Plans	1-18
Understanding Access Paths for the CBO.....	1-24

Full Table Scans.....	1-24
Rowid Scans.....	1-27
Index Scans	1-28
Cluster Scans	1-35
Hash Scans	1-36
Sample Table Scans.....	1-36
How the CBO Chooses an Access Path	1-37
Understanding Joins	1-40
How the CBO Executes Join Statements.....	1-40
How the CBO Chooses the Join Method	1-41
How the CBO Chooses Execution Plans for Join Types.....	1-42
Nested Loop Joins.....	1-45
Hash Joins	1-47
Sort Merge Joins	1-49
Cartesian Joins.....	1-50
Outer Joins	1-51
Setting Cost-Based Optimizer Parameters	1-55
Enabling CBO Features	1-55
Controlling the Behavior of the CBO	1-58
Overview of the Extensible Optimizer	1-61
Understanding User-Defined Statistics	1-62
Understanding User-Defined Selectivity	1-62
Understanding User-Defined Costs.....	1-63

2 Optimizer Operations

How the Optimizer Performs Operations	2-2
How the CBO Evaluates IN-List Iterators.....	2-3
How the CBO Evaluates Concatenation.....	2-6
How the CBO Evaluates Remote Operations	2-10
How the CBO Executes Distributed Statements	2-13
How the CBO Executes Sort Operations	2-14
How the CBO Executes Views	2-18
How the CBO Evaluates Constants.....	2-19
How the CBO Evaluates the UNION and UNION ALL Operators.....	2-20
How the CBO Evaluates the LIKE Operator	2-22

How the CBO Evaluates the IN Operator	2-22
How the CBO Evaluates the ANY or SOME Operator	2-22
How the CBO Evaluates the ALL Operator	2-23
How the CBO Evaluates the BETWEEN Operator	2-24
How the CBO Evaluates the NOT Operator.....	2-24
How the CBO Evaluates Transitivity	2-25
How the CBO Optimizes Common Subexpressions.....	2-26
How the CBO Evaluates DETERMINISTIC Functions	2-28
How the Optimizer Transforms SQL Statements.....	2-30
How the CBO Transforms ORs into Compound Queries	2-30
How the CBO Unnests Subqueries	2-33
How the CBO Merges Views	2-35
How the CBO Pushes Predicates.....	2-38
How the CBO Executes Compound Queries.....	2-48

3 Gathering Optimizer Statistics

Understanding Statistics	3-2
Generating Statistics	3-3
Getting Statistics for Partitioned Schema Objects.....	3-4
Using the DBMS_STATS Package.....	3-5
Using the ANALYZE Statement.....	3-13
Finding Data Distribution	3-13
Missing Statistics.....	3-14
Using Statistics	3-14
Managing Statistics.....	3-15
Verifying Table Statistics	3-16
Verifying Index Statistics.....	3-17
Verifying Column Statistics	3-18
Using Histograms	3-20
When to Use Histograms.....	3-21
Creating Histograms	3-22
Types of Histograms	3-23
Viewing Histograms	3-25
Verifying Histogram Statistics.....	3-25

4 Understanding Indexes and Clusters

Understanding Indexes	4-2
Tuning the Logical Structure.....	4-2
Choosing Columns and Expressions to Index.....	4-3
Choosing Composite Indexes	4-4
Writing Statements That Use Indexes.....	4-6
Writing Statements That Avoid Using Indexes.....	4-6
Re-creating Indexes	4-7
Compacting Indexes.....	4-7
Using Nonunique Indexes to Enforce Uniqueness.....	4-8
Using Enabled Novalidated Constraints.....	4-8
Using Function-based Indexes	4-10
Setting Parameters to Use Function-Based Indexes in Queries	4-10
Using Index-Organized Tables	4-12
Using Bitmap Indexes	4-12
When to Use Bitmap Indexes.....	4-12
Using Bitmap Indexes with Good Performance.....	4-15
Initialization Parameters for Bitmap Indexing.....	4-17
Using Bitmap Access Plans on Regular B-tree Indexes.....	4-18
Bitmap Index Restrictions.....	4-19
Using Bitmap Join Indexes	4-19
Using Domain Indexes	4-19
Using Clusters	4-20
Using Hash Clusters	4-21

5 Optimizer Hints

Understanding Optimizer Hints	5-2
Specifying Hints.....	5-3
Using Optimizer Hints	5-6
Hints for Optimization Approaches and Goals	5-6
Hints for Access Paths.....	5-9
Hints for Query Transformations.....	5-17
Hints for Join Orders.....	5-22
Hints for Join Operations.....	5-24
Hints for Parallel Execution	5-29

Additional Hints.....	5-34
Using Hints with Views.....	5-42

6 Optimizing SQL Statements

Goals for Tuning	6-2
Reduce the Workload.....	6-2
Balance the Workload	6-2
Parallelize the Workload	6-2
Identifying and Gathering Data on Resource-Intensive SQL	6-3
Identifying Resource-Intensive SQL.....	6-3
Gathering Data on the SQL Identified.....	6-5
Dynamic Sampling	6-6
How Dynamic Sampling Works.....	6-7
When to Use Dynamic Sampling	6-7
How to Use Dynamic Sampling to Improve Performance.....	6-7
Overview of SQL Statement Tuning	6-8
Verifying Optimizer Statistics.....	6-8
Reviewing the Execution Plan	6-9
Restructuring the SQL Statements	6-10
Controlling the Access Path and Join Order with Hints.....	6-18
Restructuring the Indexes	6-22
Modifying or Disabling Triggers and Constraints	6-23
Restructuring the Data.....	6-23
Maintaining Execution Plans Over Time	6-23
Visiting Data as Few Times as Possible.....	6-23

7 Using Plan Stability

Using Plan Stability to Preserve Execution Plans	7-2
Using Hints with Plan Stability	7-2
Storing Outlines	7-4
Enabling Plan Stability.....	7-4
Using Supplied Packages to Manage Stored Outlines.....	7-4
Creating Outlines.....	7-4
Using and Editing Stored Outlines	7-6
Viewing Outline Data	7-10

Moving Outline Tables	7-11
Using Plan Stability with the Cost-Based Optimizer	7-12
Using Outlines to Move to the Cost-Based Optimizer	7-13
Upgrading and the Cost-Based Optimizer	7-14

8 Using the Rule-Based Optimizer

Overview of the Rule-Based Optimizer (RBO)	8-2
Understanding Access Paths for the RBO	8-3
Details of the RBO Access Paths	8-4
Choosing Execution Plans for Joins with the RBO	8-15
Transforming and Optimizing Statements with the RBO	8-17
Transforming ORs into Compound Queries with the RBO	8-17
Using Alternative SQL Syntax	8-18

Part II SQL-Related Performance Tools

9 Using EXPLAIN PLAN

Understanding EXPLAIN PLAN	9-2
How Execution Plans Can Change	9-2
Minimizing Throw-Away	9-3
Looking Beyond Execution Plans	9-4
Creating the PLAN_TABLE Output Table	9-4
Running EXPLAIN PLAN	9-5
Identifying Statements for EXPLAIN PLAN	9-5
Specifying Different Tables for EXPLAIN PLAN	9-6
Displaying PLAN_TABLE Output	9-6
Reading EXPLAIN PLAN Output	9-7
EXPLAIN PLAN Examples	9-8
Viewing Bitmap Indexes with EXPLAIN PLAN	9-11
Viewing Partitioned Objects with EXPLAIN PLAN	9-12
Examples of Displaying Range and Hash Partitioning with EXPLAIN PLAN	9-12
Examples of Pruning Information with Composite Partitioned Objects	9-14
Examples of Partial Partition-wise Joins	9-16
Examples of Full Partition-wise Joins	9-17

Examples of INLIST ITERATOR and EXPLAIN PLAN	9-18
Example of Domain Indexes and EXPLAIN PLAN	9-20
Viewing Parallel Execution with EXPLAIN PLAN	9-20
CPU Costing Model.....	9-22
EXPLAIN PLAN Restrictions	9-22
PLAN_TABLE Columns	9-23
10 Using SQL Trace and TKPROF	
Understanding SQL Trace and TKPROF	10-2
Understanding the SQL Trace Facility	10-2
Understanding TKPROF	10-3
Using the SQL Trace Facility and TKPROF	10-3
Step 1: Setting Initialization Parameters for Trace File Management	10-4
Step 2: Enabling the SQL Trace Facility.....	10-5
Step 3: Formatting Trace Files with TKPROF.....	10-6
Step 4: Interpreting TKPROF Output	10-12
Step 5: Storing SQL Trace Facility Statistics.....	10-17
Avoiding Pitfalls in TKPROF Interpretation	10-20
Avoiding the Argument Trap.....	10-20
Avoiding the Read Consistency Trap.....	10-20
Avoiding the Schema Trap.....	10-21
Avoiding the Time Trap	10-22
Avoiding the Trigger Trap	10-23
Sample TKPROF Output.....	10-23
Sample TKPROF Header	10-23
Sample TKPROF Body	10-24
Sample TKPROF Summary.....	10-30
11 Using Autotrace in SQL*Plus	
Overview of the Autotrace Report.....	11-2
Configuring the Autotrace Report	11-2
Setups Required for the Autotrace Report.....	11-2
Execution Plans for SQL Statements.....	11-3
Database Statistics for SQL Statements	11-4
Tracing Statements Examples	11-5

Collecting Timing Statistics	11-7
Tracing Parallel and Distributed Queries	11-7
Monitoring Disk Reads and Buffer Gets	11-9
SYSTEM Variables Influencing SQL*Plus Performance	11-9
SET APPINFO OFF	11-9
SET ARRAYSIZE.....	11-10
SET DEFINE OFF	11-10
SET FLUSH OFF.....	11-10
SET SERVEROUTPUT	11-10
SET TRIMOUT ON	11-10
SET TRIMSPool ON	11-11
iSQL*Plus Server Statistics Report	11-11
Active Statistics	11-12
Interpreting Active Statistics.....	11-13

12 Using Oracle Trace

Overview of Oracle Trace	12-2
Event Data.....	12-2
Event Sets	12-2
Accessing Collected Data.....	12-3
Collecting Oracle Trace Data	12-3
Using the Oracle Trace Command-Line Interface	12-3
Using Initialization Parameters to Control Oracle Trace.....	12-7
Controlling Oracle Trace Collections from PL/SQL	12-10
Accessing Oracle Trace Collection Results	12-12
Formatting Oracle Trace Data to Oracle Tables	12-13
Running the Oracle Trace Reporting Utility.....	12-14
Oracle Server Events	12-15
Data Items Collected for Events	12-16
Items Associated with Each Event	12-22
Troubleshooting Oracle Trace	12-32
Oracle Trace Configuration	12-32
Formatter Tables	12-37

Part III Creating a Database for Good Performance

13 Building a Database for Performance

Initial Database Creation	13-2
Database Creation Using the Installer	13-2
Manual Database Creation	13-2
Parameters Necessary for Initial Database Creation	13-2
The CREATE DATABASE Statement	13-3
Running Data Dictionary Scripts	13-5
Sizing Redo Log Files	13-5
Creating Subsequent Tablespaces	13-6
Creating Tables for Good Performance	13-7
Data Segment Compression	13-9
Loading and Indexing Data	13-10
Using SQL*Loader for Good Performance	13-11
Efficient Index Creation	13-11
Initial Instance Configuration	13-13
Configuring Undo Space	13-15
Setting up Operating System, Database, and Network Monitoring	13-15

14 Memory Configuration and Use

Understanding Memory Allocation Issues	14-2
Oracle Memory Caches	14-2
Dynamically Changing Cache Sizes	14-2
Application Considerations	14-4
Operating System Memory Use	14-4
Iteration During Configuration	14-5
Configuring and Using the Buffer Cache	14-6
Using the Buffer Cache Effectively	14-6
Sizing the Buffer Cache	14-6
Interpreting and Using the Buffer Cache Advisory Statistics	14-11
Considering Multiple Buffer Pools	14-13
Buffer Pool Data in V\$DB_CACHE_ADVICE	14-15
Buffer Pool Hit Ratios	14-15

Determining Which Segments Have Many Buffers in the Pool.....	14-15
KEEP Pool	14-17
RECYCLE Pool.....	14-18
Configuring and Using the Shared Pool and Large Pool	14-18
Shared Pool Concepts.....	14-19
Using the Shared Pool Effectively	14-23
Sizing the Shared Pool	14-27
Interpreting Shared Pool Statistics	14-33
Using the Large Pool.....	14-35
Using CURSOR_SPACE_FOR_TIME	14-39
Caching Session Cursors.....	14-39
Configuring the Reserved Pool.....	14-40
Keeping Large Objects to Prevent Aging	14-42
CURSOR_SHARING for Existing Applications.....	14-43
Configuring and Using the Java Pool.....	14-46
Configuring and Using the Redo Log Buffer.....	14-46
Sizing the Log Buffer	14-47
Log Buffer Statistics	14-47
Configuring the PGA Working Memory	14-48
Automatic PGA Memory Management	14-50
Configuring SORT_AREA_SIZE	14-66

15 I/O Configuration and Design

Understanding I/O	15-2
Designing I/O Layouts	15-2
Disk Performance and Reliability.....	15-2
Disk Technology	15-3
What Is Disk Contention?.....	15-3
Load Balancing and Striping.....	15-4
Striping and RAID	15-4
Balancing Budget, Performance, and Availability	15-6
Basic I/O Configuration	15-6
Determining Application I/O Characteristics.....	15-6
I/O Configuration Decisions	15-10
Know Your I/O System.....	15-10

Match I/O Requirements with the I/O System.....	15-11
Lay Out the Files Using Operating System or Hardware Striping	15-12
Manually Distributing I/O	15-16
When to Separate Files.....	15-17
Three Sample Configurations	15-19
Oracle-Managed Files	15-20
Choosing Data Block Size.....	15-21

16 Understanding Operating System Resources

Understanding Operating System Performance Issues	16-2
Using Operating System Caches	16-2
Memory Usage.....	16-4
Using Process Schedulers.....	16-5
Using Operating System Resource Managers	16-5
Solving Operating System Problems.....	16-7
Performance Hints on UNIX-Based Systems	16-7
Performance Hints on NT Systems.....	16-7
Performance Hints on Midrange and Mainframe Computers	16-8
Understanding CPU	16-8
Context Switching	16-10
Finding System CPU Utilization	16-11
Checking Memory Management.....	16-12
Checking I/O Management	16-12
Checking Network Management	16-12
Checking Process Management.....	16-12

17 Configuring Instance Recovery Performance

Understanding Instance Recovery.....	17-2
Checkpointing and Cache Recovery	17-2
How Checkpoints Affect Performance.....	17-3
Reducing Checkpoint Frequency to Optimize Runtime Performance	17-3
Configuring the Duration of Cache Recovery	17-4
Initialization Parameters that Influence Cache Recovery Time.....	17-4
Use Fast-Start Checkpointing to Limit Instance Recovery Time	17-5
Set LOG_CHECKPOINT_TIMEOUT to Influence the Amount of Redo	17-7

Set LOG_CHECKPOINT_INTERVAL to Influence the Amount of Redo	17-7
Use Parallel Recovery to Speed up Redo Application	17-8
Monitoring Cache Recovery	17-9
Monitoring Estimated MTTR: Example Scenario	17-10
Calculating Performance Overhead.....	17-12
Calculating Performance Overhead: Example Scenario	17-13
Calibrating the MTTR	17-15
MTTR Advisory	17-16
How MTTR Advisory Works.....	17-16
Enabling MTTR Advisory	17-16
Viewing MTTR Advisory	17-17
Tuning Transaction Recovery	17-18
Using Fast-Start On-Demand Rollback	17-18
Using Fast-Start Parallel Rollback	17-18

18 Configuring Undo and Temporary Segments

Configuring Undo Segments	18-2
Configuring Automatic Undo Management	18-2
Configuring Rollback Segments	18-2
Configuring Temporary Tablespaces.....	18-4

19 Configuring Shared Servers

Introduction to Shared Server Performance	19-2
Configuring the Number of Shared Servers	19-2
Identifying Contention Using the Dispatcher-Specific Views	19-3
Reducing Contention for Dispatcher Processes	19-4
Reducing Contention for Shared Servers	19-5
Determining the Optimal Number of Dispatchers and Shared Servers.....	19-8

Part IV System-Related Performance Tools

20 Oracle Tools to Gather Database Statistics

Overview of Tools	20-2
Principles of Data Gathering.....	20-2

Interpreting Statistics	20-3
Oracle Enterprise Manager Diagnostics Pack	20-5
Statspack	20-7
V\$ Performance Views	20-7
Example - Saving File I/O Data	20-8

21 Using Statspack

Introduction to Statspack	21-2
Statspack Compared with BSTAT/ESTAT	21-2
How Statspack Works	21-3
Configuring Database Space Requirements for Statspack	21-4
Installing Statspack	21-4
Interactive Statspack Installation	21-4
Batch Mode Statspack Installation	21-6
Using Statspack	21-6
Taking a Statspack Snapshot.....	21-7
Automating Statistics Gathering	21-8
Running a Statspack Performance Report	21-9
Configuring the Amount of Data Captured in Statspack.....	21-15
Time Units Used for Wait Events.....	21-20
Event Timings	21-21
Managing and Sharing Statspack Performance Data.....	21-22
Oracle Real Application Clusters Considerations with Statspack	21-25
Removing Statspack	21-26
Statspack Supplied Scripts and Documentation	21-26
Scripts for Statspack Installation and Removal.....	21-27
Scripts for Statspack Reporting and Automation	21-27
Scripts for Upgrading Statspack.....	21-27
Scripts for Statspack Performance Data Maintenance	21-28
Statspack Documentation.....	21-28

Part V Optimizing Instance Performance

22 Instance Tuning

Performance Tuning Principles	22-2
Baselines.....	22-2
The Symptoms and the Problems.....	22-3
When to Tune.....	22-4
Performance Tuning Steps	22-5
Define the Problem.....	22-6
Examine the Host System.....	22-7
Examine the Oracle Statistics.....	22-10
Implement and Measure Change.....	22-14
Interpreting Oracle Statistics	22-15
Examine Load.....	22-15
Using Wait Event Statistics to Drill Down to Bottlenecks.....	22-16
Table of Wait Events and Potential Causes.....	22-19
Additional Statistics.....	22-20
Wait Events	22-24
SQL*Net.....	22-25
buffer busy waits.....	22-27
db file scattered read.....	22-29
db file sequential read.....	22-31
direct path read and direct path read (lob).....	22-33
direct path write.....	22-35
enqueue.....	22-36
free buffer waits.....	22-39
latch free.....	22-41
log buffer space.....	22-46
log file switch.....	22-46
log file sync.....	22-48
rdbms ipc reply.....	22-48
Idle Wait Events	22-49

23 Tuning Networks

Understanding Connection Models	23-2
Detecting Network Problems	23-6
Using Dynamic Performance Views for Network Performance	23-6
Understanding Latency and Bandwidth.....	23-6
Solving Network Problems	23-8
Finding Network Bottlenecks	23-8
Dissecting Network Bottlenecks.....	23-10
Using Array Interfaces.....	23-13
Adjusting Session Data Unit Buffer Size	23-13
Using TCP.NODELAY.....	23-14
Using Connection Manager	23-14

Part VI Performance-Related Reference Information

24 Dynamic Performance Views for Tuning

Dynamic Performance Tables	24-2
Current State Views.....	24-2
Counter/Accumulator Views.....	24-2
Information Views.....	24-4
Description of Dynamic Performance Views	24-5
V\$DB_OBJECT_CACHE	24-5
V\$FILESTAT.....	24-6
V\$LATCH.....	24-9
V\$LATCH_CHILDREN.....	24-13
V\$LATCHHOLDER.....	24-13
V\$LIBRARYCACHE	24-15
V\$LIBRARY_CACHE_MEMORY	24-16
V\$LOCK.....	24-17
V\$MTTR_TARGET_ADVICE.....	24-21
V\$MYSTAT.....	24-22
V\$OPEN_CURSOR	24-23
V\$PARAMETER and V\$SYSTEM_PARAMETER.....	24-25
V\$PROCESS.....	24-26

V\$ROLLSTAT.....	24-28
V\$ROWCACHE.....	24-29
V\$SEGMENT_STATISTICS.....	24-31
V\$SEGSTAT.....	24-32
V\$SEGSTAT_NAME.....	24-32
V\$SESSION.....	24-33
V\$SESSION_EVENT.....	24-36
V\$SESSION_WAIT.....	24-37
V\$SESSTAT.....	24-41
V\$SHARED_POOL_ADVICE.....	24-45
V\$SQL.....	24-45
V\$SQL_PLAN.....	24-46
V\$SQL_PLAN_STATISTICS.....	24-51
V\$SQL_PLAN_STATISTICS_ALL.....	24-53
V\$SQLAREA.....	24-57
V\$SQLTEXT.....	24-59
V\$STATISTICS_LEVEL.....	24-61
V\$SYSSTAT.....	24-61
V\$SYSTEM_EVENT.....	24-67
V\$UNDOSTAT.....	24-69
V\$WAITSTAT.....	24-70

A Schemas Used in Performance Examples

PER_ALL_PEOPLE_F Table.....	A-2
RA_CUSTOMERS Table.....	A-2
SO_HEADERS_ALL and SO_HEADERS Tables.....	A-3
MTL_SYSTEM_ITEMS Table.....	A-3
SO_LINES_ALL and SO_LINES Tables.....	A-4

Glossary

Index

Send Us Your Comments

Oracle9i Database Performance Tuning Guide and Reference, Release 2 (9.2)

Part No. A96533-02

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

- 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?

If you find any errors or have any other suggestions for improvement, please indicate the document title and part number, and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- Electronic mail: infodev_us@oracle.com
- FAX: (650) 506-7227 Attn: Server Technologies Documentation Manager

- Postal service:

Oracle Corporation
Server Technologies Documentation
500 Oracle Parkway, Mailstop 4op11
Redwood Shores, CA 94065
USA

If you would like a reply, please give your name, address, telephone number, and (optionally) electronic mail address.

If you have problems with the software, please contact your local Oracle Support Services.

Preface

This preface contains these topics:

- [Audience](#)
- [Organization](#)
- [Related Documentation](#)
- [Conventions](#)
- [Documentation Accessibility](#)

Audience

Oracle9i Database Performance Tuning Guide and Reference is an aid for people responsible for the operation, maintenance, and performance of Oracle. This book describes detailed ways to enhance Oracle performance by writing and tuning SQL properly, using performance tools, and optimizing instance performance. It also explains how to create an initial database for good performance and includes performance-related reference information.

This book could be useful for database administrators, application designers, and programmers. Readers should be familiar with Oracle9i, *Oracle9i Database Performance Planning*, the operating system, and application design before reading this manual.

Many client/server application programmers consider SQL a messaging language, because queries are issued and data is returned. However, client tools often generate inefficient SQL statements. Therefore, a good understanding of the database SQL processing engine is necessary for writing optimal SQL. This is especially true for high transaction processing systems.

Typically, SQL statements issued by OLTP applications operate on relatively few rows at a time. If an index can point to the exact rows that you want, then Oracle can construct an accurate plan to access those rows efficiently through the shortest possible path. In DSS environments, selectivity is less important, because they often access most of a table's rows. In such situations, full table scans are common, and indexes are not even used. This book is primarily focussed on OLTP-type applications. For detailed information on DSS and mixed environments, see the *Oracle9i Data Warehousing Guide*.

Before using this performance tuning reference, make sure you have read *Oracle9i Database Performance Planning*. Oracle Corporation has designed a new performance methodology, based on years of Oracle designing and performance experience. This brief book explains clear and simple activities that can dramatically improve system performance. It discusses the following topics:

- Investment Options
- Scalability
- System Architecture
- Application Design Principles
- Workload Testing, Modeling, and Implementation
- Deploying New Applications

Organization

This document contains:

Part I, "Writing and Tuning SQL"

This section provides information to help understand and manage SQL statements.

Chapter 1, "Introduction to the Optimizer"

This chapter discusses SQL processing, Oracle optimization, and how the Oracle optimizer chooses how to execute SQL statements.

Chapter 2, "Optimizer Operations"

This chapter provides details of how the CBO provides specific operations.

Chapter 3, "Gathering Optimizer Statistics"

This chapter explains why statistics are important for the cost-based optimizer and describes how to gather and use statistics.

Chapter 4, "Understanding Indexes and Clusters"

This chapter describes how to create indexes and clusters, and when to use them.

Chapter 5, "Optimizer Hints"

This chapter offers recommendations on how to use cost-based optimizer hints to enhance Oracle performance.

Chapter 6, "Optimizing SQL Statements"

This chapter describes how Oracle optimizes SQL using the cost-based optimizer (CBO).

Chapter 7, "Using Plan Stability"

This chapter describes how to use plan stability (stored outlines) to preserve performance characteristics.

Chapter 8, "Using the Rule-Based Optimizer"

This chapter discusses Oracle's rule-based optimizer (RBO).

Part II, "SQL-Related Performance Tools"

This section provides information about Oracle SQL-related performance tools.

Chapter 9, "Using EXPLAIN PLAN"

This chapter shows how to use the SQL statement `EXPLAIN PLAN` and format its output.

Chapter 10, "Using SQL Trace and TKPROF"

This chapter describes the use of the SQL trace facility and `TKPROF`, two basic performance diagnostic tools that can help you monitor and tune applications that run against the Oracle Server.

Chapter 11, "Using Autotrace in SQL*Plus"

This chapter describes the use of Autotrace, which can automatically get reports on the execution path used by the SQL optimizer and the statement execution statistics to help you monitor and tune statement performance.

Chapter 12, "Using Oracle Trace"

This chapter provides an overview of Oracle Trace usage and describes the Oracle Trace initialization parameters.

Note: Oracle Trace will be deprecated in a future release. Oracle Corporation strongly advises the use of SQL Trace and `TKPROF` instead.

Part III, "Creating a Database for Good Performance"

This section describes how to create and configure a database for good performance.

Chapter 13, "Building a Database for Performance"

This chapter describes how to design and create a database for the intended needs.

Chapter 14, "Memory Configuration and Use"

This chapter explains how to allocate memory to database structures.

Chapter 15, "I/O Configuration and Design"

This chapter introduces fundamental I/O concepts, discusses the I/O requirements of different parts of the database, and provides sample configurations for I/O subsystem design.

Chapter 16, "Understanding Operating System Resources"

This chapter explains how to tune the operating system for optimal performance of Oracle.

Chapter 17, "Configuring Instance Recovery Performance"

This chapter explains how to tune recovery performance.

Chapter 18, "Configuring Undo and Temporary Segments"

This chapter explains how to configure undo segments (using automatic undo management or using rollback segments) and how to configure temporary tablespaces.

Chapter 19, "Configuring Shared Servers"

This chapter explains how to identify and reduce contention for dispatcher processes and for shared servers.

Part IV, "System-Related Performance Tools"

This section provides information about Oracle's system-related performance tools.

Chapter 20, "Oracle Tools to Gather Database Statistics"

Oracle provides a number of tools that allow a performance engineer to gather information regarding instance and database performance. This chapter explains why performance data gathering is important, and it describes how to use available tools.

Chapter 21, "Using Statspack"

This chapter describes the use of Statspack to collect, store, and analyze system data.

Part V, "Optimizing Instance Performance"

This section describes how to tune various elements of a database system to optimize performance of an Oracle instance.

Chapter 22, "Instance Tuning"

This chapter discusses the method used for performing tuning. It also describes Oracle statistics and wait events.

Chapter 23, "Tuning Networks"

This chapter describes different connection models and networking issues that affect tuning.

Part VI, "Performance-Related Reference Information"

This section provides reference information regarding dynamic performance views and wait events.

Chapter 24, "Dynamic Performance Views for Tuning"

This chapter provides detailed information on several dynamic performance views that can help you tune your system and investigate performance problems.

Appendix

Appendix A, "Schemas Used in Performance Examples"

This appendix describes the tables used in examples in [Chapter 9, "Using EXPLAIN PLAN"](#).

Related Documentation

Before reading this manual, you should have already read *Oracle9i Database Performance Planning*, *Oracle9i Database Concepts*, the *Oracle9i Application Developer's Guide - Fundamentals*, and the *Oracle9i Database Administrator's Guide*.

For more information about Oracle Enterprise Manager and its optional applications, see *Oracle Enterprise Manager Concepts Guide*, *Oracle Enterprise Manager Administrator's Guide*, and *Database Tuning with the Oracle Tuning Pack*.

For more information about tuning data warehouse environments, see the *Oracle9i Data Warehousing Guide*.

Many of the examples in this book use the sample schemas of the seed database, which is installed by default when you install Oracle. Refer to *Oracle9i Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Printed documentation is available for sale in the Oracle Store at

<http://oraclestore.oracle.com/>

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at

<http://otn.oracle.com/admin/account/membership.html>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

<http://otn.oracle.com/docs/index.htm>

To access the database documentation search engine directly, please visit

<http://tahiti.oracle.com>

Conventions

This section describes the conventions used in the text and code examples of the this documentation set. It describes:

- [Conventions in Text](#)
- [Conventions in Code Examples](#)

Conventions in Text

We use various conventions in text to help you more quickly identify special terms. The following table describes those conventions and provides examples of their use.

Convention	Meaning	Example
Bold	Bold typeface indicates terms that are defined in the text or terms that appear in a glossary, or both.	When you specify this clause, you create an index-organized table .

Convention	Meaning	Example
<i>Italics</i>	Italic typeface indicates book titles, emphasis, syntax clauses, or placeholders.	<i>Oracle9i Database Concepts</i> You can specify the <i>parallel_clause</i> . Run <code>Uold_release.SQL</code> where <i>old_release</i> refers to the release you installed prior to upgrading.
UPPERCASE monospace (fixed-width font)	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, user names, and roles.	You can specify this clause only for a NUMBER column. You can back up the database using the BACKUP command. Query the TABLE_NAME column in the USER_TABLES data dictionary view. Specify the ROLLBACK_SEGMENTS parameter. Use the DBMS_STATS.GENERATE_STATS procedure.
lowercase monospace (fixed-width font)	Lowercase monospace typeface indicates executables and sample user-supplied elements. Such elements include computer and database names, net service names, and connect identifiers, as well as user-supplied database objects and structures, column names, packages and classes, user names and roles, program units, and parameter values.	Enter <code>sqlplus</code> to open SQL*Plus. The <code>department_id</code> , <code>department_name</code> , and <code>location_id</code> columns are in the <code>hr.departments</code> table. Set the <code>QUERY_REWRITE_ENABLED</code> initialization parameter to <code>true</code> . Connect as <code>oe</code> user.

Conventions in Code Examples

Code examples illustrate SQL, PL/SQL, SQL*Plus, or other command-line statements. They are displayed in a monospace (fixed-width) font and separated from normal text as shown in this example:

```
SELECT username FROM dba_users WHERE username = 'MIGRATE';
```

The following table describes typographic conventions used in code examples and provides examples of their use.

Convention	Meaning	Example
[]	Brackets enclose one or more optional items. Do not enter the brackets.	<code>DECIMAL (digits [, precision])</code>

Convention	Meaning	Example
{ }	Braces enclose two or more items, one of which is required. Do not enter the braces.	{ENABLE DISABLE}
	A vertical bar represents a choice of two or more options within brackets or braces. Enter one of the options. Do not enter the vertical bar.	{ENABLE DISABLE} [COMPRESS NOCOMPRESS]
...	Horizontal ellipsis points indicate either: <ul style="list-style-type: none"> That we have omitted parts of the code that are not directly related to the example That you can repeat a portion of the code 	CREATE TABLE ... AS subquery; SELECT col1, col2, ... , coln FROM employees;
.	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	SQL> SELECT NAME FROM V\$DATAFILE; NAME ----- /fsl/dbs/tbs_01.dbf /fsl/dbs/tbs_02.dbf . . . /fsl/dbs/tbs_09.dbf 9 rows selected.
Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	acctbal NUMBER(11,2); acct CONSTANT NUMBER(4) := 3;
<i>Italics</i>	Italicized text indicates placeholders or variables for which you must supply particular values.	CONNECT SYSTEM/ <i>system_password</i> DB_NAME = <i>database_name</i>
UPPERCASE	Uppercase typeface indicates elements supplied by the system. We show these terms in uppercase in order to distinguish them from terms you define. Unless terms appear in brackets, enter them in the order with the spelling shown. However, because these terms are not case sensitive, you can enter them in lowercase.	SELECT last_name, employee_id FROM employees; SELECT * FROM USER_TABLES; DROP TABLE hr.employees;

Convention	Meaning	Example
lowercase	Lowercase typeface indicates programmatic elements that you supply. For example, lowercase indicates names of tables, columns, or files. Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	<pre>SELECT last_name, employee_id FROM employees; sqlplus hr/my_hr_password CREATE USER mjones IDENTIFIED BY ty3MU9;</pre>

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle Corporation is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation JAWS, a Windows screen reader, may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, JAWS may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation This documentation may contain links to Web sites of other companies or organizations that Oracle Corporation does not own or control. Oracle Corporation neither evaluates nor makes any representations regarding the accessibility of these Web sites.

What's New in Oracle Performance?

This section describes new performance features of Oracle9i Release 2 (9.2) and provides pointers to additional information. The features and enhancements described in this section comprise the overall effort to optimize server performance.

Note: Before using this performance tuning reference, make sure you have read *Oracle9i Database Performance Planning*. Oracle Corporation has designed a new performance methodology, based on years of Oracle designing and performance experience. This brief book explains clear and simple activities that can dramatically improve system performance. It discusses the following topics:

- Investment Options
 - Scalability
 - System Architecture
 - Application Design Principles
 - Workload Testing, Modeling, and Implementation
 - Deploying New Applications
-
-

The new and updated performance features include the following:

- **Forced Rewrite**

A new setting, `FORCE`, is available for the `QUERY_REWRITE_ENABLED` session parameter.

When `QUERY_REWRITE_ENABLED` is set to `FORCE`, Oracle always uses rewrite and does not evaluate the cost before doing so. `FORCE` is useful when you know that the query will always benefit from rewrite, when reduction in compile time is important, and when you know that the optimizer may be underestimating the benefits of materialized views.

See Also: ["QUERY_REWRITE_ENABLED"](#) on page 4-10

- **Union-All Rewrite of Queries with Grouping Sets**

A new hint, `EXPAND_GSET_TO_UNION`, is available to force rewrite when using function-based indexes in queries where compilation time is important and the query always benefits from rewrite (OLAP).

The `EXPAND_GSET_TO_UNION` hint is used for queries containing grouping sets (such as queries with `GROUP BY GROUPING SET` or `GROUP BY ROLLUP`). The hint forces a query to be transformed into a corresponding query with `UNION ALL` of individual groupings.

See Also: ["EXPAND_GSET_TO_UNION"](#) on page 5-19 for details about using this hint

- **Dynamic Sampling for the Optimizer**

The purpose of dynamic sampling is to improve server performance by determining more accurate selectivity and cardinality estimates. More accurate selectivity and cardinality estimates allow the optimizer to produce better performing plans.

You can use dynamic sampling to:

- Estimate single-table predicate selectivities when collected statistics cannot be used or are likely to lead to significant errors in estimation.
- Estimate table cardinality for tables without statistics or for tables whose statistics are too out of date to trust.

The `DYNAMIC_SAMPLING` hint lets you control dynamic sampling to improve server performance by determining more accurate selectivity and cardinality

estimates. You can set the value of `DYNAMIC_SAMPLING` to a value from 0 to 10. The higher the level, the more effort the compiler puts into dynamic sampling and the more broadly it is applied. Sampling defaults to cursor level unless you specify a table.

See Also:

- ["Dynamic Sampling"](#) on page 6-6 for information about when and how to use dynamic sampling
- ["DYNAMIC_SAMPLING"](#) on page 5-39 for details about using this hint

- **Locally Managed SYSTEM Tablespace**

Oracle9i Release 2 (9.2) allows creation of a database with a locally managed SYSTEM tablespace. Use the `EXTENT MANAGEMENT LOCAL` clause of the `CREATE DATABASE` statement to create a locally managed SYSTEM tablespace.

If you specify `EXTENT MANAGEMENT LOCAL`, then you must also specify the default temporary tablespace.

See Also:

- ["Creating Temporary Tablespaces"](#) on page 13-7
- [Table 13-2, " Database Options for Initial Creation"](#) on page 13-3
- *Oracle9i SQL Reference* for detailed information about the `CREATE DATABASE` statement

- **Data Segment Compression**

Data segment compression reduces disk use and memory use (specifically, the buffer cache), often leading to a better scaleup for read-only operations. Data segment compression can also speed up query execution.

Oracle9i Release 2 (9.2) achieves a good compression ratio in many cases with no special tuning. However, if you need a better compression ratio, tuning can improve it slightly in some cases and very substantially in other cases.

See Also: ["Data Segment Compression"](#) on page 13-9

- Shared Pool Advisory Statistics

The amount of memory available for the library cache can drastically affect the parse rate of an Oracle instance. With Oracle9i, Release 2 (9.2) or higher, the shared pool advisory statistics provide a database administrator with information about library cache memory and predict how changes in the size of the shared pool can affect the parse rate.

The shared pool advisory statistics track the library cache's use of shared pool memory and predict how the library cache will behave in shared pools of different sizes. Two fixed views provide the information to determine how much memory the library cache is using, how much is currently pinned, how much is on the shared pool's LRU list, as well as how much time might be lost or gained by changing the size of the shared pool.

See Also: ["Shared Pool Advisory Statistics"](#) on page 14-31

- PGA Aggregate Target Advisory

Under automatic PGA memory management mode, the main goal of Oracle is to honor the `PGA_AGGREGATE_TARGET` limit set by the DBA, by controlling dynamically the amount of PGA memory allotted to SQL work areas. At the same time, Oracle tries to maximize the performance of all the memory-intensive SQL operators, by maximizing the number of work areas that are using an optimal amount of PGA memory (cache memory). The rest of the work areas are executed in one-pass mode, unless the PGA memory limit set by the DBA with the parameter `PGA_AGGREGATE_TARGET` is so low that multi-pass execution is required to reduce even more the consumption of PGA memory and honor the PGA target limit.

When configuring a brand new instance, it is hard to know precisely the appropriate setting for `PGA_AGGREGATE_TARGET`. You can determine this setting in three stages:

1. Make a first estimate for `PGA_AGGREGATE_TARGET`, based on a rule of thumb.
2. Run a representative workload on the instance and monitor performance, using PGA statistics collected by Oracle, to see whether the maximum PGA size is under-configured or over-configured.
3. Tune `PGA_AGGREGATE_TARGET`, using Oracle's PGA advice statistics.

See Also:

- ["Setting PGA_AGGREGATE_TARGET Initially"](#) on page 14-50
- ["Monitoring the Performance of the Automatic PGA Memory Management"](#) on page 14-51
- ["Tuning PGA_AGGREGATE_TARGET"](#) on page 14-59

- **FILESYSTEMIO_OPTIONS**

With Oracle9i Release 2 (9.2), you can use the `FILESYSTEMIO_OPTIONS` initialization parameter to enable or disable asynchronous I/O or direct I/O on file system files. This parameter is platform-specific and has a default value that is best for a particular platform. It can be dynamically changed to update the default setting.

See Also: ["FILESYSTEMIO_OPTIONS Initialization Parameter"](#) on page 16-4

- **Mean Time to Recover (MTTR) Advisory**

Starting with Oracle9i Release 2 (9.2), MTTR advisory is available to help you evaluate the effect of different MTTR settings on system performance in terms of extra physical writes.

When MTTR advisory is enabled, after the system runs a typical workload for a while, you can query `V$MTTR_TARGET_ADVICE`, which tells you the ratio of estimated number of cache writes under other MTTR settings to the number of cache writes under the current MTTR. For instance, a ratio of 1.2 indicates 20% more cache writes.

By looking at the different MTTR settings and their corresponding cache write ratio, you can decide which MTTR value fits your recovery and performance needs. `V$MTTR_TARGET_ADVICE` also gives the ratio on total physical writes, including direct writes, and the ratio on total input and output, including reads.

See Also:

- ["MTTR Advisory"](#) on page 17-16
- ["V\\$MTTR_TARGET_ADVICE"](#) on page 24-21

- **Statistics Collection Level**

Oracle9i Release 2 (9.2) provides the initialization parameter `STATISTICS_LEVEL`, which controls all major statistics collections or advisories in the database. This parameter sets the statistics collection level for the database.

`BASIC`, `TYPICAL`, or `ALL`. The default is `TYPICAL`.

See Also: ["Setting the Level of Statistics Collection"](#) on page 22-10 for details about what statistics are collected at each level

- **Segment-Level Statistics**

With Oracle9i Release 2 (9.2) and higher, you can gather segment-level statistics to help you spot performance problems associated with individual segments. Collecting and viewing segment-level statistics is a good way to effectively identify the hot table or index in an instance.

After viewing wait events or system statistics to identify the performance problem, you can use segment-level statistics to find specific tables or indexes that are causing the problem.

You can query segment-level statistics through the new dynamic views [V\\$SEGMENT_STATISTICS](#), [V\\$SEGSTAT](#), and [V\\$SEGSTAT_NAME](#).

See Also:

- ["Segment-Level Statistics"](#) on page 22-14
- [Chapter 24, "Dynamic Performance Views for Tuning"](#)

- **Runtime Row Source Statistics**

Oracle9i Release 1 (9.0.1) introduced a new dynamic performance table, `V$SQL_PLAN`, to show the execution plan for a cached cursor. Oracle9i Release 2 (9.2) introduces another dynamic performance table, `V$SQL_PLAN_STATISTICS`. This view provides, for each cached cursor, the execution statistics of each operation in the execution plan.

To view row source statistics in this view, the DBA must set the parameter `STATISTICS_LEVEL` to `ALL`.

An additional view, `V$SQL_PLAN_STATISTICS`, concatenates information from `V$SQL_PLAN` with execution statistics from `V$SQL_PLAN_STATISTICS` and `V$SQL_WORKAREA`. `V$SQL_WORKAREA` contains memory usage statistics for row sources that use SQL memory (for example, hash-join and sort).

See Also:

- ["V\\$SQL_PLAN_STATISTICS"](#) on page 24-51
- ["V\\$SQL_PLAN_STATISTICS_ALL"](#) on page 24-53

- Oracle Trace Removed from Future Releases

Oracle Trace will be deprecated in a future release. Oracle Corporation strongly advises the use of SQL Trace and TKPROF instead.

See Also: [Chapter 10, "Using SQL Trace and TKPROF"](#)

Part I

Writing and Tuning SQL

Part I provides information on understanding and managing your SQL statements for optimal performance. It is best to read these chapters in the order in which they are presented.

The chapters in this part are:

- [Chapter 1, "Introduction to the Optimizer"](#)
- [Chapter 2, "Optimizer Operations"](#)
- [Chapter 3, "Gathering Optimizer Statistics"](#)
- [Chapter 4, "Understanding Indexes and Clusters"](#)
- [Chapter 5, "Optimizer Hints"](#)
- [Chapter 6, "Optimizing SQL Statements"](#)
- [Chapter 7, "Using Plan Stability"](#)
- [Chapter 8, "Using the Rule-Based Optimizer"](#)

Introduction to the Optimizer

This chapter discusses SQL processing, optimization methods, and how the optimizer chooses a specific plan to execute SQL.

The chapter contains the following sections:

- [Overview of SQL Processing](#)
- [Overview of the Optimizer](#)
- [Choosing an Optimizer Approach and Goal](#)
- [Understanding the Cost-Based Optimizer](#)
- [Understanding Access Paths for the CBO](#)
- [Understanding Joins](#)
- [Setting Cost-Based Optimizer Parameters](#)
- [Overview of the Extensible Optimizer](#)

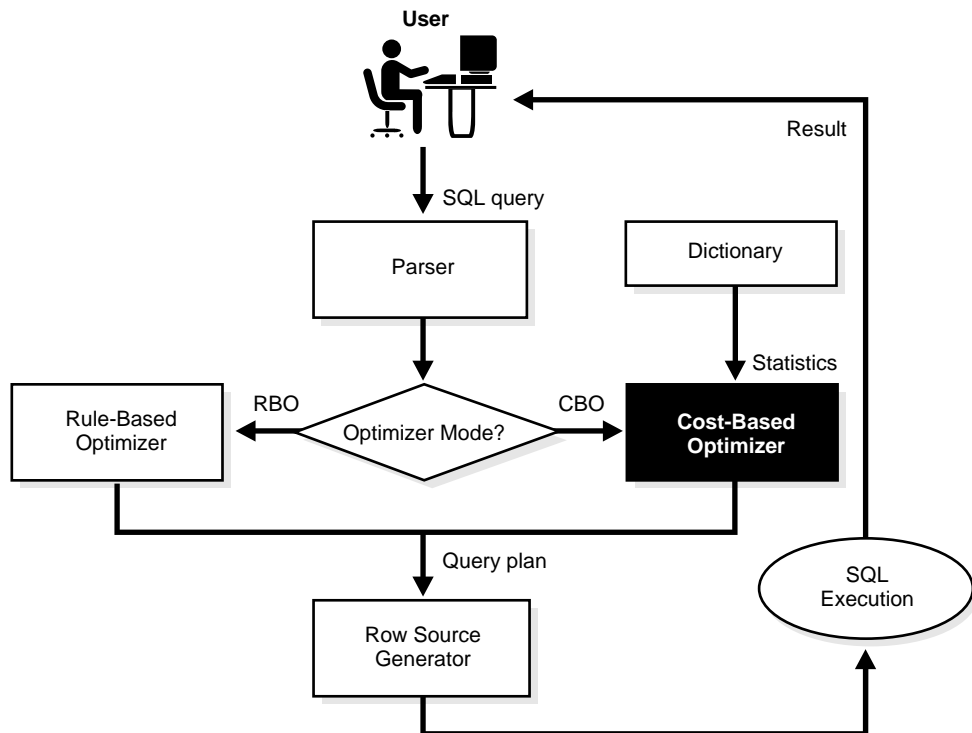
Overview of SQL Processing

SQL processing uses the following main components to execute a SQL query:

- The Parser checks both syntax and semantic analysis.
- The Optimizer uses costing methods, cost-based optimizer (CBO), or internal rules, rule-based optimizer (RBO), to determine the most efficient way of producing the result of the query.
- The Row Source Generator receives the optimal plan from the optimizer and outputs the execution plan for the SQL statement.
- The SQL Execution Engine operates on the execution plan associated with a SQL statement and then produces the results of the query.

Figure 1-1 illustrates SQL processing.

Figure 1-1 SQL Processing Overview



Overview of the Optimizer

The optimizer determines the most efficient way to execute a SQL statement after considering many factors related to the objects referenced and the conditions specified in the query. This determination is an important step in the processing of any SQL statement and can greatly affect execution time.

A SQL statement can be executed in many different ways, including the following:

- Full table scans
- Index scans
- Nested loops
- Hash joins

Note: The optimizer might not make the same decisions from one version of Oracle to the next. In recent versions, the optimizer might make different decisions, because better information is available.

The output from the optimizer is a plan that describes an optimum method of execution. The Oracle server provides the cost-based (CBO) and rule-based (RBO) optimization. In general, use the cost-based approach. Oracle Corporation is continually improving the CBO and new features require CBO.

Note: Oracle Corporation strongly advises the use of cost-based optimization. The rule-based optimization is available for backward compatibility with legacy applications and will be deprecated in a future release.

You can influence the optimizer's choices by setting the optimizer approach and goal, and by gathering representative statistics for the CBO. The optimizer goal is either throughput or response time. See "[Choosing an Optimizer Approach and Goal](#)" on page 1-5.

Sometimes, the application designer, who has more information about a particular application's data than is available to the optimizer, can choose a more effective way to execute a SQL statement. The application designer can use hints in SQL statements to specify how the statement should be executed.

See Also:

- ["Choosing an Optimizer Approach and Goal"](#) on page 1-5 for more information on optimization goals
- [Chapter 3, "Gathering Optimizer Statistics"](#) for information on gathering and using statistics
- [Chapter 5, "Optimizer Hints"](#) for more information about using hints in SQL statements

Features that Require the CBO

The following features require use of the CBO:

- Partitioned tables and indexes
- Index-organized tables
- Reverse key indexes
- Function-based indexes
- SAMPLE clauses in a SELECT statement
- Parallel query and parallel DML
- Star transformations and star joins
- Extensible optimizer
- Query rewrite with materialized views
- Enterprise Manager progress meter
- Hash joins
- Bitmap indexes and bitmap join indexes
- Index skip scans

Note: Using any of these features enables the CBO, even if the parameter `OPTIMIZER_MODE` is set to `RULE`.

Optimizer Operations

For any SQL statement processed by Oracle, the optimizer performs the operations listed in [Table 1-1](#).

Table 1-1 *Optimizer Operations*

Operation	Description
Evaluation of expressions and conditions	The optimizer first evaluates expressions and conditions containing constants as fully as possible. See "How the Optimizer Performs Operations" on page 2-2.
Statement transformation	For complex statements involving, for example, correlated subqueries or views, the optimizer might transform the original statement into an equivalent join statement. See "How the Optimizer Transforms SQL Statements" on page 2-30.
Choice of optimizer approaches	The optimizer chooses either a cost-based or rule-based approach and determines the goal of optimization. See "Choosing an Optimizer Approach and Goal" on page 1-5.
Choice of access paths	For each table accessed by the statement, the optimizer chooses one or more of the available access paths to obtain table data. See "Understanding Access Paths for the CBO" on page 1-24.
Choice of join orders	For a join statement that joins more than two tables, the optimizer chooses which pair of tables is joined first, and then which table is joined to the result, and so on. See "How the CBO Chooses Execution Plans for Join Types" on page 1-42.
Choice of join methods	For any join statement, the optimizer chooses an operation to use to perform the join. See "How the CBO Chooses the Join Method" on page 1-41.

Choosing an Optimizer Approach and Goal

By default, the goal of the CBO is the best **throughput**. This means that it chooses the least amount of resources necessary to process all rows accessed by the statement. Oracle can also optimize a statement with the goal of best **response time**. This means that it uses the least amount of resources necessary to process the first row accessed by a SQL statement.

See Also: ["How the CBO Optimizes SQL Statements for Fast Response"](#) on page 1-9

The execution plan produced by the optimizer can vary depending on the optimizer's goal. Optimizing for best throughput is more likely to result in a full table scan rather than an index scan, or a sort merge join rather than a nested loop join. Optimizing for best response time usually results in an index scan or a nested loop join.

For example, suppose you have a join statement that can be executed with either a nested loops operation or a sort-merge operation. The sort-merge operation might return the entire query result faster, while the nested loops operation might return the first row faster. If your goal is to improve throughput, then the optimizer is more likely to choose a sort merge join. If your goal is to improve response time, then the optimizer is more likely to choose a nested loop join.

Choose a goal for the optimizer based on the needs of your application:

- For applications performed in batch, such as Oracle Reports applications, optimize for best throughput. Usually, throughput is more important in batch applications, because the user initiating the application is only concerned with the time necessary for the application to complete. Response time is less important, because the user does not examine the results of individual statements while the application is running.
- For interactive applications, such as Oracle Forms applications or SQL*Plus queries, optimize for best response time. Usually, response time is important in interactive applications, because the interactive user is waiting to see the first row or first few rows accessed by the statement.

The optimizer's behavior when choosing an optimization approach and goal for a SQL statement is affected by the following factors:

- [OPTIMIZER_MODE Initialization Parameter](#)
- [CBO Statistics in the Data Dictionary](#)
- [Optimizer SQL Hints for Changing the CBO Goal](#)

OPTIMIZER_MODE Initialization Parameter

The `OPTIMIZER_MODE` initialization parameter establishes the default behavior for choosing an optimization approach for the instance. The possible values and description are listed in [Table 1-2](#).

Table 1–2 *OPTIMIZER_MODE* Parameter Values

Value	Description
CHOOSE	<p>The optimizer chooses between a cost-based approach and a rule-based approach, depending on whether statistics are available. This is the default value.</p> <ul style="list-style-type: none"> ■ If the data dictionary contains statistics for at least one of the accessed tables, then the optimizer uses a cost-based approach and optimizes with a goal of best throughput. ■ If the data dictionary contains only some statistics, then the cost-based approach is still used, but the optimizer must guess the statistics for the subjects without any statistics. This can result in suboptimal execution plans. ■ If the data dictionary contains no statistics for any of the accessed tables, then the optimizer uses a rule-based approach.
ALL_ROWS	The optimizer uses a cost-based approach for all SQL statements in the session regardless of the presence of statistics and optimizes with a goal of best throughput (minimum resource use to complete the entire statement).
FIRST_ROWS_ <i>n</i>	The optimizer uses a cost-based approach, regardless of the presence of statistics, and optimizes with a goal of best response time to return the first <i>n</i> number of rows; <i>n</i> can equal 1, 10, 100, or 1000.
FIRST_ROWS	<p>The optimizer uses a mix of cost and heuristics to find a best plan for fast delivery of the first few rows.</p> <p>Note: Using heuristics sometimes leads the CBO to generate a plan with a cost that is significantly larger than the cost of a plan without applying the heuristic. <code>FIRST_ROWS</code> is available for backward compatibility and plan stability.</p>
RULE	The optimizer chooses a rule-based approach for all SQL statements regardless of the presence of statistics.

You can change the goal of the CBO for all SQL statements in a session by changing the parameter value in initialization file or by the `ALTER SESSION SET OPTIMIZER_MODE` statement. For example:

- The following statement in an initialization parameter file changes the goal of the CBO for all sessions of the instance to best response time:

```
OPTIMIZER_MODE = FIRST_ROWS_1
```

- The following SQL statement changes the goal of the CBO for the current session to best response time:

```
ALTER SESSION SET OPTIMIZER_MODE = FIRST_ROWS_1;
```

If the optimizer uses the cost-based approach for a SQL statement, and if some tables accessed by the statement have no statistics, then the optimizer uses internal

information, such as the number of data blocks allocated to these tables, to estimate other statistics for these tables.

Optimizer SQL Hints for Changing the CBO Goal

To specify the goal of the CBO for an individual SQL statement, use one of the hints in the following list. Any of these hints in an individual SQL statement can override the `OPTIMIZER_MODE` initialization parameter for that SQL statement.

- `FIRST_ROWS(n)`, where *n* equals any positive integer
- `FIRST_ROWS`
- `ALL_ROWS`
- `CHOOSE`
- `RULE`

See Also: [Chapter 5, "Optimizer Hints"](#) for information on how to use hints

CBO Statistics in the Data Dictionary

The statistics used by the CBO are stored in the data dictionary. You can collect exact or estimated statistics about physical storage characteristics and data distribution in these schema objects by using the `DBMS_STATS` package or the `ANALYZE` statement.

Note: Oracle Corporation strongly recommends that you use the `DBMS_STATS` package rather than `ANALYZE` to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by `DBMS_STATS`. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the `ANALYZE` statement rather than `DBMS_STATS` for statistics collection not related to the cost-based optimizer, such as:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
 - To collect information on freelist blocks
-
-

To maintain the effectiveness of the CBO, you must have statistics that are representative of the data. For table columns that contain values with large variations in number of duplicates, called skewed data, you should collect histograms.

The resulting statistics provide the CBO with information about data uniqueness and distribution. Using this information, the CBO is able to compute plan costs with a high degree of accuracy. This enables the CBO to choose the best execution plan based on the least cost.

See Also: [Chapter 3, "Gathering Optimizer Statistics"](#)

How the CBO Optimizes SQL Statements for Fast Response

The CBO can optimize a SQL statement for fast response when the parameter `OPTIMIZER_MODE` is set to `FIRST_ROWS_n`, where n is 1, 10, 100, or 1000, or `FIRST_ROWS`. A hint `FIRST_ROWS(n)`, where n is any positive integer, or `FIRST_ROWS` can be used to optimize an individual SQL statement for fast response.

Fast-response optimization is suitable for online users, such as those using Oracle Forms or Web access. Typically, online users are interested in seeing the first few rows and seldom look at the entire query result, especially when the result size is large. For such users, it makes sense to optimize the query to produce the first few rows as quickly as possible, even if the time to produce the entire query result is not minimized.

With fast-response optimization, the CBO generates a plan with the lowest cost to produce the first row or the first few rows. The CBO employs two different fast-response optimizations, referred to here as the old and new methods. The old method is used with the `FIRST_ROWS` hint or parameter value. With the old method, the CBO uses a mixture of costs and rules to produce a plan. It is retained for backward compatibility reasons.

The new method is totally based on costs, and it is sensitive to the value of n . With small values of n , the CBO tends to generate plans that consist of nested loop joins with index lookups. With large values of n , the CBO tends to generate plans that consist of hash joins and full table scans.

The value of n should be chosen based on the online user requirement and depends specifically on how the result is displayed to the user. Generally, Oracle Forms users see the result one row at a time and they are typically interested in seeing the first few screens. Other online users see the result one group of rows at a time.

With the fast-response method, the CBO explores different plans and computes the cost to produce the first n rows for each. It picks the plan that produces the first n rows at lowest cost. Remember that with fast-response optimization, a plan that produces the first n rows at lowest cost might not be the optimal plan to produce the entire result. If the requirement is to obtain the entire result of a query, then fast-response optimization should not be used. Instead use the `ALL_ROWS` parameter value or hint.

Understanding the Cost-Based Optimizer

The CBO determines which execution plan is most efficient by considering available access paths and by factoring in information based on statistics for the schema objects (tables or indexes) accessed by the SQL statement. The CBO also considers hints, which are optimization suggestions placed in a comment in the statement.

See Also: [Chapter 5, "Optimizer Hints"](#) for detailed information on hints

The CBO performs the following steps:

1. The optimizer generates a set of potential plans for the SQL statement based on available access paths and hints.
2. The optimizer estimates the cost of each plan based on statistics in the data dictionary for the data distribution and storage characteristics of the tables, indexes, and partitions accessed by the statement.

The **cost** is an estimated value proportional to the expected resource use needed to execute the statement with a particular plan. The optimizer calculates the cost of access paths and join orders based on the estimated computer resources, which includes I/O, CPU, and memory.

Serial plans with higher costs take more time to execute than those with smaller costs. When using a parallel plan, however, resource use is not directly related to elapsed time.

3. The optimizer compares the costs of the plans and chooses the one with the lowest cost.

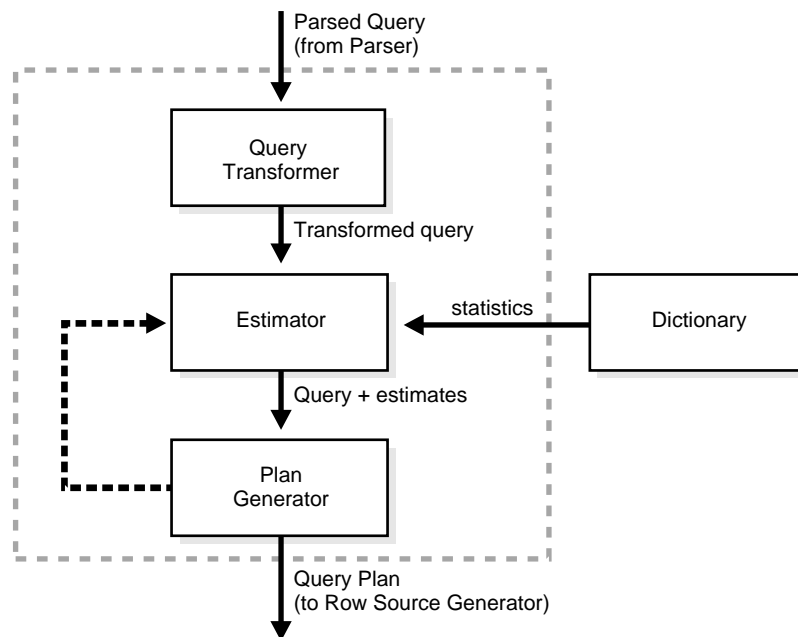
Components of the CBO

The CBO consists of the following three main components:

- Query Transformer
- Estimator
- Plan Generator

CBO components are illustrated in [Figure 1-2](#).

Figure 1-2 Cost-Based Optimizer Components



Query Transformer

The input to the query transformer is a parsed query, which is represented by a set of query blocks. The query blocks are nested or interrelated to each other. The form of the query determines how the query blocks are interrelated to each other. The main objective of the query transformer is to determine if it is advantageous to change the form of the query so that it enables generation of a better query plan.

Four different query transformation techniques are employed by the query transformer:

- [View Merging](#)
- [Predicate Pushing](#)
- [Subquery Unnesting](#)
- [Query Rewrite with Materialized Views](#)

Any combination of these transformations can be applied to a given query.

View Merging Each view referenced in a query is expanded by the parser into a separate query block. The query block essentially represents the view definition, and therefore the result of a view. One option for the optimizer is to analyze the view query block separately and generate a view subplan. The optimizer then processes the rest of the query by using the view subplan in the generation of an overall query plan. This technique usually leads to a suboptimal query plan, because the view is optimized separately from rest of the query.

The query transformer then removes the potentially suboptimal plan by merging the view query block into the query block that contains the view. Most types of views are merged. When a view is merged, the query block representing the view is merged into the containing query block. Generating a subplan is no longer necessary, because the view query block is eliminated.

Predicate Pushing For those views that are not merged, the query transformer can push the relevant predicates from the containing query block into the view query block. This technique improves the subplan of the nonmerged view, because the pushed-in predicates can be used either to access indexes or to act as filters.

Subquery Unnesting Like a view, a subquery is represented by a separate query block. Because a subquery is nested within the main query or another subquery, the plan generator is constrained in trying out different possible plans before it finds a plan with the lowest cost. For this reason, the query plan produced might not be the optimal one. The restrictions due to the nesting of subqueries can be removed by unnesting the subqueries and converting them into joins. Most subqueries are unnested by the query transformer. For those subqueries that are not unnested, separate subplans are generated. To improve execution speed of the overall query plan, the subplans are ordered in an efficient manner.

Query Rewrite with Materialized Views A materialized view is like a query with a result that is materialized and stored in a table. When a user query is found compatible

with the query associated with a materialized view, the user query can be rewritten in terms of the materialized view. This technique improves the execution of the user query, because most of the query result has been precomputed. The query transformer looks for any materialized views that are compatible with the user query and selects one or more materialized views to rewrite the user query. The use of materialized views to rewrite a query is cost-based. That is, the query is not rewritten if the plan generated without the materialized views has a lower cost than the plan generated with the materialized views.

See Also:

- ["How the Optimizer Transforms SQL Statements"](#) on page 2-30
- *Oracle9i Data Warehousing Guide* for more information on query rewrite

Estimator

The estimator generates three different types of measures:

- [Selectivity](#)
- [Cardinality](#)
- [Cost](#)

These measures are related to each other, and one is derived from another. The end goal of the estimator is to estimate the overall cost of a given plan. If statistics are available, then the estimator uses them to compute the measures. The statistics improve the degree of accuracy of the measures.

Selectivity The first measure, selectivity, represents a fraction of rows from a row set. The row set can be a base table, a view, or the result of a join or a GROUP BY operator. The selectivity is tied to a query predicate, such as `last_name = 'Smith'`, or a combination of predicates, such as `last_name = 'Smith' AND job_type = 'Clerk'`. A predicate acts as a filter that filters a certain number of rows from a row set. Therefore, the selectivity of a predicate indicates how many rows from a row set will pass the predicate test. Selectivity lies in a value range from 0.0 to 1.0. A selectivity of 0.0 means that no rows will be selected from a row set, and a selectivity of 1.0 means that all rows will be selected.

The estimator uses an internal default value for selectivity, if no statistics are available. Different internal defaults are used, depending on the predicate type. For example, the internal default for an equality predicate (`last_name = 'Smith'`) is lower than the internal default for a range predicate (`last_name > 'Smith'`). The

estimator makes this assumption because an equality predicate is expected to return a smaller fraction of rows than a range predicate.

When statistics are available, the estimator uses them to estimate selectivity. For example, for an equality predicate (`last_name = 'Smith'`), selectivity is set to the reciprocal of the number n of distinct values of `last_name`, because the query selects rows that all contain one out of n distinct values. If a histogram is available on the `last_name` column, then the estimator uses it instead of the number of distinct values. The histogram captures the distribution of different values in a column, so it yields better selectivity estimates. Having histograms on columns that contain skewed data (in other words, values with large variations in number of duplicates) greatly helps the CBO generate good selectivity estimates.

Cardinality Cardinality represents the number of rows in a row set. Here, the row set can be a base table, a view, or the result of a join or `GROUP BY` operator.

- **Base cardinality** is the number of rows in a base table. The base cardinality can be captured by analyzing the table. If table statistics are not available, then the estimator uses the number of extents occupied by the table to estimate the base cardinality.
- **Effective cardinality** is the number of rows that are selected from a base table. The effective cardinality depends on the predicates specified on different columns of a base table, with each predicate acting as a successive filter on the rows of the base table. The effective cardinality is computed as the product of the base cardinality and combined selectivity of all predicates specified on a table. When there is no predicate on a table, its effective cardinality equals its base cardinality.
- **Join cardinality** is the number of rows produced when two row sets are joined together. A **join** is a Cartesian product of two row sets, with the join predicate applied as a filter to the result. Therefore, the join cardinality is the product of the cardinalities of two row sets, multiplied by the selectivity of the join predicate.
- **Distinct cardinality** is the number of distinct values in a column of a row set. The distinct cardinality of a row set is based on the data in the column. For example, in a row set of 100 rows, if distinct column values are found in 20 rows, then the distinct cardinality is 20.
- **Group cardinality** is the number of rows produced from a row set after the `GROUP BY` operator is applied. The effect of the `GROUP BY` operator is to decrease the number of rows in a row set. The group cardinality depends on the

distinct cardinality of each of the grouping columns and on the number of rows in the row set. For an illustration of group cardinality, see [Example 1-1](#).

Example 1-1 Group Cardinality

If a row set of 100 rows is grouped by `colx`, which has a distinct cardinality of 30, then the group cardinality is 30.

However, suppose the same row set of 100 rows is grouped by `colx` and `coly`, which have distinct cardinalities of 30 and 60, respectively. In this case, the group cardinality lies between the maximum of the distinct cardinalities of `colx` and `coly`, and the lower of the product of the distinct cardinalities of `colx` and `coly`, and the number of rows in the row set.

Group cardinality in this example can be represented by the following formula:

$$\text{group cardinality lies between } \max(\text{dist. card. colx}, \text{dist. card. coly}) \\ \text{and } \min((\text{dist. card. colx} * \text{dist. card. coly}), \\ \text{num rows in row set})$$

Substituting the numbers from the example, the group cardinality is between the maximum of (30 and 60) and the minimum of (30*60 and 100). In other words, the group cardinality is between 60 and 100.

Cost The cost represents units of work or resource used. The CBO uses disk I/O, CPU usage, and memory usage as units of work. So, the cost used by the CBO represents an estimate of the number of disk I/Os and the amount of CPU and memory used in performing an operation. The operation can be scanning a table, accessing rows from a table by using an index, joining two tables together, or sorting a row set. The cost of a query plan is the number of work units that are expected to be incurred when the query is executed and its result produced.

The **access path** determines the number of units of work required to get data from a base table. The access path can be a table scan, a fast full index scan, or an index scan. During table scan or fast full index scan, multiple blocks are read from the disk in a single I/O operation. Therefore, the cost of a table scan or a fast full index scan depends on the number of blocks to be scanned and the multiblock read count value. The cost of an index scan depends on the levels in the B-tree, the number of index leaf blocks to be scanned, and the number of rows to be fetched using the rowid in the index keys. The cost of fetching rows using rowids depends on the index clustering factor.

Although the clustering factor is a property of the index, the clustering factor actually relates to the spread of similar indexed column values within data blocks in


the table. A lower clustering factor indicates that the individual rows are concentrated within fewer blocks in the table. Conversely, a high clustering factor indicates that the individual rows are scattered more randomly across blocks in the table. Therefore, a high clustering factor means that it costs more to use a range scan to fetch rows by rowid, because more blocks in the table need to be visited to return the data. [Example 1-2](#) shows how the clustering factor can affect cost.

Example 1-2 Effects of Clustering Factor on Cost

Assume the following situation:


- There is a table with 9 rows.
- There is a nonunique index on `col1` (in `tab1`).
- The `c1` column currently stores the values A, B, and C.
- The table only has three Oracle blocks.

Case 1: The index clustering factor is low for the rows as they are arranged in the following diagram.



This is because the rows that have the same indexed column values for `c1` are located within the same physical blocks in the table. The cost of using a range scan to return all of the rows that have the value A is low, because only one block in the table needs to be read.

Case 2: If the same rows in the table are rearranged so that the index values are scattered across the table blocks (rather than colocated), then the index clustering factor is higher.



This is because all three blocks in the table must be read in order to retrieve all rows with the value A in `col1`.

The **join cost** represents the combination of the individual access costs of the two row sets being joined. In a join, one row set is called inner, and the other is called outer.

- In a **nested loop join**, for every row in the outer row set, the inner row set is accessed to find all the matching rows to join. Therefore, in a nested loop join, the inner row set is accessed as many times as the number of rows in the outer row set.

$$\text{cost} = \text{outer access cost} + (\text{inner access cost} * \text{outer cardinality})$$

- In a **sort merge join**, the two row sets being joined are sorted by the join keys if they are not already in key order.

$$\text{cost} = \text{outer access cost} + \text{inner access cost} + \text{sort costs (if sort is used)}$$

- In a **hash join**, the inner row set is hashed into memory, and a hash table is built using the join key. Each row from the outer row set is then hashed, and the hash table is probed to join all matching rows. If the inner row set is very large, then only a portion of it is hashed into memory. This portion is called a hash partition.

Each row from the outer row set is hashed to probe matching rows in the hash partition. The next portion of the inner row set is then hashed into memory, followed by a probe from the outer row set. This process is repeated until all partitions of the inner row set are exhausted.

$$\text{cost} = (\text{outer access cost} * \# \text{ of hash partitions}) + \text{inner access cost}$$

See Also: ["Understanding Joins"](#) on page 1-40 for more information on joins

Plan Generator

The main function of the plan generator is to try out different possible plans for a given query and pick the one that has the lowest cost. Many different plans are possible because of the various combinations of different access paths, join methods, and join orders that can be used to access and process data in different ways and produce the same result.

A join order is the order in which different join items, such as tables, are accessed and joined together. For example, in a join order of t_1 , t_2 , and t_3 , table t_1 is accessed first. Next, t_2 is accessed, and its data is joined to t_1 data to produce a join of t_1 and t_2 . Finally, t_3 is accessed, and its data is joined to the result of the join between t_1 and t_2 .

The plan for a query is established by first generating subplans for each of the nested subqueries and nonmerged views. Each nested subquery or nonmerged view is represented by a separate query block. The query blocks are optimized

separately in a bottom-up order. That is, the innermost query block is optimized first, and a subplan is generated for it. The outermost query block, which represents the entire query, is optimized last.

The plan generator explores various plans for a query block by trying out different access paths, join methods, and join orders. The number of possible plans for a query block is proportional to the number of join items in the `FROM` clause. This number rises exponentially with the number of join items.

The plan generator uses an internal cutoff to reduce the number of plans it tries when finding the one with the lowest cost. The cutoff is based on the cost of the current best plan. If the current best cost is large, then the plan generator tries harder (in other words, explores more alternate plans) to find a better plan with lower cost. If the current best cost is small, then the plan generator ends the search swiftly, because further cost improvement will not be significant.

The cutoff works well if the plan generator starts with an initial join order that produces a plan with cost close to optimal. Finding a good initial join order is a difficult problem. The plan generator uses a simple heuristic for the initial join order. It orders the join items by their effective cardinalities. The join item with the smallest effective cardinality goes first, and the join item with the largest effective cardinality goes last.

Understanding Execution Plans

To execute a SQL statement, Oracle might need to perform many steps. Each of these steps either retrieves rows of data physically from the database or prepares them in some way for the user issuing the statement. The combination of the steps Oracle uses to execute a statement is called an **execution plan**. An execution plan includes an **access path** for each table that the statement accesses and an ordering of the tables (the **join order**) with the appropriate **join method**.

See Also:

- ["Understanding Access Paths for the CBO"](#) on page 1-24
- ["Understanding Access Paths for the RBO"](#) on page 8-3
- [Chapter 9, "Using EXPLAIN PLAN"](#)

Overview of EXPLAIN PLAN

You can examine the execution plan chosen by the optimizer for a SQL statement by using the `EXPLAIN PLAN` statement. When the statement is issued, the optimizer

chooses an execution plan and then inserts data describing the plan into a database table. Simply issue the EXPLAIN PLAN statement and then query the output table.

These are the basics of using the EXPLAIN PLAN statement:

- Use the SQL script `UTLXPLAN.SQL` to create a sample output table called `PLAN_TABLE` in your schema. See ["Creating the PLAN_TABLE Output Table"](#) on page 9-4.
- Include the `EXPLAIN PLAN FOR` clause prior to the SQL statement. See ["Running EXPLAIN PLAN"](#) on page 9-5.
- After issuing the `EXPLAIN PLAN` statement, use one of the scripts provided by Oracle to display the most recent plan table output. See ["Displaying PLAN_TABLE Output"](#) on page 9-6.
- The execution order in `EXPLAIN PLAN` output begins with the line that is the furthest indented to the right. The next step is the parent of that line. If two lines are indented equally, then the top line is executed first.

Notes:

- The `EXPLAIN PLAN` output tables in this chapter were displayed with the `utlxlpls.sql` script.
 - The steps in the `EXPLAIN PLAN` output tables in this chapter may be different on your system. Depending on database configurations, the optimizer may choose different execution plans.
-
-

Example 1-3 uses `EXPLAIN PLAN` to examine a SQL statement that selects the `employee_id`, `job_title`, `salary`, and `department_name` for the employees whose IDs are less than 103.

Example 1-3 Using EXPLAIN PLAN

```
EXPLAIN PLAN FOR
SELECT e.employee_id, j.job_title, e.salary, d.department_name
FROM employees e, jobs j, departments d
WHERE e.employee_id < 103
      AND e.job_id = j.job_id
      AND e.department_id = d.department_id;
```

The resulting output table in [Example 1-4](#) shows the execution plan chosen by the optimizer to execute the SQL statement in the example:

Example 1-4 EXPLAIN PLAN Output

```
-----
```

Id	Operation	Name	Rows	Bytes	Cost (%CPU)
0	SELECT STATEMENT		3	189	10 (10)
1	NESTED LOOPS		3	189	10 (10)
2	NESTED LOOPS		3	141	7 (15)
* 3	TABLE ACCESS FULL	EMPLOYEES	3	60	4 (25)
4	TABLE ACCESS BY INDEX ROWID	JOBS	19	513	2 (50)
* 5	INDEX UNIQUE SCAN	JOB_ID_PK	1		
6	TABLE ACCESS BY INDEX ROWID	DEPARTMENTS	27	432	2 (50)
* 7	INDEX UNIQUE SCAN	DEPT_ID_PK	1		

```
-----
```

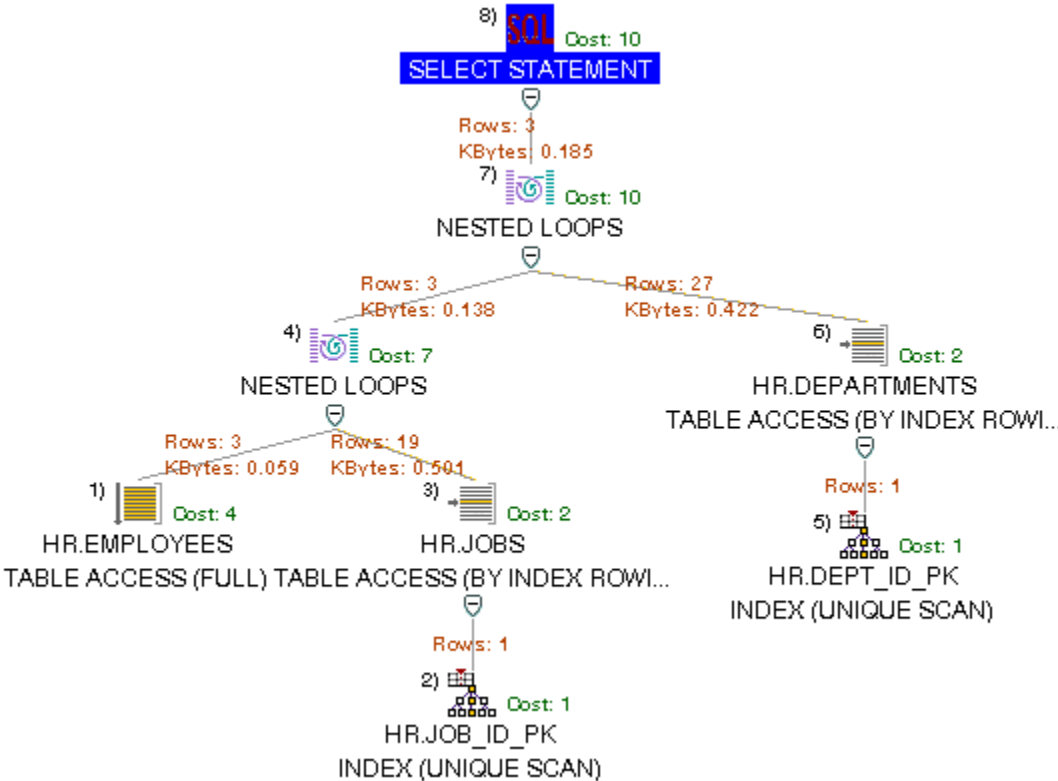
Predicate Information (identified by operation id):

- ```

```
- 3 - filter("E"."EMPLOYEE\_ID"<103)
  - 5 - access("E"."JOB\_ID"="J"."JOB\_ID")
  - 7 - access("E"."DEPARTMENT\_ID"="D"."DEPARTMENT\_ID")

Oracle provides additional graphical tools for displaying EXPLAIN PLAN output. [Figure 6-1, "Oracle SQL Analyze"](#) on page 6-4 is an example of the SQL statement displayed in Oracle SQL Analyze. [Figure 1-3](#) shows a graphical EXPLAIN PLAN for the SQL statement in [Example 1-3](#) which has been generated in an Oracle Enterprise Manager SQL Scratchpad window.

Figure 1-3 Graphical View of SQL Explain Plan in SQL Scratchpad



Note that the execution steps in Figure 1-3 are identified by the order of execution, rather than an ID as in Example 1-4.

**See Also:** For more information about Oracle Enterprise Manager and its optional applications, see *Oracle Enterprise Manager Concepts Guide*, *Oracle Enterprise Manager Administrator's Guide*, and *Database Tuning with the Oracle Tuning Pack*.

### Steps in the Execution Plan

Each row in the output table corresponds to a single step in the execution plan. Note that the step IDs with asterisks are listed in the Predicate Information section.

**See Also:**

- ["Steps in the Execution Plan"](#) on page 1-21
- [Chapter 9, "Using EXPLAIN PLAN"](#)

Each step of the execution plan returns a set of rows that either is used by the next step or, in the last step, is returned to the user or application issuing the SQL statement. A set of rows returned by a step is called a **row set**.

The numbering of the step Ids reflects the order in which they are displayed in response to the EXPLAIN PLAN statement. Each step of the execution plan either retrieves rows from the database or accepts rows from one or more row sources as input.

- The following steps in [Example 1-4](#) physically retrieve data from an object in the database:
  - Step 3 reads all rows of the `employees` table.
  - Step 5 looks up each `job_id` in `JOB_ID_PK` index and finds the rowids of the associated rows in the `jobs` table.
  - Step 4 retrieves the rows with rowids that were returned by Step 5 from the `jobs` table.
  - Step 7 looks up each `department_id` in `DEPT_ID_PK` index and finds the rowids of the associated rows in the `departments` table.
  - Step 6 retrieves the rows with rowids that were returned by Step 7 from the `departments` table.
- The following steps in [Example 1-4](#) operate on rows returned by the previous row source:
  - Step 2 performs the nested loop operation on `job_id` in the `jobs` and `employees` tables, accepting row sources from Steps 3 and 4, joining each row from Step 3 source to its corresponding row in Step 4, and returning the resulting rows to Step 2.
  - Step 1 performs the nested loop operation, accepting row sources from Step 2 and Step 6, joining each row from Step 2 source to its corresponding row in Step 6, and returning the resulting rows to Step 1.

**See Also:**

- ["Understanding Access Paths for the CBO"](#) on page 1-24 and ["Understanding Access Paths for the RBO"](#) on page 8-3 for more information on access paths
- ["Understanding Joins"](#) on page 1-40 for more information on the methods by which Oracle joins row sources

**Understanding Execution Order** The steps of the execution plan are not performed in the order in which they are numbered in [Example 1-3](#) on page 1-19. Oracle first performs the steps that appear indented most to the right in the EXPLAIN PLAN output. In [Figure 1-3](#) on page 1-21, the steps are numbered in the order that they are performed. The rows returned by each step become the row sources of its parent step. Oracle then performs the parent steps.

Oracle performs the following steps in [Example 1-4](#) on page 1-20 to execute the statement in [Example 1-3](#) on page 1-19:

- Oracle performs Step 3 and Step 4, returning the resulting rows to Step 2. For each row returned to Step 2, Oracle performs Step 5, returning resulting rowid to Step 4.
- Oracle performs Step 1, joining the single row from Step 2 with a single row from Step 6, returning the resulting rows, if any, to the user issuing the SQL statement. For each row returned to Step 6, Oracle performs Step 7, returning resulting rowid to Step 4.

If a parent step requires only a single row from its child step before it can be executed, then Oracle performs the parent step as soon as a single row has been returned from the child step. If the parent of that parent step also can be activated by the return of a single row, then it is executed as well.

Statement execution can cascade up the tree, possibly to encompass the rest of the execution plan. Oracle performs the parent step and all cascaded steps once for each row retrieved by the child step. The parent steps that are triggered for each row returned by a child step include table accesses, index accesses, nested loop joins, and filters.

If a parent step requires all rows from its child step before it can be executed, then Oracle cannot perform the parent step until all rows have been returned from the child step. Such parent steps include sorts, sort merge joins, and aggregate functions.

## Understanding Access Paths for the CBO

Access paths are ways in which data is retrieved from the database. In general, index access paths should be used for statements that retrieve a small subset of table rows, while full scans are more efficient when accessing a large portion of the table. Online transaction processing (OLTP) applications, which consist of short-running SQL statements with high selectivity, often are characterized by the use of index access paths. Decision support systems, on the other hand, tend to use partitioned tables and perform full scans of the relevant partitions.

This section describes the data access paths that can be used to locate and retrieve any row in any table.

- [Full Table Scans](#)
- [Rowid Scans](#)
- [Index Scans](#)
- [Cluster Scans](#)
- [Hash Scans](#)
- [Sample Table Scans](#)
- [How the CBO Chooses an Access Path](#)

### Full Table Scans

This type of scan reads all rows from a table and filters out those that do not meet the selection criteria. During a full table scan, all blocks in the table that are under the high water mark are scanned. Each row is examined to determine whether it satisfies the statement's `WHERE` clause.

When Oracle performs a full table scan, the blocks are read sequentially. Because the blocks are adjacent, I/O calls larger than a single block can be used to speed up the process. The size of the read calls range from one block to the number of blocks indicated by the initialization parameter `DB_FILE_MULTIBLOCK_READ_COUNT`. Using multiblock reads means a full table scan can be performed very efficiently. Each block is read only once.

**Example 1-4, "EXPLAIN PLAN Output"** on page 1-20 contains an example of a full table scan on the `employees` table.



## Why a Full Table Scan Is Faster for Accessing Large Amounts of Data

Full table scans are cheaper than index range scans when accessing a large fraction of the blocks in a table. This is because full table scans can use larger I/O calls, and making fewer large I/O calls is cheaper than making many smaller calls.

## When the Optimizer Uses Full Table Scans

The optimizer uses a full table scan in any of the following cases:

**Lack of Index** If the query is unable to use any existing indexes, then it uses a full table scan. For example, if there is a function used on the indexed column in the query, the optimizer is unable to use the index and instead uses a full table scan as in [Example 1-5](#).

### Example 1-5 Full Table Scan

```
SELECT last_name, first_name
 FROM employees
 WHERE UPPER(last_name) LIKE :b1
```

If you need to use the index for case-independent searches, then either do not permit mixed-case data in the search columns or create a function-based index, such as `UPPER(last_name)`, on the search column. See ["Using Function-based Indexes"](#) on page 4-10.

**Large Amount of Data** If the optimizer thinks that the query will access most of the blocks in the table, then it uses a full table scan, even though indexes might be available.

**Small Table** If a table contains less than `DB_FILE_MULTIBLOCK_READ_COUNT` blocks under the high water mark, which can be read in a single I/O call, then a full table scan might be cheaper than an index range scan, regardless of the fraction of tables being accessed or indexes present.

**High Degree of Parallelism** A high degree of parallelism for a table skews the optimizer toward full table scans over range scans. Examine the `DEGREE` column in `ALL_TABLES` for the table to determine the degree of parallelism.

## Full Table Scan Hints

Use the hint `FULL(table alias)` if you want to force the use of a full table scan. For more information on the `FULL` hint, see ["FULL"](#) on page 5-10.

[Example 1-6](#) shows a query that uses an index range scan. [Example 1-7](#) shows the same query using the FULL hint to force a full table scan.

**Example 1-6 Before Using the FULL Hint**

```
SELECT employee_id, last_name
 FROM employees
 WHERE last_name LIKE :b1;
```

```

```

| Id  | Operation                   | Name        | Rows | Bytes | Cost (%CPU) |
|-----|-----------------------------|-------------|------|-------|-------------|
| 0   | SELECT STATEMENT            |             | 5    | 95    | 3 (34)      |
| 1   | TABLE ACCESS BY INDEX ROWID | EMPLOYEES   | 5    | 95    | 3 (34)      |
| * 2 | INDEX RANGE SCAN            | EMP_NAME_IX | 2    |       | 3 (34)      |

```

```

Predicate Information (identified by operation id):

```

```

```
2 - access("EMPLOYEES"."LAST_NAME" LIKE :Z)
 filter("EMPLOYEES"."LAST_NAME" LIKE :Z)
```

[Example 1-7](#) shows the [Example 1-6](#) query using the FULL hint to force a full table scan.

**Example 1-7 After Using the FULL Hint**

```
SELECT /*+ FULL(e) */ employee_id, last_name
 FROM employees e
 WHERE last_name LIKE :b1;
```

```

```

| Id  | Operation         | Name      | Rows | Bytes | Cost (%CPU) |
|-----|-------------------|-----------|------|-------|-------------|
| 0   | SELECT STATEMENT  |           | 5    | 95    | 4 (25)      |
| * 1 | TABLE ACCESS FULL | EMPLOYEES | 5    | 95    | 4 (25)      |

```

```

Predicate Information (identified by operation id):

```

```

```
1 - filter("E"."LAST_NAME" LIKE :Z)
```

**Assessing I/O for Blocks, not Rows**

Oracle does I/O by blocks. Therefore, the optimizer's decision to use full table scans is influenced by the percentage of blocks accessed, not rows. This is called the index

clustering factor. If blocks contain single rows, then rows accessed and blocks accessed are the same.

However, most tables have multiple rows in each block. Consequently, the desired number of rows could be clustered together in a few blocks, or they could be spread out over a larger number of blocks.

**See Also:** ["Estimator"](#) on page 1-13 for more information on the index clustering factor

### High Water Mark in DBA\_TABLES

The data dictionary keeps track of the blocks that have been populated with rows. The high water mark is used as the end marker during a full table scan. The high water mark is stored in `DBA_TABLES.BLOCKS`. It is reset when the table is dropped or truncated.

For example, consider a table that had a large number of rows in the past. Most of the rows have been deleted, and now most of the blocks under the high water mark are empty. A full table scan on this table exhibits poor performance because all the blocks under the high water mark are scanned.

### Parallel Query Execution

When a full table scan is required, response time can be improved by using multiple parallel execution servers for scanning the table. Parallel queries are used generally in low-concurrency data warehousing environments, because of the potential resource usage.

**See Also:** *Oracle9i Data Warehousing Guide*

## Rowid Scans

The rowid of a row specifies the datafile and data block containing the row and the location of the row in that block. Locating a row by specifying its rowid is the fastest way to retrieve a single row, because the exact location of the row in the database is specified.

To access a table by rowid, Oracle first obtains the rowids of the selected rows, either from the statement's `WHERE` clause or through an index scan of one or more of the table's indexes. Oracle then locates each selected row in the table based on its rowid.

In [Example 1-4, "EXPLAIN PLAN Output"](#) on page 1-20, an index scan is performed the `jobs` and `departments` tables. The rowids retrieved are used to return the row data.

### When the Optimizer Uses Rowids

This is generally the second step after retrieving the rowid from an index. The table access might be required for any columns in the statement not present in the index.

Access by rowid does not need to follow every index scan. If the index contains all the columns needed for the statement, then table access by rowid might not occur.

---

---

**Note:** Rowids are an internal Oracle representation of where data is stored. They can change between versions. Accessing data based on position is not recommended, because rows can move around due to row migration and chaining and also after export and import. Foreign keys should be based on primary keys. For more information on rowids, see *Oracle9i Application Developer's Guide - Fundamentals*.

---

---

## Index Scans

In this method, a row is retrieved by traversing the index, using the indexed column values specified by the statement. An index scan retrieves data from an index based on the value of one or more columns in the index. To perform an index scan, Oracle searches the index for the indexed column values accessed by the statement. If the statement accesses only columns of the index, then Oracle reads the indexed column values directly from the index, rather than from the table.

The index contains not only the indexed value, but also the rowids of rows in the table having that value. Therefore, if the statement accesses other columns in addition to the indexed columns, then Oracle can find the rows in the table by using either a table access by rowid or a cluster scan.

An index scan can be one of the following types:

- [Index Unique Scans](#)
- [Index Range Scans](#)
- [Index Range Scans Descending](#)
- [Index Skip Scans](#)
- [Full Scans](#)

- [Fast Full Index Scans](#)
- [Index Joins](#)
- [Bitmap Joins](#)

### Index Unique Scans

This scan returns, at most, a single rowid. Oracle performs a unique scan if a statement contains a `UNIQUE` or a `PRIMARY KEY` constraint that guarantees that only a single row is accessed.

In [Example 1-4, "EXPLAIN PLAN Output"](#) on page 1-20, an index scan is performed on the `jobs` and `departments` tables, using the `job_id_pk` and `dept_id_pk` indexes respectively.

**When the Optimizer Uses Index Unique Scans** This access path is used when all columns of a unique (B-tree) index are specified with equality conditions.

**See Also:** *Oracle9i Database Concepts* for more details on index structures and for detailed information on how a B-tree is searched

**Index Unique Scan Hints** In general, you should not need to use a hint to do a unique scan. There might be cases where the table is across a database link and being accessed from a local table, or where the table is small enough for the optimizer to prefer a full table scan.

The hint `INDEX(alias index_name)` specifies the index to use, but not an access path (range scan or unique scan). For more information on the `INDEX` hint, see ["INDEX"](#) on page 5-12.

### Index Range Scans

An index range scan is a common operation for accessing selective data. It can be bounded (bounded on both sides) or unbounded (on one or both sides). Data is returned in the ascending order of index columns. Multiple rows with identical values are sorted in ascending order by rowid.

If data must be sorted by order, then use the `ORDER BY` clause, and do not rely on an index. If an index can be used to satisfy an `ORDER BY` clause, then the optimizer uses this option and avoids a sort.

In [Example 1-8](#), the order has been imported from a legacy system, and you are querying the order by the reference used in the legacy system. Assume this reference is the `order_date`.

**Example 1–8 Index Range Scan**

```
SELECT order_status, order_id
FROM orders
WHERE order_date = :b1;
```

```

```

| Id  | Operation                   | Name              | Rows | Bytes | Cost (%CPU) |
|-----|-----------------------------|-------------------|------|-------|-------------|
| 0   | SELECT STATEMENT            |                   | 1    | 20    | 3 (34)      |
| 1   | TABLE ACCESS BY INDEX ROWID | ORDERS            | 1    | 20    | 3 (34)      |
| * 2 | INDEX RANGE SCAN            | ORD_ORDER_DATE_IX | 1    |       | 2 (50)      |

```

```

Predicate Information (identified by operation id):

```

2 - access("ORDERS"."ORDER_DATE"=:Z)
```

This should be a highly selective query, and you should see the query using the index on the column to retrieve the desired rows. The data returned is sorted in ascending order by the rowids for the `order_date`. Because the index column `order_date` is identical for the selected rows here, the data is sorted by rowid.

**When the Optimizer Uses Index Range Scans** The optimizer uses a range scan when it finds one or more leading columns of an index specified in conditions, such as the following:

- `col1 = :b1`
- `col1 < :b1`
- `col1 > :b1`
- AND combination of the preceding conditions for leading columns in the index
- `col1 like '%ASD'` Wild-card searches should not be in a leading position. The condition `col1 like '%ASD'` does not result in a range scan.

Range scans can use unique or nonunique indexes. Range scans avoid sorting when index columns constitute the `ORDER BY/GROUP BY` clause.

**Index Range Scan Hints** A hint might be required if the optimizer chooses some other index or uses a full table scan. The hint `INDEX(table_alias index_name)` specifies the index to use. For more information on the `INDEX` hint, see ["INDEX"](#) on page 5-12.

Suppose that `order_id` has a skewed distribution. The column has histograms, so the optimizer knows about the distribution. However, with a bind variable, the

optimizer does not know the value and could choose a full table scan. You have two options:

- Use literals rather than bind variables, which can cause problems due to nonsharing of SQL statements.
- Use hints in order to share the statements.

[Example 1-9](#) shows a query before using the `INDEX` hint.

#### **Example 1-9 Before Using the INDEX Hint**

```
SELECT l.line_item_id, order_id, l.unit_price * l.quantity
 FROM order_items l
 WHERE l.order_id = :b1;
```

| Id  | Operation         | Name        | Rows | Bytes | Cost (%CPU) |
|-----|-------------------|-------------|------|-------|-------------|
| 0   | SELECT STATEMENT  |             | 6    | 90    | 4 (25)      |
| * 1 | TABLE ACCESS FULL | ORDER_ITEMS | 6    | 90    | 4 (25)      |

Predicate Information (identified by operation id):

```
1 - filter("L"."ORDER_ID"=TO_NUMBER(:Z))
```

[Example 1-10](#) shows the [Example 1-9](#) query using the `INDEX` hint.

#### **Example 1-10 Using Bind Variables and INDEX Hint:**

```
SELECT /*+ INDEX(l item_order_ix) */ l.line_item_id, order_id,
 l.unit_price * l.quantity
 FROM order_items l
 WHERE l.order_id = :b1;
```

| Id  | Operation                   | Name          | Rows | Bytes | Cost (%CPU) |
|-----|-----------------------------|---------------|------|-------|-------------|
| 0   | SELECT STATEMENT            |               | 6    | 90    | 10 (10)     |
| 1   | TABLE ACCESS BY INDEX ROWID | ORDER_ITEMS   | 6    | 90    | 10 (10)     |
| * 2 | INDEX RANGE SCAN            | ITEM_ORDER_IX | 6    |       | 2 (50)      |

Predicate Information (identified by operation id):

```
2 - access("L"."ORDER_ID"=TO_NUMBER(:Z))
```

## Index Range Scans Descending

An index range scan descending is identical to an index range scan, except that the data is returned in descending order. Indexes, by default, are stored in ascending order. Usually, this scan is used when ordering data in a descending order to return the most recent data first, or when seeking a value less than a specified value.

[Example 1-11](#) uses a two-column unique index on `order_id`, `line_item_id`.

### **Example 1-11** Index Range Scan Descending Using Two-Column Unique Index

```
SELECT line_item_id, order_id
 FROM order_items
 WHERE order_id < :b1
 ORDER BY order_id DESC;
```

| Id  | Operation                   | Name          | Rows | Bytes | Cost (%CPU) |
|-----|-----------------------------|---------------|------|-------|-------------|
| 0   | SELECT STATEMENT            |               | 33   | 231   | 3 (34)      |
| 1   | TABLE ACCESS BY INDEX ROWID | ORDER_ITEMS   | 33   | 231   | 3 (34)      |
| * 2 | INDEX RANGE SCAN DESCENDING | ITEM_ORDER_IX | 6    |       | 3 (34)      |

Predicate Information (identified by operation id):

```
2 - access("ORDER_ITEMS"."ORDER_ID"<TO_NUMBER(:Z))
 filter("ORDER_ITEMS"."ORDER_ID"<TO_NUMBER(:Z))
```

The data is sorted in descending order by the `order_id`, `line_item_id`, `rowid` of the selected rows. However, because there is only one row for each `order_id`, `line_item_id` (`item_order_ix` is a unique index on the two columns), the rows are sorted by `order_id`, `line_item_id`.

**When the Optimizer Uses Index Range Scans Descending** The optimizer uses index range scan descending when an order by descending clause can be satisfied by an index.

**Index Range Scan Descending Hints** The hint `INDEX_DESC(table_alias index_name)` is used for this access path. For more information on the `INDEX_DESC` hint, see ["INDEX\\_DESC"](#) on page 5-15.

## Index Skip Scans

Index skip scans improve index scans by nonprefix columns. Often, scanning index blocks is faster than scanning table data blocks.



Skip scanning lets a composite index be split logically into smaller subindexes. In skip scanning, the initial column of the composite index is not specified in the query. In other words, it is skipped.

The number of logical subindexes is determined by the number of distinct values in the initial column. Skip scanning is advantageous if there are few distinct values in the leading column of the composite index and many distinct values in the nonleading key of the index.

**Example 1–12 Index Skip Scan**

Consider, for example, a table `employees` (`sex`, `employee_id`, `address`) with a composite index on (`sex`, `employee_id`). Splitting this composite index would result in two logical subindexes, one for M and one for F.


For this example, suppose you have the following index data:

```
('F' , 98)
('F' , 100)
('F' , 102)
('F' , 104)
('M' , 101)
('M' , 103)
('M' , 105)
```

The index is split logically into the following two subindexes:

- The first subindex has the keys with the value F.
- The second subindex has the keys with the value M.

**Figure 1–4 Index Skip Scan Illustration**



The column `sex` is skipped in the following query:

```
SELECT *
 FROM employees
 WHERE employee_id = 101;
```

A complete scan of the index is not performed, but the subindex with the value `F` is searched first, followed by a search of the subindex with the value `M`.

## Full Scans

A full scan is available if a predicate references one of the columns in the index. The predicate does not need to be an index driver. A full scan is also available when there is no predicate, if both the following conditions are met:

- All of the columns in the table referenced in the query are included in the index.
- At least one of the index columns is not null.

A full scan can be used to eliminate a sort operation, because the data is ordered by the index key. It reads the blocks singly.

## Fast Full Index Scans

Fast full index scans are an alternative to a full table scan when the index contains all the columns that are needed for the query, and at least one column in the index key has the `NOT NULL` constraint. A fast full scan accesses the data in the index itself, without accessing the table. It cannot be used to eliminate a sort operation, because the data is *not* ordered by the index key. It reads the entire index using multiblock reads, unlike a full index scan, and can be parallelized.

Fast full scan is available only with the CBO. You can specify it with the initialization parameter `OPTIMIZER_FEATURES_ENABLE` or the `INDEX_FFS` hint. Fast full index scans cannot be performed against bitmap indexes.

A fast full scan is faster than a normal full index scan in that it can use multiblock I/O and can be parallelized just like a table scan.

**Fast Full Index Scan Hints** The fast full scan has a special index hint, `INDEX_FFS`, which has the same format and arguments as the regular `INDEX` hint. For more information on the `INDEX_FFS` hint, see "[INDEX\\_FFS](#)" on page 5-16.

**Fast Full Index Scan Restrictions** Fast full index scans have the following restrictions:

- At least one indexed column of the table must have the `NOT NULL` constraint.

- There must be a parallel clause on the index if you want to perform a fast full index scan in parallel. The parallel degree of the index is set independently. The index does *not* inherit the degree of parallelism of the table.
- You must have analyzed the index; otherwise, the optimizer might decide not to use it.

### Index Joins

An index join is a hash join of several indexes that together contain all the table columns that are referenced in the query. If an index join is used, then no table access is needed, because all the relevant column values can be retrieved from the indexes. An index join cannot be used to eliminate a sort operation. The index join is available only with the CBO.

**Index Join Hints** You can specify an index join with the initialization parameter `OPTIMIZER_FEATURES_ENABLE` or the `INDEX_JOIN` hint. For more information on the `INDEX_JOIN` hint, see "[INDEX\\_JOIN](#)" on page 5-15.

### Bitmap Joins

A bitmap join uses a bitmap for key values and a mapping function that converts each bit position to a rowid. Bitmaps can efficiently merge indexes that correspond to several conditions in a `WHERE` clause, using Boolean operations to resolve `AND` and `OR` conditions.

Bitmap access is available only with the CBO.

---

---

**Note:** Bitmap indexes and bitmap join indexes are available only if you have purchased the Oracle9i Enterprise Edition.

---

---

**See Also:** *Oracle9i Data Warehousing Guide* for more information about bitmap indexes

## Cluster Scans

A cluster scan is used to retrieve, from a table stored in an indexed cluster, all rows that have the same cluster key value. In an indexed cluster, all rows with the same cluster key value are stored in the same data block. To perform a cluster scan, Oracle first obtains the rowid of one of the selected rows by scanning the cluster index. Oracle then locates the rows based on this rowid.

## Hash Scans

A hash scan is used to locate rows in a hash cluster, based on a hash value. In a hash cluster, all rows with the same hash value are stored in the same data block. To perform a hash scan, Oracle first obtains the hash value by applying a hash function to a cluster key value specified by the statement. Oracle then scans the data blocks containing rows with that hash value.

## Sample Table Scans

A sample table scan retrieves a random sample of data from a table. This access path is used when a statement's FROM clause includes the SAMPLE clause or the SAMPLE BLOCK clause. To perform a sample table scan when sampling by rows (the SAMPLE clause), Oracle reads a specified percentage of rows in the table. To perform a sample table scan when sampling by blocks (the SAMPLE BLOCK clause), Oracle reads a specified percentage of table blocks.

Oracle does not support sample table scans when the query involves a join or a remote table. However, you can perform an equivalent operation by using a CREATE TABLE AS SELECT query to materialize a sample of an underlying table. You then rewrite the original query to refer to the newly created table sample. Additional queries can be written to materialize samples for other tables. Sample table scans require the CBO.

[Example 1-13](#) uses a sample table scan to access 1% of the employees table, sampling by blocks.

### **Example 1-13 Sample Table Scan**

```
SELECT *
 FROM employees SAMPLE BLOCK (1);
```

The EXPLAIN PLAN output for this statement might look like this:

| Id | Operation           | Name      | Rows | Bytes | Cost (%CPU) |
|----|---------------------|-----------|------|-------|-------------|
| 0  | SELECT STATEMENT    |           | 1    | 68    | 3 (34)      |
| 1  | TABLE ACCESS SAMPLE | EMPLOYEES | 1    | 68    | 3 (34)      |

## How the CBO Chooses an Access Path

The CBO chooses an access path based on the following factors:

- The available access paths for the statement
- The estimated cost of executing the statement, using each access path or combination of paths

To choose an access path, the optimizer first determines which access paths are available by examining the conditions in the statement's `WHERE` clause and its `FROM` clause for the `SAMPLE` or `SAMPLE BLOCK` clause. The optimizer then generates a set of possible execution plans using available access paths and estimates the cost of each plan, using the statistics for the index, columns, and tables accessible to the statement. Finally, the optimizer chooses the execution plan with the lowest estimated cost.

When choosing an access path, the CBO is influenced by the following:

- Optimizer Hints

The optimizer's choice among available access paths can be overridden with hints, except when the statement's `FROM` clause contains `SAMPLE` or `SAMPLE BLOCK`.

**See Also:** [Chapter 5, "Optimizer Hints"](#) for information about hints in SQL statements

- Old Statistics

For example, if a table has not been analyzed since it was created, and if it has less than `DB_FILE_MULTIBLOCK_READ_COUNT` blocks under the high water mark, then the optimizer thinks that the table is small and uses a full table scan. Review the `LAST_ANALYZED` and `BLOCKS` columns in the `ALL_TABLES` table to see what the statistics reflect.

### Choosing an Access Path Examples

This section discusses how the optimizer chooses an access path.

**Case 1** In [Example 1-14](#) the query uses an equality condition in its `WHERE` clause to select all employees named `Jackson`.

**Example 1–14 Access Path Selection**

```
SELECT *
 FROM employees
 WHERE last_name = 'JACKSON';
```

If the `last_name` column is a unique or primary key, then the optimizer determines that there is only one employee named Jackson, and the query returns only one row. In this case, the query is very selective, and the optimizer is most likely to access the table using a unique scan on the index that enforces the unique or primary key.

**Case 2** Consider again the query in [Example 1–14](#). If the `last_name` column is not a unique or primary key, then the optimizer can use the following statistics to estimate the query's selectivity:

- `USER_TAB_COLUMNS.NUM_DISTINCT`, which is the number of values for each column in the table
- `USER_TABLES.NUM_ROWS`, which is the number of rows in each table

By dividing the number of rows in the `employees` table by the number of distinct values in the `last_name` column, the optimizer estimates what percentage of employees have the same name. By assuming that the `last_name` values are distributed uniformly, the optimizer uses this percentage as the estimated selectivity of the query.

**Case 3** The following query selects all employees with employee ID numbers less than 7500:

```
SELECT *
 FROM employees
 WHERE employee_id < 7500;
```

To estimate the selectivity of the query, the optimizer uses the boundary value of 7500 in the `WHERE` clause condition and the values of the `HIGH_VALUE` and `LOW_VALUE` statistics for the `employee_id` column, if available. These statistics can be found in the `USER_TAB_COL_STATISTICS` view or the `USER_TAB_COLUMNS` view. The optimizer assumes that `employee_id` values are distributed evenly in the range between the lowest value and highest value. The optimizer then determines what percentage of this range is less than the value 7500 and uses this value as the estimated selectivity of the query.

**Case 4** The following query uses a bind variable rather than a literal value for the boundary value in the `WHERE` clause condition:

```
SELECT *
 FROM employees
 WHERE employee_id < :e1;
```

The optimizer does not know the value of the bind variable `e1`. The value of `e1` might be different each time the query is executed. For this reason, the optimizer cannot use the means described in the previous example to determine selectivity of this query. In this case, the optimizer heuristically guesses a small value for the selectivity, using an internal default value. The optimizer makes this assumption whenever a bind variable is used as a boundary value in a condition with one of the following operators: `<`, `>`, `<=`, or `>=`.

The optimizer's treatment of bind variables can cause it to choose different execution plans for SQL statements that differ only in the use of bind variables rather than constants. For example, the optimizer might choose different execution plans for an embedded SQL statement with a bind variable in an Oracle precompiler program and the same SQL statement with a constant in SQL\*Plus.

**Case 5** The following query uses two bind variables as boundary values in the condition with the `BETWEEN` operator:

```
SELECT *
 FROM employees
 WHERE employee_id BETWEEN :low_e AND :high_e;
```

The optimizer rewrites the `BETWEEN` condition as the following two conditions:

```
employee_id >= :low_e
employee_id <= :high_e
```

The optimizer heuristically estimates a small selectivity (an internal default value) for indexed columns in order to favor the use of the index.

**Case 6** The following query uses the `BETWEEN` operator to select all employees with employee ID numbers between 7500 and 7800:

```
SELECT *
 FROM employees
 WHERE employee_id BETWEEN 7500 AND 7800;
```

To determine the selectivity of this query, the optimizer rewrites the `WHERE` clause condition into the following two conditions:

```
employee_id >= 7500
employee_id <= 7800
```

The optimizer estimates the individual selectivity of each condition, using the means described in [Case 4](#) on page 1-38. The optimizer then uses these selectivities (S1 and S2) and the absolute value function (ABS) to estimate the selectivity (S) of the BETWEEN condition, as follows:

$$S = \text{ABS}(S1 + S2 - 1)$$

## Understanding Joins

Joins are statements that retrieve data from more than one table. A join is characterized by multiple tables in the FROM clause, and the relationship between the tables is defined through the existence of a join condition in the WHERE clause.

This section discusses:

- [How the CBO Executes Join Statements](#)
- [How the CBO Chooses the Join Method](#)
- [How the CBO Chooses Execution Plans for Join Types](#)
- Join Methods:
  - [Nested Loop Joins](#)
  - [Hash Joins](#)
  - [Sort Merge Joins](#)
  - [Cartesian Joins](#)
  - [Outer Joins](#)

**See Also:** A discussion of joins in the *Oracle9i SQL Reference*

### How the CBO Executes Join Statements

To choose an execution plan for a join statement, the optimizer must make these interrelated decisions:

**Access Paths** As for simple statements, the optimizer must choose an access path to retrieve data from each table in the join statement.

**Join Method** To join each pair of row sources, Oracle must perform a join operation. Join methods include nested loop, sort merge, cartesian, and hash joins.



**Join Order** To execute a statement that joins more than two tables, Oracle joins two of the tables and then joins the resulting row source to the next table. This process is continued until all tables are joined into the result.

**See Also:**

["Understanding Access Paths for the RBO" on page 8-3](#)

["Understanding Access Paths for the CBO" on page 1-24](#)

## How the CBO Chooses the Join Method

The optimizer estimates the cost of each join method and chooses the method with the least cost. If a join returns many rows, then the optimizer considers the following three factors:

- A nested loop join is inefficient when a join returns a large number of rows (typically, more than 10,000 rows is considered large), and the optimizer might choose not to use it. The cost of a nested loop join is calculated by the following formula:

`cost= access cost of A + (access cost of B * number of rows from A)`

- If you are using the CBO, then a hash join is the most efficient join when a join returns a large number of rows. The cost of a hash join is calculated by the following formula:

`cost= (access cost of A * number of hash partitions of B) + access cost of B`

- If you are using the RBO, then a merge join is the most efficient join when a join returns a large number of rows. The cost of a merge join is calculated by the following formula:

`cost= access cost of A + access cost of B +(sort cost of A + sort cost of B)`

When the data is presorted, the sort costs are both zero.

---

---

**Note:** Oracle Corporation strongly advises the use of cost-based optimization. Rule-based optimization will be deprecated in a future release.

---

---

## How the CBO Chooses Execution Plans for Join Types

The following considerations apply to both the cost-based and rule-based approaches:

- The optimizer first determines whether joining two or more tables definitely results in a row source containing at most one row. The optimizer recognizes such situations based on `UNIQUE` and `PRIMARY KEY` constraints on the tables. If such a situation exists, then the optimizer places these tables first in the join order. The optimizer then optimizes the join of the remaining set of tables.
- For join statements with outer join conditions, the table with the outer join operator must come after the other table in the condition in the join order. The optimizer does not consider join orders that violate this rule.

With the CBO, the optimizer generates a set of execution plans, according to possible join orders, join methods, and available access paths. The optimizer then estimates the cost of each plan and chooses the one with the lowest cost. The optimizer estimates costs in the following ways:

- The cost of a nested loops operation is based on the cost of reading each selected row of the outer table and each of its matching rows of the inner table into memory. The optimizer estimates these costs using the statistics in the data dictionary.
- The cost of a sort merge join is based largely on the cost of reading all the sources into memory and sorting them.

The optimizer also considers other factors when determining the cost of each operation. For example:

- A smaller sort area size is likely to increase the cost for a sort merge join because sorting takes more CPU time and I/O in a smaller sort area. Sort area size is specified by the initialization parameter `SORT_AREA_SIZE`.

---

---

**Note:** Oracle does not recommend using the `SORT_AREA_SIZE` parameter unless the instance is configured with the shared server option. Oracle recommends instead that you enable automatic sizing of SQL work areas by setting `PGA_AGGREGATE_TARGET`. `SORT_AREA_SIZE` is retained only for backward compatibility.

---

---

- A larger multiblock read count is likely to decrease the cost for a sort merge join in relation to a nested loop join. If a large number of sequential blocks can be read from disk in a single I/O, then an index on the inner table for the nested

loop join is less likely to improve performance over a full table scan. The multiblock read count is specified by the initialization parameter `DB_FILE_MULTIBLOCK_READ_COUNT`.

With the CBO, the optimizer's choice of join orders can be overridden with the `ORDERED` hint. If the `ORDERED` hint specifies a join order that violates the rule for an outer join, then the optimizer ignores the hint and chooses the order. Also, you can override the optimizer's choice of join method with hints.

**See Also:** [Chapter 5, "Optimizer Hints"](#) for more information about optimizer hints

### How the CBO Executes Anti-joins

An anti-join returns rows from the left side of the predicate for which there are no corresponding rows on the right side of the predicate. That is, it returns rows that fail to match (`NOT IN`) the subquery on the right side. For example, an anti-join can select a list of employees who are not in a particular set of departments:

#### *Example 1–15 Anti-join on Employees and Departments*

```
SELECT * FROM employees
WHERE department_id NOT IN
 (SELECT department_id FROM departments
 WHERE location_id = 1700);
```

The optimizer uses a nested loops algorithm for `NOT IN` subqueries by default. However, if the `MERGE_AJ`, `HASH_AJ`, or `NL_AJ` hint is used and various required conditions are met, the `NOT IN` uncorrelated subquery can be changed into a sort merge or hash antijoin.

### How the CBO Executes Semi-joins

A semi-join returns rows that match an `EXISTS` subquery without duplicating rows from the left side of the predicate when multiple rows on the right side satisfy the criteria of the subquery. For example:

#### *Example 1–16 Semi-join on Employees and Departments*

```
SELECT * FROM departments
WHERE EXISTS
 (SELECT * FROM employees
 WHERE departments.department_id = employees.department_id
 AND employees.salary > 2500);
```

In this query, only one row needs to be returned from `departments`, even though many rows in `employees` might match the subquery. If there is no index on the `salary` column in `employees`, then a semi-join can be used to improve query performance.

The optimizer uses a nested loops algorithm by default for `IN` or `EXISTS` subqueries that cannot be merged with the containing query. However, if the `MERGE_SJ`, `HASH_SJ`, or `NL_SJ` hint is used and various required conditions are met, the subquery can be changed into a sort merge or hash semi-join.

**See Also:** [Chapter 5, "Optimizer Hints"](#) for information about optimizer hints

---

---

**Note:** Semi-join transformation cannot be done if the subquery is on an `OR` branch. This also applies to anti-joins.

---

---

### How the CBO Executes Star Queries

Some data warehouses are designed around a star schema, which includes a large fact table and several small dimension (lookup) tables. The fact table stores primary information. Each dimension table stores information about an attribute in the fact table.

A star query is a join between a fact table and a number of lookup tables. Each lookup table is joined by its primary keys to the corresponding foreign keys of the fact table, but the lookup tables are not joined to each other.

The CBO recognizes star queries and generates efficient execution plans for them. Star queries are not recognized by the RBO.

A typical fact table contains keys and measures. For example, a simple fact table might contain the measure `Sales`, and the keys `Time`, `Product`, and `Market`. In this case there would be corresponding dimension tables for `Time`, `Product`, and `Market`. The `Product` dimension table, for example, typically contains information about each product number that appears in the fact table.

A star join uses a join of foreign keys in a fact table to the corresponding primary keys in dimension tables. The fact table normally has a concatenated index on the foreign key columns to facilitate this type of join, or it has a separate bitmap index on each foreign key column.

**See Also:** *Oracle9i Data Warehousing Guide* for more information about tuning star queries

## Nested Loop Joins

Nested loop joins are useful when small subsets of data are being joined and if the join condition is an efficient way of accessing the second table.

It is very important to ensure that the inner table is driven from (dependent on) the outer table. If the inner table's access path is independent of the outer table, then the same rows are retrieved for every iteration of the outer loop, degrading performance considerably. In such cases, hash joins joining the two independent row sources perform better.

**See Also:** "[Cartesian Joins](#)" on page 1-50

A nested loop join involves the following steps:

1. The optimizer determines the driving table and designates it as the outer table.
2. The other table is designated as the inner table.
3. For every row in the outer table, Oracle accesses all the rows in the inner table. The outer loop is for every row in outer table and the inner loop is for every row in the inner table. The outer loop appears before the inner loop in the execution plan, as follows:

```

NESTED LOOPS
 outer_loop
 inner_loop

```

### Nested Loop Example

This section discusses the outer and inner loops for the following nested loops in the query in [Example 1-3](#).

```

...
2	NESTED LOOPS		3	141	7 (15)
* 3	TABLE ACCESS FULL	EMPLOYEES	3	60	4 (25)
4	TABLE ACCESS BY INDEX ROWID	JOBS	19	513	2 (50)
* 5	INDEX UNIQUE SCAN	JOB_ID_PK	1		
...

```

In this example, the outer loop retrieves all the rows of the `employees` table. For every employee retrieved by the outer loop, the inner loop retrieves the associated row in the `jobs` table.

**Outer loop** In the execution plan in [Example 1-4](#), the outer loop and the equivalent statement are as follows:

```
3 | TABLE ACCESS FULL | EMPLOYEES

SELECT e.employee_id, e.salary
 FROM employees e
 WHERE e.employee_id < 103
```

**Inner loop** The execution plan in [Example 1-4](#) shows the inner loop being iterated for every row fetched from the outer loop, as follows:

```
4 | TABLE ACCESS BY INDEX ROWID | JOBS
5 | INDEX UNIQUE SCAN | JOB_ID_PK

SELECT j.job_title
 FROM jobs j
 WHERE e.job_id = j.job_id
```

### When the Optimizer Uses Nested Loop Joins

The optimizer uses nested loop joins when joining small number of rows, with a good driving condition between the two tables. You drive from the outer loop to the inner loop, so the order of tables in the execution plan is important.

The outer loop is the driving row source. It produces a set of rows for driving the join condition. The row source can be a table accessed using an index scan or a full table scan. Also, the rows can be produced from any other operation. For example, the output from a nested loop join can be used as a row source for another nested loop join.

The inner loop is iterated for every row returned from the outer loop, ideally by an index scan. If the access path for the inner loop is not dependent on the outer loop, then you can end up with a Cartesian product; for every iteration of the outer loop, the inner loop produces the same set of rows. Therefore, you should use other join methods when two independent row sources are joined together.

### Nested Loop Join Hints

If the optimizer is choosing to use some other join method, you can use the `USE_NL(table1 table2)` hint, where `table1` and `table2` are the aliases of the tables being joined.

For some SQL examples, the data is small enough for the optimizer to prefer full table scans and use hash joins. This is the case for the SQL example shown in [Example 1-17, "Hash Joins"](#) on page 1-48. However, you can add a `USE_NL` hint that

changes the join method to nested loop. For more information on the `USE_NL` hint, see "[USE\\_NL](#)" on page 5-24.

## Nesting Nested Loops

The outer loop of a nested loop can be a nested loop itself. You can nest two or more outer loops together to join as many tables as needed. Each loop is a data access method, as follows:

```
SELECT STATEMENT
 NESTED LOOP 3
 NESTED LOOP 2 (OUTER LOOP 3.1)
 NESTED LOOP 1 (OUTER LOOP 2.1)
 OUTER LOOP 1.1 - #1
 INNER LOOP 1.2 - #2
 INNER LOOP 2.2 - #3
 INNER LOOP 3.2 - #4
```

## Hash Joins

Hash joins are used for joining large data sets. The optimizer uses the smaller of two tables or data sources to build a hash table on the join key in memory. It then scans the larger table, probing the hash table to find the joined rows.

This method is best used when the smaller table fits in available memory. The cost is then limited to a single read pass over the data for the two tables.

However, if the hash table grows too big to fit into the memory, then the optimizer breaks it up into different partitions. As the partitions exceed allocated memory, parts are written to temporary segments on disk. Larger temporary extent sizes lead to improved I/O when writing the partitions to disk; the recommended temporary extent is about 1 MB. Temporary extent size is specified by `INITIAL` and `NEXT` for permanent tablespaces and by `UNIFORM SIZE` for temporary tablespaces.

After the hash table is complete, the following processes occur:

1. The second, larger table is scanned.
2. It is broken up into partitions like the smaller table.
3. The partitions are written to disk.

When the hash table build is complete, it is possible that an entire hash table partition is resident in memory. Then, you do not need to build the corresponding partition for the second (larger) table. When that table is scanned, rows that hash to the resident hash table partition can be joined and returned immediately.

Each hash table partition is then read into memory, and the following processes occur:

1. The corresponding partition for the second table is scanned.
2. The hash table is probed to return the joined rows.

This process is repeated for the rest of the partitions. The cost can increase to two read passes over the data and one write pass over the data.

If the hash table does not fit in the memory, it is possible that parts of it may need to be swapped in and out, depending on the rows retrieved from the second table. Performance for this scenario can be extremely poor.

### When the Optimizer Uses Hash Joins

The optimizer uses a hash join to join two tables if they are joined using an equijoin and if either of the following conditions are true:

- A large amount of data needs to be joined.
- A large fraction of the table needs to be joined.

In [Example 1-17](#), the table `orders` is used to build the hash table, and `order_items` is the larger table, which is scanned later.

#### **Example 1-17 Hash Joins**

```
SELECT o.customer_id, l.unit_price * l.quantity
 FROM orders o ,order_items l
 WHERE l.order_id = o.order_id;
```

| Id  | Operation         | Name        | Rows | Bytes | Cost (%CPU) |
|-----|-------------------|-------------|------|-------|-------------|
| 0   | SELECT STATEMENT  |             | 665  | 13300 | 8 (25)      |
| * 1 | HASH JOIN         |             | 665  | 13300 | 8 (25)      |
| 2   | TABLE ACCESS FULL | ORDERS      | 105  | 840   | 4 (25)      |
| 3   | TABLE ACCESS FULL | ORDER_ITEMS | 665  | 7980  | 4 (25)      |

Predicate Information (identified by operation id):

```
1 - access("L"."ORDER_ID"="O"."ORDER_ID")
```



## Hash Join Hints

Apply the `USE_HASH` hint to advise the optimizer to use a hash join when joining two tables together. If you are having trouble getting the optimizer to use hash joins, investigate the values for the `HASH_AREA_SIZE` and `HASH_JOIN_ENABLED` parameters.

---

---

**Note:** Oracle does not recommend using the `HASH_AREA_SIZE` parameter unless the instance is configured with the shared server option. Oracle recommends instead that you enable automatic sizing of SQL work areas by setting `PGA_AGGREGATE_TARGET`. `HASH_AREA_SIZE` is retained only for backward compatibility.

---

---

For more information on the `USE_HASH` hint, see "[USE\\_HASH](#)" on page 5-26.

## Sort Merge Joins

Sort merge joins can be used to join rows from two independent sources. Hash joins generally perform better than sort merge joins. On the other hand, sort merge joins can perform better than hash joins if both of the following conditions exist:

- The row sources are sorted already.
- A sort operation does not have to be done.

However, if a sort merge join involves choosing a slower access method (an index scan as opposed to a full table scan), then the benefit of using a sort merge might be lost.

Sort merge joins are useful when the join condition between two tables is an inequality condition (but not a nonequality) like `<`, `<=`, `>`, or `>=`. Sort merge joins perform better than nested loop joins for large data sets. You cannot use hash joins unless there is an equality condition.

In a merge join, there is no concept of a driving table. The join consists of two steps:

1. Sort join operation: Both the inputs are sorted on the join key.
2. Merge join operation: The sorted lists are merged together.

If the input is already sorted by the join column, then a sort join operation is not performed for that row source.

### When the Optimizer Uses Sort Merge Joins

The optimizer can choose a sort merge join over a hash join for joining large amounts of data if any of the following conditions are true:

- The join condition between two tables is not an equi-join.
- `OPTIMIZER_MODE` is set to `RULE`.
- `HASH_JOIN_ENABLED` is false.
- Because of sorts already required by other operations, the optimizer finds it is cheaper to use a sort merge than a hash join.
- The optimizer thinks that the cost of a hash join is higher, based on the settings of `HASH_AREA_SIZE` and `SORT_AREA_SIZE`.

---

---

**Note:** Oracle does not recommend using the `HASH_AREA_SIZE` and `SORT_AREA_SIZE` parameters unless the instance is configured with the shared server option. Oracle recommends instead that you enable automatic sizing of SQL work areas by setting `PGA_AGGREGATE_TARGET`. `HASH_AREA_SIZE` and `SORT_AREA_SIZE` are retained only for backward compatibility.

---

---

### Sort Merge Join Hints

To advise the optimizer to use a sort merge join, apply the `USE_MERGE` hint. You might also need to give hints to force an access path.

There are situations where it is better to override the optimizer with the `USE_MERGE` hint. For example, the optimizer can choose a full scan on a table and avoid a sort operation in a query. However, there is an increased cost because a large table is accessed through an index and single block reads, as opposed to faster access through a full table scan.

For more information on the `USE_MERGE` hint, see "[USE\\_MERGE](#)" on page 5-25.

## Cartesian Joins

A Cartesian join is used when one or more of the tables does not have any join conditions to any other tables in the statement. The optimizer joins every row from one data source with every row from the other data source, creating the Cartesian product of the two sets.

### When the Optimizer Uses Cartesian Joins

The optimizer uses Cartesian joins when it is asked to join two tables with no join conditions. In some cases, a common filter condition between the two tables could be picked up by the optimizer as a possible join condition. This is even more dangerous, because the joins are not flagged in the execution plan as being a Cartesian product.

Cartesian joins generally result from poorly written SQL. You should specify join criteria and avoid Cartesian products. In some situations, a query can involve a large number of tables and an extra table is contained in the `FROM` clause, but not in the `WHERE` clause. With such queries, a `DISTINCT` clause can weed out multiple rows. However, while the `DISTINCT` clause can be used to remove the extra tuples generated by the Cartesian product, the performance can be severely degraded.

In a SQL query where the inner table of a nested loop operation is not driven from the outer table, but from an independent row source, then the rows accessed can be the same as in a Cartesian product. Because the join condition is present but is applied after accessing the table, the result is not a Cartesian product. However, the cost of accessing the table (rows accessed) is about the same.

### Cartesian Join Hints

Applying the `ORDERED` hint, causes the optimizer uses a Cartesian join. By specifying a table before its join table is specified, the optimizer does a Cartesian join. For more information on the `ORDERED` hint, see "[ORDERED](#)" on page 5-23.

## Outer Joins

An outer join extends the result of a simple join. An outer join returns all rows that satisfy the join condition and also returns some or all of those rows from one table for which no rows from the other satisfy the join condition.

### Nested Loop Outer Joins

This operation is used when an outer join is used between two tables. The outer join returns the outer (preserved) table rows, even when there are no corresponding rows in the inner (optional) table.

In a regular outer join, the optimizer chooses the order of tables (driving and driven) based on the cost. However, in a nested loop outer join, the order of tables is determined by the join condition. The outer table, with rows that are being preserved, is used to drive to the inner table.

The optimizer uses nested loop joins to process an outer join in the following circumstances:

- It is possible to drive from the outer table to inner table.
- Data volume is low enough to make the nested loop method efficient.

For an example of a nested loop outer join, you can add the `USE_NL` hint to [Example 1-18](#) to ensure that a nested loop is used. For example:

```
SELECT /*+ USE_NL(c o) */ cust_last_name, sum(nvl2(o.customer_id,0,1)) "Count"
```

### Hash Join Outer Joins

The optimizer uses hash joins for processing an outer join if the data volume is high enough to make the hash join method efficient or if it is not possible to drive from the outer table to inner table.

Like an outer join, the order of tables is not determined by the cost, but by the join condition. The outer table (with preserved rows) is used to build the hash table, and the inner table is used to probe the hash table.

[Example 1-18](#) shows a typical hash join outer join query. In this example, all the customers with credit limits greater than 1000 are queried. An outer join is needed so that you do not miss the customers who do not have any orders.

#### **Example 1-18 Hash Join Outer Joins**

```
SELECT cust_last_name, sum(nvl2(o.customer_id,0,1)) "Count"
FROM customers c, orders o
WHERE c.credit_limit > 1000
AND c.customer_id = o.customer_id(+)
GROUP BY cust_last_name;
```

| Id  | Operation         | Name      | Rows | Bytes | Cost (%CPU) |
|-----|-------------------|-----------|------|-------|-------------|
| 0   | SELECT STATEMENT  |           | 168  | 3192  | 11 (28)     |
| 1   | SORT GROUP BY     |           | 168  | 3192  | 11 (28)     |
| * 2 | HASH JOIN OUTER   |           | 260  | 4940  | 10 (20)     |
| * 3 | TABLE ACCESS FULL | CUSTOMERS | 260  | 3900  | 6 (17)      |
| * 4 | TABLE ACCESS FULL | ORDERS    | 105  | 420   | 4 (25)      |

Predicate Information (identified by operation id):  
 -----

```

2 - access("C"."CUSTOMER_ID"="O"."CUSTOMER_ID"(+))
3 - filter("C"."CREDIT_LIMIT">1000)
4 - filter("O"."CUSTOMER_ID"(+)>0)

```

The query looks for customers which satisfy various conditions. An outer join returns NULL for the inner table columns along with the outer (preserved) table rows when it does not find any corresponding rows in the inner table. This operation finds all the `customers` rows that do not have any `orders` rows.

In this case, the outer join condition is the following:

```
customers.customer_id = orders.customer_id(+)
```

The components of this condition represent the following:

- The outer table is `customers`.
- The inner table is `orders`.
- The join preserves the `customers` rows, including those rows without a corresponding row in `orders`.
- The hash table is built using `customers`.
- The hash table is probed using `orders`.

You could use a `NOT EXISTS` subquery to return the rows. However, because you are querying all the rows in the table, the hash join performs better (unless the `NOT EXISTS` subquery is not nested).

In [Example 1–19](#), the outer join is to a multitable view. The optimizer cannot drive into the view like in a normal join or push the predicates, so it builds the entire row set of the view.

#### **Example 1–19 Outer Join to a Multitable View**

```

SELECT c.cust_last_name, sum(revenue)
 FROM customers c, v_orders o
 WHERE c.credit_limit > 2000
 AND o.customer_id(+) = c.customer_id
 GROUP BY c.cust_last_name;

```

```

| Id | Operation | Name | Rows | Bytes | Cost (%CPU)|

0	SELECT STATEMENT		1	44	16 (32)
1	SORT GROUP BY		1	44	16 (32)
* 2	HASH JOIN OUTER		663	21216	15 (27)
* 3	TABLE ACCESS FULL	CUSTOMERS	195	2925	6 (17)
4	VIEW	V_ORDERS	665	11305	
5	SORT GROUP BY		665	15960	9 (34)
* 6	HASH JOIN		665	15960	8 (25)
* 7	TABLE ACCESS FULL	ORDERS	105	840	4 (25)
8	TABLE ACCESS FULL	ORDER_ITEMS	665	10640	4 (25)

```

Predicate Information (identified by operation id):

- ```

-----
2 - access("O"."CUSTOMER_ID"(+)="C"."CUSTOMER_ID")
3 - filter("C"."CREDIT_LIMIT">2000)
6 - access("O"."ORDER_ID"="L"."ORDER_ID")
7 - filter("O"."CUSTOMER_ID">0)

```

The view definition is as follows:

```

CREATE OR REPLACE view v_orders AS
SELECT l.product_id, SUM(l.quantity*unit_price) revenue,
       o.order_id, o.customer_id
FROM orders o, order_items l
WHERE o.order_id = l.order_id
GROUP BY l.product_id, o.order_id, o.customer_id;

```

Sort Merge Outer Joins

When an outer join cannot drive from the outer (preserved) table to the inner (optional) table, it cannot use a hash join or nested loop joins. Then it uses the sort merge outer join for performing the join operation.

The optimizer uses sort merge for an outer join if a nested loop join is inefficient. A nested loop join can be inefficient because of data volumes; or if, because of sorts already required by other operations, the optimizer finds it is cheaper to use a sort merge over a hash join.

Full Outer Joins

A full outer join acts like a combination of the left and right outer joins. In addition to the inner join, rows from both tables that have not been returned in the result of the inner join are preserved and extended with nulls. In other words, full outer joins

let you join tables together, yet still show rows that do not have corresponding rows in the joined tables.

The query in [Example 1–20](#) retrieves all departments and all employees in each department, but also includes:

- Any employees without departments
- Any departments without employees

Example 1–20 Full Outer Join

```
SELECT d.department_id, e.employee_id
   FROM employees e
  FULL OUTER JOIN departments d
     ON e.department_id = d.department_id
  ORDER BY d.department_id;
```

The statement produces the following output:

DEPARTMENT_ID	EMPLOYEE_ID
10	200
20	201
20	202
30	114
30	115
30	116
...	
270	
280	
	178
	207

125 rows selected.

Setting Cost-Based Optimizer Parameters

This section contains some of the parameters specific to the optimizer. The following sections are especially useful when tuning Oracle applications.

Enabling CBO Features

You enable optimizer features by setting the `OPTIMIZER_FEATURES_ENABLE` initialization parameter.

OPTIMIZER_FEATURES_ENABLE Parameter

The `OPTIMIZER_FEATURES_ENABLE` parameter acts as an umbrella parameter for the CBO. This parameter can be used to enable a series of CBO-related features, depending on the release. It accepts one of a list of valid string values corresponding to the release numbers, such as 8.0.4, 8.1.5, and so on. For example, the following statement enables the use of the optimizer features in Oracle9i, Release 2 (9.2).

```
OPTIMIZER_FEATURES_ENABLE=9.2.0;
```

This statement causes the Release 2 (9.2) optimizer features to be used in generating query plans. For example, you can use `ALL_ROWS` or `FIRST_ROWS` optimizer mode for recursive user SQL generated by PL/SQL procedures. Prior to Oracle8i Release 8.1.6, only `RULE` or `CHOOSE` optimizer mode was used for such recursive SQL, and when the user explicitly set the `OPTIMIZER_MODE` parameter to `FIRST_ROWS` or `ALL_ROWS`, a `CHOOSE` mode was used instead.

The `OPTIMIZER_FEATURES_ENABLE` parameter was introduced in Oracle8, release 8.0.4. The main goal was to allow customers the ability to upgrade the Oracle server, yet preserve the old behavior of the CBO after the upgrade. For example, when you upgrade the Oracle server from release 8.1.5 to release 8.1.6, the default value of the `OPTIMIZER_FEATURES_ENABLE` parameter changes from 8.1.5 to 8.1.6. This upgrade results in the CBO enabling optimization features based on 8.1.6, as opposed to 8.1.5.

For plan stability or backward compatibility reasons, you might not want the query plans to change because of new optimizer features in a new release. In such a case, you can set the `OPTIMIZER_FEATURES_ENABLE` parameter to an earlier version. For example, to preserve the behavior of the CBO to release 8.1.5, set the parameter as follows:

```
OPTIMIZER_FEATURES_ENABLE=8.1.5;
```

This statement disables all new optimizer features that were added in releases following release 8.1.5.

Note: If you upgrade to a new release and you want to enable the features available with that release, then you do not need to explicitly set the `OPTIMIZER_FEATURES_ENABLE` parameter.

Oracle Corporation does not recommend explicitly setting the `OPTIMIZER_FEATURES_ENABLE` parameter to an earlier release. Instead, execution plan or query performance issues should be resolved on a case-by-case basis.

Table 1–3 describes some of the optimizer features that are enabled when you set the `OPTIMIZER_FEATURES_ENABLE` parameter to each of the following release values.

Table 1–3 Features Included with the `OPTIMIZER_FEATURES_ENABLE` Parameter

Set to Value	New Features include
8.0.4	Index fast full scan Ordered nested loop join method
8.0.5	No new features
8.0.6	Improved outer join cardinality estimation
8.1.4	No new features
8.1.5	Improved verification of NULLs inclusion in B-tree indexes
8.1.6	Use of <code>FIRST_ROWS</code> or <code>ALL_ROWS</code> mode for user recursive SQL Random distribution of left input of nested loop join Improved row-length calculation Improved method of computing selectivity based on histogram Partition pruning based on predicates in a subquery
8.1.7	Common subexpression optimization Statistics of a column imbedded in some selected functions such as <code>TO_CHAR</code> to compute selectivity Improved partition statistics aggregation
9.0.1	Peeking of user-defined bind variables Complex view merging Push-join predicate Consideration of bitmap access paths for tables with only B-tree indexes Subquery unnesting Index joins
9.0.2	Parallel broadcast in parallel query using joins to small reference or lookup tables

Peeking of User-Defined Bind Variables

The CBO peeks at the values of user-defined bind variables on the first invocation of a cursor. This feature lets the optimizer determine the selectivity of any `WHERE` clause condition, as well as if literals have been used instead of bind variables. On subsequent invocations of the cursor, no peeking takes place, and the cursor is shared, based on the standard cursor-sharing criteria, even if subsequent invocations use different bind values.

When bind variables are used in a statement, it is assumed that cursor sharing is intended and that different invocations are supposed to use the same execution plan. If different invocations of the cursor would significantly benefit from different execution plans, then bind variables may have been used inappropriately in the SQL statement.

Controlling the Behavior of the CBO

This section lists some initialization parameters that can be used to control the behavior of the cost-based optimizer. These parameters can be used to enable various optimizer features in order to improve the performance of SQL execution.

CURSOR_SHARING

This parameter converts literal values in SQL statements to bind variables. Converting the values improves cursor sharing and can affect the execution plans of SQL statements. The optimizer generates the execution plan based on the presence of the bind variables and not the actual literal values.

DB_FILE_MULTIBLOCK_READ_COUNT

This parameter specifies the number of blocks that are read in a single I/O during a full table scan or index fast full scan. The optimizer uses the value of `DB_FILE_MULTIBLOCK_READ_COUNT` to cost full table scans and index fast full scans. Larger values result in a cheaper cost for full table scans and can result in the optimizer choosing a full table scan over an index scan.

HASH_AREA_SIZE

This parameter specifies the amount of memory (in bytes) to be used for hash joins. The CBO uses this parameter to cost a hash join operation. Larger values for `HASH_AREA_SIZE` reduce the cost of hash joins.

Note: Oracle does not recommend using the `HASH_AREA_SIZE` parameter unless the instance is configured with the shared server option. Oracle recommends instead that you enable automatic sizing of SQL work areas by setting `PGA_AGGREGATE_TARGET`. `HASH_AREA_SIZE` is retained only for backward compatibility.

HASH_JOIN_ENABLED

This parameter can be used to enable or disable the use of hash joins as a join method chosen by the optimizer. When set to `true`, the optimizer considers hash joins as a possible join method. The CBO chooses a hash join if the cost is better than other join methods, such as nested loops or sort merge joins.

OPTIMIZER_INDEX_CACHING

This parameter controls the costing of an index probe in conjunction with a nested loop. The range of values 0 to 100 for `OPTIMIZER_INDEX_CACHING` indicates percentage of index blocks in the buffer cache, which modifies the optimizer's assumptions about index caching for nested loops and IN-list iterators. A value of 100 infers that 100% of the index blocks are likely to be found in the buffer cache and the optimizer adjusts the cost of an index probe or nested loop accordingly. Use caution when using this parameter because execution plans can change in favor of index caching.

OPTIMIZER_INDEX_COST_ADJ

This parameter can be used to adjust the cost of index probes. The range of values is 1 to 10000. The default value is 100, which means that indexes are evaluated as an access path based on the normal costing model. A value of 10 means that the cost of an index access path is one-tenth the normal cost of an index access path.

OPTIMIZER_MAX_PERMUTATIONS

This parameter controls the maximum number of permutations that the CBO considers when generating execution plans for SQL statements with joins. The range of values is 4 to 80000. A value of 80000 corresponds to no limit. Setting this parameter to a value less than 1000 normally ensures parse times of a few seconds or less.

The `OPTIMIZER_MAX_PERMUTATIONS` parameter can be used to reduce parse times for complex SQL statements that join a large number of tables. However, reducing its value can result in the optimizer missing an optimal join permutation.

OPTIMIZER_MODE

This initialization parameter sets the mode of the optimizer at instance startup. The possible values are `RULE`, `CHOOSE`, `ALL_ROWS`, `FIRST_ROWS_n`, and `FIRST_ROWS`. For description of these parameter values, see "[OPTIMIZER_MODE Initialization Parameter](#)" on page 1-6.

Note: Oracle Corporation strongly advises the use of cost-based optimization. Rule-based optimization will be deprecated in a future release.

PARTITION_VIEW_ENABLED

This parameter enables the partition view pruning feature. If set to `true`, then the CBO scans only the required partitions, based on the view predicates or filters.

QUERY_REWRITE_ENABLED

This parameter enables the query rewrite feature, which works in conjunction with materialized views. If set to `true`, then the CBO considers query rewrites, using materialized views to satisfy the original query. This parameter also controls whether or not function-based indexes are used.

SORT_AREA_SIZE

This parameter specifies the amount of memory (in bytes) that will be used to perform sorts. If a sort operation is performed, and if the amount of data to be sorted exceeds the value of `SORT_AREA_SIZE`, then data beyond the value of `SORT_AREA_SIZE` is written to the temporary tablespace. The CBO uses the value of `SORT_AREA_SIZE` to cost sort operations including sort merge joins. Larger values for `SORT_AREA_SIZE` result in cheaper CBO costs for sort operations.

Note: Oracle does not recommend using the `SORT_AREA_SIZE` parameter unless the instance is configured with the shared server option. Oracle recommends instead that you enable automatic sizing of SQL work areas by setting `PGA_AGGREGATE_TARGET`. `SORT_AREA_SIZE` is retained only for backward compatibility.

STAR_TRANSFORMATION_ENABLED

This parameter, if set to `true`, enables the CBO to cost a star transformation for star queries. The star transformation combines the bitmap indexes on the various fact table columns rather than using a Cartesian approach.

See Also: *Oracle9i Database Reference* for complete information about each parameter

Overview of the Extensible Optimizer

The extensible optimizer is part of the CBO. It allows the authors of user-defined functions and domain indexes to control the three main components that the CBO uses to select an execution plan: statistics, selectivity, and cost evaluation.

The extensible optimizer lets you:

- Associate cost function and default costs with domain indexes, indextypes, packages, and standalone functions
- Associate selectivity function and default selectivity with methods of object types, package functions, and standalone functions
- Associate statistics collection functions with domain indexes and columns of tables
- Order predicates with functions based on cost
- Select a user-defined access path (domain index) for a table based on access cost
- Use the `DBMS_STATS` package or the `ANALYZE` statement to invoke user-defined statistics collection and deletion functions

Note: Oracle Corporation strongly recommends that you use the `DBMS_STATS` package rather than `ANALYZE` to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by `DBMS_STATS`. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the `ANALYZE` statement rather than `DBMS_STATS` for statistics collection not related to the cost-based optimizer, such as:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
 - To collect information on freelist blocks
-
-
- Use new data dictionary views to include information about the statistics collection, cost, or selectivity functions associated with columns, domain indexes, indextypes, or functions
 - Add a hint to preserve the order of evaluation for function predicates

See Also: *Oracle9i Data Cartridge Developer's Guide* for details about the extensible optimizer

Understanding User-Defined Statistics

You can define statistics collection functions for domain indexes, individual columns of a table, and user-defined datatypes.

Whenever a domain index is analyzed to gather statistics, Oracle calls the associated statistics collection function. Whenever a column of a table is analyzed, Oracle collects the standard statistics for that column and calls any associated statistics collection function. If a statistics collection function exists for a datatype, then Oracle calls it for each column that has that datatype in the table being analyzed.

Understanding User-Defined Selectivity

The selectivity of a predicate in a SQL statement is used to estimate the cost of a particular access path; it is also used to determine the optimal join order. The optimizer cannot compute an accurate selectivity for predicates that contain

user-defined operators, because it does not have any information about these operators.

You can define selectivity functions for predicates containing user-defined operators, standalone functions, package functions, or type methods. The optimizer calls the user-defined selectivity function whenever it encounters a predicate that contains the operator, function, or method in one of the following relations with a constant: `<`, `<=`, `=`, `>=`, `>`, or `LIKE`.

Understanding User-Defined Costs

The optimizer cannot compute an accurate estimate of the cost of a domain index because it does not know the internal storage structure of the index. Also, the optimizer might underestimate the cost of a user-defined function that invokes PL/SQL, uses recursive SQL, accesses a `BFILE`, or is CPU-intensive.

You can define costs for domain indexes and user-defined standalone functions, package functions, and type methods. These user-defined costs can be in the form of default costs that the optimizer simply looks up, or they can be full-fledged cost functions that the optimizer calls to compute the cost.

Optimizer Operations

This chapter expands on the ideas introduced in [Chapter 1, "Introduction to the Optimizer"](#) and explains optimizer actions in greater detail for specific cases. This chapter describes how the cost-based optimizer evaluates expressions and performs specific operations. It also explains how the CBO transforms some SQL statements into others to achieve the same goal more efficiently.

The chapter contains the following sections:

- [How the Optimizer Performs Operations](#)
- [How the Optimizer Transforms SQL Statements](#)

How the Optimizer Performs Operations

The optimizer fully evaluates expressions whenever possible and translates certain syntactic constructs into equivalent constructs. Constructs are translated in the following cases:

- Oracle can more quickly evaluate the resulting expression than the original expression.
- The original expression is a syntactic equivalent of the resulting expression. Sometimes, different SQL constructs can operate identically. For example, Oracle maps `= ANY (subquery)` and `IN (subquery)` to a single construct.

This section contains the following discussions:

- [How the CBO Evaluates IN-List Iterators](#)
- [How the CBO Evaluates Concatenation](#)
- [How the CBO Evaluates Remote Operations](#)
- [How the CBO Executes Distributed Statements](#)
- [How the CBO Executes Sort Operations](#)
- [How the CBO Executes Views](#)
- [How the CBO Evaluates Constants](#)
- [How the CBO Evaluates the UNION and UNION ALL Operators](#)
- [How the CBO Evaluates the LIKE Operator](#)
- [How the CBO Evaluates the IN Operator](#)
- [How the CBO Evaluates the ANY or SOME Operator](#)
- [How the CBO Evaluates the ALL Operator](#)
- [How the CBO Evaluates the BETWEEN Operator](#)
- [How the CBO Evaluates the NOT Operator](#)
- [How the CBO Evaluates Transitivity](#)
- [How the CBO Optimizes Common Subexpressions](#)
- [How the CBO Evaluates DETERMINISTIC Functions](#)

How the CBO Evaluates IN-List Iterators

The IN-list iterator is used when a query contains an IN clause with values. The execution plan is identical to what would result for a statement with an equality clause instead of IN except for one additional step. That extra step occurs when the IN-list iterator feeds the equality clause with unique values from the IN-list.

Both of the statements in [Example 2-1](#) and [Example 2-1](#) are equivalent and produce the same plan.

Example 2-1 IN-List Iterators Initial Statement

```
SELECT header_id, line_id, revenue_amount
   FROM so_lines_all
  WHERE header_id IN (1011,1012,1013);
```

```
SELECT header_id, line_id, revenue_amount
   FROM so_lines_all
  WHERE header_id = 1011
     OR header_id = 1012
     OR header_id = 1013;
```

Plan

```
-----
SELECT STATEMENT
  INLIST ITERATOR
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
```

An equivalent statement in [Example 2-1](#) binds :b1 to the different unique values by the IN-list iterator.

Example 2-2 Alternate Statement with Equality

```
SELECT header_id, line_id, revenue_amount
   FROM so_lines_all l
  WHERE header_id = :b1;
```

Plan

```
-----
SELECT STATEMENT
  TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N1
```

Example 2-3 uses a unique index. Because there is a sort involved on the IN-list, even with complete keys of unique indexes, there is still a range scan.

Example 2-3 IN-List Iterator with Unique Index

```
SELECT header_id, line_id, revenue_amount
   FROM so_lines_all
  WHERE line_id IN (1011,1012,1013);
```

Plan

```
-----
SELECT STATEMENT
  INLIST ITERATOR
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_U1
```

In **Example 2-4**, the IN-list operator can be used when driving into a table with a nested loop operation.

Example 2-4 IN-List Iterator with a Nested Loop

```
SELECT h.header_id, l.line_id, l.revenue_amount
   FROM so_headers_all h, so_lines_all l
  WHERE l.inventory_item_id = :b1
        AND h.order_number = l.header_id
        AND h.order_type_id IN (1,2,3);
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N5
    INLIST ITERATOR
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_U2
```

The IN-list operator is especially useful if there is an expensive first step that you do not want to repeat for every IN-list element. In **Example 2-5**, even though there are three IN-list elements, the full scan on so_lines_all happens only once.

Example 2-5 Using the IN-List Iterator to Avoid an Expensive First Step

```

SELECT h.header_id, l.line_id, l.revenue_amount
  FROM so_headers_all h, so_lines_all l
 WHERE l.s7 = :b1
        AND h.order_number = l.header_id
        AND h.order_type_id IN (1,2,3);

```

Plan

```

-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS FULL SO_LINES_ALL
    INLIST ITERATOR
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_U2

```

When the Optimizer Uses IN-List Iterators

The optimizer uses an IN-list iterator when an IN clause is specified with values, and the optimizer finds a selective index for that column. If there are multiple OR clauses using the same index, then the optimizer chooses this operation rather than CONCATENATION or UNION ALL, because it is more efficient.

IN-List Iterator Hints

There are no hints for this operation. You can provide a hint to use the relevant index, which can cause this operation. [Example 2-6](#) shows a query without an INDEX hint, and the resulting execution plan.

Example 2-6 Calling the IN-List Iterator without the INDEX Hint

```

SELECT h.customer_id, l.line_id, l.revenue_amount
  FROM so_lines_all l, so_headers_all h
 WHERE l.s7 = 20
        AND h.original_system_reference = l.attribute5
        AND h.original_system_source_code IN (1013,1014);

```

Plan

```

-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS FULL SO_LINES_ALL
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
      INDEX RANGE SCAN SO_HEADERS_N5

```

[Example 2-7](#) shows a query with an INDEX hint, and the resulting execution plan.

Example 2-7 Calling the IN-List Iterator with the INDEX Hint

```
SELECT /*+INDEX(h so_headers_n9 */ h.customer_id, l.line_id, l.revenue_amount
      FROM so_lines_all l, so_headers_all h
      WHERE l.s7 = 20
             AND h.original_system_reference = l.attribute5
             AND h.original_system_source_code IN (1013,1014);
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS FULL SO_LINES_ALL
      INLIST ITERATOR
        TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
          INDEX RANGE SCAN SO_HEADERS_N9
```

How the CBO Evaluates Concatenation

Concatenation is useful for statements with different conditions combined with an OR clause. With concatenation, you get a good execution plan with appropriate indexes. [Examples 2-8](#) and [2-9](#) show two plans, each accessing the table through the appropriate index and combined using concatenation.

The plan in [Example 2-8](#) does not return duplicate rows, so for each component it appends a negation of the previous components.

Example 2-8 How the CBO Evaluates Concatenation

```
SELECT l.header_id, l.line_id, l.revenue_amount
      FROM so_lines_all l
      WHERE l.parent_line_id = :b1
             OR l.service_parent_line_id = :b1;
```

Plan

```
-----
SELECT STATEMENT
  CONCATENATION
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N20
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N17
```

In [Example 2-8](#), the components are the following:

- `l.parent_line_id = :b1`
- `l.service_parent_line_id = :b1` and `l.parent_line_id != :b1`

[Example 2-9](#) shows how the optimizer uses concatenation to optimize a nested OR statement.

Example 2-9 Concatenation of a Query with a Nested OR Statement

```
SELECT p.header_id, l.line_id, l.revenue_amount
   FROM so_lines_all p, so_lines_all l
  WHERE p.header_id = :b1
        AND (l.parent_line_id = p.line_id
            OR l.service_parent_line_id = p.line_id);
```

Plan

```
-----
SELECT STATEMENT
CONCATENATION
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N1
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N20
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N1
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N17
```

[Example 2-10](#) shows the resulting execution plan for the same query shown in [Example 2-9](#), but with concatenation specifically disabled by the `NO_EXPAND` hint.

Example 2-10 Query with Nested OR with Concatenation Disabled

```
SELECT /*+NO_EXPAND */ p.header_id, l.line_id, l.revenue_amount
   FROM so_lines_all p, so_lines_all l
  WHERE p.header_id = :b1
        AND (l.parent_line_id = p.line_id
            OR l.service_parent_line_id = p.line_id);
```

```
Plan
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
    TABLE ACCESS FULL SO_LINES_ALL
```

Trying to execute the statement in a single query produces a poor execution plan. Because the optimizer has two paths to follow and has been instructed not to decompose the query, it needs to access all the rows in the second table to see if any rows match one of the conditions.

Concatenation Hints

Use the hint `USE_CONCAT` for this operation.

When Not to Use Concatenation

Concatenation is expensive in the following cases and should not be used:

- When OR conditions are on same column and can use the IN-list operator, which is more efficient than concatenation
- If an expensive step gets repeated for every concatenation

[Example 2-11](#) illustrates this point.

Example 2-11 When Not to Use Concatenation

Consider the following statement:

```
SELECT h.customer_id, l.line_id, l.revenue_amount
  FROM so_lines_all l, so_headers_all h
 WHERE l.s7 = 20
        AND h.original_system_reference = l.attribute5
        AND h.original_system_source_code IN (1013,1014);
```

```
Plan
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS FULL SO_LINES_ALL
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
      INDEX RANGE SCAN SO_HEADERS_N5
```


The initial statement calls for a full scan on `so_lines_all` as the first stage. The optimizer chooses to use a single column index for the second table, but we want it to use a two-column index.

You could use a hint to force concatenation as in shown in [Example 2-12](#), but the initial full scan is still repeated which is not desirable.

Example 2-12 Using the Concatenation Hint

```
SELECT /*+USE_CONCAT*/ h.customer_id, l.line_id, l.revenue_amount
  FROM so_lines_all l, so_headers_all h
 WHERE l.s7 = 20
       AND h.original_system_reference = l.attribute5
       AND h.original_system_source_code IN (1013,1014);
```

Plan

```
-----
SELECT STATEMENT
  CONCATENATION
    NESTED LOOPS
      TABLE ACCESS FULL SO_LINES_ALL
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_N9
    NESTED LOOPS
      TABLE ACCESS FULL SO_LINES_ALL
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_N9
```

If, instead, you provide a hint to use the two-column index, then the optimizer switches to that with an `IN-list` iterator. The initial scan is not repeated, and a better execution plan results, as follows:

Example 2-13 Using the Index Hint

```
SELECT /*+INDEX(h so_headers_n9 */ h.customer_id, l.line_id, l.revenue_amount
  FROM so_lines_all l, so_headers_all h
 WHERE l.s7 = 20
       AND h.original_system_reference = l.attribute5
       AND h.original_system_source_code IN (1013,1014);
```

```
Plan
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS FULL SO_LINES_ALL
    INLIST ITERATOR
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_N9
```

How the CBO Evaluates Remote Operations

The remote operation indicates that there is a table from another database being accessed through a database link. [Example 2-14](#) has a remote driving table:

Example 2-14 How the CBO Evaluates a Query with a Remote Driving Table

```
SELECT c.customer_name, count(*)
  FROM ra_customers c, so_headers_all@oe h
 WHERE c.customer_id = h.customer_id
        AND h.order_number = :b1
 GROUP BY c.customer_name;
```

```
Plan
-----
SELECT STATEMENT
  SORT GROUP BY
    NESTED LOOPS
      REMOTE
        TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
          INDEX UNIQUE SCAN RA_CUSTOMERS_U1
```

Remote Database Query Obtained from the Library Cache

```
SELECT "ORDER_NUMBER", "CUSTOMER_ID"
  FROM "SO_HEADERS_ALL" "H"
 WHERE "ORDER_NUMBER"=: "SYS_B_0";
```

[Example 2-15](#) has a local driving table.

Example 2-15 How the CBO Evaluates a Query with a Local Driving Table

```
SELECT c.customer_name, h.order_number
  FROM ra_customers c, so_headers_all@oe h
 WHERE c.customer_id = h.customer_id
        AND c.customer_name LIKE :b1;
```

Plan

```
-----
SELECT STATEMENT
  MERGE JOIN
    REMOTE
      SORT JOIN
        TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
          INDEX RANGE SCAN RA_CUSTOMERS_N1
```

Remote Database Query Obtained from the Library Cache

```
SELECT "ORDER_NUMBER", "CUSTOMER_ID"
  FROM "SO_HEADERS_ALL" "H"
 WHERE "CUSTOMER_ID" IS NOT NULL
 ORDER BY "CUSTOMER_ID";
```

A couple of factors influence the execution plan:

- Network round trips can be several orders of magnitude more expensive than the physical and logical I/Os.
- The optimizer does not have any statistics on the remote database, which might not be an Oracle database.

In general, the optimizer chooses to access the remote tables first, before accessing the local tables. This works well for cases like [Example 2-14](#), where the driving table is the remote table. However, if the driving table is the local table, then there might not be any selective way of accessing the remote table without first accessing the local tables. In such cases, you might need to provide appropriate hints to avoid performance problems.

Example 2-16 *How the CBO Evaluates a Query with a Local Driving Table and a Nested Loops Hint*

```
SELECT /*+USE_NL(c h) */ c.customer_name, h.order_number
  FROM ra_customers c, so_headers_all@oe h
 WHERE c.customer_id = h.customer_id
        AND c.customer_name LIKE :b1;
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
      INDEX RANGE SCAN RA_CUSTOMERS_N1
    REMOTE
```

Remote Database Query Obtained from the Library Cache

```
SELECT /*+ USE_NL("H") */ "ORDER_NUMBER", "CUSTOMER_ID"  
  FROM "SO_HEADERS_ALL" "H" WHERE :1="CUSTOMER_ID"  
FILTER;
```

The construct used by the optimizer is applying a filter condition to filter out rows; this filter could not be applied when the table was accessed. [Example 2-17](#) uses no filter.

Example 2-17 Querying a Remote Table Using No Filter

```
SELECT h.order_number  
  FROM so_headers_all h  
 WHERE h.open_flag = 'Y'  
       AND attribute1 IS NOT NULL;
```

Plan

```
-----  
SELECT STATEMENT  
  TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL  
    INDEX RANGE SCAN SO_HEADERS_N2
```

Besides the conditions used in the access path, a table might have additional conditions to filter rows when the table is visited. Conditions that get applied when the table is accessed, like `attribute1 IS NOT NULL`, do not show up as `FILTER`. [Example 2-18](#) and [Example 2-19](#) illustrate such conditions.

[Example 2-18](#) shows a query with a `GROUP BY` condition that creates a filter.

Example 2-18 Querying with a Filter Created by a GROUP BY Condition

```
SELECT h.order_number, count(*)  
  FROM so_headers_all h  
 WHERE h.open_flag = 'Y'  
       AND attribute1 IS NOT NULL  
 GROUP BY h.order_number  
HAVING COUNT(*) = 1 § Filter condition;
```

Plan

```
-----  
SELECT STATEMENT  
  FILTER  
    SORT GROUP BY  
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL  
        INDEX RANGE SCAN SO_HEADERS_N2
```

Example 2–19 shows a query with a subquery that creates a filter.

Example 2–19 Querying with a Filter Created by a Subquery

```
SELECT h.order_number
   FROM so_headers_all h
  WHERE h.open_flag = 'Y'
     AND EXISTS (SELECT null FROM so_lines_all l
                 WHERE l.header_id = h.header_id
                 AND l.revenue_amount > 10000);
```

Plan

```
-----
SELECT STATEMENT
  FILTER
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
      INDEX RANGE SCAN SO_HEADERS_N2
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
```

In this example, for every row meeting the condition of the outer query, the correlated EXISTS subquery is executed. If a row meeting the condition is found in the so_lines_all table, then the row from so_headers_all is returned.

How the CBO Executes Distributed Statements

The optimizer chooses execution plans for SQL statements that access data on remote databases in much the same way that it chooses execution plans for statements that access only local data:

- If all the tables accessed by a SQL statement are located together on the same remote database, then Oracle sends the SQL statement to that remote database. The remote Oracle instance executes the statement and sends only the results back to the local database.
- If a SQL statement accesses tables that are located on different databases, then Oracle decomposes the statement into individual fragments, each of which accesses tables on a single database. Oracle then sends each fragment to the database that it accesses. The remote Oracle instance for each of these databases executes its fragment and returns the results to the local database, where the local Oracle instance can perform any additional processing that the statement requires.

When choosing a cost-based execution plan for a distributed statement, the optimizer considers the available indexes on remote databases, just as it considers indexes on the local database. The optimizer also considers statistics on remote databases for the CBO. Furthermore, the optimizer considers the location of data when estimating the cost of accessing it. For example, a full scan of a remote table has a greater estimated cost than a full scan of an identical local table.

For a rule-based execution plan, the optimizer does not consider indexes on remote tables.

See Also: [Chapter 6, "Optimizing SQL Statements"](#) for more information on tuning distributed queries

How the CBO Executes Sort Operations

Sort operations result when users specify some operation that requires a sort. Commonly encountered operations include the following:

- [SORT UNIQUE](#)
- [SORT AGGREGATE](#)
- [SORT GROUP BY](#)
- [SORT JOIN](#)
- [SORT ORDER BY](#)

SORT UNIQUE

`SORT UNIQUE` occurs if a user specifies a `DISTINCT` clause ([Example 2-20](#)) or if an operation requires unique values for the next step ([Example 2-21](#)).

Example 2-20 DISTINCT Clause Causing SORT UNIQUE

```
SELECT DISTINCT last_name, first_name
FROM per_all_people_f
WHERE full_name LIKE :b1;
```

Plan

```
-----
SELECT STATEMENT
  SORT UNIQUE
    TABLE ACCESS BY INDEX ROWID PER_ALL_PEOPLE_F
      INDEX RANGE SCAN PER_PEOPLE_F_N54
```

Example 2-21 IN Subquery Causing SORT UNIQUE

`SORT UNIQUE` provides the outer query with a unique list of `header_ids`. The plan shows that the `IN` subquery has been un-nested and transformed into `VW_NSO_1`.

```
SELECT c.customer_name, h.order_number
       FROM ra_customers c, so_headers_all h
       WHERE c.customer_id = h.customer_id
              AND h.header_id in
              (SELECT l.header_id FROM so_lines_all l
               WHERE l.inventory_item_id = :b1
                  AND ordered_quantity > 10);
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    NESTED LOOPS
      VIEW VW_NSO_1
        SORT UNIQUE
          TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
            INDEX RANGE SCAN SO_LINES_N5
          TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
            INDEX UNIQUE SCAN SO_HEADERS_U1
        TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
          INDEX UNIQUE SCAN RA_CUSTOMERS_U1
```

If the optimizer can guarantee (with unique keys) that duplicate values will not be passed, then a sort can be avoided, as in [Example 2-22](#).

Example 2-22 IN Subquery That Does Not Cause SORT UNIQUE

```
UPDATE so_lines_all l
       SET line_status = 'HOLD'
       WHERE l.header_id IN
              ( SELECT h.header_id FROM so_headers_all h
                WHERE h.customer_id = :b1);
```

Plan

```
-----
UPDATE STATEMENT
  UPDATE SO_LINES_ALL
    NESTED LOOPS
      TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
        INDEX RANGE SCAN SO_HEADERS_N1
      INDEX RANGE SCAN SO_LINES_N1
```

SORT AGGREGATE

SORT AGGREGATE does not actually involve a sort. It is used when aggregates are being computed across the whole set of rows, as shown in [Example 2-23](#).

Example 2-23 Query Causing SORT AGGREGATE

```
SELECT SUM(l.revenue_amount)
  FROM so_lines_all l
 WHERE l.header_id = :b1;
```

Plan

```
-----
SELECT STATEMENT
  SORT AGGREGATE
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
```

SORT GROUP BY

SORT GROUP BY is used when aggregates are being computed for different groups in the data. The sort is required to separate the rows into different groups, as shown in [Example 2-24](#).

Example 2-24 Query That Causes SORT GROUP BY

```
SELECT created_by, SUM(l.revenue_amount)
  FROM so_lines_all l
 WHERE header_id > :b1
 GROUP BY created_by;
```

Plan

```
-----
SELECT STATEMENT
  SORT GROUP BY
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
```


SORT JOIN

SORT JOIN happens during a SORT MERGE JOIN, if the rows need to be sorted by the join key, as shown in [Example 2–25](#).

Example 2–25 Query Causing SORT JOIN

```
SELECT SUM(l.revenue_amount), l2.creation_date
   FROM so_lines_all l, so_lines_All l2
  WHERE l.creation_date < l2.creation_date
        AND l.header_id <> l2.header_id
 GROUP BY l2.creation_date, l2.line_id;
```

Plan

```
-----
SELECT STATEMENT
  SORT GROUP BY
    MERGE JOIN
      SORT JOIN
        TABLE ACCESS FULL SO_LINES_ALL
      FILTER
        SORT JOIN
          TABLE ACCESS FULL SO_LINES_ALL
```

SORT ORDER BY

SORT ORDER BY is required when the statement specifies an ORDER BY that cannot be satisfied by one of the indexes, as shown in [Example 2–26](#).

Example 2–26 Query Causing SORT ORDER BY

```
SELECT h.order_number
   FROM so_headers_all h
  WHERE h.customer_id = :b1
 ORDER BY h.creation_date DESC;
```

Plan

```
-----
SELECT STATEMENT
  SORT ORDER BY
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
    INDEX RANGE SCAN SO_HEADERS_N1
```

How the CBO Executes Views

Either of the following can result in creation of a view by the CBO:

- A complex view has not been decomposed ([Example 2-27](#)).
- A temporary or inline view is being used ([Examples 2-28](#) and [2-29](#)).

Example 2-27 Query Causing a View

```
SELECT order_id
FROM orders
  WHERE customer_id = :b1
     AND revenue > :b2;
```

Plan

```
-----
SELECT STATEMENT
  VIEW ORDERS
  FILTER
    SORT GROUP BY
      NESTED LOOPS
        TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
          INDEX RANGE SCAN SO_HEADERS_N1
        TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
          INDEX RANGE SCAN SO_LINES_N1
```

In [Example 2-28](#), a view is created because the `IN` subquery requires a `SORT UNIQUE` on the values being selected. This view would be unnecessary if the columns being selected were unique, not requiring a sort.

Example 2-28 IN Subquery Causing a View

```
SELECT c.customer_name, h.order_number
FROM ra_customers c, so_headers_all h
  WHERE c.customer_id = h.customer_id
     AND h.header_id IN
        (SELECT l.header_id FROM so_lines_all l
          WHERE l.inventory_item_id = :b1
             AND ordered_quantity > 10);
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    NESTED LOOPS
      VIEW VW_NSO_1
```

```

SORT UNIQUE
  TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
    INDEX RANGE SCAN SO_LINES_N5
  TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
    INDEX UNIQUE SCAN SO_HEADERS_U1
  TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
    INDEX UNIQUE SCAN RA_CUSTOMERS_U1

```

The query in [Example 2-29](#) examines the distribution of orders and revenue by the number of lines in each order. The CBO uses temporary inline views to do the double grouping.

Example 2-29 Inline Query Causing a View

```

SELECT COUNT(*) "Orders", cnt "Lines", sum(rev) "Revenue"
FROM (SELECT header_id, COUNT(*) cnt, SUM(revenue_amount) rev
      FROM so_lines_all
      GROUP BY header_id)
GROUP BY cnt;

```

Plan

```

-----
SELECT STATEMENT
  SORT GROUP BY
    VIEW
      SORT GROUP BY
        TABLE ACCESS FULL SO_LINES_ALL

```

How the CBO Evaluates Constants

Computation of constants is performed only once, when the statement is optimized, rather than each time the statement is executed.

For example, the following conditions all test for monthly salaries greater than 2000:

```

salary > 24000/12
salary > 2000
salary*12 > 24000

```

If a SQL statement contains the first condition, then the optimizer simplifies it into the second condition.

Note: The optimizer does not simplify expressions across comparison operators: In the preceding examples, the optimizer does not simplify the third expression into the second. For this reason, application developers write conditions that compare columns with constants whenever possible, rather than conditions with expressions involving columns.

How the CBO Evaluates the UNION and UNION ALL Operators

This operator is useful for combining OR clauses into one compound statement or for breaking up a complex statement into a compound statement containing simpler select statements that are easier to optimize and understand.

As with concatenation, you do not want to duplicate expensive operations by using UNION ALL.

When the Optimizer Uses UNION or UNION ALL

The optimizer uses UNION or UNION ALL when the SQL statement contains UNION or UNION ALL clauses.

Example 2-30 shows a query without a UNION clause. The query finds customers who are new or have open orders.

Example 2-30 Query Without UNION Clause

```
SELECT c.customer_name, c.creation_date
FROM ra_customers c
WHERE c.creation_date > SYSDATE - 30
OR customer_id IN
  (SELECT customer_id FROM so_headers_all h
   WHERE h.open_flag = 'Y');
```

Plan

```
-----
SELECT STATEMENT
  FILTER
    TABLE ACCESS FULL RA_CUSTOMERS
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
    INDEX RANGE SCAN SO_HEADERS_N1
```

Because the driving conditions come from different tables, you cannot execute the query effectively in a single statement.

With a `UNION` clause, you can break the query into two statements:

- New customers
- Customers with open orders

These two statements can be optimized easily. Because you do not want duplicates (some customers meet both criteria), use `UNION`, which eliminates duplicates by using a sort. If you use two disjoint sets, then you can use `UNION ALL`, eliminating the sort. The query from [Example 2-30](#) is shown using `UNION` in [Example 2-31](#).

Example 2-31 Query Using the UNION Clause

```
SELECT c.customer_name, c.creation_date
   FROM ra_customers c
  WHERE c.creation_date > SYSDATE - 30
UNION ALL
SELECT c.customer_name, c.creation_date
   FROM ra_customers c
  WHERE customer_id IN
      (SELECT customer_id FROM so_headers_all h
       WHERE h.open_flag = 'Y');
```

Plan

```
-----
SELECT STATEMENT
  SORT UNIQUE
    UNION-ALL
      TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
        INDEX RANGE SCAN RA_CUSTOMERS_N2
      NESTED LOOPS
        TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
          INDEX RANGE SCAN SO_HEADERS_N2
        TABLE ACCESS BY INDEX ROWID RA_CUSTOMERS
          INDEX UNIQUE SCAN RA_CUSTOMERS_U1
```

UNION and UNION ALL Hints

There are no hints for this operation.

How the CBO Evaluates the LIKE Operator

The optimizer simplifies conditions that use the `LIKE` comparison operator to compare an expression with no wildcard characters into an equivalent condition that uses an equality operator instead.

In the following example, the optimizer simplifies the first condition into the second:

```
last_name LIKE 'SMITH'
```

is transformed into

```
last_name = 'SMITH'
```

The optimizer can simplify these expressions only when the comparison involves variable-length datatypes. For example, if `last_name` was of type `CHAR(10)`, then the optimizer cannot transform the `LIKE` operation into an equality operation due to the equality operator following blank-padded semantics and `LIKE` not following blank-padded semantics.

How the CBO Evaluates the IN Operator

The optimizer expands a condition that uses the `IN` comparison operator to an equivalent condition that uses equality comparison operators and `OR` logical operators.

In the following example, the optimizer expands the first condition into the second:

```
last_name IN ('SMITH', 'KING', 'JONES')
```

is transformed into

```
last_name = 'SMITH' OR last_name = 'KING' OR last_name = 'JONES'
```

See Also: ["How the CBO Merges an IN Subquery"](#) on page 2-38

How the CBO Evaluates the ANY or SOME Operator

The optimizer expands a condition that uses the `ANY` or `SOME` comparison operator followed by a parenthesized list of values into an equivalent condition that uses equality comparison operators and `OR` logical operators.

In the following example, the optimizer expands the first condition into the second:

```
salary > ANY (:first_sal, :second_sal)
```

is transformed into

```
salary > :first_sal OR salary > :second_sal
```

The optimizer transforms a condition that uses the `ANY` or `SOME` operator followed by a subquery into a condition containing the `EXISTS` operator and a correlated subquery.

In the following example, the optimizer transforms the first condition into the second:

```
x > ANY (SELECT salary
        FROM employees
        WHERE job_id = 'IT_PROG')
```

is transformed into

```
EXISTS (SELECT salary
        FROM employees
        WHERE job_id = 'IT_PROG'
        AND x > salary)
```

How the CBO Evaluates the ALL Operator

The optimizer expands a condition that uses the `ALL` comparison operator followed by a parenthesized list of values into an equivalent condition that uses equality comparison operators and `AND` logical operators.

In the following example, the optimizer expands the first condition into the second:

```
salary > ALL (:first_sal, :second_sal)
```

is transformed into

```
salary > :first_sal AND salary > :second_sal
```

The optimizer transforms a condition that uses the `ALL` comparison operator followed by a subquery into an equivalent condition that uses the `ANY` comparison operator and a complementary comparison operator. In the following example, the optimizer transforms the first condition into the second:

```
x > ALL (SELECT salary
        FROM employees
        WHERE department_id = 50)
```

is transformed into

```
NOT (x <= ANY (SELECT salary
              FROM employees
              WHERE department_id = 50) )
```

The optimizer then further transforms the second query into the following query using the rule for transforming conditions with the ANY comparison operator, followed by a correlated subquery:

```
NOT EXISTS (SELECT salary
            FROM employees
            WHERE department_id = 50
            AND x <= salary)
```

How the CBO Evaluates the BETWEEN Operator

The optimizer always replaces a condition that uses the BETWEEN comparison operator with an equivalent condition that uses the >= and <= comparison operators. In the following example, the optimizer replaces the first condition with the second:

```
salary BETWEEN 2000 AND 3000
```

is transformed into

```
salary >= 2000 AND salary <= 3000
```

How the CBO Evaluates the NOT Operator

The optimizer simplifies a condition to eliminate the NOT logical operator. The simplification involves removing the NOT logical operator and replacing a comparison operator with its opposite comparison operator. In [Example 2-32](#), the optimizer simplifies the first condition into the second one:

Example 2-32 *How the CBO Evaluates the NOT Operator*

```
NOT department_id = (SELECT department_id FROM employees WHERE last_name = 'Taylor')
```

is transformed into

```
department_id <> (SELECT department_id FROM employees WHERE last_name = 'Taylor')
```

Often, a condition containing the NOT logical operator can be written in many different ways. The optimizer attempts to transform such a condition so that the subconditions negated by NOTs are as simple as possible, even if the resulting

condition contains more NOTs. In [Example 2-33](#), the optimizer simplifies the first condition into the second, and then into the third.

Example 2-33 How the CBO Simplifies a NOT Statement

```
NOT (salary < 1000 OR commission_pct IS NULL)
```

is transformed into

```
NOT salary < 1000 AND commission_pct IS NOT NULL
```

which is further transformed into

```
salary >= 1000 AND commission_pct IS NOT NULL
```

How the CBO Evaluates Transitivity

If two conditions in the WHERE clause involve a common column, then the optimizer sometimes can infer a third condition, using the transitivity principle. The optimizer can then use the inferred condition to optimize the statement. The inferred condition can make available an index access path that was not made available by the original conditions.

Note: Transitivity is used only by the CBO.

Consider a WHERE clause containing two conditions of these forms:

```
WHERE column1 comp_oper constant
      AND column1 = column2
```

In this case, the optimizer infers the condition:

```
column2 comp_oper constant
```

where:

- comp_oper is any of the comparison operators =, !=, ^=, <, <>, >, <=, or >=
- constant is any constant expression involving operators, SQL functions, literals, bind variables, and correlation variables

In [Example 2-34](#), the WHERE clause contains two conditions, each of which uses the employees.department_id column.

Example 2–34 Query Causing Transitivity Evaluation

```
SELECT *
  FROM employees, departments
 WHERE employees.department_id = 20
        AND employees.department_id = departments.department_id;
```

Using transitivity, the optimizer infers the following condition:

```
departments.department_id = 20
```

If an index exists on the `departments.department_id` column, then this condition makes available access paths using that index.

The optimizer only infers conditions that relate columns to constant expressions, rather than columns to other columns. Consider a `WHERE` clause containing two conditions of these forms:

```
WHERE column1 comp_oper column3
      AND column1 = column2
```

In this case, the optimizer does not infer the condition, `column2 comp_oper column3`.

How the CBO Optimizes Common Subexpressions

Common subexpression optimization is an optimization heuristic that identifies, removes, and collects common subexpressions from disjunctive (`OR`) branches of a query. In most cases, it reduces the number of joins to be performed. This optimization is enabled with the `OPTIMIZER_FEATURES_ENABLE` initialization parameter.

A query is considered valid for common subexpression optimization if its `WHERE` clause is in the following form:

- The top-level is a disjunction (a list of `OR`ed logs).
- Each disjunct is either a simple predicate or a conjunction (a list of `AND`ed logs).
- Each conjunct is either a simple predicate or a disjunction of simple predicates. (A predicate is considered simple if it does not contain `AND` or `OR`.)
- The expression appears in all the disjunctive branches of the query.

Examples of Common Subexpression Optimization

The query in [Example 2-35](#) finds names of employees who work in a department located in London and who make more than 60K or who are accountants. The query contains common subexpressions in its two disjunctive branches.

Example 2-35 Optimization of Common Subexpressions in Two OR Branches

Initial Query

```
SELECT employees.last_name
FROM employees E, departments D
WHERE (D.department_id = E.department_id AND E.job_id = 'AC_ACCOUNT' AND
D.location = 2400)
OR
E.department_id = D.department_id AND E.salary > 60000 AND D.location = 2400);
```

The elimination of the common subexpressions transforms this query into the following query, reducing the number of joins from two to one.

Optimized Query

```
SELECT employees.last_name
FROM employees E, departments D
WHERE (D.department_id = E.department_id AND D.location = 2400)
AND (E.job_id = 'AC_ACCOUNT' OR E.salary > 60000);
```

The query in [Example 2-36](#) contains common subexpression in three disjunctive branches.

Example 2-36 Optimization of Common Subexpressions in Three OR Branches

Initial Query

```
SELECT SUM (l_extendedprice* (1 - l_discount))
FROM PARTS, LINEITEM
WHERE (p_partkey = l_partkey
AND p_brand = 'Brand#12'
AND p_container IN ('SM CASE', 'SM BOX', 'SM PACK', 'SM PKG')
AND l_quantity >= 1 AND l_quantity <= 1 + 10
AND p_size >= 1 AND p_size <= 5
AND l_shipmode IN ('AIR', 'REG AIR')
AND l_shipinstruct = 'DELIVER IN PERSON')
OR (l_partkey = p_partkey)
AND p_brand = 'Brand#23'
AND p_container IN ('MED BAG', 'MED BOX', 'MED PKG', 'MED PACK')
AND l_quantity >= 10 AND l_quantity <= 10 + 10
```

```
AND p_size >= 1 AND p_size <= 10 AND p_size BETWEEN 1 AND 10
AND l_shipmode IN ('AIR', 'REG AIR')
AND l_shipinstruct = 'DELIVER IN PERSON')
OR (p_partkey = l_partkey
AND p_brand = 'Brand#34'
AND p_container IN ('LG CASE', 'LG BOX', 'LG PACK', 'LG PKG')
AND l_quantity >= 20 AND l_quantity <= 20 + 10
AND p_size >= 1 AND p_size <= 15
AND l_shipmode IN ('AIR', 'REG AIR')
AND l_shipinstruct = 'DELIVER IN PERSON');
```

This query is transformed by common subexpression optimization to the following query, reducing the number joins from three to one.

Optimized Query

```
SELECT SUM (l_extendedprice* (1 - l_discount))
FROM PARTS, LINEITEM
WHERE (p_partkey = l_partkey /* these are the four common subexpressions */
AND p_size >= 1
AND l_shipmode IN ('AIR', 'REG AIR')
AND l_shipinstruct = 'DELIVER IN PERSON')
AND
((p_brand = 'Brand#12'
AND p_container IN ( 'SM CASE', 'SM BOX', 'SM PACK', 'SM PKG')
AND l_quantity >= 1 AND l_quantity <= 1 + 10
AND p_size <= 5)
OR (p_brand = 'Brand#23'
AND p_container IN ('MED BAG', 'MED BOX', 'MED PKG', 'MED PACK')
AND l_quantity >= 10 AND l_quantity <= 10 + 10
AND p_size <= 10)
OR (p_brand = 'Brand#34'
AND p_container IN ('LG CASE', 'LG BOX', 'LG PACK', 'LG PKG')
AND l_quantity >= 20 AND l_quantity <= 20 + 10
AND p_size <= 15));
```

How the CBO Evaluates DETERMINISTIC Functions

In some cases, the optimizer can use a previously calculated value rather than executing a user-written function. This is only safe for functions that behave in a restricted manner. The function must return the same output return value for any given set of input argument values.

The function's result must not differ because of differences in the content of package variables or the database, or session parameters such as the globalization support

parameters. Furthermore, if the function is redefined in the future, then its output return value must be the same as that calculated with the prior definition for any given set of input argument values. Finally, there must be no meaningful side effects to using a precalculated value instead of executing the function again.

The creator of a function can promise to the Oracle server that the function behaves according to these restrictions by using the keyword `DETERMINISTIC` when declaring the function with a `CREATE FUNCTION` statement or in a `CREATE PACKAGE` or `CREATE TYPE` statement. The server does not attempt to verify this declaration—even a function that obviously manipulates the database or package variables can be declared `DETERMINISTIC`. It is the programmer's responsibility to use this keyword only when appropriate.

Calls to a `DETERMINISTIC` function might be replaced by the use of an already calculated value when the function is called multiple times within the same query, or if there is a function-based index or a materialized view defined that includes a relevant call to the function.

See Also:

- *Oracle9i Application Developer's Guide - Fundamentals* for more information on `DETERMINISTIC` functions
- *Oracle9i SQL Reference* for descriptions of `CREATE FUNCTION`, `CREATE INDEX`, and `CREATE MATERIALIZED VIEW`
- *Oracle9i Database Concepts* for a description of function-based indexes
- *Oracle9i Data Warehousing Guide* for detailed information about materialized views

How the Optimizer Transforms SQL Statements

SQL is a very flexible query language; there are often many statements you could use to achieve the same goal. Sometimes, the optimizer (Query Transformer) transforms one such statement into another that achieves the same goal if the second statement can be executed more efficiently.

This section discusses the following topics:

- [How the CBO Transforms ORs into Compound Queries](#)
- [How the CBO Unnests Subqueries](#)
- [How the CBO Merges Views](#)
- [How the CBO Pushes Predicates](#)
- [How the CBO Executes Compound Queries](#)

See Also: "[Understanding Joins](#)" on page 1-40 for additional information about optimizing statements that contain joins, semi-joins, or anti-joins

How the CBO Transforms ORs into Compound Queries

If a query contains a `WHERE` clause with multiple conditions combined with `OR` operators, then the optimizer transforms it into an equivalent compound query that uses the `UNION ALL` set operator, if this makes the query execute more efficiently:

- If each condition individually makes an index access path available, then the optimizer can make the transformation. The optimizer chooses an execution plan for the resulting statement that accesses the table multiple times using the different indexes and then puts the results together.
- If any condition requires a full table scan because it does not make an index available, then the optimizer does not transform the statement. The optimizer chooses a full table scan to execute the statement, and Oracle tests each row in the table to determine whether it satisfies any of the conditions.
- For statements that use the CBO, the optimizer might use statistics to determine whether to make the transformation, by estimating and then comparing execution costs of the original statement and the resulting statement.
- The CBO does not use the `OR` transformation for `IN`-lists or `OR`s on the same column; instead, it uses the `INLIST` iterator operator.

See Also: ["Understanding Access Paths for the RBO"](#) on page 8-3 and ["How the CBO Chooses an Access Path"](#) on page 1-37 for information on access paths and how indexes make them available

[Example 2-37](#) shows how an OR query is transformed into a compound query. In this example, the WHERE clause contains two conditions combined with an OR operator.

Example 2-37 Transforming an OR Query into a Compound Query

Initial Query

```
SELECT *
  FROM employees
 WHERE job_id = 'ST_CLERK'
        OR department_id = 50;
```

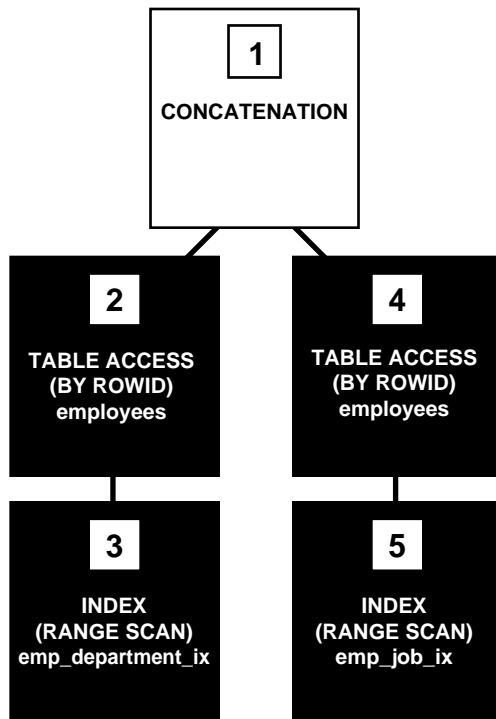
If there are indexes on both the `job_id` and `department_id` columns, then the optimizer might transform this query into the following equivalent query:

Optimized Query

```
SELECT *
  FROM employees
 WHERE job_id = 'ST_CLERK'
UNION ALL
SELECT *
  FROM employees
 WHERE department_id = 50
        AND job_id <> 'ST_CLERK';
```

When the CBO is deciding whether to make a transformation, the optimizer compares the cost of executing the original query using a full table scan with that of executing the resulting query.

The execution plan for the transformed statement might look like the illustration in [Figure 2-1](#). The shaded boxes indicate steps that physically retrieve data and the clear boxes indicate steps that operate on data returned from the previous step.

Figure 2–1 Execution Plan for a Transformed Query Containing OR

To execute the transformed query, Oracle performs the steps in [Figure 2–1](#) in the following order:

1. Steps 3 and 5 scan the indexes on the `job_id` and `department_id` columns using the conditions of the component queries. These steps obtain rowids of the rows that satisfy the component queries.
2. Steps 2 and 4 use the rowids from steps 3 and 5 to locate the rows that satisfy each component query.
3. Step 1 puts together the row sources returned by steps 2 and 4.

If either the `job_id` column or the `department_id` column is not indexed, then the optimizer does not even consider the transformation, because the resulting compound query would require a full table scan to execute one of its component queries. Executing the compound query with a full table scan in addition to an index scan could not possibly be faster than executing the original query with a full table scan.

[Example 2-38](#) shows the kind of query that is not transformed. The query in this example assumes that there is an index on the `last_name` column only:

Example 2-38 Query Containing OR That Is Not Transformed

```
SELECT *
  FROM employees
 WHERE last_name = 'Smith'
        OR salary > commission_pct*100000;
```

Transforming this query would result in the compound query in [Example 2-39](#).

Example 2-39 Compound Query Containing OR That Is Transformed

```
SELECT *
  FROM employees
 WHERE last_name = 'Smith'
UNION ALL
SELECT *
  FROM employees
 WHERE salary > commission_pct*100000;
```

Because the condition in the `WHERE` clause of the second component query (`salary > commission_pct`) does not make an index available, the compound query requires a full table scan. For this reason, the optimizer does not make the transformation, but chooses a full table scan to execute the original statement.

How the CBO Unnests Subqueries

To optimize a complex statement, the optimizer chooses to do one of the following:

- Transform the complex statement into an equivalent join statement, and then optimize the join statement
- Optimize the complex statement as it is

The optimizer transforms a complex statement into a join statement whenever the resulting join statement is guaranteed to return exactly the same rows as the complex statement. This transformation allows Oracle to execute the statement by taking advantage of join optimizer techniques.

[Example 2-40](#) shows how the optimizer uses a join to unnest a subquery. The complex statement in this example selects all rows from the `orders` table for which owners appear in the `customers` table.

Example 2–40 How the CBO Unnests Subqueries

Initial Query

```
SELECT *
  FROM orders
 WHERE customer_id IN
    (SELECT customer_id FROM customers);
```

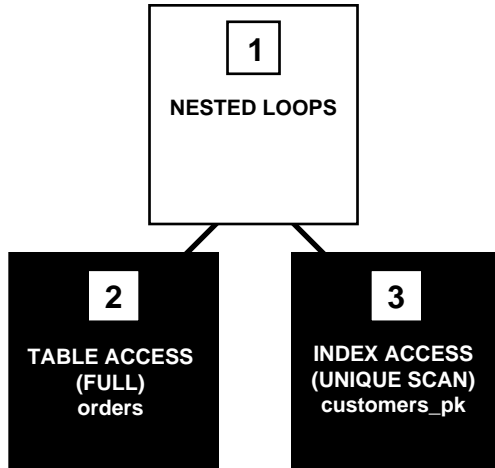
If the `customer_id` of the `customers` table is a primary key or has a `UNIQUE` constraint, then the optimizer can transform the complex query into the following join statement that is guaranteed to return the same data:

Optimized Query

```
SELECT orders.*
  FROM orders, customers
 WHERE orders.customer_id = customers.customer_id;
```

To execute this statement, Oracle performs a nested-loops join operation. The execution plan for this statement might look like [Figure 2–2](#).

Figure 2–2 Execution Plan for a Nested Loops Join



If the optimizer cannot transform a complex statement into a join statement, then it chooses execution plans for the parent statement and the subquery as though they were separate statements. Oracle then executes the subquery and uses the rows returned to execute the parent query.

[Example 2-41](#) shows the kind of statement is not transformed. The complex statement in this example returns all rows from the `customers` table that have credit limits greater than the average credit limit.

Example 2-41 Complex Statement That Is Not Transformed

```
SELECT *
  FROM customers
 WHERE credit_limit >
       (SELECT AVG(credit_limit) FROM customers);
```

No join statement can perform the function of this statement, so the optimizer does not transform the statement.

Note: Complex queries whose subqueries contain aggregate functions such as `AVG` cannot be transformed into join statements.

See Also: ["Understanding Joins"](#) on page 1-40 for information on nested loops joins

How the CBO Merges Views

To merge the view's query into a referencing query block in the accessing statement, the optimizer replaces the name of the view with the names of its base tables in the query block and adds the condition of the view's query's `WHERE` clause to the accessing query block's `WHERE` clause.

This optimization applies to **select-project-join** views, which are views that contain only selections, projections, and joins. This means that these views do not contain items such as set operators, aggregate functions, `DISTINCT`, `GROUP BY`, and `CONNECT BY`. See ["Mergeable and Nonmergeable Views"](#) on page 2-36.

[Example 2-42](#) shows how the CBO merges views. The view in this example is of all employees who work in department 10:

Example 2-42 How the CBO Merges Views

```
CREATE VIEW emp_10
  AS SELECT employee_id, last_name, job_id, manager_id, hire_date, salary,
 commission_pct, department_id
  FROM employees
  WHERE department_id = 10;
```

The following query accesses the view. The query selects the Ids greater than 170 of employees who work in department 10:

```
SELECT employee_id
FROM emp_10
WHERE employee_id > 170;
```

The optimizer transforms the query into the following query that accesses the view's base table:

```
SELECT employee_id
FROM employees
WHERE department_id = 10
AND employee_id > 170;
```

If there are indexes on the `department_id` or `employee_id` columns, then the resulting `WHERE` clause makes them available.

Mergeable and Nonmergeable Views The optimizer can merge a view into a referencing query block when the view has one or more base tables, provided the view does not contain any of the following:

- Set operators (UNION, UNION ALL, INTERSECT, MINUS)
- A `CONNECT BY` clause
- A `ROWNUM` pseudocolumn
- Aggregate functions (AVG, COUNT, MAX, MIN, SUM) in the select list

When a view contains one of the following structures, it can be merged into a referencing query block only if [Complex View Merging](#) is enabled:

- A `GROUP BY` clause
- A `DISTINCT` operator in the select list

View merging is not possible for a view that has multiple base tables, if it is on the right side of an outer join. However, if a view on the right side of an outer join has only one base table, then the optimizer can use complex view merging, even if an expression in the view can return a nonnull value for a `NULL`.

If a query has a `CURSOR` expression, then no view merging will take place, even for views which would normally be mergeable. An example is the following:

```
CREATE VIEW emp_v AS
SELECT last_name, employee_id FROM employees;
SELECT CURSOR(select * from sys.dual), last_name, employee_id from emp_v;
```

This query will not merge the emp_v view.

See Also: ["Understanding Joins"](#) on page 1-40

Complex View Merging If a view's query contains a GROUP BY clause or DISTINCT operator in the select list, then the optimizer can merge the view's query into the accessing statement only if complex view merging is enabled. Complex merging of a view with a GROUP BY clause is illustrated in [Example 2-43](#).

Complex merging can also be used to merge an IN subquery into the accessing statement if the subquery is uncorrelated, as shown in [Example 2-44](#). Complex merging is not cost-based; it must be enabled with the initialization parameter OPTIMIZER_FEATURES_ENABLE or the MERGE hint. Without this hint or parameter setting, the optimizer uses another approach, pushing predicates.

See Also:

- ["How the CBO Pushes Predicates"](#) on page 2-38 for an alternate approach
- [Chapter 5, "Optimizer Hints"](#) for details about the MERGE and NO_MERGE hints

Example 2-43 *How the CBO Merges a View with a GROUP BY Clause*

The view avg_salary_view contains the average salaries for each department:

```
CREATE VIEW avg_salary_view AS
  SELECT department_id, AVG(salary) AS avg_sal_dept,
         FROM employees
         GROUP BY department_id;
```

If complex view merging is enabled, then the optimizer can transform the following query, which finds the average salaries of departments in London:

```
SELECT departments.location_id, avg_sal_dept
  FROM departments, avg_salary_view
 WHERE departments.department_id = avg_salary_view.department_id
    AND departments.location_id = 2400;
```

into the following query:

```
SELECT departments.loc, AVG(salary)
  FROM departments, employees
 WHERE departments.department_id = employees.department_id
```

```
    AND departments.location_id = 2400
GROUP BY departments.rowid, departments.location_id;
```

The transformed query accesses the view's base table, selecting only the rows of employees who work in London and grouping them by department.

Example 2-44 *How the CBO Merges an IN Subquery*

Complex merging can be used for an `IN` clause with a noncorrelated subquery, as well as for views. The view `min_salary_view` contains the minimum salaries for each department:

```
CREATE VIEW min_salary_view AS
SELECT department_id, MIN(salary) min_sal
   FROM employees
  GROUP BY department_id;
```

If complex merging is enabled, then the optimizer can transform the following query, which finds all employees who earn the minimum salary for their department in London:

```
SELECT employees.last_name, employees.salary
   FROM employees, departments
  WHERE (employees.department_id, employees.salary) IN
        (select department_id, min_sal from min_salary_view)
 AND employees.department_id = departments.department_id
 AND departments.location_id = 2400;
```

into the following query (where `e1` and `e2` represent the `employees` table as it is referenced in the accessing query block and the view's query block, respectively):

```
SELECT e1.last_name, e1.salary
   FROM employees e1, departments, employees e2
  WHERE e1.department_id = departments.department_id
 AND departments.location_id = 2400
 AND e1.department_id = e2.department_id
  GROUP BY e1.rowid, departments.rowid, e1.last_name, e1.salary
 HAVING e1.salary = MIN(e2.salary);
```

How the CBO Pushes Predicates

The optimizer can transform a query block that accesses a nonmergeable view by pushing the query block's predicates inside the view's query. [Examples 2-45](#) and [2-46](#) illustrate this process.

In [Example 2-45](#), the `two_emp_tables` view is the union of two employee tables. The view is defined with a compound query that uses the `UNION` set operator.

Example 2-45 How the CBO Pushes Predicates - the UNION Set Operator

```
CREATE VIEW two_emp_tables
  (employee_id, last_name, job_id, manager_id, hire_date, salary,
   commission_pct, department_id) AS
  SELECT employee_id, last_name, job_id, manager_id, hire_date, salary,
  commission_pct, department_id
    FROM emp1
   UNION
  SELECT employee_id, last_name, job_id, manager_id, hire_date,
  salary, commission_pct, department_id
    FROM emp2;
```

The following query accesses the view. The query selects the IDs and names of all employees in either table who work in department 50:

```
SELECT employee_id, last_name
  FROM two_emp_tables
 WHERE department_id = 50;
```

Because the view is defined as a compound query, the optimizer cannot merge the view's query into the accessing query block. Instead, the optimizer can transform the accessing statement by pushing its predicate, the `WHERE` clause condition (`department_id = 50`), into the view's compound query.

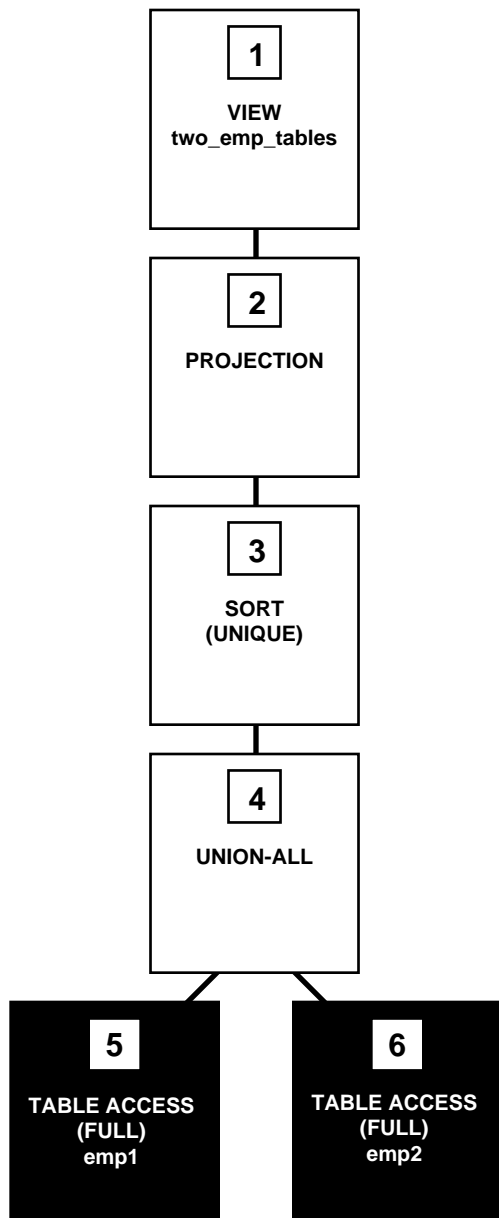
The resulting statement looks like the following:

```
SELECT employee_id, last_name
  FROM ( SELECT employee_id, last_name, job_id, manager_id, hire_date,
  salary, commission_pct, department_id
    FROM emp1
   WHERE department_id = 50
   UNION
  SELECT employee_id, last_name, job_id, manager_id, hire_date,
  salary, commission_pct, department_id
    FROM emp2
   WHERE department_id = 50 );
```

If there is an index on the `department_id` column, then the resulting `WHERE` clauses make it available.

[Figure 2-3](#) shows the execution plan of the resulting statement.

Figure 2–3 Accessing a View Defined with the UNION Set Operator



To execute this statement, Oracle performs the steps in [Figure 2-3](#) in the following order:

1. Steps 5 and 6 perform full scans of the `emp1` and `emp2` tables.
2. Step 4 performs a `UNION-ALL` operation returning all rows returned by either step 5 or step 6, including all copies of duplicates.
3. Step 3 sorts the result of step 4, eliminating duplicate rows.
4. Step 2 extracts the desired columns from the result of step 3.
5. Step 1 indicates that the view's query was not merged into the accessing query.

In [Example 2-46](#), the view `emp_group_by_deptno` contains the department number, average salary, minimum salary, and maximum salary of all departments that have employees.

Example 2-46 How the CBO Pushes Predicates - the GROUP BY Clause

```
CREATE VIEW emp_group_by_deptno
AS SELECT department_id,
        AVG(salary) avg_sal,
        MIN(salary) min_sal,
        MAX(salary) max_sal
FROM employees
GROUP BY department_id;
```

The following query selects the average, minimum, and maximum salaries of department 50 from the `emp_group_by_deptno` view:


```
SELECT *
FROM emp_group_by_deptno
WHERE department_id = 50;
```

The optimizer transforms the statement by pushing its predicate (the `WHERE` clause condition) into the view's query. The resulting statement looks like the following:

```
SELECT department_id,
        AVG(salary) avg_sal,
        MIN(salary) min_sal,
        MAX(salary) max_sal,
FROM employees
WHERE department_id = 50
GROUP BY department_id;
```

If there is an index on the `department_id` column, then the resulting `WHERE` clause makes it available. [Figure 2-4](#) shows the execution plan for the resulting statement. The execution plan uses an index on the `department_id` column.

Figure 2-4 Accessing a View Defined with a *GROUP BY* Clause



To execute this statement, Oracle performs the steps in [Figure 2-4](#) in the following order:

1. Step 4 performs a range scan on the index `emp_department_ix` (an index on the `department_id` column of the `employees` table) to retrieve the rowids of all rows in the `employees` table with a `department_id` value of 50.
2. Step 3 accesses the `employees` table using the rowids retrieved by step 4.
3. Step 2 sorts the rows returned by step 3 to calculate the average, minimum, and maximum `salary` values.
4. Step 1 indicates that the view's query was not merged into the accessing query.

How the CBO Applies an Aggregate Function to the View The optimizer can transform a query that contains an aggregate function (`AVG`, `COUNT`, `MAX`, `MIN`, `SUM`) by applying the function to the view's query.

The query in [Figure 2-47](#) accesses the `emp_group_by_deptno` view defined in [Figure 2-46](#). The query derives the averages for the average department salary, the minimum department salary, and the maximum department salary from the employee table.

Example 2-47 How the CBO Applies an Aggregate Function to a View


```
SELECT AVG(avg_sal), AVG(min_sal), AVG(max_sal)
       FROM emp_group_by_deptno;
```

The optimizer transforms this statement by applying the `AVG` aggregate function to the select list of the view's query:

```
SELECT AVG(AVG(salary)), AVG(MIN(salary)), AVG(MAX(salary))
       FROM employees
       GROUP BY department_id;
```

[Figure 2-5](#) shows the execution plan of the resulting statement.

Figure 2-5 *Applying Aggregate Functions to a View Defined with GROUP BY Clause*



To execute this statement, Oracle performs the steps in [Figure 2-5](#) in the following order:

1. Step 4 performs a full scan of the `employees` table.
2. Step 3 sorts the rows returned by step 4 into groups based on their `department_id` values and calculates the average, minimum, and maximum salary value of each group.
3. Step 2 indicates that the view's query was not merged into the accessing query.
4. Step 1 calculates the averages of the values returned by step 2.

How the CBO Executes Views in Outer Joins

For a view that is on the right side of an outer join, the optimizer can use one of two methods, depending on how many base tables the view accesses:

- If the view has only one base table, then the optimizer can use view merging.
- If the view has multiple base tables, then the optimizer can push the join predicate into the view.

How the CBO Accesses the View's Rows with the Original Statement

The optimizer cannot transform all statements that access views into equivalent statements that access base table(s). For example, if a query accesses a `ROWNUM` pseudocolumn in a view, then the view cannot be merged into the query, and the query's predicate cannot be pushed into the view.

To execute a statement that cannot be transformed into one that accesses base tables, Oracle issues the view's query, collects the resulting set of rows, and then accesses this set of rows with the original statement as though it were a table. [Example 2-48](#) illustrates this process.

Example 2-48 *How the CBO Accesses the View's Rows*

Consider the `emp_group_by_deptno` view defined in [Figure 2-46](#):

```
CREATE VIEW emp_group_by_deptno
AS SELECT department_id,
         AVG(salary) avg_sal,
         MIN(salary) min_sal,
         MAX(salary) max_sal
FROM employees
GROUP BY department_id;
```

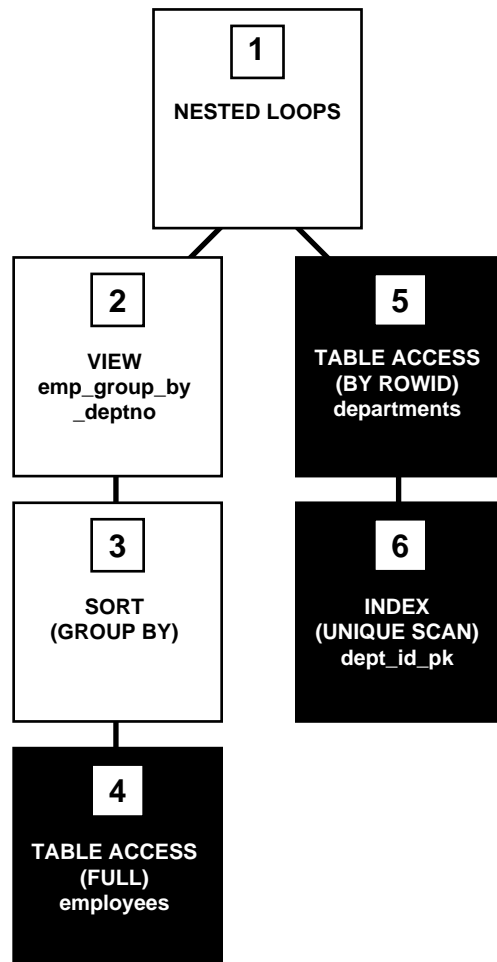
The following query accesses this view. The query joins the average, minimum, and maximum salaries from each department represented in the view to the name and location of the department in the `departments` table.

```
SELECT emp_group_by_deptno.department_id, avg_sal, min_sal,
       max_sal, department_name, location_id
FROM emp_group_by_deptno, departments
WHERE emp_group_by_deptno.department_id = departments.department_id;
```

Because there is no equivalent statement that accesses only base tables, the optimizer cannot transform this statement. Instead, the optimizer chooses an execution plan that issues the view's query and then uses the resulting set of rows as it would the rows resulting from a table access.

See Also: ["Understanding Joins"](#) on page 1-40 for more information on how Oracle performs a nested loops join operation

[Figure 2-6](#) shows the execution plan for this statement.

Figure 2-6 *Joining a View Defined with a GROUP BY Clause to a Table*

To execute this statement, Oracle performs the steps in [Figure 2-6](#) in the following order:

1. Step 4 performs a full scan of the `employees` table.
2. Step 3 sorts the results of step 4 and calculates the average, minimum, and maximum `salary` values selected by the query for the `emp_group_by_deptno` view.

3. Step 2 used the data from the previous two steps for a view.
4. For each row returned by step 2, step 6 uses the `department_id` value to perform a unique scan of the `dept_id_pk` index.
5. Step 5 uses each rowid returned by step 6 to locate the row in the `departments` table with the matching `department_id` value.
6. Oracle combines each row returned by step 2 with the matching row returned by step 5 and returns the result.

How the CBO Executes Compound Queries

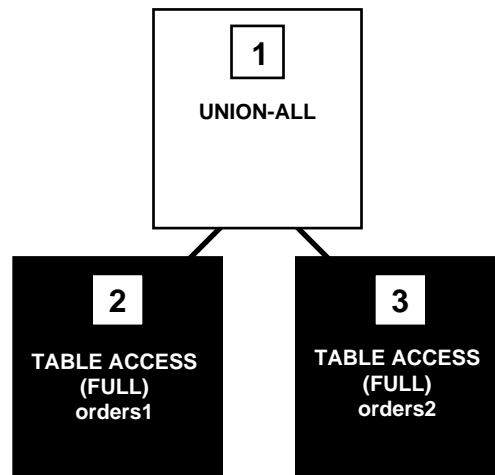
To choose the execution plan for a compound query, the optimizer chooses an execution plan for each of its component queries, and then combines the resulting row sources with the union, intersection, or minus operation, depending on the set operator used in the compound query. [Examples 2-49, 2-50](#) and [2-51](#) illustrate the process.

The query in [Example 2-49](#) uses the `UNION ALL` operator to select all occurrences of all parts in either the `orders1` table or the `orders2` table.

Example 2-49 *How the CBO Executes Compound Queries with `UNION ALL`*

```
SELECT part FROM orders1
UNION ALL
SELECT part FROM orders2;
```

[Figure 2-7](#) shows the execution plan for this statement:

Figure 2–7 Compound Query with UNION ALL Set Operator

To execute this statement, Oracle performs the steps in [Figure 2–7](#) in the following order:

1. Steps 2 and 3 perform full table scans on the `orders1` and `orders2` tables.
2. Step 1 performs a `UNION-ALL` operation returning all rows that are returned by either step 2 or step 3 including all copies of duplicates.

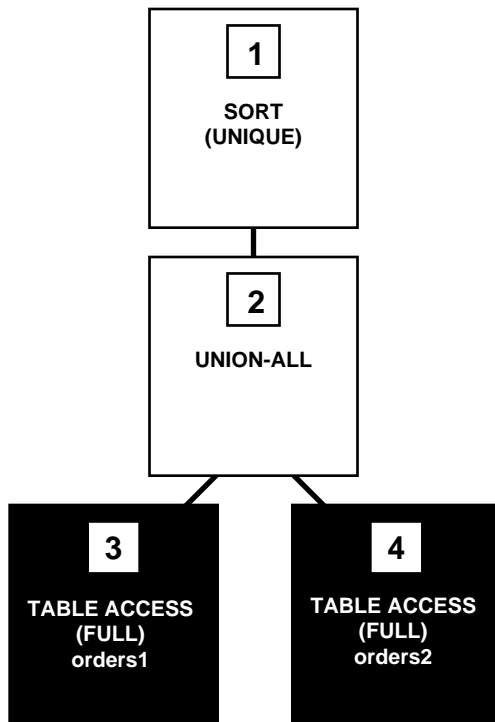
The query in [Example 2–50](#) uses the `UNION` operator to select all parts that appear in either the `orders1` or `orders2` table.

Example 2–50 How the CBO Executes Compound Queries with UNION

```

SELECT part FROM orders1
UNION
SELECT part FROM orders2;
  
```

[Figure 2–8](#) shows the execution plan for this statement:

Figure 2–8 Compound Query with UNION Set Operator

This execution plan is identical to the one for the UNION ALL operator shown in [Figure 2–7](#), except that in this case, Oracle uses the SORT operation to eliminate the duplicates returned by the UNION ALL operation.

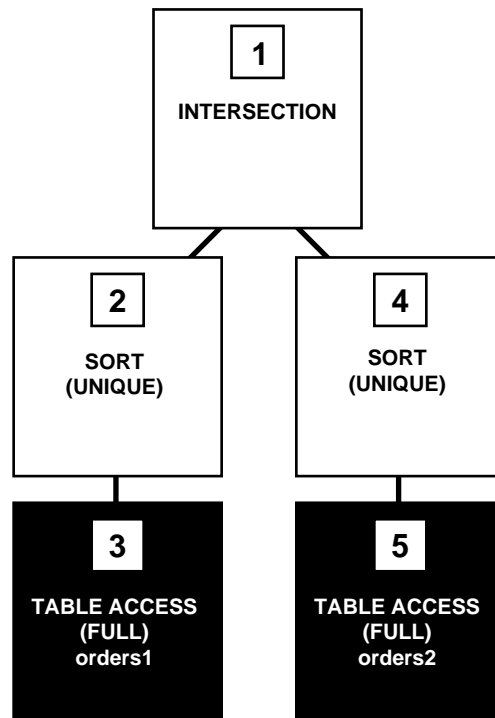
The query in [Example 2–51](#) uses the INTERSECT operator to select only those parts that appear in both the orders1 and orders2 tables.

Example 2–51 How the CBO Executes Compound Queries with INTERSECT

```

SELECT part FROM orders1
INTERSECT
SELECT part FROM orders2;
  
```

[Figure 2–9](#) shows the execution plan for this statement:

Figure 2–9 Compound Query with INTERSECT Set Operator

To execute this statement, Oracle performs the steps in [Figure 2–9](#) in the following order:

1. Steps 3 and 5 perform full table scans of the `orders1` and `orders2` tables.
2. Steps 2 and 4 sort the results of steps 3 and 5, eliminating duplicates in each row source.
3. Step 1 performs an `INTERSECT` operation that returns only rows that are returned by both steps 2 and 4.

Gathering Optimizer Statistics

This chapter explains why statistics are important for the cost-based optimizer and how to gather and use statistics.

The chapter contains the following sections:

- [Understanding Statistics](#)
- [Generating Statistics](#)
- [Using Statistics](#)
- [Using Histograms](#)

Understanding Statistics

As database administrator, you can generate statistics that quantify the data distribution and storage characteristics of tables, columns, indexes, and partitions. The cost-based optimization approach uses these statistics to calculate the selectivity of predicates and to estimate the cost of each execution plan. **Selectivity** is the fraction of rows in a table that the SQL statement's predicate chooses. The optimizer uses the selectivity of a predicate to estimate the cost of a particular access method and to determine the optimal join order and join method.

The statistics are stored in the data dictionary and can be exported from one database and imported into another. For example, you might want to transfer your statistics to a test system to simulate your production environment.

Note: The statistics mentioned in this section are CBO statistics, not instance performance statistics visible through `V$` views.

You should gather statistics periodically for objects where the statistics become stale over time because of changing data volumes or changes in column values. New statistics should be gathered after a schema object's data or structure are modified in ways that make the previous statistics inaccurate. For example, after loading a significant number of rows into a table, collect new statistics on the number of rows. After updating data in a table, you do not need to collect new statistics on the number of rows, but you might need new statistics on the average row length.

Use the `DBMS_STATS` package to generate statistics.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference*

Statistics generated include the following:

- Table statistics
 - Number of rows
 - Number of blocks
 - Average row length
- Column statistics
 - Number of distinct values (NDV) in column
 - Number of nulls in column

- Data distribution (histogram)
- Index statistics
 - Number of leaf blocks
 - Levels
 - Clustering factor
- System statistics
 - I/O performance and utilization
 - CPU performance and utilization

Generating Statistics

Because the cost-based approach relies on statistics, you should generate statistics for all tables and clusters and all indexes accessed by your SQL statements before using the cost-based approach. If the size and data distribution of the tables change frequently, then regenerate these statistics regularly to ensure the statistics accurately represent the data in the tables.

Oracle generates statistics using the following techniques:

- Estimation based on random data sampling
- Exact computation
- User-defined statistics collection methods

To perform an exact computation, Oracle requires enough space to perform a scan and sort of the table. If there is not enough space in memory, then temporary space might be required. For estimations, Oracle requires enough space to perform a scan and sort of only the rows in the requested sample of the table. For indexes, computation does not take up as much time or space.

Some statistics are computed exactly, such as the number of data blocks currently containing data in a table or the depth of an index from its root block to its leaf blocks.

Oracle Corporation recommends setting the `ESTIMATE_PERCENT` parameter of the `DBMS_STATS` gathering procedures to `DBMS_STATS.AUTO_SAMPLE_SIZE` to maximize performance gains while achieving necessary statistical accuracy. `AUTO_SAMPLE_SIZE` lets Oracle determine the best sample size for good statistics. For example, to collect table and column statistics for all tables in the `OE` schema with auto-sampling:

```
EXECUTE DBMS_STATS.GATHER_SCHEMA_STATS('OE',DBMS_STATS.AUTO_SAMPLE_SIZE);
```

To estimate statistics, Oracle selects a random sample of data. You can specify the sampling percentage and whether sampling should be based on rows or blocks. Oracle Corporation recommends using `DBMS_STATS.AUTO_SAMPLE_SIZE` for the sampling percentage. When in doubt, choose row sampling.

- **Row sampling** reads rows without regard to their physical placement on disk. This provides the most random data for estimates, but it can result in reading more data than necessary. For example, in the worst case a row sample might select one row from each block, requiring a full scan of the table or index.
- **Block sampling** reads a random sample of blocks and uses all of the rows in those blocks for estimates. This reduces the amount of I/O activity for a given sample size, but it can reduce the randomness of the sample if rows are not randomly distributed on disk. Block sampling is not available for index statistics.

When you generate statistics for a table, column, or index, if the data dictionary already contains statistics for the object, then Oracle updates the existing statistics. Oracle also invalidates any currently parsed SQL statements that access the object.

The next time such a statement executes, the optimizer automatically chooses a new execution plan based on the new statistics. Distributed statements issued on remote databases that access the analyzed objects use the new statistics the next time Oracle parses them.

When you associate a statistics type with a column or domain index, Oracle calls the statistics collection method in the statistics type, if you analyze the column or domain index.

Getting Statistics for Partitioned Schema Objects

Partitioned schema objects can contain multiple sets of statistics. They can have statistics that refer to any of the following:

- The entire schema object as a whole (global statistics)
- An individual partition
- An individual subpartition of a composite partitioned object

Unless the query predicate narrows the query to a single partition, the optimizer uses the global statistics. Because most queries are not likely to be this restrictive, it is most important to have accurate global statistics. Intuitively, it can seem that generating global statistics from partition-level statistics is straightforward;

however, this is true only for some of the statistics. For example, it is very difficult to figure out the number of distinct values for a column from the number of distinct values found in each partition, because of the possible overlap in values. Therefore, actually gathering global statistics with the `DBMS_STATS` package is highly recommended, rather than calculating them with the `ANALYZE` statement.

Note: Oracle Corporation strongly recommends that you use the `DBMS_STATS` package rather than `ANALYZE` to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by `DBMS_STATS`. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the `ANALYZE` statement rather than `DBMS_STATS` for statistics collection not related to the cost-based optimizer, such as:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
 - To collect information on freelist blocks
-
-

Using the `DBMS_STATS` Package

The PL/SQL package `DBMS_STATS` lets you generate and manage statistics for cost-based optimization. You can use this package to gather, modify, view, export, import, and delete statistics. You can also use this package to identify or name statistics gathered.

The `DBMS_STATS` package can gather statistics on indexes, tables, columns, and partitions, as well as statistics on all schema objects in a schema or database. It does not gather cluster statistics—you can use `DBMS_STATS` to gather statistics on the individual tables instead of the whole cluster.

The statistics-gathering operations can run either serially or in parallel. Index statistics are not gathered in parallel.

For partitioned tables and indexes, `DBMS_STATS` can gather separate statistics for each partition, as well as global statistics for the entire table or index. Similarly, for composite partitioning, `DBMS_STATS` can gather separate statistics for subpartitions, partitions, and the entire table or index. Depending on the SQL

statement being optimized, the optimizer can choose to use either the partition (or subpartition) statistics or the global statistics.

DBMS_STATS gathers only statistics needed for cost-based optimization; it does not gather other statistics. For example, the table statistics gathered by DBMS_STATS include the number of rows, number of blocks currently containing data, and average row length, but not the number of chained rows, average free space, or number of unused data blocks.

See Also:

- *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information about the DBMS_STATS package
- *Oracle9i Data Cartridge Developer's Guide* for more information about user-defined statistics

Gathering Statistics with the DBMS_STATS Package

Table 3-1 lists the procedures in the DBMS_STATS package for gathering statistics:

Table 3-1 Statistics Gathering Procedures in the DBMS_STATS Package

Procedure	Collects
GATHER_INDEX_STATS	Index statistics
GATHER_TABLE_STATS	Table, column, and index statistics
GATHER_SCHEMA_STATS	Statistics for all objects in a schema
GATHER_DATABASE_STATS	Statistics for all objects in a database
GATHER_SYSTEM_STATS	CPU and I/O statistics for the system

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for syntax and examples of all DBMS_STATS procedures

Gathering System Statistics

System statistics enable the optimizer to consider a system's I/O and CPU performance and utilization. For each plan candidate, the optimizer computes estimates for I/O and CPU costs. It is important to know the system characteristics to pick the most efficient plan with optimal proportion between I/O and CPU cost.

System I/O characteristics depend on many factors and do not stay constant all the time. Using system statistics management routines, database administrators can capture statistics in the interval of time when the system has the most common workload. For example, database applications can process OLTP transactions during the day and run OLAP reports at night. Administrators can gather statistics for both states and activate appropriate OLTP or OLAP statistics when needed. This enables the optimizer to generate relevant costs with respect to available system resource plans.

When Oracle generates system statistics, it analyzes system activity in a specified period of time. Unlike table, index, or column statistics, Oracle does not invalidate already parsed SQL statements when system statistics get updated. All new SQL statements are parsed using new statistics. Oracle Corporation highly recommends that you gather system statistics.

The `DBMS_STATS.GATHER_SYSTEM_STATS` routine collects system statistics in a user-defined timeframe. You can also set system statistics values explicitly using `DBMS_STATS.SET_SYSTEM_STATS`. Use `DBMS_STATS.GET_SYSTEM_STATS` to verify system statistics.

Note: You must have DBA privileges to update dictionary system statistics.

[Example 3-1](#) shows database applications processing OLTP transactions during the day and running reports at night. First, system statistics must be collected. The values in this example are user-defined; in other words, you must determine an appropriate time interval and name for your environment.

Example 3-1 *Generating System Statistics*

Gather statistics during the day. Gathering ends after 720 minutes and is stored in the `mystats` table:

```
BEGIN
DBMS_STATS.GATHER_SYSTEM_STATS(
    gathering_mode => 'interval',
    interval => 720,
    stattab => 'mystats',
    statid => 'OLTP');

END;
/
```

Gather statistics during the night. Gathering ends after 720 minutes and is stored in the `mystats` table:

```
BEGIN
DBMS_STATS.GATHER_SYSTEM_STATS(
    gathering_mode => 'interval',
    interval => 720,
    stattab => 'mystats',
    statid => 'OLAP');
END;
/
```

If appropriate, you can switch between the statistics gathered. It is possible to automate this process by submitting a job to update the dictionary with appropriate statistics.

During the day, the following jobs import the OLTP statistics for the daytime run:

```
VARIABLE jobno number;
BEGIN
    DBMS_JOB.SUBMIT(:jobno,
        'DBMS_STATS.IMPORT_SYSTEM_STATS(''mystats'', ''OLTP'');'
        SYSDATE, 'SYSDATE + 1');
    COMMIT;
END;
/
```

During the night, the following jobs import the OLAP statistics for the nighttime run:

```
BEGIN
    DBMS_JOB.SUBMIT(:jobno,
        'DBMS_STATS.IMPORT_SYSTEM_STATS(''mystats'', ''OLAP'');'
        SYSDATE + 0.5, 'SYSDATE + 1');
    COMMIT;
END;
/
```

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for detailed information on the procedures in the `DBMS_STATS` package for implementing system statistics

Gathering Index Statistics

Oracle can gather some statistics automatically while creating or rebuilding a B-tree or bitmap index. The `COMPUTE STATISTICS` option of `CREATE INDEX` or `ALTER INDEX ... REBUILD` enables this gathering of statistics.

The statistics that Oracle gathers for the `COMPUTE STATISTICS` option depend on whether the index is partitioned or nonpartitioned.

- For a nonpartitioned index, Oracle gathers index, table, and column statistics while creating or rebuilding the index. In a concatenated-key index, the column statistics refer only to the leading column of the key.
- For a partitioned index, Oracle does not gather any table or column statistics while creating the index or rebuilding its partitions.
 - While creating a partitioned index, Oracle gathers index statistics for each partition and for the entire index. If the index uses composite partitioning, then Oracle also gathers statistics for each subpartition.
 - While rebuilding a partition or subpartition of an index, Oracle gathers index statistics only for that partition or subpartition.

To ensure correctness of the statistics, Oracle always uses base tables when creating an index with the `COMPUTE STATISTICS` option, even if another index is available that could be used to create the index.

If you do not use the `COMPUTE STATISTICS` clause, or if you have made significant changes to the data, then use the `DBMS_STATS.GATHER_INDEX_STATS` procedure to collect index statistics.

See Also: *Oracle9i SQL Reference* for more information about the `COMPUTE STATISTICS` clause

Gathering Statistics for Function-Based Indexes You should analyze the table after creating a function-based index, to allow Oracle to collect column statistics equivalent information for the expression. Optionally, you can collect histograms for the index expressions by specifying `for all hidden columns size number_of_buckets` in the `METHOD_OPT` argument to the `DBMS_STATS` procedures.

See Also:

- ["Using Histograms"](#) on page 3-24
- *Oracle9i Supplied PL/SQL Packages and Types Reference* for the syntax of the `METHOD_OPT` argument

Gathering New Optimizer Statistics

Before gathering new statistics for a particular schema, use the `DBMS_STATS.EXPORT_SCHEMA_STATS` procedure to extract and save existing statistics. Then use `DBMS_STATS.GATHER_SCHEMA_STATS` to gather new statistics. You can implement both of these with a single call to the `GATHER_SCHEMA_STATS` procedure, by specifying additional parameters.

See Also: ["Preserving Versions of Statistics"](#) on page 3-12

If key SQL statements experience significant performance degradation, then you can either gather statistics again using a larger sample size or perform the following steps:

1. Use `DBMS_STATS.EXPORT_SCHEMA_STATS` to save the new statistics in a different statistics table or a statistics table with a different statistics identifier.
2. Use `DBMS_STATS.IMPORT_SCHEMA_STATS` to restore the old statistics. The application is now ready to run again.

You might want to use the new statistics if they result in improved performance for the majority of SQL statements and if the number of problem SQL statements is small. In this case, perform the following steps:

1. Create a stored outline for each problematic SQL statement, using the old statistics. Stored outlines are precompiled execution plans that Oracle can use to mimic proven application performance characteristics.

See Also: [Chapter 7, "Using Plan Stability"](#)

2. Use `DBMS_STATS.IMPORT_SCHEMA_STATS` to restore the new statistics. The application is now ready to run with the new statistics. However, you will continue to achieve the previous performance levels for the problem SQL statements.

Gathering Automated Statistics

You can automatically gather statistics or create lists of tables that have stale or no statistics. To automatically gather statistics, run the `DBMS_STATS.GATHER_SCHEMA_STATS` and `DBMS_STATS.GATHER_DATABASE_STATS` procedures with the `OPTIONS` and `objlist` parameters. Use the values listed in [Table 3-2](#) for the `options` parameter.

Table 3–2 Options Parameters for Gathering Statistics

Value	Meaning
GATHER STALE	Gathers statistics on tables with stale statistics.
GATHER	Gathers statistics on all tables. (default)
GATHER EMPTY	Gathers statistics only on tables without statistics.
LIST STALE	Creates a list of tables with stale statistics.
LIST EMPTY	Creates a list of tables that do not have statistics.
GATHER AUTO	Gathers <i>all</i> the statistics for the objects of a specific schema (or database with <code>DBMS_STATS.GATHER_DATABASE_STATS()</code>) that are not up-to-date.

The `objlist` parameter identifies an output parameter for the `LIST STALE` and `LIST EMPTY` options. The `objlist` parameter is of type `DBMS_STATS.OBJECTTAB`.

Designating Tables for Monitoring and Automated Statistics Gathering Before you can utilize automated statistics gathering for a particular table, you must bring either the tables of a specific schema or a complete database into the monitoring mode. Do this with the `DBMS_STATS.ALTER_SCHEMA_TAB_MONITORING` or `DBMS_STATS.ALTER_DATABASE_TAB_MONITORING` procedures. Alternatively, you can enable the monitoring attribute using the `MONITORING` keyword. This keyword is part of the `CREATE TABLE` and `ALTER TABLE` statement syntax. Monitoring tracks the approximate number of `INSERTS`, `UPDATES`, and `DELETES` for that table since the last time statistics were gathered. Oracle uses this data to identify tables with stale statistics. Then, you can enable automated statistics gathering by setting up a recurring job (perhaps by using job queues) that invokes `DBMS_STATS.GATHER_TABLE_STATS` with the `GATHER STALE` option at an appropriate interval for your application.

Objects are considered stale when 10% of the total rows have been changed. When you issue `GATHER_TABLE_STATS` with `GATHER STALE`, the procedure checks the `USER_TAB_MODIFICATIONS` view. If a monitored table has been modified more than 10%, then statistics are gathered again. The information about changes of tables, as shown in the `USER_TAB_MODIFICATIONS` view, can be flushed from the SGA into the data dictionary with the `DBMS_STATS.FLUSH_DATABASE_MONITORING_INFO` procedure.

Note: There can be a few minutes delay while Oracle propagates information to this view.

To disable monitoring, use the `DBMS_STATS.ALTER_SCHEMA_TAB_MONITORING` or `DBMS_STATS.ALTER_DATABASE_TAB_MONITORING` procedures, or use the `NOMONITORING` keyword.

See Also: *Oracle9i SQL Reference* for more information about the `CREATE TABLE` and `ALTER TABLE` syntax and the `MONITORING` and `NOMONITORING` keywords

Enabling Automated Statistics Gathering The `GATHER STALE` option gathers statistics only for tables that have stale statistics and for which you have enabled monitoring.

The `GATHER STALE` option maintains up-to-date statistics for the cost-based optimizer. Using this option at regular intervals also avoids the overhead associated with gathering statistics on all tables at one time. The `GATHER` option can incur much more overhead, because this option generally gathers statistics for a greater number of tables than `GATHER STALE`.

Use a script or job scheduling tool for the `GATHER_SCHEMA_STATS` and `GATHER_DATABASE_STATS` procedures to establish a frequency of statistics collection that is appropriate for the application. The frequency of collection intervals should balance the task of providing accurate statistics for the optimizer against the processing overhead incurred by the statistics collection process.

Creating Lists of Tables with Stale or No Statistics You can use the `GATHER_SCHEMA_STATS` and `GATHER_DATABASE_STATS` procedures to create a list of tables with stale statistics. You can also use these procedures to create a list of tables with no statistics. Use the lists to identify tables for which you want to gather manual statistics.

Preserving Versions of Statistics

You can preserve versions of statistics for tables by specifying the `stattab`, `statid`, and `statown` parameters in the `DBMS_STATS` package. Use `stattab` to identify a destination table for archiving previous versions of statistics. Further identify these versions using `statid` to denote the date and time the version was made. Use `statown` to identify a destination schema, if it is different from the schema(s) of the actual tables. You must first create such a table, using the `CREATE_STAT_TABLE` procedure of the `DBMS_STATS` package.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on DBMS_STATS procedures and parameters

Using the ANALYZE Statement

You can use the ANALYZE statement to generate statistics for cost-based optimization.

Note: Oracle Corporation strongly recommends that you use the DBMS_STATS package rather than ANALYZE to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by DBMS_STATS. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the ANALYZE statement rather than DBMS_STATS for statistics collection not related to the cost-based optimizer, such as:

- To use the VALIDATE or LIST CHAINED ROWS clauses
 - To collect information on freelist blocks
-
-

Finding Data Distribution

The statistics gathered help you determine how the data is distributed across the tables. The optimizer assumes that the data is uniformly distributed. You can analyze the actual data distribution in the tables by viewing the appropriate dictionary table: DBA_TABLES for tables and DBA_TAB_COL_STATISTICS for column statistics.

You can use histograms to determine attribute skew.

See Also:

- ["Understanding Statistics"](#) on page 3-2
- ["Using Histograms"](#) on page 3-20

Missing Statistics

When statistics do not exist, the optimizer uses the default values shown in [Table 3-4](#) and [Table 3-4](#).

Table 3-3 Default Table Values When Statistics are Missing

Table Statistic	Default Value Used by Optimizer
Cardinality	$\text{num_of_blocks} * (\text{block_size} - \text{cache_layer}) / \text{avg_row_len}$
Average row length	100 bytes
Number of blocks	
Remote cardinality	2000 rows
Remote average row length	100 bytes

Table 3-4 Default Index Values When Statistics are Missing

Index Statistic	Default Value Used by Optimizer
Levels	1
Leaf blocks	25
Leaf blocks/key	1
Data blocks/key	1
Distinct keys	100
Clustering factor	800 (8 * number of blocks)

Using Statistics

This section provides guidelines on how to use and view statistics. This includes:

- [Managing Statistics](#)
- [Verifying Table Statistics](#)
- [Verifying Index Statistics](#)
- [Verifying Column Statistics](#)
- [Using Histograms](#)

Managing Statistics

This section describes statistics tables and lists the views that display information about statistics stored in the data dictionary.

Using Statistics Tables

The `DBMS_STATS` package lets you store statistics in a statistics table. You can transfer the statistics for a column, table, index, or schema into a statistics table and subsequently restore those statistics to the data dictionary. The optimizer does not use statistics that are stored in a statistics table.

Statistics tables enable you to experiment with different sets of statistics. For example, you can back up a set of statistics before you delete them, modify them, or generate new statistics. You can then compare the performance of SQL statements optimized with different sets of statistics, and if the statistics stored in a table give the best performance, you can restore them to the data dictionary.

A statistics table can keep multiple distinct sets of statistics, or you can create multiple statistics tables to store distinct sets of statistics separately.

Viewing Statistics

Use the `DBMS_STATS` package to view the statistics stored in the data dictionary or in a statistics table. [Example 3-2](#) queries a statistics table.

Example 3-2 Viewing Statistics in a Statistics Table

```

DECLARE
    num_rows NUMBER;
    num_blocks NUMBER;
    avg_row_len NUMBER;

BEGIN
    -- retrieve the values of table statistics on OE.ORDERS
    -- statistics table name: OE.SAVESTATS    statistics ID: TEST1

    DBMS_STATS.GET_TABLE_STATS('OE', 'ORDERS', null,
        'SAVESTATS', 'TEST1',
        num_rows, num_blocks, avg_row_len);

    -- print the values
    DBMS_OUTPUT.PUT_LINE('num_rows=' || num_rows || ', num_blocks=' || num_blocks ||
        ', avg_row_len=' || avg_row_len);
END;

```

Note: Statistics held in a statistics table are held in a form that is only understood by using `DBMS_STATS` package.

To view statistics in the data dictionary, query the appropriate data dictionary view (`USER_`, `ALL_`, or `DBA_`). The following list shows the `DBA_` views:

- `DBA_TABLES`
- `DBA_TAB_COL_STATISTICS`
- `DBA_INDEXES`
- `DBA_CLUSTERS`
- `DBA_TAB_PARTITIONS`
- `DBA_TAB_SUBPARTITIONS`
- `DBA_IND_PARTITIONS`
- `DBA_IND_SUBPARTITIONS`
- `DBA_PART_COL_STATISTICS`
- `DBA_SUBPART_COL_STATISTICS`

See Also: *Oracle9i Database Reference* for information on the statistics in these views

Verifying Table Statistics

To verify that the table statistics are available, query the data dictionary view `DBA_TABLES`, using a statement like the one in [Example 3-3](#):

Example 3-3 *Verifying Table Statistics*

```
SQL> SELECT TABLE_NAME, NUM_ROWS, BLOCKS, AVG_ROW_LEN,  
        TO_CHAR(LAST_ANALYZED, 'MM/DD/YYYY HH24:MI:SS')  
        FROM DBA_TABLES  
        WHERE TABLE_NAME IN ('SO_LINES_ALL', 'SO_HEADERS_ALL', 'SO_LAST_ALL');
```

This returns the following typical data:

TABLE_NAME	NUM_ROWS	BLOCKS	AVG_ROW_LEN	LAST_ANALYZED
SO_HEADERS_ALL	1632264	207014	449	07/29/1999 00:59:51
SO_LINES_ALL	10493845	1922196	663	07/29/1999 01:16:09
SO_LAST_ALL ...				

Note: Because the table SO_LAST_ALL has no statistics, it returns blanks for all columns.

Verifying Index Statistics

To verify that index statistics are available and decide which are the best indexes to use in an application, query the data dictionary view DBA_INDEXES, using a statement like the one in [Example 3-4](#):

Example 3-4 Verifying Index Statistics

```
SQL> SELECT INDEX_NAME "NAME", NUM_ROWS, DISTINCT_KEYS "DISTINCT",
1 LEAF_BLOCKS, CLUSTERING_FACTOR "CF", BLEVEL "LEVEL",
2 AVG_LEAF_BLOCKS_PER_KEY "ALFBPKEY"
3 FROM DBA_INDEXES
4 WHERE OWNER = 'SH'
5* ORDER BY INDEX_NAME;
```

Typical output is:

NAME	NUM_ROWS	DISTINCT	LEAF_BLOCKS	CF	LEVEL	ALFBPKEY
CUSTOMERS_PK	50000	50000	454	4405	2	1
PRODUCTS_PK	10000	10000	90	1552	1	1
PRODUCTS_PROD_CAT_IX	10000	4	99	4422	1	24
PRODUCTS_PROD_SUBCAT_IX	10000	37	170	6148	2	4
SALES_PROD_BIX	6287	909	1480	6287	1	1
SALES_PROMO_BIX	4727	459	570	4727	1	1

6 rows selected.

Optimizer Index Determination Criteria

The optimizer uses the following criteria when determining which index to use:

- Number of rows in the index (cardinality).

- Number of distinct keys. These define the selectivity of the index.
- Level or height of the index. This indicates how deeply the data probe must search in order to find the data.
- Number of leaf blocks in the index. This is the number of I/Os needed to find the desired rows of data.
- Clustering factor (CF). This is the collocation amount of the index block relative to data blocks. The higher the CF, the less likely the optimizer is to select this index.
- Average leaf blocks for each key (ALFBKEY). Average number of leaf blocks in which each distinct value in the index appears, rounded to the nearest integer. For indexes that enforce `UNIQUE` and `PRIMARY KEY` constraints, this value is always one.

Determining if You Have Chosen the Right Index

Use the following notes to assist you in deciding whether you have chosen an appropriate index for a table, data, and query:

DISTINCT Consider index `ap_invoices_n3`, having two distinct keys. The resulting selectivity based on index `ap_invoices_n3` is poor, and the optimizer is not likely to use this index. Using this index fetches 50% of the data in the table. In this case, a full table scan is cheaper than using index `ap_invoices_n3`.

Index Cost Tie The optimizer uses alphabetic determination. If the optimizer determines that the selectivity, cost, and cardinality of two finalist indexes is the same, then it looks at the names of the indexes and chooses the name that begins with the lower alphabetic letter or number.

Verifying Column Statistics

To verify that column statistics are available, query the data dictionary view `DBA_TAB_COL_STATISTICS`, using a statement like the one in [Example 3-5](#):

Example 3-5 Verifying Column Statistics

```
SQL> SELECT COLUMN_NAME, NUM_DISTINCT, NUM_NULLS, NUM_BUCKETS, DENSITY
       FROM DBA_TAB_COL_STATISTICS
       WHERE TABLE_NAME = "PA_EXPENDITURE_ITEMS_ALL"
       ORDER BY COLUMN_NAME;
```

This returns the following data:

COLUMN_NAME	NUM_DISTINCT	NUM_NULLS	NUM_BUCKETS	DENSITY
BURDEN_COST	4300	71957	1	.000232558
BURDEN_COST_RATE	675	7376401	1	.001481481
CONVERTED_FLAG	1	16793903	1	1
COST_BURDEN_DISTRIBUTED_FLAG	2	15796	1	.5
COST_DISTRIBUTED_FLAG	2	0	1	.5
COST_IND_COMPILED_SET_ID	87	6153143	1	.011494253
EXPENDITURE_ID	1171831	0	1	8.5337E-07
TASK_ID	8648	0	1	.000115634
TRANSFERRED_FROM_EXP_ITEM_ID	1233787	15568891	1	8.1051E-07

Verifying column statistics are important for the following conditions:

- Join conditions
- When the WHERE clause includes a column(s) with a bind variable; for example:
column x = :variable_y

In these cases, you can use the stored column statistics to get a representative cardinality estimation for the given expression.

The following subsections examine the data returned by the query in [Example 3-5](#).

NUM_DISTINCT Column Statistic

NUM_DISTINCT indicates the number of distinct values for a column.

Low In [Example 3-5](#), the number of distinct values for the column CONVERTED_FLAG is 1. In this case, this column has only one value. If there is a bind variable on column CONVERTED_FLAG in the WHERE clause (for example, CONVERTED_FLAG = :variable_y), then this leads to poor selectivity, and CONVERTED_FLAG is a poor candidate to be used as the index.

Column COST_BURDEN_DISTRIBUTED_FLAG: NUM_DISTINCT = 2. Likewise, this value is low. COST_BURDEN_DISTRIBUTED_FLAG is not a good candidate for an index unless there is much skew or there are a lot of nulls. If there is data skew of, say, 90%, then 90% of the data has one particular value and 10% of the data has another value. If the query only needs to access the 10%, then a histogram is needed on that column in order for the optimizer to recognize the skew and use an index on this column.

High In [Example 3-5](#), NUM_DISTINCT is more than 1 million for column EXPENDITURE_ID. If there is a bind variable on column EXPENDITURE_ID, then

this leads to high selectivity (implying high density of data on this column). In other words, `EXPENDITURE_ID` is a good candidate to be used as the index.

NUM_NULL Column Statistic

`NUM_NULLS` indicates the number of null values for that column.

Low If a single column index has few nulls, such as the `COST_DISTRIBUTED_FLAG` column in [Example 3-5](#), and if this column is used as the index, then the resulting data set is large.

High If there are many nulls on a particular column, such as the `CONVERTED_FLAG` column in [Example 3-5](#), and if this column is used as the index, then the resulting data set is small. This means that `COST_DISTRIBUTED_FLAG` is a more appropriate column to index.

DENSITY Column Statistic

This indicates the density of the values of that column. This is calculated as 1 divided by `NUM_DISTINCT`.

Column Statistics and Join Methods

Column statistics are useful to help determine the most efficient join method, which, in turn, is also based on the number of rows returned.

Using Histograms

The cost-based optimizer can use data value histograms to get accurate estimates of the distribution of column data. A **histogram** partitions the values in the column into bands, so that all column values in a band fall within the same range. Histograms provide improved selectivity estimates in the presence of data skew, resulting in optimal execution plans with nonuniform data distributions.

One of the fundamental tasks of the cost-based optimizer is determining the selectivity of predicates that appear in queries. Selectivity estimates are used to decide when to use an index and the order in which to join tables. Some attribute domains (a table's columns) are not uniformly distributed.


The cost-based optimizer uses height-based histograms on specified attributes to describe the distributions of nonuniform domains. In a height-based histogram, the column values are divided into bands so that each band contains approximately the

same number of values. The useful information that the histogram provides, then, is where in the range of values the endpoints fall.

See Also: ["Types of Histograms"](#) on page 3-23

Consider a column C with values between 1 and 100 and a histogram with 10 buckets. If the data in C is uniformly distributed, then the histogram looks similar to [Figure 3-1](#), where the numbers are the endpoint values.


Figure 3-1 *Histogram with Uniform Distribution*



The number of rows in each bucket is one tenth the total number of rows in the table. Four-tenths of the rows have values between 60 and 100 in this example of uniform distribution.

If the data is not uniformly distributed, then the histogram might look similar to [Figure 3-2](#).

Figure 3-2 *Histogram with Non-Uniform Distribution*



In this case, most of the rows have the value 5 for the column; only 1/10 of the rows have values between 60 and 100.

When to Use Histograms

Histograms can affect performance and should be used only when they substantially improve query plans. Histogram statistics data is persistent, so the space required to save the data depends on the sample size. In general, create histograms on columns that are used frequently in `WHERE` clauses of queries and have a highly skewed data distribution. For uniformly distributed data, the cost-based optimizer can make fairly accurate guesses about the cost of executing a particular statement without the use of histograms.

Histograms, like all other optimizer statistics, are static. They are useful only when they reflect the current data distribution of a given column. (The data in the column

can change as long as the *distribution* remains constant.) If the data distribution of a column changes frequently, you must recompute its histogram frequently.

Histograms are not useful for columns with the following characteristics:

- All predicates on the column use bind variables.
- The column data is uniformly distributed.
- The column is unique and is used only with equality predicates.

Creating Histograms

You generate histograms by using the `DBMS_STATS` package. You can generate histograms for columns of a table or partition. For example, to create a 10-bucket histogram on the `SAL` column of the `emp` table, issue the following statement:

```
EXECUTE DBMS_STATS.GATHER_TABLE_STATS
('scott','emp', METHOD_OPT => 'FOR COLUMNS SIZE 10 sal');
```

The `SIZE` keyword declares the maximum number of buckets for the histogram. You would create a histogram on the `SAL` column if there were an unusually high number of employees with the same salary and few employees with other salaries. You can also collect histograms for a single partition of a table.

Oracle Corporation recommends using the `DBMS_STATS` package to have the database automatically decide which columns need histograms. This is done by specifying `SIZE AUTO`.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on the `DBMS_STATS` package

Choosing the Number of Buckets for a Histogram

If the number of frequently occurring distinct values in a column is relatively small, then set the number of buckets to be greater than that number. The default number of buckets for a histogram is 75. This value provides an appropriate level of detail for most data distributions. However, because the number of buckets and the data distribution both affect a histogram's usefulness, you might need to experiment with different numbers of buckets to obtain optimal results.

Types of Histograms

There are two types of histograms:

- [Understanding Height-Based Histograms](#)
- [Understanding Value-Based Histograms](#)

Understanding Height-Based Histograms

Height-based histograms place approximately the same number of values into each range, so that the endpoints of the range are determined by how many values are in that range. Only the last (largest) values in each bucket appear as bucket (end point) values.

Consider that a table's query results in the following four sample values: 4, 18, 30, and 35.

For a height-based histogram, each of these values occupies a portion of one bucket, in proportion to their size. The resulting selectivity is computed with the following formula:

$$S = \text{Height}(35) / \text{Height}(4 + 18 + 30 + 35)$$

Understanding Value-Based Histograms

Value-based histograms are created when the number of distinct values is less than or equal to the number of histogram buckets specified. In value-based histograms, all the values in the column have corresponding buckets, and the bucket number reflects the repetition count of each value. These can also be known as frequency histograms.

Consider the same four sample values in the previous example. In a value-based histogram, a bucket is used to represent each of the four distinct values. In other words, one bucket represents 4, one bucket represents 18, another represents 30, and another represents 35. The resulting selectivity is computed with the following formula:

$$S = [\#\text{rows}(35) / (\#\text{rows}(4) + \#\text{rows}(18) + \#\text{rows}(30) + \#\text{rows}(35))] / \#\text{buckets}$$

If there are many different values anticipated for a particular column of table, it is preferable to use the value-based histogram rather than the height-based histogram. This is because if there is much data skew in the height, then the skew can offset the selectivity calculation and give a nonrepresentative selectivity value.

Using Histograms

Example 3–6 illustrates the use of a histogram in order to improve the execution plan and demonstrate the skewed behavior of the `s6` indexed column.

Example 3–6 Using a Histogram to Improve an Execution Plan

```
UPDATE so_lines l
SET open_flag=null,
    s6=10,
    s6_date=sysdate,
WHERE l.line_type_code in ('REGULAR','DETAIL','RETURN') AND
      l.open_flag = 'Y' AND NVL(l.shipped_quantity, 0)=0 OR
      NVL(l.shipped_quantity, 0) != 0 AND
      l.shipped_quantity +NVL(l.cancelled_quantity, 0)= l.ordered_quantity)) AND
      l.s6=18
```

This query shows the skewed distribution of data values for `s6`. In this case, there are two distinct non-null values: 10 and 18. The majority of the rows consists of `s6 = 10` (1,589,464), while a small number of rows consist of `s6 = 18` (13,091).

```
S6:          COUNT(*)
=====
10          1,589,464
18           13,091
NULL        21,889
```

The selectivity of column `s6`, where `s6 = 18`:

$$S = 13,091 / (13,091 + 1,589,464) = 0.008$$

If No Histogram is Used: The selectivity of column `s6` is assumed to be 50%, uniformly distributed across 10 and 18. This is not selective; therefore, `s6` is not an ideal choice for use as an index.

If a Histogram is Used: The data distribution information is stored in the dictionary. This allows the optimizer to use this information and compute the correct selectivity based on the data distribution. In **Example 3–6**, the selectivity, based on the histogram data, is 0.008. This is a relatively high, or good, selectivity, which leads the optimizer to use an index on column `s6` in the execution plan.

Viewing Histograms

To view histogram information, query the appropriate data dictionary view (USER_, ALL_, or DBA_). The following list shows the DBA_ views:

- DBA_HISTOGRAMS
- DBA_PART_HISTOGRAMS
- DBA_SUBPART_HISTOGRAMS
- DBA_TAB_COL_STATISTICS

Number of Rows View the DBA_HISTOGRAMS dictionary table for the number of buckets (in other words, the number of rows) for each column:

- ENDPOINT_NUMBER
- ENDPOINT_VALUE

See Also: *Oracle9i Database Reference* for column descriptions of data dictionary views, as well as histogram use and restrictions

Verifying Histogram Statistics

To verify that histogram statistics are available, query the data dictionary's DBA_HISTOGRAMS table, using a statement similar to [Example 3-7](#).

Example 3-7 Verifying Histogram Statistics

```
SQL> SELECT ENDPOINT_NUMBER, ENDPOINT_VALUE
       FROM DBA_HISTOGRAMS
       WHERE TABLE_NAME = "SO_LINES_ALL" AND COLUMN_NAME="S2"
       ORDER BY ENDPOINT_NUMBER;
```

This query returns the following typical data:

ENDPOINT_NUMBER	ENDPOINT_VALUE
1365	4
1370	5
2124	8
2228	18

One row corresponds to one bucket in the histogram. Consider the differences between ENDPOINT_NUMBER values in [Example 3-7](#) listed in [Table 3-5](#).

Table 3–5 *ENDPOINT_NUMBER Differences*

Bucket (values)	ENDPOINT_NUMBER Difference	Number of Values in Bucket
1 (0 to 4)	N/A	N/A
2 (4 to 5)	1370 - 1365	5
3 (5 to 8)	2124 - 1370	754
4 (8 to 18)	2228 - 2124	104

Table 3–5 shows that the buckets hold very different numbers of values. The data is skewed: 754 values are between 5 and 8, but only 104 are between 8 and 18. More buckets should be used.

Understanding Indexes and Clusters

This chapter provides an overview of data access methods using indexes and clusters that can enhance or degrade performance.

The chapter contains the following sections:

- [Understanding Indexes](#)
- [Using Function-based Indexes](#)
- [Using Index-Organized Tables](#)
- [Using Bitmap Indexes](#)
- [Using Bitmap Join Indexes](#)
- [Using Domain Indexes](#)
- [Using Clusters](#)
- [Using Hash Clusters](#)

Understanding Indexes

This section describes the following:

- [Tuning the Logical Structure](#)
- [Choosing Columns and Expressions to Index](#)
- [Choosing Composite Indexes](#)
- [Writing Statements That Use Indexes](#)
- [Writing Statements That Avoid Using Indexes](#)
- [Re-creating Indexes](#)
- [Using Nonunique Indexes to Enforce Uniqueness](#)
- [Using Enabled Novalidated Constraints](#)

Tuning the Logical Structure

Although cost-based optimization helps avoid the use of nonselective indexes within query execution, the SQL engine must continue to maintain all indexes defined against a table, regardless of whether they are used. Index maintenance can present a significant CPU and I/O resource demand in any write-intensive application. In other words, do not build indexes unless necessary.

To maintain optimal performance, drop indexes that an application is not using. You can find indexes that are not being used by using the `ALTER INDEX MONITORING USAGE` functionality over a period of time that is representative of your workload. This monitoring feature records whether or not an index has been used. If you find that an index has not been used, then drop it. Be careful to select a representative workload to monitor.

See Also: *Oracle9i SQL Reference*

Indexes within an application sometimes have uses that are not immediately apparent from a survey of statement execution plans. An example of this is a foreign key index on a parent table, which prevents share locks from being taken out on a child table.

If you are deciding whether to create new indexes to tune statements, then you can also use the `EXPLAIN PLAN` statement to determine whether the optimizer will choose to use these indexes when the application is run. If you create new indexes to tune a statement that is currently parsed, then Oracle invalidates the statement.

When the statement is next executed, the optimizer automatically chooses a new execution plan that could potentially use the new index. If you create new indexes on a remote database to tune a distributed statement, then the optimizer considers these indexes when the statement is next parsed.

Also keep in mind that the way you tune one statement can affect the optimizer's choice of execution plans for other statements. For example, if you create an index to be used by one statement, then the optimizer can choose to use that index for other statements in the application as well. For this reason, reexamine the application's performance and rerun the SQL trace facility after you have tuned those statements that you initially identified for tuning.

Note: You can use the Oracle Index Tuning Wizard to detect tables with inefficient indexes. The Oracle Index Tuning wizard is an Oracle Enterprise Manager integrated application available with the Oracle Tuning Pack. Similar functionality is available from the Virtual Index Advisor (a feature of SQL Analyze) and Oracle Expert.

See Also: *Database Tuning with the Oracle Tuning Pack*

Choosing Columns and Expressions to Index

A key is a column or expression on which you can build an index. Follow these guidelines for choosing keys to index:

- Consider indexing keys that are used frequently in `WHERE` clauses.
- Consider indexing keys that are used frequently to join tables in SQL statements. For more information on optimizing joins, see the section "[Using Hash Clusters](#)" on page 4-21.
- Index keys that have high selectivity. The selectivity of an index is the percentage of rows in a table having the same value for the indexed key. An index's selectivity is optimal if few rows have the same value.

Note: Oracle automatically creates indexes, or uses existing indexes, on the keys and expressions of unique and primary keys that you define with integrity constraints.

Indexing low selectivity columns can be helpful if the data distribution is skewed so that one or two values occur much less often than other values.

- Do not use standard B-tree indexes on keys or expressions with few distinct values. Such keys or expressions usually have poor selectivity and therefore do not optimize performance unless the frequently selected key values appear less frequently than the other key values. You can use bitmap indexes effectively in such cases, unless a high concurrency OLTP application is involved where the index is modified frequently.
- Do not index columns that are modified frequently. `UPDATE` statements that modify indexed columns and `INSERT` and `DELETE` statements that modify indexed tables take longer than if there were no index. Such SQL statements must modify data in indexes as well as data in tables. They also generate additional undo and redo.
- Do not index keys that appear only in `WHERE` clauses with functions or operators. A `WHERE` clause that uses a function, other than `MIN` or `MAX`, or an operator with an indexed key does not make available the access path that uses the index except with function-based indexes.
- Consider indexing foreign keys of referential integrity constraints in cases in which a large number of concurrent `INSERT`, `UPDATE`, and `DELETE` statements access the parent and child tables. Such an index allows `UPDATES` and `DELETES` on the parent table without share locking the child table.
- When choosing to index a key, consider whether the performance gain for queries is worth the performance loss for `INSERTS`, `UPDATES`, and `DELETES` and the use of the space required to store the index. You might want to experiment by comparing the processing times of the SQL statements with and without indexes. You can measure processing time with the SQL trace facility.

See Also: *Oracle9i Application Developer's Guide - Fundamentals* for more information on the effects of foreign keys on locking

Choosing Composite Indexes

A composite index contains more than one key column. Composite indexes can provide additional advantages over single-column indexes:

- Improved selectivity
Sometimes two or more columns or expressions, each with poor selectivity, can be combined to form a composite index with higher selectivity.
- Reduced I/O

If all columns selected by a query are in a composite index, then Oracle can return these values from the index without accessing the table.

A SQL statement can use an access path involving a composite index if the statement contains constructs that use a leading portion of the index.

Note: This is no longer the case with index skip scans. See "[Index Skip Scans](#)" on page 1-32.

A leading portion of an index is a set of one or more columns that were specified first and consecutively in the list of columns in the `CREATE INDEX` statement that created the index. Consider this `CREATE INDEX` statement:

```
CREATE INDEX comp_ind
ON table1(x, y, z);
```

- `x`, `xy`, and `xyz` combinations of columns are leading portions of the index
- `yz`, `y`, and `z` combinations of columns are *not* leading portions of the index

Choosing Keys for Composite Indexes

Follow these guidelines for choosing keys for composite indexes:

- Consider creating a composite index on keys that are used together frequently in `WHERE` clause conditions combined with `AND` operators, especially if their combined selectivity is better than the selectivity of either key individually.
- If several queries select the same set of keys based on one or more key values, then consider creating a composite index containing all of these keys.

Of course, consider the guidelines associated with the general performance advantages and trade-offs of indexes described in the previous sections.

Ordering Keys for Composite Indexes

Follow these guidelines for ordering keys in composite indexes:

- Create the index so the keys used in `WHERE` clauses make up a leading portion.
- If some keys are used in `WHERE` clauses more frequently, then be sure to create the index so that the more frequently selected keys make up a leading portion to allow the statements that use only these keys to use the index.

- If all keys are used in `WHERE` clauses equally often, then ordering these keys from most selective to least selective in the `CREATE INDEX` statement best improves query performance.
- If all keys are used in the `WHERE` clauses equally often but the data is physically ordered on one of the keys, then place that key first in the composite index.

Writing Statements That Use Indexes

Even after you create an index, the optimizer cannot use an access path that uses the index simply because the index exists. The optimizer can choose such an access path for a SQL statement only if it contains a construct that makes the access path available. To allow the CBO the option of using an index access path, ensure that the statement contains a construct that makes such an access path available.

Writing Statements That Avoid Using Indexes

In some cases, you might want to prevent a SQL statement from using an access path that uses an existing index. You might want to do this if you know that the index is not very selective and that a full table scan would be more efficient. If the statement contains a construct that makes such an index access path available, then you can force the optimizer to use a full table scan through one of the following methods:

- Use the `NO_INDEX` hint to give the CBO maximum flexibility while disallowing the use of a certain index.
- Use the `FULL` hint to force the optimizer to choose a full table scan instead of an index scan.
- Use the `INDEX`, `INDEX_COMBINE`, or `AND_EQUAL` hints to force the optimizer to use one index or a set of listed indexes instead of another.

See Also: [Chapter 5, "Optimizer Hints"](#) for more information on the `NO_INDEX`, `FULL`, `INDEX`, `INDEX_COMBINE`, and `AND_EQUAL` hints

Parallel execution uses indexes effectively. It does not perform parallel index range scans, but it does perform parallel index lookups for parallel nested loop join execution. If an index is very selective (there are few rows for each index entry), then it might be better to use sequential index lookup rather than parallel table scan.

Re-creating Indexes

You might want to re-create an index to compact it and minimize fragmented space, or to change the index's storage characteristics. When creating a new index that is a subset of an existing index or when rebuilding an existing index with new storage characteristics, Oracle might use the existing index instead of the base table to improve the performance of the index build.

Note: To avoid calling `DBMS_STATS` after the index creation or rebuild, include the `COMPUTE STATISTICS` statement on the `CREATE` or `REBUILD`. You can use the Oracle Enterprise Manager Reorg Wizard to identify indexes that require rebuilding. The Reorg Wizard can also be used to rebuild the indexes.

However, there are cases where it can be beneficial to use the base table instead of the existing index. Consider an index on a table on which a lot of DML has been performed. Because of the DML, the size of the index can increase to the point where each block is only 50% full, or even less. If the index refers to most of the columns in the table, then the index could actually be *larger* than the table. In this case, it is faster to use the base table rather than the index to re-create the index.

Use the `ALTER INDEX ... REBUILD` statement to reorganize or compact an existing index or to change its storage characteristics. The `REBUILD` statement uses the existing index as the basis for the new one. All index storage statements are supported, such as `STORAGE` (for extent allocation), `TABLESPACE` (to move the index to a new tablespace), and `INITRANS` (to change the initial number of entries).

Usually, `ALTER INDEX ... REBUILD` is faster than dropping and re-creating an index, because this statement uses the fast full scan feature. It reads all the index blocks using multiblock I/O, then discards the branch blocks. A further advantage of this approach is that the old index is still available for queries while the rebuild is in progress.

See Also: *Oracle9i SQL Reference* for more information about the `CREATE INDEX` and `ALTER INDEX` statements, as well as restrictions on rebuilding indexes

Compacting Indexes

You can coalesce leaf blocks of an index by using the `ALTER INDEX` statement with the `COALESCE` option. This option lets you combine leaf levels of an index to free blocks for reuse. You can also rebuild the index online.

See Also: *Oracle9i SQL Reference* and *Oracle9i Database Administrator's Guide* for more information about the syntax for this statement

Using Nonunique Indexes to Enforce Uniqueness

You can use an existing nonunique index on a table to enforce uniqueness, either for `UNIQUE` constraints or the unique aspect of a `PRIMARY KEY` constraint. The advantage of this approach is that the index remains available and valid when the constraint is disabled. Therefore, enabling a disabled `UNIQUE` or `PRIMARY KEY` constraint does not require rebuilding the unique index associated with the constraint. This can yield significant time savings on enable operations for large tables.

Using a nonunique index to enforce uniqueness also lets you eliminate redundant indexes. You do not need a unique index on a primary key column if that column already is included as the prefix of a composite index. You can use the existing index to enable and enforce the constraint. You also save significant space by not duplicating the index. However, if the existing index is partitioned, then the partitioning key of the index must also be a subset of the `UNIQUE` key; otherwise, Oracle creates an additional unique index to enforce the constraint.

Using Enabled Novalidated Constraints

An enabled novalidated constraint behaves similarly to an enabled validated constraint for new data. Placing a constraint in the enabled novalidated state signifies that any new data entered into the table must conform to the constraint. Existing data is not checked. By placing a constraint in the enabled novalidated state, you enable the constraint without locking the table.

If you change a constraint from disabled to enabled, then the table must be locked. No new DML, queries, or DDL can occur, because there is no mechanism to ensure that operations on the table conform to the constraint during the enable operation. The enabled novalidated state prevents operations violating the constraint from being performed on the table.

An enabled novalidated constraint can be validated with a parallel, consistent-read query of the table to determine whether any data violates the constraint. No locking is performed, and the enable operation does not block readers or writers to the table. In addition, enabled novalidated constraints can be validated in parallel: Multiple constraints can be validated at the same time and each constraint's validity check can be determined using parallel query.

Use the following approach to create tables with constraints and indexes:

1. Create the tables with the constraints. NOT NULL constraints can be unnamed and should be created enabled and validated. All other constraints (CHECK, UNIQUE, PRIMARY KEY, and FOREIGN KEY) should be named and created disabled.

Note: By default, constraints are created in the ENABLED state.

2. Load old data into the tables.
3. Create all indexes, including indexes needed for constraints.
4. Enable novalidate all constraints. Do this to primary keys before foreign keys.
5. Allow users to query and modify data.
6. With a separate ALTER TABLE statement for each constraint, validate all constraints. Do this to primary keys before foreign keys. For example,

```
CREATE TABLE t (a NUMBER CONSTRAINT apk PRIMARY KEY DISABLE,
b NUMBER NOT NULL);
CREATE TABLE x (c NUMBER CONSTRAINT afk REFERENCES t DISABLE);
```

Now you can use Import or Fast Loader to load data into table t.

```
CREATE UNIQUE INDEX tai ON t (a);
CREATE INDEX tci ON x (c);
ALTER TABLE t MODIFY CONSTRAINT apk ENABLE NOVALIDATE;
ALTER TABLE x MODIFY CONSTRAINT afk ENABLE NOVALIDATE;
```

At this point, users can start performing INSERTS, UPDATES, DELETES, and SELECTS on table t.

```
ALTER TABLE t ENABLE CONSTRAINT apk;
ALTER TABLE x ENABLE CONSTRAINT afk;
```

Now the constraints are enabled and validated.

See Also: *Oracle9i Database Concepts* for a complete discussion of integrity constraints

Using Function-based Indexes

A function-based index includes columns that are either transformed by a function, such as the `UPPER` function, or included in an expression, such as `col1 + col2`.

Defining a function-based index on the transformed column or expression allows that data to be returned using the index when that function or expression is used in a `WHERE` clause or an `ORDER BY` clause. Therefore, a function-based index can be beneficial when frequently-executed SQL statements include transformed columns, or columns in expressions, in a `WHERE` or `ORDER BY` clause.

Function-based indexes defined with the `UPPER(column_name)` or `LOWER(column_name)` keywords allow case-insensitive searches. For example, the following index:

```
CREATE INDEX uppercase_idx ON employees (UPPER(last_name));
```

facilitates processing queries such as:

```
SELECT * FROM employees
      WHERE UPPER(last_name) = 'MARKSON';
```

Setting Parameters to Use Function-Based Indexes in Queries

To use function-based indexes in queries, you need to set the `QUERY_REWRITE_ENABLED` and `QUERY_REWRITE_INTEGRITY` parameters.

QUERY_REWRITE_ENABLED

To enable function-based indexes for queries, set the `QUERY_REWRITE_ENABLED` session parameter to `TRUE`. `QUERY_REWRITE_ENABLED` can be set to the following values:

- `TRUE`: cost - based rewrite
- `FALSE`: no rewrite
- `FORCE`: forced rewrite

When `QUERY_REWRITE_ENABLED` is set to `FALSE`, then function-based indexes are not used for obtaining the values of an expression in the function-based index. However, function-based indexes can still be used for obtaining values in real columns.

When `QUERY_REWRITE_ENABLED` is set to `FORCE`, Oracle always uses rewrite and does not evaluate the cost before doing so. `FORCE` is useful when you know that the query will always benefit from rewrite, when reduction in compile time is

important, and when you know that the optimizer may be underestimating the benefits of materialized views.

QUERY_REWRITE_ENABLED is a session-level and also an instance-level parameter.

QUERY_REWRITE_INTEGRITY

Setting the value of the QUERY_REWRITE_INTEGRITY parameter determines how function-based indexes are used,

- If the QUERY_REWRITE_INTEGRITY parameter is set to ENFORCED (the default), then Oracle uses function-based indexes to derive values of SQL expressions only. This also includes SQL functions.
- If QUERY_REWRITE_INTEGRITY is set to any value other than ENFORCED, then Oracle uses the function-based index, even if it is based on a user-defined, rather than SQL, function.

Function-based indexes are an efficient mechanism for evaluating statements that contain functions in WHERE clauses. You can create a function-based index to store computation-intensive expressions in the index. This permits Oracle to bypass computing the value of the expression when processing SELECT and DELETE statements. When processing INSERT and UPDATE statements, however, Oracle evaluates the function to process the statement.

For example, if you create the following index:

```
CREATE INDEX idx ON table_1 (a + b * (c - 1), a, b);
```

then Oracle can use it when processing queries such as:

```
SELECT a
   FROM table_1
  WHERE a + b * (c - 1) < 100;
```

You can also use function-based indexes for linguistic sort indexes that provide efficient linguistic collation in SQL statements.

Oracle treats descending indexes as function-based indexes. The columns marked DESC are sorted in descending order.

See Also:

- *Oracle9i Application Developer's Guide - Fundamentals* and *Oracle9i Database Administrator's Guide* for more information on using function-based indexes
- *Oracle9i SQL Reference* for more information on the `CREATE INDEX` statement

Using Index-Organized Tables

An index-organized table differs from an ordinary table in that the data for the table is held in its associated index. Changes to the table data, such as adding new rows, updating rows, or deleting rows, result only in updating the index. Because data rows are stored in the index, index-organized tables provide faster key-based access to table data for queries that involve exact match or range search or both.

See Also: *Oracle9i Database Concepts*

Using Bitmap Indexes

This section describes:

- [When to Use Bitmap Indexes](#)
- [Using Bitmap Indexes with Good Performance](#)
- [Initialization Parameters for Bitmap Indexing](#)
- [Using Bitmap Access Plans on Regular B-tree Indexes](#)
- [Bitmap Index Restrictions](#)

See Also: *Oracle9i Database Concepts* and *Oracle9i Data Warehousing Guide* for more information on bitmap indexing

When to Use Bitmap Indexes

This section describes three aspects of indexing that you must evaluate when deciding whether to use bitmap indexing on a given table:

- [Performance Considerations for Bitmap Indexes](#)
- [Comparing B-tree Indexes to Bitmap Indexes](#)

- [Maintenance Considerations for Bitmap Indexes](#)

Performance Considerations for Bitmap Indexes

Bitmap indexes can substantially improve performance of queries that have all of the following characteristics:

- The `WHERE` clause contains multiple predicates on low- or medium-cardinality columns.
- The individual predicates on these low- or medium-cardinality columns select a large number of rows.
- Bitmap indexes have been created on some or all of these low- or medium-cardinality columns.
- The tables being queried contain many rows.

You can use multiple bitmap indexes to evaluate the conditions on a single table. Bitmap indexes are thus highly advantageous for complex *ad hoc* queries that contain lengthy `WHERE` clauses. Bitmap indexes can also provide optimal performance for aggregate queries and for optimizing joins in star schemas.

See Also: *Oracle9i Database Concepts* for more information on optimizing anti-joins and semi-joins

Comparing B-tree Indexes to Bitmap Indexes

Bitmap indexes can provide considerable storage savings over the use of B-tree indexes. In databases containing only B-tree indexes, you must anticipate the columns that are commonly accessed together in a single query, and create a composite B-tree index on these columns.

Not only would this B-tree index require a large amount of space, it would also be ordered. That is, a B-tree index on `(marital_status, region, gender)` is useless for queries that only access `region` and `gender`. To completely index the database, you must create indexes on the other permutations of these columns. For the simple case of three low-cardinality columns, there are six possible composite B-tree indexes. You must consider the trade-offs between disk space and performance needs when determining which composite B-tree indexes to create.

Bitmap indexes solve this dilemma. Bitmap indexes can be efficiently combined during query execution, so three small single-column bitmap indexes can do the job of six three-column B-tree indexes.

Bitmap indexes are much more efficient than B-tree indexes, especially in data warehousing environments. Bitmap indexes are created not only for efficient space usage but also for efficient execution, and the latter is somewhat more important.

Do not create bitmap indexes on unique key columns. However, for columns where each value is repeated hundreds or thousands of times, a bitmap index typically is less than 25% of the size of a regular B-tree index. The bitmaps themselves are stored in compressed format.

Simply comparing the relative sizes of B-tree and bitmap indexes is not an accurate measure of effectiveness, however. Because of their different performance characteristics, keep B-tree indexes on high-cardinality columns, while creating bitmap indexes on low-cardinality columns.

Maintenance Considerations for Bitmap Indexes

Bitmap indexes benefit data warehousing applications, but they are not appropriate for OLTP applications with a heavy load of concurrent `INSERTS`, `UPDATES`, and `DELETES`. In a data warehousing environment, data is maintained usually by way of bulk inserts and updates. Index maintenance is deferred until the end of each DML operation. For example, when you insert 1000 rows, the inserted rows are placed into a sort buffer and then the updates of all 1000 index entries are batched. (This is why `SORT_AREA_SIZE` must be set properly for good performance with inserts and updates on bitmap indexes.) Thus, each bitmap segment is updated only once in each DML operation, even if more than one row in that segment changes.

Note: The sorts described previously are regular sorts and use the regular sort area, determined by `SORT_AREA_SIZE`. The `BITMAP_MERGE_AREA_SIZE` and `CREATE_BITMAP_AREA_SIZE` initialization parameters described in ["Initialization Parameters for Bitmap Indexing"](#) on page 4-17 affect only the specific operations indicated by the parameter names.

DML and DDL statements, such as `UPDATE`, `DELETE`, and `DROP TABLE`, affect bitmap indexes the same way they do traditional indexes; the consistency model is the same. A compressed bitmap for a key value is made up of one or more bitmap segments, each of which is at most half a block in size (although it can be smaller). The locking granularity is one such bitmap segment. This can affect performance in environments where many transactions make simultaneous updates. If numerous DML operations have caused increased index size and decreasing performance for

queries, then you can use the `ALTER INDEX ... REBUILD` statement to compact the index and restore efficient performance.

A B-tree index entry contains a single rowid. Therefore, when the index entry is locked, a single row is locked. With bitmap indexes, an entry can potentially contain a range of rowids. When a bitmap index entry is locked, the entire range of rowids is locked. The number of rowids in this range affects concurrency. As the number of rowids increases in a bitmap segment, concurrency decreases.

Locking issues affect DML operations and can affect heavy OLTP environments. Locking issues do not, however, affect query performance. As with other types of indexes, updating bitmap indexes is a costly operation. Nonetheless, for bulk inserts and updates where many rows are inserted or many updates are made in a single statement, performance with bitmap indexes can be better than with regular B-tree indexes.

Using Bitmap Indexes with Good Performance

This section discusses performance issues with bitmap indexes.

Using Hints with Bitmap Indexes

The `INDEX` hint works with bitmap indexes in the same way as with traditional indexes.

The `INDEX_COMBINE` hint identifies the most cost effective indexes for the optimizer. The optimizer recognizes all indexes that can potentially be combined, given the predicates in the `WHERE` clause. However, it might not be cost effective to use all of them. Oracle recommends using `INDEX_COMBINE` rather than `INDEX` for bitmap indexes, because it is a more versatile hint.

In deciding which of these hints to use, the optimizer includes nonhinted indexes that appear cost effective, as well as indexes named in the hint. If certain indexes are given as arguments for the hint, then the optimizer tries to use some combination of those particular bitmap indexes.

If the hint does not name indexes, then all indexes are considered hinted. Hence, the optimizer tries to combine as many as possible, given the `WHERE` clause, without regard to cost effectiveness. The optimizer always tries to use hinted indexes in the plan, regardless of whether it considers them cost effective.

See Also: [Chapter 5, "Optimizer Hints"](#) for more information on the `INDEX_COMBINE` hint

Performance Tips for Bitmap Indexes

When creating bitmap indexes, Oracle needs to consider the theoretical maximum number of rows that will fit in a data block. For this reason, to get optimal performance and disk space usage with bitmap indexes, consider the following tips:

- To make compressed bitmaps as small as possible, declare `NOT NULL` constraints on all columns that cannot contain null values.
- Fixed-length datatypes are more amenable to a compact bitmap representation than variable length datatypes.

See Also: [Chapter 9, "Using EXPLAIN PLAN"](#) for more information about bitmap `EXPLAIN PLAN` output

Mapping Bitmaps to Rowids Efficiently

Use SQL statements with the `ALTER TABLE` syntax to optimize the mapping of bitmaps to rowids. The `MINIMIZE RECORDS_PER_BLOCK` clause enables this optimization, and the `NOMINIMIZE RECORDS_PER_BLOCK` clause disables it.

When `MINIMIZE RECORDS_PER_BLOCK` is enabled, Oracle scans the table and determines the maximum number of records in any block and restricts this table to this maximum number. This enables bitmap indexes to allocate fewer bits for each block and results in smaller bitmap indexes. The block and record allocation restrictions that this statement places on the table are beneficial only to bitmap indexes. Therefore, Oracle does not recommend using this mapping on tables that are not heavily indexed with bitmap indexes.

See Also:

- ["Using Bitmap Indexes"](#) on page 4-12 for more information
- *Oracle9i SQL Reference* for syntax on `MINIMIZE` and `NOMINIMIZE`

Using Bitmap Indexes on Index-Organized Tables

The rowids used in bitmap indexes on index-organized tables are in a mapping table, not in the base table. The mapping table maintains a mapping of logical rowids (needed to access the index-organized table) to physical rowids (needed by the bitmap index code). Each index-organized table has one mapping table, used by all the bitmap indexes created on that table.

Note: Moving rows in an index-organized table does not make the bitmap indexes built on that index-organized table unusable.

See Also: *Oracle9i Database Concepts* for information on bitmap indexes and index-organized tables

Indexing Null Values

Bitmap indexes index nulls, whereas all other index types do not. Consider, for example, a table with `STATE` and `PARTY` columns, on which you want to perform the following query:

```
SELECT COUNT(*)
FROM people
WHERE state='CA'
      AND party !='D';
```

Indexing nulls enables a bitmap minus plan where bitmaps for party equal to D and NULL are subtracted from state bitmaps equal to CA. The `EXPLAIN PLAN` output looks like the following:

```
SELECT STATEMENT
  SORT AGGREGATE
    BITMAP CONVERSION COUNT
      BITMAP MINUS
        BITMAP MINUS
          BITMAP INDEX SINGLE VALUE STATE_BM
          BITMAP INDEX SINGLE VALUE PARTY_BM
          BITMAP INDEX SINGLE VALUE PARTY_BM
```

If a NOT NULL constraint exists on party, then the second minus operation (where party is null) is left out, because it is not needed.

Initialization Parameters for Bitmap Indexing

The following initialization parameters affect the use of bitmap indexes:

- `CREATE_BITMAP_AREA_SIZE` affects memory allocated for bitmap creation.
- `BITMAP_MERGE_AREA_SIZE` affects memory used to merge bitmaps from an index range scan.
- `SORT_AREA_SIZE` affects memory used when inserting or updating bitmap indexes.

See Also: *Oracle9i Database Reference* for more information on these parameters

Using Bitmap Access Plans on Regular B-tree Indexes

If there is at least one bitmap index on the table, then the optimizer considers using a bitmap access path using regular B-tree indexes for that table. This access path can involve combinations of B-tree and bitmap indexes, but might not involve any bitmap indexes at all. However, the optimizer does not generate a bitmap access path using a single B-tree index unless instructed to do so by a hint.

To use bitmap access paths for B-tree indexes, the rowids stored in the indexes must be converted to bitmaps. After such a conversion, the various Boolean operations available for bitmaps can be used. As an example, consider the following query, where there is a bitmap index on column *c1*, and regular B-tree indexes on columns *c2* and *c3*.

```
EXPLAIN PLAN FOR
SELECT COUNT(*)
FROM t
WHERE c1 = 2 AND c2 = 6
OR c3 BETWEEN 10 AND 20;

SELECT STATEMENT
  SORT AGGREGATE
    BITMAP CONVERSION COUNT
      BITMAP OR
        BITMAP AND
          BITMAP INDEX c1_ind SINGLE VALUE
          BITMAP CONVERSION FROM ROWIDS
            INDEX c2_ind RANGE SCAN
        BITMAP CONVERSION FROM ROWIDS
          SORT ORDER BY
            INDEX c3_ind RANGE SCAN
```

Note: This statement is executed by accessing indexes only, so no table access is necessary.

Here, a **COUNT** option for the **BITMAP CONVERSION** row source counts the number of rows matching the query. There are also conversions from rowids in the plan to generate bitmaps from the rowids retrieved from the B-tree indexes. The **ORDER BY**

sort appears in the plan because the conditions on column `c3` result in the return of more than one list of rowids from the B-tree index. These lists are sorted before being converted into a bitmap.

Bitmap Index Restrictions

Bitmap indexes have the following restrictions:

- For bitmap indexes with direct load, the `SORTED_INDEX` flag does not apply.
- Bitmap indexes are not considered by the rule-based optimizer.
- Bitmap indexes cannot be used for referential integrity checking.

Using Bitmap Join Indexes

In addition to a bitmap index on a single table, you can create a bitmap join index, which is a bitmap index for the join of two or more tables. A bitmap join index is a space-saving way to reduce the volume of data that must be joined, by performing restrictions in advance. For each value in a column of a table, a bitmap join index stores the rowids of corresponding rows in another table. In a data warehousing environment, the join condition is an equi-inner join between the primary key column(s) of the dimension tables and the foreign key column(s) in the fact table.

Bitmap join indexes are much more efficient in storage than materialized join views, an alternative for materializing joins in advance. This is because the materialized join views do not compress the rowids of the fact tables.

See Also: *Oracle9i Data Warehousing Guide* for examples and restrictions of bitmap join indexes

Using Domain Indexes

Domain indexes are built using the indexing logic supplied by a user-defined indextype. An indextype provides an efficient mechanism to access data that satisfy certain operator predicates. Typically, the user-defined indextype is part of an Oracle option, like the `Spatial` option. For example, the `SpatialIndextype` allows efficient search and retrieval of spatial data that overlap a given bounding box.

The cartridge determines the parameters you can specify in creating and maintaining the domain index. Similarly, the performance and storage characteristics of the domain index are presented in the specific cartridge documentation.

Refer to the appropriate cartridge documentation for information such as the following:

- What datatypes can be indexed?
- What indextypes are provided?
- What operators does the indextype support?
- How can the domain index be created and maintained?
- How do we efficiently use the operator in queries?
- What are the performance characteristics?

Note: You can also create index types with the `CREATE INDEXTYPE` statement.

See Also: *Oracle Spatial User's Guide and Reference* for information about the `SpatialIndextype`

Using Clusters

Clusters are groups of one or more tables that are physically stored together because they share common columns and usually are used together. Because related rows are physically stored together, disk access time improves.

To create a cluster, use the `CREATE CLUSTER` statement.

See Also: *Oracle9i Database Concepts* for more information on clusters

Follow these guidelines when deciding whether to cluster tables:

- Cluster tables that are accessed frequently by the application in join statements.
- Do not cluster tables if the application joins them only occasionally or modifies their common column values frequently. Modifying a row's cluster key value takes longer than modifying the value in an unclustered table, because Oracle might need to migrate the modified row to another block to maintain the cluster.
- Do not cluster tables if the application often performs full table scans of only one of the tables. A full table scan of a clustered table can take longer than a full

table scan of an unclustered table. Oracle is likely to read more blocks, because the tables are stored together.

- Cluster master-detail tables if you often select a master record and then the corresponding detail records. Detail records are stored in the same data block(s) as the master record, so they are likely still to be in memory when you select them, requiring Oracle to perform less I/O.
- Store a detail table alone in a cluster if you often select many detail records of the same master. This measure improves the performance of queries that select detail records of the same master, but does not decrease the performance of a full table scan on the master table. An alternative is to use an index organized table.
- Do not cluster tables if the data from all tables with the same cluster key value exceeds more than one or two Oracle blocks. To access a row in a clustered table, Oracle reads all blocks containing rows with that value. If these rows take up multiple blocks, then accessing a single row could require more reads than accessing the same row in an unclustered table.
- Do not cluster tables when the number of rows for each cluster key value varies significantly. This causes waste of space for the low cardinality key value; it causes collisions for the high cardinality key values. Collisions degrade performance.

Consider the benefits and drawbacks of clusters with respect to the needs of the application. For example, you might decide that the performance gain for join statements outweighs the performance loss for statements that modify cluster key values. You might want to experiment and compare processing times with the tables both clustered and stored separately.

See Also: *Oracle9i Database Administrator's Guide* for more information on creating clusters

Using Hash Clusters

Hash clusters group table data by applying a hash function to each row's cluster key value. All rows with the same cluster key value are stored together on disk. Consider the benefits and drawbacks of hash clusters with respect to the needs of the application. You might want to experiment and compare processing times with a particular table as it is stored in a hash cluster, and as it is stored alone with an index.

Follow these guidelines for choosing when to use hash clusters:

- Use hash clusters to store tables accessed frequently by SQL statements with `WHERE` clauses, if the `WHERE` clauses contain equality conditions that use the same column or combination of columns. Designate this column or combination of columns as the cluster key.
- Store a table in a hash cluster if you can determine how much space is required to hold all rows with a given cluster key value, including rows to be inserted immediately as well as rows to be inserted in the future.
- Do not store a table in a hash cluster if the application often performs full table scans and if you must allocate a great deal of space to the hash cluster in anticipation of the table growing. Such full table scans must read all blocks allocated to the hash cluster, even though some blocks might contain few rows. Storing the table alone reduces the number of blocks read by full table scans.
- Do not store a table in a hash cluster if the application frequently modifies the cluster key values. Modifying a row's cluster key value can take longer than modifying the value in an unclustered table, because Oracle might need to migrate the modified row to another block to maintain the cluster.

Storing a single table in a hash cluster can be useful, regardless of whether the table is joined frequently with other tables, as long as hashing is appropriate for the table based on the points in this list.

Optimizer Hints

Optimizer hints can be used with SQL statements to alter execution plans. This chapter explains how to use hints to force various approaches.

The chapter contains the following sections:

- [Understanding Optimizer Hints](#)
- [Using Optimizer Hints](#)

Understanding Optimizer Hints

Hints let you make decisions usually made by the optimizer. As an application designer, you might know information about your data that the optimizer does not know.

For example, you might know that a certain index is more selective for certain queries. Based on this information, you might be able to choose a more efficient execution plan than the optimizer. In such a case, use hints to force the optimizer to use the optimal execution plan.

You can use hints to specify the following:

- The optimization approach for a SQL statement
- The goal of the cost-based optimizer for a SQL statement
- The access path for a table accessed by the statement
- The join order for a join statement
- A join operation in a join statement

Note: The use of hints involves extra code that must be managed, checked, and controlled.

Hints provide a mechanism to direct the optimizer to choose a certain query execution plan based on the following criteria:

- Join order
- Join method
- Access path
- Parallelization

Hints (except for the `RULE` hint) invoke the cost-based optimizer (CBO). If you have not gathered statistics, then defaults are used.

See Also: [Chapter 3, "Gathering Optimizer Statistics"](#) for more information on default values

Specifying Hints

Hints apply only to the optimization of the statement block in which they appear. A statement block is any one of the following statements or parts of statements:

- A simple `SELECT`, `UPDATE`, or `DELETE` statement.
- A parent statement or subquery of a complex statement.
- A part of a compound query.

For example, a compound query consisting of two component queries combined by the `UNION` operator has two statement blocks, one for each component query. For this reason, hints in the first component query apply only to its optimization, not to the optimization of the second component query.

You can send hints for a SQL statement to the optimizer by enclosing them in a comment within the statement.

See Also: *Oracle9i SQL Reference* for more information on comments

A statement block can have only one comment containing hints. This comment can only follow the `SELECT`, `UPDATE`, or `DELETE` keyword.

Exception: The `APPEND` hint always follows the `INSERT` keyword, and the `PARALLEL` hint can follow the `INSERT` keyword.

The following syntax shows hints contained in both styles of comments that Oracle supports within a statement block.

```
{DELETE|INSERT|SELECT|UPDATE} /*+ hint [text] [hint[text]]... */
```

or

```
{DELETE|INSERT|SELECT|UPDATE} --+ hint [text] [hint[text]]...
```

where:

- `DELETE`, `INSERT`, `SELECT`, and `UPDATE` are keywords that begin a statement block. Comments containing hints can appear only after these keywords.
- `+` causes Oracle to interpret the comment as a list of hints. The plus sign must immediately follow the comment delimiter; no space is permitted.

- `hint` is one of the hints discussed in this section. If the comment contains multiple hints, then each hint must be separated from the others by at least one space.
- `text` is other commenting text that can be interspersed with the hints.

If you specify hints incorrectly, then Oracle ignores them but does not return an error:

- Oracle ignores hints if the comment containing them does not follow a `DELETE`, `SELECT`, or `UPDATE` keyword.
- Oracle ignores hints containing syntax errors, but considers other correctly specified hints within the same comment.
- Oracle ignores combinations of conflicting hints, but considers other hints within the same comment.
- Oracle ignores hints in all SQL statements in those environments that use PL/SQL version 1, such as Forms version 3 triggers, Oracle Forms 4.5, and Oracle Reports 2.5. These hints can be passed to the server, but the server ignores them.

The optimizer recognizes hints only when using the cost-based approach. If you include a hint (except the `RULE` hint) in a statement block, then the optimizer automatically uses the cost-based approach.

See Also:

- ["Using Optimizer Hints"](#) on page 5-6 for the syntax of each hint
- ["INDEX"](#) on page 5-12 and following sections, for conditions specific to index type

Specifying a Full Set of Hints

When using hints, in some cases, you might need to specify a full set of hints in order to ensure the optimal execution plan. For example, if you have a very complex query, which consists of many table joins, and if you specify only the `INDEX` hint for a given table, then the optimizer needs to determine the remaining access paths to be used, as well as the corresponding join methods. Therefore, even though you gave the `INDEX` hint, the optimizer might not necessarily use that hint, because the optimizer might have determined that the requested index cannot be used due to the join methods and access paths selected by the optimizer.

In [Example 5-1](#), the `ORDERED` hint specifies the exact join order to be used; the join methods to be used on the different tables are also specified.

Example 5–1 Specifying a Full Set of Hints

```

SELECT /*+ ORDERED INDEX (b, jl_br_balances_n1) USE_NL (j b)
        USE_NL (glcc glf) USE_MERGE (gp gsb) */
  b.application_id ,
  b.set_of_books_id ,
  b.personnel_id,
  p.vendor_id Personnel,
  p.segment1 PersonnelNumber,
  p.vendor_name Name
FROM   jl_br_journals j,
       jl_br_balances b,
       gl_code_combinations glcc,
       fnd_flex_values_vl glf,
       gl_periods gp,
       gl_sets_of_books gsb,
       po_vendors p
WHERE  ...

```

Note that the hints could have been also been in this format:

```

SELECT --- ORDERED INDEX (b, jl_br_balances_n1) USE_NL (j b)
        USE_NL (glcc glf) USE_MERGE (gp gsb)

```

Using Hints Against Views

By default, hints do not propagate inside a complex view. For example, if you specify a hint in a query that selects against a complex view, then that hint is not honored, because it is not pushed inside the view.

Note: If the view is a single-table, then the hint is not propagated.

Unless the hints are inside the base view, they might not be honored from a query against the view.

Local Hints Compared with Global Hints

Table hints (in other words, hints that specify a table) generally refer to tables in the DELETE, SELECT, or UPDATE statement in which the hint occurs, not to tables inside any views referenced by the statement. When you want to specify hints for tables that appear inside views, Oracle recommends using global hints instead of embedding the hint in the view. Any table hint described in this chapter can be transformed into a global hint by using an extended syntax for the table name.

Note: The SQL Analyze tool that is available with the Oracle Tuning Pack, provides a graphical user interface for working with optimizer hints. The Hint Wizard, a feature of SQL Analyze, helps you easily add or modify hints in SQL statements.

See Also:

- ["Global Hints"](#) on page 5-44 for information on how to create global hints
- *Database Tuning with the Oracle Tuning Pack* for more information on Oracle SQL Analyze

Using Optimizer Hints

Optimizer hints can be categorized as follows:

- [Hints for Optimization Approaches and Goals](#)
- [Hints for Access Paths](#)
- [Hints for Query Transformations](#)
- [Hints for Join Orders](#)
- [Hints for Join Operations](#)
- [Hints for Parallel Execution](#)
- [Additional Hints](#)

Hints for Optimization Approaches and Goals

The hints described in this section let you choose between the cost-based and the rule-based optimization approaches. With the cost-based approach, this also includes the goal of best throughput or best response time.

- [ALL_ROWS](#)
- [FIRST_ROWS\(n\)](#)
- [CHOOSE](#)
- [RULE](#)


If a SQL statement has a hint specifying an optimization approach and goal, then the optimizer uses the specified approach regardless of the presence or absence of statistics, the value of the `OPTIMIZER_MODE` initialization parameter, and the `OPTIMIZER_MODE` parameter of the `ALTER SESSION` statement.

Note: The optimizer goal applies only to queries submitted directly. Use hints to determine the access path for any SQL statements submitted from within PL/SQL. The `ALTER SESSION... SET OPTIMIZER_MODE` statement does not affect SQL that is run from within PL/SQL.

ALL_ROWS

The `ALL_ROWS` hint explicitly chooses the cost-based approach to optimize a statement block with a goal of best throughput (that is, minimum total resource consumption).

`all_rows_hint::=`



For example, the optimizer uses the cost-based approach to optimize this statement for best throughput:

```
SELECT /*+ ALL_ROWS */ employee_id, last_name, salary, job_id
FROM employees
WHERE employee_id = 7566;
```

FIRST_ROWS(n)

The hints `FIRST_ROWS(n)` (where *n* is any positive integer) or `FIRST_ROWS` instruct Oracle to optimize an individual SQL statement for fast response. `FIRST_ROWS(n)` affords greater precision, because it instructs Oracle to choose the plan that returns the first *n* rows most efficiently. The `FIRST_ROWS` hint, which optimizes for the best plan to return the first single row, is retained for backward compatibility and plan stability.

`first_rows_hint::=`



For example, the optimizer uses the cost-based approach to optimize this statement for best response time:

```
SELECT /*+ FIRST_ROWS(10) */ employee_id, last_name, salary, job_id
FROM employees
WHERE department_id = 20;
```

In this example each department contains many employees. The user wants the first 10 employees of department #20 to be displayed as quickly as possible.

The optimizer ignores this hint in DELETE and UPDATE statement blocks and in SELECT statement blocks that contain any of the following syntax:

- Set operators (UNION, INTERSECT, MINUS, UNION ALL)
- GROUP BY clause
- FOR UPDATE clause
- Aggregate functions
- DISTINCT operator
- ORDER BY clauses, when there is no index on the ordering columns

These statements cannot be optimized for best response time, because Oracle must retrieve all rows accessed by the statement before returning the first row. If you specify this hint in any of these statements, then the optimizer uses the cost-based approach and optimizes for best throughput.

If you specify either the ALL_ROWS or the FIRST_ROWS hint in a SQL statement, and if the data dictionary does not have statistics about tables accessed by the statement, then the optimizer uses default statistical values (such as allocated storage for such tables) to estimate the missing statistics and, subsequently, to choose an execution plan.


These estimates might not be as accurate as those gathered by the DBMS_STATS package. Therefore, use the DBMS_STATS package to gather statistics. If you specify hints for access paths or join operations along with either the ALL_ROWS or FIRST_ROWS hint, then the optimizer gives precedence to the access paths and join operations specified by the hints.

See Also: ["How the CBO Optimizes SQL Statements for Fast Response"](#) on page 1-9 for an explanation of the difference between FIRST_ROWS(*n*) and FIRST_ROWS

CHOOSE

The `CHOOSE` hint causes the optimizer to choose between the rule-based and cost-based approaches for a SQL statement. The optimizer bases its selection on the presence of statistics for the tables accessed by the statement. If the data dictionary has statistics for at least one of these tables, then the optimizer uses the cost-based approach and optimizes with the goal of best throughput. If the data dictionary does not have statistics for these tables, then it uses the rule-based approach.

choose_hint::=



For example:

```
SELECT /*+ CHOOSE */ employee_id, last_name, salary, job_id
FROM employees
WHERE employee_id = 7566;
```

RULE

rule_hint::=



For example:

```
SELECT /*+ RULE */
employee_id, last_name, salary, job_id
FROM employees
WHERE employee_id = 7566;
```

Note: Oracle Corporation strongly advises the use of cost-based optimization. Rule-based optimization will be deprecated in a future release.

Hints for Access Paths

Each hint described in this section suggests an access path for a table.

- `FULL`
- `ROWID`
- `CLUSTER`

- `HASH`
- `INDEX`
- `INDEX_ASC`
- `INDEX_COMBINE`
- `INDEX_JOIN`
- `INDEX_DESC`
- `INDEX_FFS`
- `NO_INDEX`
- `AND_EQUAL`

Specifying one of these hints causes the optimizer to choose the specified access path only if the access path is available based on the existence of an index or cluster and on the syntactic constructs of the SQL statement. If a hint specifies an unavailable access path, then the optimizer ignores it.

You must specify the table to be accessed exactly as it appears in the statement. If the statement uses an alias for the table, then use the alias rather than the table name in the hint. The table name within the hint should not include the schema name if the schema name is present in the statement.

Note: For access path hints, Oracle ignores the hint if you specify the `SAMPLE` option in the `FROM` clause of a `SELECT` statement.

See Also: *Oracle9i SQL Reference* for more information on the `SAMPLE` option

FULL

The `FULL` hint explicitly chooses a full table scan for the specified table.

full_hint::=

→ (/*+ → FULL → (→ table →) → */ →

where `table` specifies the name or alias of the table on which the full table scan is to be performed. If the statement does not use aliases, then the table name is the default alias.

For example:

```
SELECT /*+ FULL(e) */ employee_id, last_name
  FROM employees e
 WHERE last_name LIKE :b1;
```

Oracle performs a full table scan on the `employees` table to execute this statement, even if there is an index on the `last_name` column that is made available by the condition in the `WHERE` clause.

Note: Because the `employees` table has alias `e` the hint must refer to the table by its alias rather than by its name. Also, do not specify schema names in the hint even if they are specified in the `FROM` clause.

ROWID

The `ROWID` hint explicitly chooses a table scan by rowid for the specified table.

`rowid_hint::=`

→ (/*+ → ROWID → (→ table →) → */ →

where `table` specifies the name or alias of the table on which the table access by rowid is to be performed.

For example:

```
SELECT /*+ROWID(employees)*/ *
  FROM employees
 WHERE rowid > 'AAAAtkAABAAAFNTAAA' AND employee_id = 155;
```

CLUSTER

The `CLUSTER` hint explicitly chooses a cluster scan to access the specified table. It applies only to clustered objects.

`cluster_hint::=`

→ (/*+ → CLUSTER → (→ table →) → */ →

where `table` specifies the name or alias of the table to be accessed by a cluster scan.


For example:

```
SELECT /*+ CLUSTER */
employees.last_name, department_id
FROM employees, departments
WHERE department_id = 10
AND employees.department_id = departments.department_id;
```

HASH

The `HASH` hint explicitly chooses a hash scan to access the specified table. It applies only to tables stored in a cluster.

`hash_hint::=`




where `table` specifies the name or alias of the table to be accessed by a hash scan.

INDEX

The `INDEX` hint explicitly chooses an index scan for the specified table. You can use the `INDEX` hint for domain, B-tree, bitmap, and bitmap join indexes. However, Oracle recommends using `INDEX_COMBINE` rather than `INDEX` for bitmap indexes, because it is a more versatile hint.

`index_hint::=`



where:

- `table` specifies the name or alias of the table associated with the index to be scanned.
- `index` specifies an index on which an index scan is to be performed.

This hint can optionally specify one or more indexes:

- If this hint specifies a single available index, then the optimizer performs a scan on this index. The optimizer does not consider a full table scan or a scan on another index on the table.

- If this hint specifies a list of available indexes, then the optimizer considers the cost of a scan on each index in the list and then performs the index scan with the lowest cost. The optimizer can also choose to scan multiple indexes from this list and merge the results, if such an access path has the lowest cost. The optimizer does not consider a full table scan or a scan on an index not listed in the hint.
- If this hint specifies no indexes, then the optimizer considers the cost of a scan on each available index on the table and then performs the index scan with the lowest cost. The optimizer can also choose to scan multiple indexes and merge the results, if such an access path has the lowest cost. The optimizer does not consider a full table scan.

For example, consider this query that selects the name, height, and weight of all male patients in a hospital:

```
SELECT name, height, weight
FROM patients
WHERE sex = 'm';
```

Assume that there is an index on the `SEX` column and that this column contains the values `m` and `f`. If there are equal numbers of male and female patients in the hospital, then the query returns a relatively large percentage of the table's rows, and a full table scan is likely to be faster than an index scan. However, if a very small percentage of the hospital's patients are male, then the query returns a relatively small percentage of the table's rows, and an index scan is likely to be faster than a full table scan.

Barring the use of frequency histograms, the number of occurrences of each distinct column value is not available to the optimizer. The cost-based approach assumes that each value has an equal probability of appearing in each row. For a column having only two distinct values, the optimizer assumes each value appears in 50% of the rows, so the cost-based approach is likely to choose a full table scan rather than an index scan.

If you know that the value in the `WHERE` clause of the query appears in a very small percentage of the rows, then you can use the `INDEX` hint to force the optimizer to choose an index scan. In this statement, the `INDEX` hint explicitly chooses an index scan on the `sex_index`, the index on the `sex` column:


```
SELECT /*+ INDEX(patients sex_index) use sex_index because there are few
      male patients */ name, height, weight
FROM patients
WHERE sex = 'm';
```

The `INDEX` hint applies to `IN`-list predicates; it forces the optimizer to use the hinted index, if possible, for an `IN`-list predicate. Multicolumn `IN`-lists will not use an index.

INDEX_ASC

The `INDEX_ASC` hint explicitly chooses an index scan for the specified table. If the statement uses an index range scan, then Oracle scans the index entries in ascending order of their indexed values.

`index_asc_hint::=`




Each parameter serves the same purpose as in the `INDEX` hint.

Because Oracle's default behavior for a range scan is to scan index entries in ascending order of their indexed values, this hint does not specify anything more than the `INDEX` hint. However, you might want to use the `INDEX_ASC` hint to specify ascending range scans explicitly should the default behavior change.

INDEX_COMBINE

The `INDEX_COMBINE` hint explicitly chooses a bitmap access path for the table. If no indexes are given as arguments for the `INDEX_COMBINE` hint, then the optimizer uses whatever Boolean combination of bitmap indexes has the best cost estimate for the table. If certain indexes are given as arguments, then the optimizer tries to use some Boolean combination of those particular bitmap indexes.

`index_combine_hint::=`




For example:

```
SELECT /*+INDEX_COMBINE(employees salary_bmi hire_date_bmi)*/ *
FROM employees
WHERE salary < 50000 AND hire_date < '01-JAN-1990';
```

INDEX_JOIN

The `INDEX_JOIN` hint explicitly instructs the optimizer to use an index join as an access path. For the hint to have a positive effect, a sufficiently small number of indexes must exist that contain all the columns required to resolve the query.

`index_join_hint::=`



where:

- `table` specifies the name or alias of the table associated with the index to be scanned.
- `index` specifies an index on which an index scan is to be performed.


For example, the following query uses an index join to access the `employee_id` and `department_id` columns, both of which are indexed in the `employees` table.

```
SELECT /*+index_join(employees emp_emp_id_pk emp_department_ix)*/
       employee_id, department_id
FROM   employees
WHERE  department_id > 50;
```

INDEX_DESC

The `INDEX_DESC` hint explicitly chooses an index scan for the specified table. If the statement uses an index range scan, then Oracle scans the index entries in descending order of their indexed values. In a partitioned index, the results are in descending order within each partition.

`index_desc_hint::=`




Each parameter serves the same purpose as in the `INDEX` hint. For example:

```
SELECT /*+ INDEX_DESC(a ord_order_date_ix) */
       a.order_date, a.promotion_id, a.order_id
FROM   orders a
WHERE  a.order_date = :b1;
```

INDEX_FFS

The `INDEX_FFS` hint causes a fast full index scan to be performed rather than a full table scan.

`index_ffs_hint::=`



For example:


```
SELECT /*+ INDEX_FFS ( o order_pk ) */ COUNT(*)
FROM order_items l, orders o
WHERE l.order_id > 50
      AND l.order_id = o.order_id;
```

See Also: ["Full Scans"](#) on page 1-34

NO_INDEX

The `NO_INDEX` hint explicitly disallows a set of indexes for the specified table.

`no_index_hint::=`



- If this hint specifies a single available index, then the optimizer does not consider a scan on this index. Other indexes not specified are still considered.
- If this hint specifies a list of available indexes, then the optimizer does not consider a scan on any of the specified indexes. Other indexes not specified in the list are still considered.
- If this hint specifies no indexes, then the optimizer does not consider a scan on any index on the table. This behavior is the same as a `NO_INDEX` hint that specifies a list of all available indexes for the table.

The `NO_INDEX` hint applies to function-based, B-tree, bitmap, cluster, or domain indexes. If a `NO_INDEX` hint and an index hint (`INDEX`, `INDEX_ASC`, `INDEX_DESC`, `INDEX_COMBINE`, or `INDEX_FFS`) both specify the same indexes, then both the `NO_INDEX` hint and the index hint are ignored for the specified indexes and the optimizer considers the specified indexes.


For example:

```
SELECT /*+NO_INDEX(employees emp_empid)*/ employee_id
FROM employees
WHERE employee_id > 200;
```

AND_EQUAL

The `AND_EQUAL` hint explicitly chooses an execution plan that uses an access path that merges the scans on several single-column indexes.

`and_equal_hint::=`



where:

- `table` specifies the name or alias of the table associated with the indexes to be merged.
- `index` specifies an index on which an index scan is to be performed. You must specify at least two indexes. You cannot specify more than five.

Hints for Query Transformations

Each hint described in this section suggests a SQL query transformation.


- `USE_CONCAT`
- `NO_EXPAND`
- `REWRITE`
- `EXPAND_GSET_TO_UNION`
- `NOREWRITE`
- `MERGE`
- `NO_MERGE`
- `STAR_TRANSFORMATION`
- `FACT`
- `NO_FACT`

USE_CONCAT

The `USE_CONCAT` hint forces combined `OR` conditions in the `WHERE` clause of a query to be transformed into a compound query using the `UNION ALL` set operator. Generally, this transformation occurs only if the cost of the query using the concatenations is cheaper than the cost without them.

The `USE_CONCAT` hint turns off `IN`-list processing and `OR`-expands all disjunctions, including `IN`-lists.

`use_concat_hint::=`




For example:

```
SELECT /*+USE_CONCAT*/ *
FROM employees
WHERE employee_id > 50 OR salary < 50000;
```

NO_EXPAND

The `NO_EXPAND` hint prevents the cost-based optimizer from considering `OR`-expansion for queries having `OR` conditions or `IN`-lists in the `WHERE` clause. Usually, the optimizer considers using `OR` expansion and uses this method if it decides that the cost is lower than not using it.

`no_expand_hint::=`



For example:


```
SELECT /*+NO_EXPAND*/ *
FROM employees
WHERE employee_id = 50 OR employee_id = 100;
```

REWRITE

The `REWRITE` hint forces the cost-based optimizer to rewrite a query in terms of materialized views, when possible, without cost consideration. Use the `REWRITE` hint with or without a view list. If you use `REWRITE` with a view list and the list contains an eligible materialized view, then Oracle uses that view regardless of its cost.

Oracle does not consider views outside of the list. If you do not specify a view list, then Oracle searches for an eligible materialized view and always uses it regardless of its cost.

rewrite_hint::=



See Also:

- *Oracle9i Database Concepts* and *Oracle9i Advanced Replication* for more information on materialized views
- *Oracle9i Data Warehousing Guide* for more information on using REWRITE with materialized views

EXPAND_GSET_TO_UNION

The EXPAND_GSET_TO_UNION hint is used for queries containing grouping sets (such as queries with GROUP BY GROUPING SET or GROUP BY ROLLUP). The hint forces a query to be transformed into a corresponding query with UNION ALL of individual groupings.

expand_gset_to_union_hint::=



For example:

```
SELECT year, quarter, month, sum(sales)
FROM T
GROUP BY year, rollup(quarter, month)
```

is first transformed to

```
SELECT year, quarter, month, sum(sales)
FROM T
GROUP BY year, quarter, month UNION ALL
SELECT year, quarter, null, sum(sales)
FROM T
GROUP BY year, quarter UNION ALL
SELECT year, null, null, sum(sales)
FROM T
GROUP BY year
```

Then, for each branch of the `UNION ALL`, Oracle tries a rewrite with a materialized view. The rewrite may do a joinback and rollup of the materialized view. Finally, Oracle looks at the branches not rewritten and tries to represent them as a single query block with grouping sets. So for example, if only the last branch of the `UNION ALL` was rewritten with materialized view `MV`, Oracle replaces the first two branches with a the equivalent `GROUPING SET` query, as follows:

```
SELECT year, quarter, month, sum(sales)
FROM T
GROUP BY grouping set ( (year, quarter, month), (year, quarter) ) UNION ALL
SELECT year, null, null, sum_sales
FROM MV
```

NOREWRITE

The `NOREWRITE` hint disables query rewrite for the query block, overriding the setting of the parameter `QUERY_REWRITE_ENABLED`. Use the `NOREWRITE` hint on any query block of a request.

Note: The `NOREWRITE` hint disables the use of function-based indexes.

norewrite_hint::=



MERGE

The `MERGE` hint lets you merge a view for each query.

If a view's query contains a `GROUP BY` clause or `DISTINCT` operator in the `SELECT` list, then the optimizer can merge the view's query into the accessing statement only if complex view merging is enabled. Complex merging can also be used to merge an `IN` subquery into the accessing statement if the subquery is uncorrelated.

Complex merging is not cost-based; that is, the accessing query block must include the `MERGE` hint. Without this hint, the optimizer uses another approach.

merge_hint::=



For example:


```
SELECT /*+MERGE(v)*/ e1.last_name, e1.salary, v.avg_salary
FROM employees e1,
     (SELECT department_id, avg(salary) avg_salary
      FROM employees e2
      GROUP BY department_id) v
WHERE e1.department_id = v.department_id AND e1.salary > v.avg_salary;
```

Note: This example requires that complex view merging be enabled.

NO_MERGE

The `NO_MERGE` hint causes Oracle not to merge mergeable views.

`no_merge_hint::=`



This hint lets the user have more influence over the way in which the view is accessed.

For example:

```
SELECT /*+NO_MERGE(dallas_dept)*/ e1.last_name, dallas_dept.dname
FROM employees e1,
     (SELECT department_id, dname
      FROM departments
      WHERE loc = 'DALLAS') dallas_dept
WHERE e1.department_id = dallas_dept.department_id;
```

This causes view `dallas_dept` not to be merged.

When the `NO_MERGE` hint is used without an argument, it should be placed in the view query block. When `NO_MERGE` is used with the view name as an argument, it should be placed in the surrounding query.

STAR_TRANSFORMATION

The `STAR_TRANSFORMATION` hint makes the optimizer use the best plan in which the transformation has been used. Without the hint, the optimizer could make a cost-based decision to use the best plan generated without the transformation, instead of the best plan for the transformed query.

Even if the hint is given, there is no guarantee that the transformation will take place. The optimizer only generates the subqueries if it seems reasonable to do so. If no subqueries are generated, then there is no transformed query, and the best plan for the untransformed query is used, regardless of the hint.

star_transformation_hint::=




See Also:

- *Oracle9i Database Concepts* for a full discussion of star transformation.
- *Oracle9i Database Reference* for more information on the `STAR_TRANSFORMATION_ENABLED` initialization parameter.

FACT

The `FACT` hint is used in the context of the star transformation to indicate to the transformation that the hinted table should be considered as a fact table.


fact_hint::=



NO_FACT

The `NO_FACT` hint is used in the context of the star transformation to indicate to the transformation that the hinted table should not be considered as a fact table.

no_fact_hint::=



Hints for Join Orders

The hints in this section suggest join orders:

- **ORDERED**
- **STAR**

ORDERED

The `ORDERED` hint causes Oracle to join tables in the order in which they appear in the `FROM` clause.

If you omit the `ORDERED` hint from a SQL statement performing a join, then the optimizer chooses the order in which to join the tables. You might want to use the `ORDERED` hint to specify a join order if you know something about the number of rows selected from each table that the optimizer does not. Such information lets you choose an inner and outer table better than the optimizer could.

ordered_hint::=



The following query is an example of the use of the `ORDERED` hint:

```
SELECT /*+ORDERED */ o.order_id, c.customer_id, l.unit_price * l.quantity
  FROM customers c, order_items l, orders o
 WHERE c.cust_last_name = :b1
       AND o.customer_id = c.customer_id
       AND o.order_id = l.order_id;
```

STAR

The `STAR` hint forces a star query plan to be used, if possible. A star plan has the largest table in the query last in the join order and joins it with a nested loops join on a concatenated index. The `STAR` hint applies when there are at least three tables, the large table's concatenated index has at least three columns, and there are no conflicting access or join method hints. The optimizer also considers different permutations of the small tables.

star_hint::=



Usually, if you analyze the tables, then the optimizer selects an efficient star plan. You can also use hints to improve the plan. The most precise method is to order the tables in the `FROM` clause in the order of the keys in the index, with the large table last. Then use the following hints:

```
/*+ ORDERED USE_NL(FACTS) INDEX(facts fact_concat) */
```

where `facts` is the table and `fact_concat` is the index. A more general method is to use the `STAR` hint.

Hints for Join Operations

Each hint described in this section suggests a join operation for a table.

- `USE_NL`
- `USE_MERGE`
- `USE_HASH`
- `DRIVING_SITE`
- `LEADING`
- `HASH_AJ, MERGE_AJ, and NL_AJ`
- `HASH_SJ, MERGE_SJ, and NL_SJ`


In the hint you must specify a table exactly the same way it appears in the statement. If the statement uses an alias for the table, then you must use the alias rather than the table name in the hint. However, the table name within the hint should *not* include the schema name, if the schema name is present in the statement.

Use of the `USE_NL` and `USE_MERGE` hints is recommended with the `ORDERED` hint. Oracle uses these hints when the referenced table is forced to be the inner table of a join; the hints are ignored if the referenced table is the outer table.

USE_NL

The `USE_NL` hint causes Oracle to join each specified table to another row source with a nested loops join, using the specified table as the inner table.

`use_nl_hint::=`



where `table` is the name or alias of a table to be used as the inner table of a nested loops join.

For example, consider this statement, which joins the `accounts` and `customers` tables. Assume that these tables are not stored together in a cluster:

```
SELECT accounts.balance, customers.last_name, customers.first_name
   FROM accounts, customers
   WHERE accounts.customer_id = customers.customer_id;
```

Because the default goal of the cost-based approach is best throughput, the optimizer chooses either a nested loops operation, a sort-merge operation, or a hash

operation to join these tables, depending on which is likely to return all the rows selected by the query more quickly.

However, you might want to optimize the statement for best response time or the minimal elapsed time necessary to return the first row selected by the query, rather than best throughput. If so, then you can force the optimizer to choose a nested loops join by using the `USE_NL` hint. In this statement, the `USE_NL` hint explicitly chooses a nested loops join with the `customers` table as the inner table:

```
SELECT /*+ ORDERED USE_NL(customers) to get first row faster */
      accounts.balance, customers.last_name, customers.first_name
FROM accounts, customers
WHERE accounts.customer_id = customers.customer_id;
```

In many cases, a nested loops join returns the first row faster than a sort merge join. A nested loops join can return the first row after reading the first selected row from one table and the first matching row from the other and combining them, while a sort merge join cannot return the first row until after reading and sorting all selected rows of both tables and then combining the first rows of each sorted row source.

In the following statement where a nested loop is forced through a hint, `orders` is accessed through a full table scan and the filter condition `l.order_id = h.order_id` is applied to every row. For every row that meets the filter condition, `order_items` is accessed through the index `order_id`.


```
SELECT /*+ USE_NL(l h) */ h.customer_id, l.unit_price * l.quantity
      FROM orders h ,order_items l
      WHERE l.order_id = h.order_id;
```

Adding an `INDEX` hint to the query could avoid the full table scan on `orders`, resulting in an execution plan similar to one used on larger systems, even though it might not be particularly efficient here.

USE_MERGE

The `USE_MERGE` hint causes Oracle to join each specified table with another row source, using a sort-merge join.

`use_merge_hint::=`



where `table` is a table to be joined to the row source resulting from joining the previous tables in the join order using a sort merge join.

For example:

```
SELECT /*+USE_MERGE(employees departments)*/ *
FROM employees, departments
WHERE employees.department_id = departments.department_id;
```

The following query shows an inventory usage report in which the optimizer avoids a sort for the `GROUP BY` operation by using the sort merge operation specified by the `USE_MERGE` hint.

```
SELECT /*+ USE_MERGE(inv l) */inv.product_id, SUM(l.quantity)
FROM inventories inv, order_items l
WHERE inv.product_id = l.product_id(+)
GROUP BY inv.product_id;
```


The following is a query applying the `USE_MERGE` hint with the `FULL` hint.

```
SELECT /*+USE_MERGE(h l) FULL(h l) */ h.customer_id, l.unit_price * l.quantity
FROM orders h ,order_items l
WHERE l.order_id = h.order_id;
```

USE_HASH

The `USE_HASH` hint causes Oracle to join each specified table with another row source, using a hash join.

`use_hash_hint::=`



where `table` is a table to be joined to the row source resulting from joining the previous tables in the join order using a hash join.

For example:

```
SELECT /*+USE_HASH(l l2) */ l.order_date, l.order_id, l2.product_id,
SUM(l2.unit_price*quantity)
FROM orders l, order_items l2
WHERE l.order_id = l2.order_id
GROUP BY l2.product_id, l.order_date, l.order_id;
```


Another example:

```
SELECT /*+use_hash(employees departments)*/ *
FROM employees, departments
WHERE employees.department_id = departments.department_id;
```

DRIVING_SITE

The `DRIVING_SITE` hint forces query execution to be done at a different site than that selected by Oracle. This hint can be used with either rule-based or cost-based optimization.

driving_site_hint:=



where `table` is the name or alias for the table at which site the execution should take place.

For example:

```
SELECT /*+DRIVING_SITE(departments)*/ *
FROM employees, departments@rsite
WHERE employees.department_id = departments.department_id;
```

If this query is executed without the hint, then rows from `departments` are sent to the local site, and the join is executed there. With the hint, the rows from `employees` are sent to the remote site, and the query is executed there, returning the result to the local site.


This hint is useful if you are using distributed query optimization.

LEADING

The `LEADING` hint causes Oracle to use the specified table as the first table in the join order.

If you specify two or more `LEADING` hints on different tables, then all of them are ignored. If you specify the `ORDERED` hint, then it overrides all `LEADING` hints.

leading_hint:=



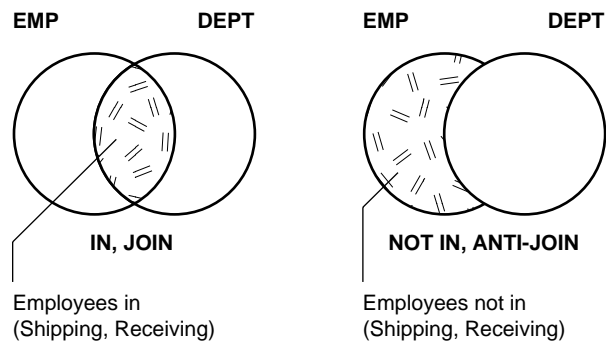
where `table` is the name or alias of a table to be used as the first table in the join order.

HASH_AJ, MERGE_AJ, and NL_AJ

For a specific query, place the `MERGE_AJ`, `HASH_AJ`, or `NL_AJ` hint into the `NOT IN` subquery. `MERGE_AJ` uses a sort-merge anti-join, `HASH_AJ` uses a hash anti-join, and `NL_AJ` uses a nested loop anti-join.

As illustrated in [Figure 5–1](#), the SQL `IN` predicate can be evaluated using a join to intersect two sets. Thus, `employees.department_id` can be joined to `departments.department_id` to yield a list of employees in a set of departments.

Figure 5–1 Parallel Hash Anti-join



Alternatively, the SQL predicate `NOT IN` can be evaluated using an anti-join to subtract two sets. Thus, `employees.department_id` can be anti-joined to `departments.department_id` to select all employees who are not in a set of departments, and you can get a list of all employees who are not in the shipping or receiving departments.

HASH_SJ, MERGE_SJ, and NL_SJ

For a specific query, place the `HASH_SJ`, `MERGE_SJ`, or `NL_SJ` hint into the `EXISTS` subquery. `HASH_SJ` uses a hash semi-join, `MERGE_SJ` uses a sort merge semi-join, and `NL_SJ` uses a nested loop semi-join.

For example:

```
SELECT * FROM departments
WHERE exists (SELECT /*+HASH_SJ*/ *
             FROM employees
             WHERE employees.department_id = departments.department_id
                AND salary > 200000);
```


This converts the subquery into a special type of join between t_1 and t_2 that preserves the semantics of the subquery. That is, even if there is more than one matching row in t_2 for a row in t_1 , the row in t_1 is returned only once.

A subquery is evaluated as a semi-join only with these limitations:

- There can only be one table in the subquery.
- The outer query block must not itself be a subquery.
- The subquery must be correlated with an equality predicate.
- The subquery must have no GROUP BY, CONNECT BY, or ROWNUM references.

See Also: *Oracle9i SQL Reference* for more information about joins

Hints for Parallel Execution

The hints described in this section determine how statements are parallelized or not parallelized when using parallel execution.

- [PARALLEL](#)
- [NOPARALLEL](#)
- [PQ_DISTRIBUTE](#)
- [PARALLEL_INDEX](#)
- [NOPARALLEL_INDEX](#)

See Also: *Oracle9i Data Warehousing Guide* for more information on parallel execution


PARALLEL

The `PARALLEL` hint lets you specify the desired number of concurrent servers that can be used for a parallel operation. The hint applies to the `SELECT`, `INSERT`, `UPDATE`, and `DELETE` portions of a statement, as well as to the table scan portion.

Note: The number of servers that can be used is twice the value in the `PARALLEL` hint, if sorting or grouping operations also take place.

If any parallel restrictions are violated, then the hint is ignored.

parallel_hint::=



The `PARALLEL` hint must use the table alias, if an alias is specified in the query. The hint can then take two values, separated by commas after the table name. The first value specifies the degree of parallelism for the given table, and the second value specifies how the table is to be split among the Oracle Real Application Clusters instances. Specifying `DEFAULT` or no value signifies that the query coordinator should examine the settings of the initialization parameters to determine the default degree of parallelism. In the following example, the `PARALLEL` hint overrides the degree of parallelism specified in the `employees` table definition:

```
SELECT /*+ FULL(hr_emp) PARALLEL(hr_emp, 5) */ last_name
FROM hr.employees hr_emp;
```


In the next example, the `PARALLEL` hint overrides the degree of parallelism specified in the `employees` table definition and tells the optimizer to use the default degree of parallelism determined by the initialization parameters. This hint also specifies that the table should be split among all of the available instances, with the of parallelism on each instance.

```
SELECT /*+ FULL(hr_emp) PARALLEL(hr_emp, DEFAULT,DEFAULT) */ last_name
FROM hr.employees hr_emp;
```

NOPARALLEL

The `NOPARALLEL` hint overrides a `PARALLEL` specification in the table clause. In general, hints take precedence over table clauses.

noparallel_hint::=



The following example illustrates the `NOPARALLEL` hint:


```
SELECT /*+ NOPARALLEL(hr_emp) */ last_name
FROM hr.employees hr_emp;
```

PQ_DISTRIBUTE

The `PQ_DISTRIBUTE` hint improves the performance of parallel join operations. Do this by specifying how rows of joined tables should be distributed among producer and consumer query servers. Using this hint overrides decisions the optimizer would normally make.

Use the `EXPLAIN PLAN` statement to identify the distribution chosen by the optimizer. The optimizer ignores the distribution hint, if both tables are serial.

`pq_distribute_hint::=`



where:

- `table_name` is the name or alias of a table to be used as the inner table of a join.
- `outer_distribution` is the distribution for the outer table.
- `inner_distribution` is the distribution for the inner table.

See Also: *Oracle9i Database Concepts* for more information on how Oracle parallelizes join operations

There are six combinations for table distribution. Only a subset of distribution method combinations for the joined tables is valid, as explained in [Table 5-1](#).

Table 5-1 *Distribution Hint Combinations*

Distribution	Interpretation
Hash, Hash	Maps the rows of each table to consumer query servers, using a hash function on the join keys. When mapping is complete, each query server performs the join between a pair of resulting partitions. This hint is recommended when the tables are comparable in size and the join operation is implemented by hash-join or sort merge join.
Broadcast, None	All rows of the outer table are broadcast to each query server. The inner table rows are randomly partitioned. This hint is recommended when the outer table is very small compared to the inner table. As a general rule, use the Broadcast/None hint when <i>inner table size * number of query servers > outer table size</i> .

Table 5–1 (Cont.) Distribution Hint Combinations

Distribution	Interpretation
None, Broadcast	All rows of the inner table are broadcast to each consumer query server. The outer table rows are randomly partitioned. This hint is recommended when the inner table is very small compared to the outer table. As a general rule, use the None/Broadcast hint when <i>inner table size * number of query servers < outer table size</i> .
Partition, None	Maps the rows of the outer table, using the partitioning of the inner table. The inner table must be partitioned on the join keys. This hint is recommended when the number of partitions of the outer table is equal to or nearly equal to a multiple of the number of query servers; for example, 14 partitions and 15 query servers. Note: The optimizer ignores this hint if the inner table is not partitioned or not equijoined on the partitioning key.
None, Partition	Maps the rows of the inner table using the partitioning of the outer table. The outer table must be partitioned on the join keys. This hint is recommended when the number of partitions of the outer table is equal to or nearly equal to a multiple of the number of query servers; for example, 14 partitions and 15 query servers. Note: The optimizer ignores this hint if the outer table is not partitioned or not equijoined on the partitioning key.
None, None	Each query server performs the join operation between a pair of matching partitions, one from each table. Both tables must be equipartitioned on the join keys.

For example: Given two tables, *r* and *s*, that are joined using a hash-join, the following query contains a hint to use hash distribution:

```
SELECT column_list /*+ORDERED PQ_DISTRIBUTE(s HASH, HASH) USE_HASH (s) */
FROM r,s
WHERE r.c=s.c;
```


To broadcast the outer table *r*, the query is:

```
SELECT column list /*+ORDERED PQ_DISTRIBUTE(s BROADCAST, NONE) USE_HASH (s) */
FROM r,s
WHERE r.c=s.c;
```

PARALLEL_INDEX

The `PARALLEL_INDEX` hint specifies the desired number of concurrent servers that can be used to parallelize index range scans for partitioned indexes.

`parallel_index_hint::=`



where:

- `table` is the name or alias of the table associated with the index to be scanned.
- `index` is the index on which an index scan is to be performed (optional).

The hint can take two values, separated by commas after the table name. The first value specifies the degree of parallelism for the given table. The second value specifies how the table is to be split among the Oracle Real Application Clusters instances. Specifying `DEFAULT` or no value signifies the query coordinator should examine the settings of the initialization parameters to determine the default degree of parallelism.

For example:


```
SELECT /*+ PARALLEL_INDEX(table1, index1, 3, 2) +/
```

In this example, there are three parallel execution processes to be used on each of two instances.

NOPARALLEL_INDEX

The `NOPARALLEL_INDEX` hint overrides a `PARALLEL` attribute setting on an index to avoid a parallel index scan operation.

`noparallel_index_hint::=`



Additional Hints

Several additional hints are included in this section:

- APPEND
- NOAPPEND
- CACHE
- NOCACHE
- UNNEST
- NO_UNNEST
- PUSH_PRED
- NO_PUSH_PRED
- PUSH_SUBQ
- NO_PUSH_SUBQ
- ORDERED_PREDICATES
- CURSOR_SHARING_EXACT
- DYNAMIC_SAMPLING

APPEND

The `APPEND` hint lets you enable direct-path `INSERT` if your database is running in serial mode. Your database is in serial mode if you are not using Enterprise Edition. Conventional `INSERT` is the default in serial mode, and direct-path `INSERT` is the default in parallel mode.

In direct-path `INSERT`, data is appended to the end of the table, rather than using existing space currently allocated to the table. As a result, direct-path `INSERT` can be considerably faster than conventional `INSERT`.

`append_hint::=`



NOAPPEND

The `NOAPPEND` hint enables conventional `INSERT` by disabling parallel mode for the duration of the `INSERT` statement. (Conventional `INSERT` is the default in serial mode, and direct-path `INSERT` is the default in parallel mode).


noappend_hint::=



CACHE

The **CACHE** hint specifies that the blocks retrieved for the table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This option is useful for small lookup tables.

cache_hint::=




In the following example, the **CACHE** hint overrides the table's default caching specification:

```
SELECT /*+ FULL (hr_emp) CACHE(hr_emp) */ last_name
FROM hr.employees hr_emp;
```

NOCACHE

The **NOCACHE** hint specifies that the blocks retrieved for the table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed. This is the normal behavior of blocks in the buffer cache.

nocache_hint::=



For example:

```
SELECT /*+ FULL(hr_emp) NOCACHE(hr_emp) */ last_name
FROM hr.employees hr_emp;
```

Note: The **CACHE** and **NOCACHE** hints affect system statistics "table scans(long tables)" and "table scans(short tables)", as shown in the **V\$SYSSTAT** view.

Automatic Caching of Small Tables Starting with Oracle9i, Release 2 (9.2), small tables are automatically cached, according to the criteria in [Table 5-2](#).

Table 5–2 Table Caching Criteria

Table Size	Size Criteria	Caching
Small	Number of blocks < 20 or 2% of total cached blocks, whichever is larger	Always cached
Medium	Larger than a small table, but < 10% of total cached blocks	Oracle decides whether to cache a table on the basis of its table scan and workload history. It caches the table only if a future table scan is likely to find the cached blocks.
Large	> 10% of total cached blocks	Not cached

Automatic caching of small tables is disabled for tables that are created or altered with the `CACHE` attribute.

UNNEST

The `UNNEST` hint specifies subquery unnesting. Subquery unnesting unnests and merges the body of the subquery into the body of the statement that contains it, allowing the optimizer to consider them together when evaluating access paths and joins.

If the `UNNEST` hint is used, Oracle first verifies if the statement is valid. If the statement is not valid, then subquery unnesting cannot proceed. The statement must then must pass a heuristic test.

The `UNNEST` hint tells Oracle to check the subquery block for validity only. If the subquery block is valid, then subquery unnesting is enabled without Oracle's checking the heuristics.

See Also:

- *Oracle9i SQL Reference* for more information on unnesting nested subqueries and the conditions that make a subquery block valid
- "[Subquery Unnesting](#)" on page 1-12 and "[How the CBO Unnests Subqueries](#)" on page 2-33 for more information on the use of unnesting


unnest_hint::=



NO_UNNEST

Use of the `NO_UNNEST` hint turns off unnesting for specific subquery blocks.


no_unnest_hint::=



PUSH_PRED

The `PUSH_PRED` hint forces pushing of a join predicate into the view.

push_pred_hint::=



For example:

```

SELECT /*+ PUSH_PRED(v) */ t1.x, v.y
FROM t1
      (SELECT t2.x, t3.y
       FROM t2, t3
       WHERE t2.x = t3.x) v
WHERE t1.x = v.x and t1.y = 1;
  
```

NO_PUSH_PRED

The `NO_PUSH_PRED` hint prevents pushing of a join predicate into the view.

no_push_pred_hint::=




PUSH_SUBQ

The `PUSH_SUBQ` hint causes non-merged subqueries to be evaluated at the earliest possible step in the execution plan. Generally, subqueries that are not merged are executed as the last step in the execution plan. If the subquery is relatively inexpensive and reduces the number of rows significantly, then it improves performance to evaluate the subquery earlier.

This hint has no effect if the subquery is applied to a remote table or one that is joined using a merge join.


`push_subq_hint::=`



NO_PUSH_SUBQ

The `NO_PUSH_SUBQ` hint causes non-merged subqueries to be evaluated as the last step in the execution plan. If the subquery is relatively expensive or does not reduce the number of rows significantly, then it improves performance to evaluate the subquery last.

`no_push_subq_hint::=`



ORDERED_PREDICATES

The `ORDERED_PREDICATES` hint forces the optimizer to preserve the order of predicate evaluation, except for predicates used as index keys. Use this hint in the `WHERE` clause of `SELECT` statements.

If you do not use the `ORDERED_PREDICATES` hint, then Oracle evaluates all predicates in the following order:

1. Predicates without user-defined functions, type methods, or subqueries are evaluated first, in the order specified in the `WHERE` clause.
2. Predicates with user-defined functions and type methods that have user-computed costs are evaluated next, in increasing order of their cost.
3. Predicates with user-defined functions and type methods without user-computed costs are evaluated next, in the order specified in the `WHERE` clause.
4. Predicates not specified in the `WHERE` clause (for example, predicates transitively generated by the optimizer) are evaluated next.
5. Predicates with subqueries are evaluated last, in the order specified in the `WHERE` clause.

Note: Remember, you cannot use the `ORDERED_PREDICATES` hint to preserve the order of predicate evaluation on index keys.

`ordered_predicates_hint::=`




See Also: *Oracle9i Database Concepts*

CURSOR_SHARING_EXACT

Oracle can replace literals in SQL statements with bind variables, if it is safe to do so. This is controlled with the `CURSOR_SHARING` startup parameter. The `CURSOR_SHARING_EXACT` hint causes this behavior to be switched off. In other words, Oracle executes the SQL statement without any attempt to replace literals by bind variables.


`cursor_sharing_exact_hint::=`



DYNAMIC_SAMPLING

The `DYNAMIC_SAMPLING` hint lets you control dynamic sampling to improve server performance by determining more accurate selectivity and cardinality estimates. You can set the value of `DYNAMIC_SAMPLING` to a value from 0 to 10. The higher the level, the more effort the compiler puts into dynamic sampling and the more broadly it is applied. Sampling defaults to cursor level unless you specify a table.

`dynamic_sampling_hint::=`



where:

- `table` specifies the name or alias of the table on which the dynamic sampling is to be performed.
- `integer` is a value from 0 to 10 indicating the degree of sampling. If the statement does not use aliases, then the table name is the default alias.

For example:

```
SELECT /*+ dynamic_sampling(1) */ *  
FROM ...
```

enables dynamic sampling if all of the following conditions are true:

- There is more than one table in the query.
- Some table has not been analyzed and has no indexes.
- The optimizer determines that a relatively expensive table scan would be required for this table that has not been analyzed.

The sampling levels are as follows if the dynamic sampling level used is from a cursor hint or from the `optimizer_dynamic_sampling` parameter:

- Level 0: Do not use dynamic sampling.
- Level 1: Sample all tables that have not been analyzed if the following criteria are met: (1) there is at least 1 unanalyzed table in the query; (2) this unanalyzed table is joined to another table or appears in a subquery or non-mergeable view; (3) this unanalyzed table has no indexes; (4) this unanalyzed table has more blocks than the number of blocks that would be used for dynamic sampling of this table. The number of blocks sampled is the default number of dynamic sampling blocks (32).
- Level 2: Apply dynamic sampling to all unanalyzed tables. The number of blocks sampled is the default number of dynamic sampling blocks.
- Level 3: Apply dynamic sampling to all tables that meet Level 2 criteria, plus all tables for which standard selectivity estimation used a guess for some predicate that is a potential dynamic sampling predicate. The number of blocks sampled is the default number of dynamic sampling blocks.
- Level 4: Apply dynamic sampling to all tables that meet Level 3 criteria, plus all tables that have single-table predicates that reference 2 or more columns. The number of blocks sampled is the default number of dynamic sampling blocks.
- Level 5: Apply dynamic sampling to all tables that meet the Level 4 criteria using 2 times the default number of dynamic sampling blocks.
- Level 6: Apply dynamic sampling to all tables that meet the Level 5 criteria using 4 times the default number of dynamic sampling blocks.
- Level 7: Apply dynamic sampling to all tables that meet the Level 6 criteria using 8 times the default number of dynamic sampling blocks.

- Level 8: Apply dynamic sampling to all tables that meet the Level 7 criteria using 32 times the default number of dynamic sampling blocks.
- Level 9: Apply dynamic sampling to all tables that meet the Level 8 criteria using 128 times the default number of dynamic sampling blocks.
- Level 10: Apply dynamic sampling to all tables that meet the Level 9 criteria using all blocks in the table.

The sampling levels are as follows if the dynamic sampling level used is from a table hint:

- Level 0: Do not use dynamic sampling.
- Level 1: The number of blocks sampled is the default number of dynamic sampling blocks (32).
- Level 2: The number of blocks sampled is 2 times the default number of dynamic sampling blocks.
- Level 3: The number of blocks sampled is 4 times the default number of dynamic sampling blocks.
- Level 4: The number of blocks sampled is 8 times the default number of dynamic sampling blocks.
- Level 5: The number of blocks sampled is 16 times the default number of dynamic sampling blocks.
- Level 6: The number of blocks sampled is 32 times the default number of dynamic sampling blocks.
- Level 7: The number of blocks sampled is 64 times the default number of dynamic sampling blocks.
- Level 8: The number of blocks sampled is 128 times the default number of dynamic sampling blocks.
- Level 9: The number of blocks sampled is 256 times the default number of dynamic sampling blocks.
- Level 10: Read all blocks in the table.

To apply dynamic sampling to a specific table, use the following form of the hint:

```
SELECT /*+ dynamic_sampling(employees 1) */ *  
FROM employees  
WHERE ...
```

If there is a table hint, dynamic sampling is used unless the table is analyzed and there are no predicates on the table. For example, the following query will not result in any dynamic sampling if `employees` is analyzed:

```
SELECT /*+ dynamic_sampling(e 1) */ count(*)  
FROM employees e;
```

The cardinality statistic is used, if it exists. If there is a predicate, dynamic sampling is done with a table hint and cardinality is not estimated.

To force cardinality estimation even for an analyzed table, you can use a further hint, `dynamic_sampling_est_cdn`, as in the following example:

```
SELECT /*+ dynamic_sampling(e 1) dynamic_sampling_est_cdn(t) */ count(*)  
FROM employees e;
```

This forces cardinality estimation for `employees`, even if the table is analyzed. The following query does both selectivity and cardinality estimation for `employees`:

```
SELECT /*+ dynamic_sampling(e 1) dynamic_sampling_est_cdn(e) */ count(*)  
FROM employees e  
WHERE cols > 3;
```

Using Hints with Views

Oracle does not encourage the use of hints inside or on views (or subqueries). This is because you can define views in one context and use them in another. However, such hints can result in unexpected execution plans. In particular, hints inside views or on views are handled differently, depending on whether the view is mergeable into the top-level query.

If you decide, nonetheless, to use hints with views, the following sections describe the behavior in each case.

- [Hints and Mergeable Views](#)
- [Hints and Nonmergeable Views](#)

If you want to specify a hint for a table in a view or subquery, then the global hint syntax is recommended. The following section describes this in detail.

- [Global Hints](#)

Hints and Mergeable Views

This section describes hint behavior with mergeable views.

Optimization Approaches and Goal Hints Optimization approach and goal hints can occur in a top-level query or inside views.

- If there is such a hint in the top-level query, then that hint is used regardless of any such hints inside the views.
- If there is no top-level optimizer mode hint, then mode hints in referenced views are used as long as all mode hints in the views are consistent.
- If two or more mode hints in the referenced views conflict, then all mode hints in the views are discarded and the session mode is used, whether default or user-specified.

Access Path and Join Hints on Views Access path and join hints on referenced views are ignored, unless the view contains a single table (or references an [Additional Hints](#) view with a single table). For such single-table views, an access path hint or a join hint on the view applies to the table inside the view.

Access Path and Join Hints Inside Views Access path and join hints can appear in a view definition.

- If the view is a subquery (that is, if it appears in the `FROM` clause of a `SELECT` statement), then all access path and join hints inside the view are preserved when the view is merged with the top-level query.
- For views that are not subqueries, access path and join hints in the view are preserved only if the top-level query references no other tables or views (that is, if the `FROM` clause of the `SELECT` statement contains only the view).

Parallel Execution Hints on Views `PARALLEL`, `NOPARALLEL`, `PARALLEL_INDEX`, and `NOPARALLEL_INDEX` hints on views are applied recursively to all the tables in the referenced view. Parallel execution hints in a top-level query override such hints inside a referenced view.

Parallel Execution Hints Inside Views `PARALLEL`, `NOPARALLEL`, `PARALLEL_INDEX`, and `NOPARALLEL_INDEX` hints inside views are preserved when the view is merged with the top-level query. Parallel execution hints on the view in a top-level query override such hints inside a referenced view.

Hints and Nonmergeable Views

With nonmergeable views, optimization approach and goal hints inside the view are ignored; the top-level query decides the optimization mode.

Because nonmergeable views are optimized separately from the top-level query, access path and join hints inside the view are preserved. For the same reason, access path hints on the view in the top-level query are ignored.

However, join hints on the view in the top-level query are preserved because, in this case, a nonmergeable view is similar to a table.

Global Hints

Table hints (hints that specify a table) normally refer to tables in the `DELETE`, `SELECT`, or `UPDATE` statement in which the hint occurs, not to tables inside any views referenced by the statement. When you want to specify hints for tables that appear inside views, use global hints instead of embedding the hint in the view. You can transform any table hint in this chapter into a global hint by using an extended syntax for the table name, described as follows.

Consider the following view definitions and `SELECT` statement:

```
CREATE OR REPLACE VIEW v1 AS
  SELECT *
  FROM employees
  WHERE employee_id < 150;

CREATE OR REPLACE VIEW v2 AS
  SELECT v1.employee_id employee_id, departments.department_id department_id
  FROM v1, departments
  WHERE v1.department_id = departments.department_id;

SELECT /*+ INDEX( v2.v1.employees emp_emp_id_pk ) FULL(v2.departments) */ *
  FROM v2
  WHERE department_id = 30;
```

The view `V1` retrieves all employees whose employee number is less than 150. The view `V2` performs a join between the view `V1` and the department table. The `SELECT` statement retrieves rows from the view `V2` restricting it to the department whose number is 30.

There are two global hints in the `SELECT` statement. The first hint specifies an index scan for the employee table referenced in the view `V1`, which is referenced in the view `V2`. The second hint specifies a full table scan for the department table referenced in the view `V2`. Note the dotted syntax for the view tables.

A hint such as:

```
INDEX(employees emp_emp_id_pk)
```


in the `SELECT` statement is ignored because the employee table does not appear in the `FROM` clause of the `SELECT` statement.

The global hint syntax also applies to unmergeable views. Consider the following `SELECT` statement:

```
SELECT /*+ NO_MERGE(v2) INDEX(v2.v1.employees emp_emp_id_pk)
        FULL(v2.departments) */ *
FROM v2
WHERE department_id = 30;
```

It causes `V2` not to be merged and specifies access path hints for the employee and department tables. These hints are pushed down into the (nonmerged) view `V2`.

If a global hint references a `UNION` or `UNION ALL` view, then the hint is applied to the first branch that contains the hinted table. Consider the `INDEX` hint in the following `SELECT` statement:

```
SELECT /*+ INDEX(v.employees emp_emp_id_pk) */ *
FROM (SELECT *
      FROM employees
      WHERE employee_id < 150
      UNION ALL
      SELECT *
      FROM employees
      WHERE employee_id > 175) v
WHERE department_id = 30;
```

The `INDEX` hint applies to the employee table in the first branch of the `UNION ALL` view `v`, not to the employee table in the second branch.

Optimizing SQL Statements

This chapter describes how to identify high-resource SQL statements, explains what should be collected, and provides tuning suggestions.

This chapter contains the following sections:

- [Goals for Tuning](#)
- [Identifying and Gathering Data on Resource-Intensive SQL](#)
- [Dynamic Sampling](#)
- [Overview of SQL Statement Tuning](#)

Goals for Tuning

The objective of tuning a system is either to reduce the response time for end users of the system, or to reduce the resources used to process the same work. You can accomplish both of these objectives in several ways:

- [Reduce the Workload](#)
- [Balance the Workload](#)
- [Parallelize the Workload](#)

Reduce the Workload

This is what commonly constitutes SQL tuning: finding more efficient ways to process the same workload. It is possible to change the execution plan of the statement without altering the functionality to reduce the resource consumption.

Two examples of how resource usage can be reduced are:

1. If a commonly executed query needs to access a small percentage of data in the table, then it can be executed more efficiently by using an index. By creating such an index, you reduce the amount of resources used.
2. If a user is looking at the first twenty rows of the 10,000 rows returned in a specific sort order, and if the query (and sort order) can be satisfied by an index, then the user does not need to access and sort the 10,000 rows to see the first 20 rows.

Balance the Workload

Systems often tend to have peak usage in the daytime when real users are connected to the system, and low usage in the nighttime. If noncritical reports and batch jobs can be scheduled to run in the nighttime and their concurrency during day time reduced, then it frees up resources for the more critical programs in the day.

Parallelize the Workload

Queries that access large amounts of data (typical data warehouse queries) often can be parallelized. This is extremely useful for reducing the response time in low concurrency data warehouse. However, for OLTP environments, which tend to be high concurrency, this can adversely impact other users by increasing the overall resource usage of the program.

Identifying and Gathering Data on Resource-Intensive SQL

This section describes the steps involved in identifying and gathering data on poorly-performing SQL statements.

Identifying Resource-Intensive SQL

The first step in identifying resource-intensive SQL is to categorize the problem you are attempting to fix: is the problem specific to a single program (or small number of programs), or is the problem generic over the application?

Tuning a Specific Program

If you are tuning a specific program (GUI or 3GL), then identifying the SQL to examine is a simple matter of looking at the SQL executed within the program.

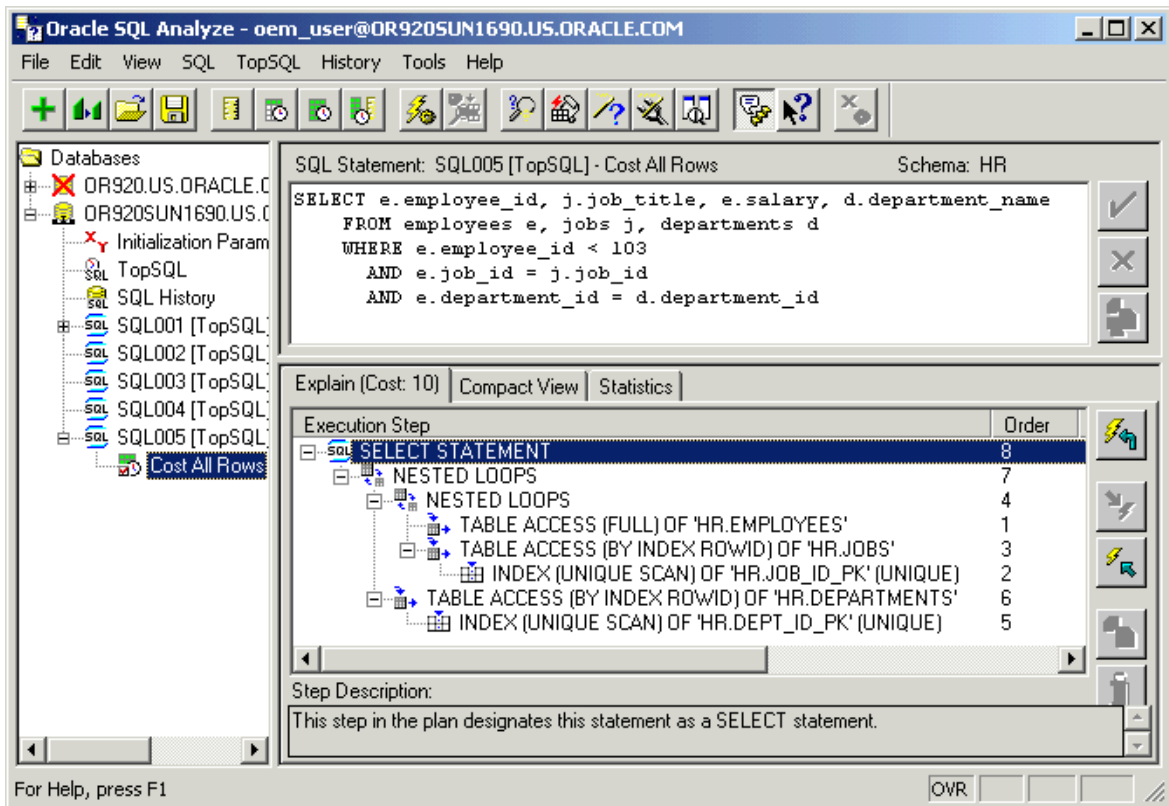
If it is not possible to identify the SQL (for example, the SQL is generated dynamically), then use `SQL_TRACE` to generate a trace file that contains the SQL executed, then use `TKPROF` to generate an output file.

The SQL statements in the `TKPROF` output file can be ordered by various parameters, such as the execution elapsed time (`exeela`), which usually assists in the identification by ordering the SQL statements by elapsed time (with highest elapsed time SQL statements at the top of the file). This makes the job of identifying the poorly performing SQL easier if there are many SQL statements in the file.

See Also: [Chapter 10, "Using SQL Trace and TKPROF"](#)

Oracle SQL Analyze can be used for identifying resource intensive SQL statements, generating explain plans, and evaluating SQL performance. [Figure 6-1](#) is an illustration of SQL Analyze displaying the SQL statement used in [Example 1-3, "Using EXPLAIN PLAN"](#) on page 1-19.

Figure 6–1 Oracle SQL Analyze



See Also: For more information on Oracle SQL Analyze, see the *Database Tuning with the Oracle Tuning Pack* manual

Tuning an Application / Reducing Load

If your whole application is performing suboptimally, or if you are attempting to reduce the overall CPU or I/O load on the database server, then identifying resource-intensive SQL involves the following steps:

1. Determine which period in the day you would like to examine; typically this is the application's peak processing time.
2. Gather operating system and Oracle statistics over that period. The minimum of Oracle statistics gathered should be file I/O (V\$FILESTAT), system statistics

(V\$SYSSTAT), and SQL statistics (V\$SQLAREA or V\$SQL, V\$SQLTEXT and V\$SQL_PLAN).

See Also: [Chapter 21, "Using Statspack"](#) for information on how to use Statspack to gather Oracle instance performance data for you

3. Using the data collected in step two, identify the SQL statements using the most resources. A good way to identify candidate SQL statements is to query V\$SQLAREA. V\$SQLAREA contains resource usage information for all SQL statements in the shared pool. The data in V\$SQLAREA should be ordered by resource usage. The most common resources are:
 - Buffer gets (V\$SQLAREA.BUFFER_GETS, for high CPU using statements)
 - Disk reads (V\$SQLAREA.DISK_READS, for high I/O statements)
 - Sorts (V\$SQLAREA.SORTS, for many sorts)

One method to identify which SQL statements are creating the highest load is to compare the resources used by a SQL statement to the total amount of that resource used in the period. For BUFFER_GETS, divide each SQL statement's BUFFER_GETS by the total number of buffer gets during the period. The total number of buffer gets in the system is available in the V\$SYSSTAT table, for the statistic session logical reads.

Similarly, it is possible to apportion the percentage of disk reads a statement performs out of the total disk reads performed by the system by dividing V\$SQLAREA.DISK_READS by the value for the V\$SYSSTAT statistic physical reads. The SQL sections of the Statspack report include this data, so you do not need to perform the percentage calculations manually.

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#) for more information on V\$SQLAREA and V\$SQL

After you have identified the candidate SQL statements, the next stage is to gather information that enables you to examine the statements and tune them.

Gathering Data on the SQL Identified

If you are most concerned with CPU, then examine the top SQL statements that performed the most BUFFER_GETS during that interval. Otherwise, start with the SQL statement that performed the most DISK_READS.

Information to Gather During Tuning

The tuning process begins by determining the structure of the underlying tables and indexes.

Information gathered includes the following:

1. Complete SQL text from `V$SQLTEXT`
2. Structure of the tables referenced in the SQL statement (usually by describing the table in SQL*Plus)
3. Definitions of any indexes (columns, column orderings), and whether the indexes are unique or nonunique
4. CBO statistics for the segments (including the number of rows each table, selectivity of the index columns), including the date when the segments were last analyzed
5. Definitions of any views referred to in the SQL statement
6. Repeat steps two and three for any tables referenced in the view definitions found in step four
7. Optimizer plan for the SQL statement (either from `EXPLAIN PLAN`, `V$SQL_PLAN`, or the `TKPROF` output)
8. Any previous optimizer plans for that SQL statement

Note: It is important to generate and review execution plans for all of the key SQL statements in your application. Doing so lets you compare the optimizer execution plans of a SQL statement when the statement performed well to the plan when that the statement is not performing well. Having the comparison, along with information such as changes in data volumes, can assist in identifying the cause of performance degradation.

Dynamic Sampling

The purpose of dynamic sampling is to improve server performance by determining more accurate selectivity and cardinality estimates. More accurate selectivity and cardinality estimates allow the optimizer to produce better performing plans.

You can use dynamic sampling to:

- Estimate single-table predicate selectivities when collected statistics cannot be used or are likely to lead to significant errors in estimation.
- Estimate table cardinality for tables without statistics or for tables whose statistics are too out of date to trust.

How Dynamic Sampling Works

The primary performance attribute is compile time. Oracle determines at compile time whether a query would benefit from dynamic sampling. If so, a recursive SQL statement is issued to scan a small random sample of the table's blocks, and to apply the relevant single table predicates to estimate predicate selectivities. The sample cardinality can also be used, in some cases, to estimate table cardinality.

Depending on the value of the `OPTIMIZER_DYNAMIC_SAMPLING` initialization parameter, a certain number of blocks are read by the dynamic sampling query.

When to Use Dynamic Sampling

For a query that normally completes quickly (in less than a few seconds), you will not want to incur the cost of dynamic sampling. However, dynamic sampling can be beneficial under any of the following conditions:

- A better plan can be found using dynamic sampling.
- The sampling time is a small fraction of total execution time for the query.
- The query will be executed many times.

Dynamic sampling can be applied to a subset of a single table's predicates and combined with standard selectivity estimates of predicates for which dynamic sampling is not done.

How to Use Dynamic Sampling to Improve Performance

You control dynamic sampling with the `OPTIMIZER_DYNAMIC_SAMPLING` parameter, which can be set to a value from 0 to 10.

- A value of 0 means dynamic sampling will not be done.
- A value of 1 (the default) means dynamic sampling will be performed if all of the following conditions are true:
 - There is more than one table in the query.
 - Some table has not been analyzed and has no indexes.

- The optimizer determines that a relatively expensive table scan would be required for this unanalyzed table.
- Increasing the value of the parameter results in more aggressive application of dynamic sampling, in terms of both the type of tables sampled (analyzed or unanalyzed) and the amount of I/O spent on sampling.

Dynamic sampling is repeatable if no rows have been inserted, deleted, or updated in the table being sampled.

The parameter `OPTIMIZER_FEATURES_ENABLE` turns off dynamic sampling if set to a version prior to 9.0.2.

See Also: ["DYNAMIC_SAMPLING"](#) on page 5-39 for details about using this hint

Overview of SQL Statement Tuning

This section describes ways you can improve SQL statement efficiency:

- [Verifying Optimizer Statistics](#)
- [Reviewing the Execution Plan](#)
- [Restructuring the SQL Statements](#)
- [Restructuring the Indexes](#)
- [Modifying or Disabling Triggers and Constraints](#)
- [Restructuring the Data](#)
- [Maintaining Execution Plans Over Time](#)
- [Visiting Data as Few Times as Possible](#)

Note: The guidelines described in this section are oriented to production SQL that will be executed frequently. Most of the techniques that are discouraged here can legitimately be employed in ad hoc statements or in applications run infrequently where performance is not critical.

Verifying Optimizer Statistics

The CBO uses statistics gathered on tables and indexes when determining the optimal execution plan. If these statistics have not been gathered, or if the statistics

are no longer representative of the data stored within the database, then the optimizer does not have sufficient information to generate the best plan.

Things to check:

- If you gather statistics for some tables in your database, then it is probably best to gather statistics for all tables. This is especially true if your application includes SQL statements that perform joins.
- If the optimizer statistics in the data dictionary are no longer representative of the data in the tables and indexes, then gather new statistics. One way to check whether the dictionary statistics are stale is to compare the real cardinality (row count) of a table to the value of `DBA_TABLES.NUM_ROWS`. If there is significant data skew on predicate columns, then consider using histograms.

Reviewing the Execution Plan

When tuning (or writing) a SQL statement in an OLTP environment, the goal is to drive from the table that has the most selective filter. This means that there are fewer rows passed to the next step. If the next step is a join, then this means that fewer rows are joined. Check to see whether the access paths are optimal.

When examining the optimizer execution plan, look for the following:

- The plan is such that the driving table has the best filter.
- The join order in each step means that the fewest number of rows are being returned to the next step (that is, the join order should reflect, where possible, going to the best not-yet-used filters).
- The join method is appropriate for the number of rows being returned. For example, nested loop joins through indexes may not be optimal when many rows are being returned.
- Views are used efficiently. Look at the `SELECT` list to see whether access to the view is necessary.
- There are any unintentional Cartesian products (even with small tables).
- Each table is being accessed efficiently:

Consider the predicates in the SQL statement and the number of rows in the table. Look for suspicious activity, such as a full table scans on tables with large number of rows, which have predicates in the where clause. Determine why an index is not used for such a selective predicate.

A full table scan does not mean inefficiency. It might be more efficient to perform a full table scan on a small table, or to perform a full table scan to leverage a better join method (for example, `hash_join`) for the number of rows returned.

If any of these conditions are not optimal, then consider restructuring the SQL statement or the indexes available on the tables.

Restructuring the SQL Statements

Often, rewriting an inefficient SQL statement is easier than repairing it. If you understand the purpose of a given statement, then you might be able to quickly and easily write a new statement that meets the requirement.

Compose Predicates Using AND and =

To improve SQL efficiency, use equijoins whenever possible. Statements that perform equijoins on untransformed column values are the easiest to tune.

Avoid Transformed Columns in the WHERE Clause

Use untransformed column values. For example, use:

```
WHERE a.order_no = b.order_no
```

rather than:

```
WHERE TO_NUMBER (SUBSTR(a.order_no, INSTR(b.order_no, '.') - 1))  
= TO_NUMBER (SUBSTR(a.order_no, INSTR(b.order_no, '.') - 1))
```

Do not use SQL functions in predicate clauses or `WHERE` clauses. Any expression using a column, such as a function having the column as its argument, causes the optimizer to ignore the possibility of using an index on that column, even a unique index, unless there is a function-based index defined that can be used.

Avoid mixed-mode expressions, and beware of implicit type conversions. When you want to use an index on the `VARCHAR2` column `charcol`, but the `WHERE` clause looks like this:

```
AND charcol = numexpr
```

where `numexpr` is an expression of number type (for example, 1, `USERENV('SESSIONID')`, `numcol`, `numcol+0`,...), Oracle translates that expression into:

```
AND TO_NUMBER(charcol) = numexpr
```

Avoid the following kinds of complex expressions:

- `col1 = NVL (:b1,col1)`
- `NVL (col1, -999) =`
- `TO_DATE(), TO_NUMBER(),` and so on

These expressions prevent the optimizer from assigning valid cardinality or selectivity estimates and can in turn affect the overall plan and the join method.

Add the predicate versus using `NVL()` technique.

For example:

```
SELECT employee_num, full_name Name, employee_id
FROM mtl_employees_current_view
WHERE (employee_num = NVL (:b1,employee_num)) AND (organization_id=:1)
ORDER BY employee_num;
```

Also:

```
SELECT employee_num, full_name Name, employee_id
FROM mtl_employees_current_view
WHERE (employee_num = :b1) AND (organization_id=:1)
ORDER BY employee_num;
```

When you need to use SQL functions on filters or join predicates, do not use them on the columns on which you want to have an index; rather, use them on the opposite side of the predicate, as in the following statement:

```
TO_CHAR(numcol) = varcol
```

rather than

```
varcol = TO_CHAR(numcol)
```

See Also: [Chapter 4, "Understanding Indexes and Clusters"](#) for more information on function-based indexes

Write Separate SQL Statements for Specific Tasks

SQL is not a procedural language. Using one piece of SQL to do many different things usually results in a less-than-optimal result for each task. If you want SQL to accomplish different things, then write various statements, rather than writing one statement to do different things depending on the parameters you give it.

Note: Oracle Forms and Reports are powerful development tools that allow application logic to be coded using PL/SQL (triggers or program units). This helps reduce the complexity of SQL by allowing complex logic to be handled in the Forms or Reports. You can also invoke a server side PL/SQL package that performs the few SQL statements in place of a single large complex SQL statement. Because the package is a server-side unit, there are no issues surrounding client to database round-trips and network traffic.

It is always better to write separate SQL statements for different tasks, but if you must use one SQL statement, then you can make a very complex statement slightly less complex by using the UNION ALL operator.

Optimization (determining the execution plan) takes place before the database knows what values will be substituted into the query. An execution plan cannot, therefore, depend on what those values are. For example:

```
SELECT info
FROM tables
WHERE ...
      AND somecolumn BETWEEN DECODE(:loval, 'ALL', somecolumn, :loval)
      AND DECODE(:hival, 'ALL', somecolumn, :hival);
```

Written as shown, the database cannot use an index on the `somecolumn` column, because the expression involving that column uses the same column on both sides of the `BETWEEN`.

This is not a problem if there is some other highly selective, indexable condition you can use to access the driving table. Often, however, this is not the case. Frequently, you might want to use an index on a condition like that shown but need to know the values of `:loval`, and so on, in advance. With this information, you can rule out the `ALL` case, which should *not* use the index.

If you want to use the index whenever real values are given for `:loval` and `:hival` (if you expect narrow ranges, even ranges where `:loval` often equals `:hival`), then you can rewrite the example in the following logically equivalent form:

```
SELECT /* change this half of UNION ALL if other half changes */ info
FROM tables
WHERE ...
      AND somecolumn BETWEEN :loval AND :hival
      AND (:hival != 'ALL' AND :loval != 'ALL')
```

```
UNION ALL
SELECT /* Change this half of UNION ALL if other half changes. */ info
FROM tables
WHERE ...
      AND (:hival = 'ALL' OR :loval = 'ALL');
```

If you run `EXPLAIN PLAN` on the new query, then you seem to get both a desirable and an undesirable execution plan. However, the first condition the database evaluates for either half of the `UNION ALL` is the combined condition on whether `:hival` and `:loval` are `ALL`. The database evaluates this condition before actually getting any rows from the execution plan for that part of the query.

When the condition comes back false for one part of the `UNION ALL` query, that part is not evaluated further. Only the part of the execution plan that is optimum for the values provided is actually carried out. Because the final conditions on `:hival` and `:loval` are guaranteed to be mutually exclusive, only one half of the `UNION ALL` actually returns rows. (The `ALL` in `UNION ALL` is logically valid because of this exclusivity. It allows the plan to be carried out without an expensive sort to rule out duplicate rows for the two halves of the query.)

Use of EXISTS versus IN for Subqueries

In certain circumstances, it is better to use `IN` rather than `EXISTS`. In general, if the selective predicate is in the subquery, then use `IN`. If the selective predicate is in the parent query, then use `EXISTS`.

Note: This discussion is most applicable in an OLTP environment, where the access paths either to the parent SQL or subquery are through indexed columns with high selectivity. In a DSS environment, there can be low selectivity in the parent SQL or subquery, and there might not be any indexes on the join columns. In a DSS environment, consider using semi-joins for the `EXISTS` case.

See Also:

- ["How the CBO Executes Anti-joins"](#) on page 1-43
- ["HASH_AJ, MERGE_AJ, and NL_AJ"](#) on page 5-28 and ["HASH_SJ, MERGE_SJ, and NL_SJ"](#) on page 5-28
- *Oracle9i Data Warehousing Guide*

Sometimes, Oracle can rewrite a subquery when used with an `IN` clause to take advantage of selectivity specified in the subquery. This is most beneficial when the most selective filter appears in the subquery and there are indexes on the join columns. Conversely, using `EXISTS` is beneficial when the most selective filter is in the parent query. This allows the selective predicates in the parent query to be applied before filtering the rows against the `EXISTS` criteria.

Note: You should verify the CBO cost of the statement with the actual number of resources used (`BUFFER_GETS`, `DISK_READS`, `CPU_TIME` from `V$SQL` or `V$SQLAREA`). Situations such as data skew (without the use of histograms) can adversely affect the optimizer's estimated cost for an operation.

Below are two examples that demonstrate the benefits of `IN` and `EXISTS`. Both examples use the same schema with the following characteristics:

- There is a unique index on the `employees.employee_id` field.
- There is an index on the `orders.customer_id` field.
- There is an index on the `employees.department_id` field.
- The `employees` table has 27,000 rows.
- The `orders` table has 10,000 rows.
- The `OE` and `HR` schemas, which own these segments, were both analyzed with `COMPUTE`.

Example 1: Using `IN` - Selective Filters in the Subquery This example demonstrates how rewriting a query to use `IN` can improve performance. This query identifies all employees who have placed orders on behalf of customer 144.

The following SQL statement uses `EXISTS`:

```
SELECT /* EXISTS example */
       e.employee_id
       , e.first_name
       , e.last_name
       , e.salary
FROM   employees e
WHERE  EXISTS (SELECT 1 FROM orders o
              WHERE e.employee_id = o.sales_rep_id
              AND o.customer_id = 144);
/* Note 1 */
/* Note 2 */
/* Note 3 */
```


Notes:

- Note 1: This shows the line containing EXISTS.
- Note 2: This shows the line that makes the subquery a correlated subquery.
- Note 3: This shows the line where the correlated subqueries include the highly selective predicate `customer_id = number`.

Below is the execution plan (from `V$SQL_PLAN`) for the preceding statement. The plan requires a full table scan of the `employees` table, returning many rows. Each of these rows is then filtered against the `orders` table (through an index).

ID	OPERATION	OPTIONS	OBJECT_NAME	OPT	COST
0	SELECT STATEMENT			CHO	
1	FILTER				
2	TABLE ACCESS	FULL	EMPLOYEES	ANA	155
3	TABLE ACCESS	BY INDEX ROWID	ORDERS	ANA	3
4	INDEX	RANGE SCAN	ORD_CUSTOMER_IX	ANA	1

Rewriting the statement using `IN` results in significantly fewer resources used.

The SQL statement using `IN`:

```
SELECT /* IN example */
       e.employee_id
       , e.first_name
       , e.last_name
       , e.salary
FROM   employees e
WHERE  e.employee_id IN (SELECT o.sales_rep_id          /* Note 4 */
                        FROM orders o
                        WHERE o.customer_id = 144); /* Note 3 */
```

Note:

- Note 3: This shows the line where the correlated subqueries include the highly selective predicate `customer_id = number`
 - Note 4: This indicates that an `IN` is being used. The subquery is no longer correlated, because the `IN` clause replaces the join in the subquery.
-
-

Below is the execution plan (from `V$SQL_PLAN`) for the preceding statement. The optimizer rewrites the subquery into a view, which is then joined through a unique index to the `employees` table. This results in a significantly better plan, because the view (that is, subquery) has a selective predicate, thus returning only a few `employee_ids`. These few `employee_ids` are then used to access the `employees` table through the unique index.

ID	OPERATION	OPTIONS	OBJECT_NAME	OPT	COST
0	SELECT STATEMENT			CHO	
1	NESTED LOOPS				5
2	VIEW				3
3	SORT	UNIQUE			3
4	TABLE ACCESS	FULL	ORDERS	ANA	1
5	TABLE ACCESS	BY INDEX ROWID	EMPLOYEES	ANA	1
6	INDEX	UNIQUE SCAN	EMP_EMP_ID_PK	ANA	

Example 2: Using EXISTS - Selective Predicate in the Parent This example demonstrates how rewriting a query to use `EXISTS` can improve performance. This query identifies all employees from department 80 who are sales reps who have placed orders.

The following SQL statement uses `IN`:

```

SELECT /* IN example */
       e.employee_id
       , e.first_name
       , e.last_name
       , e.department_id
       , e.salary
FROM   employees  e
WHERE  e.department_id = 80                                /* Note 5 */
       AND e.job_id      = 'SA_REP'                       /* Note 6 */
       AND e.employee_id IN (SELECT o.sales_rep_id FROM orders o); /* Note 4 */

```

Note:

- Note 4: This indicates that an `IN` is being used. The subquery is no longer correlated, because the `IN` clause replaces the join in the subquery.
- Note 5 and 6: These are the selective predicates in the parent SQL.

Below is the execution plan (from `V$SQL_PLAN`) for the preceding statement. The SQL statement was rewritten by the optimizer to use a view on the `orders` table, which requires sorting the data to return all unique `employee_ids` existing in the `orders` table. Because there is no predicate, many `employee_ids` are returned. The large list of resulting `employee_ids` are then used to access the `employees` table through the unique index.

ID	OPERATION	OPTIONS	OBJECT_NAME	OPT	COST
0	SELECT STATEMENT			CHO	
1	NESTED LOOPS				125
2	VIEW				116
3	SORT	UNIQUE			116
4	TABLE ACCESS	FULL	ORDERS	ANA	40
5	TABLE ACCESS	BY INDEX ROWID	EMPLOYEES	ANA	1
6	INDEX	UNIQUE SCAN	EMP_EMP_ID_PK	ANA	

The following SQL statement uses `EXISTS`:

```
SELECT /* EXISTS example */
       e.employee_id
       , e.first_name
       , e.last_name
       , e.salary
FROM   employees e
WHERE  e.department_id = 80                /* Note 5 */
       AND e.job_id      = 'SA_REP'        /* Note 6 */
       AND EXISTS (SELECT 1
                   FROM orders o
                   WHERE e.employee_id = o.sales_rep_id); /* Note 2 */
```

Note:

- Note 1: This shows the line containing EXISTS.
 - Note 2: This shows the line that makes the subquery a correlated subquery.
 - Note 5 & 6: These are the selective predicates in the parent SQL.
-
-

Below is the execution plan (from V\$SQL_PLAN) for the preceding statement. The cost of the plan is reduced by rewriting the SQL statement to use an EXISTS. This plan is more effective, because two indexes are used to satisfy the predicates in the parent query, thus returning only a few employee_ids. The employee_ids are then used to access the orders table through an index.

ID	OPERATION	OPTIONS	OBJECT_NAME	OPT	COST
0	SELECT STATEMENT			CHO	
1	FILTER				
2	TABLE ACCESS	BY INDEX ROWID	EMPLOYEES	ANA	98
3	AND-EQUAL				
4	INDEX	RANGE SCAN	EMP_JOB_IX	ANA	
5	INDEX	RANGE SCAN	EMP_DEPARTMENT_IX	ANA	
6	INDEX	RANGE SCAN	ORD_SALES_REP_IX	ANA	8

Note: An even more efficient approach is to have a concatenated index on department_id and job_id. This eliminates the need to access two indexes and reduces the resources used.

Controlling the Access Path and Join Order with Hints

You can influence the optimizer's choices by setting the optimizer approach and goal, and by gathering representative statistics for the CBO. Sometimes, the application designer, who has more information about a particular application's data than is available to the optimizer, can choose a more effective way to execute a SQL statement. You can use hints in SQL statements to specify how the statement should be executed.

Hints, such as /*+FULL */ control access paths. For example:

```
SELECT /*+ FULL(e) */ e.ename
FROM emp e
WHERE e.job = 'CLERK';
```

See Also: [Chapter 1, "Introduction to the Optimizer"](#) and [Chapter 5, "Optimizer Hints"](#)

Join order can have a significant effect on performance. The main objective of SQL tuning is to avoid performing unnecessary work to access rows that do not affect the result. This leads to three general rules:

- Avoid a full-table scan if it is more efficient to get the required rows through an index.
- Avoid using an index that fetches 10,000 rows from the driving table if you could instead use another index that fetches 100 rows.
- Choose the join order so as to join fewer rows to tables later in the join order.

The following example shows how to tune join order effectively:

```
SELECT info
FROM taba a, tabb b, tabc c
WHERE a.acol BETWEEN 100 AND 200
      AND b.bcol BETWEEN 10000 AND 20000
      AND c.ccol BETWEEN 10000 AND 20000
      AND a.key1 = b.key1
      AND a.key2 = c.key2;
```

1. Choose the driving table and the driving index (if any).

The first three conditions in the previous example are filter conditions applying to only a single table each. The last two conditions are join conditions.

Filter conditions dominate the choice of driving table and index. In general, the driving table is the one containing the filter condition that eliminates the highest percentage of the table. Thus, because the range of 100 to 200 is narrow compared with the range of `acol`, but the ranges of 10000 and 20000 are relatively large, `taba` is the driving table, all else being equal.

With nested loop joins, the joins all happen through the join indexes, the indexes on the primary or foreign keys used to connect that table to an earlier table in the join tree. Rarely do you use the indexes on the nonjoin conditions, except for the driving table. Thus, after `taba` is chosen as the driving table, use the indexes on `b.key1` and `c.key2` to drive into `tabb` and `tabc`, respectively.

2. Choose the best join order, driving to the best unused filters earliest.

The work of the following join can be reduced by first joining to the table with the best still-unused filter. Thus, if "bcol BETWEEN ..." is more restrictive (rejects a higher percentage of the rows seen) than "ccol BETWEEN ...", the last join can be made easier (with fewer rows) if tabb is joined before tabc.

3. You can use the ORDERED or STAR hint to force the join order.

See Also: ["Hints for Join Orders"](#) on page 5-22

Use Caution When Managing Views

Be careful when joining views, when performing outer joins to views, and when reusing an existing view for a new purpose.

Use Caution When Joining Complex Views Joins to complex views are not recommended, particularly joins from one complex view to another. Often this results in the entire view being instantiated, and then the query is run against the view data.

For example, the following statement creates a view that lists employees and departments:

```
CREATE OR REPLACE VIEW emp_dept
AS
SELECT d.department_id
      , d.department_name
      , d.location_id
      , e.employee_id
      , e.last_name
      , e.first_name
      , e.salary
      , e.job_id
FROM   departments d
      , employees e
WHERE  e.department_id (+) = d.department_id
/
```

The following query finds employees in a specified state:

```
SELECT v.last_name, v.first_name, l.state_province
FROM   locations l, emp_dept v
WHERE  l.state_province = 'California'
AND    v.location_id = l.location_id (+)
/
```

In the following plan, note that the emp_dept view is instantiated:

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT						
FILTER						
NESTED LOOPS OUTER						
VIEW	EMP_DEPT					
NESTED LOOPS OUTER						
TABLE ACCESS FULL	DEPARTMEN					
TABLE ACCESS BY INDEX	EMPLOYEES					
INDEX RANGE SCAN	EMP_DEPAR					
TABLE ACCESS BY INDEX R	LOCATIONS					
INDEX UNIQUE SCAN	LOC_ID_PK					

Do Not Recycle Views Beware of writing a view for one purpose and then using it for other purposes to which it might be ill-suited. Querying from a view requires all tables from the view to be accessed for the data to be returned. Before reusing a view, determine whether all tables in the view need to be accessed to return the data. If not, then do not use the view. Instead, use the base table(s), or if necessary, define a new view. The goal is to refer to the minimum number of tables and views necessary to return the required data.

Consider the following example:

```
SELECT dname
FROM emp_dept
WHERE deptno=10;
```

The entire view is first instantiated by performing a join of the `emp` and `dept` tables and then aggregating the data. However, you can obtain `dname` and `deptno` directly from the `dept` table. It is inefficient to obtain this information by querying the `dx` view (which was declared in the earlier example).

Use Caution When Unnesting Subqueries Subquery unnesting merges the body of the subquery into the body of the statement that contains it, allowing the optimizer to consider them together when evaluating access paths and joins.

See Also: *Oracle9i Data Warehousing Guide* for an explanation of the dangers with subquery unnesting

Use Caution When Performing Outer Joins to Views In the case of an outer join to a multitable view, the CBO (in Release 8.1.6 and later) can drive from an outer join column, if an equality predicate is defined on it.

An outer join *within* a view is problematic because the performance implications of the outer join are not visible.

Store Intermediate Results

Intermediate, or staging, tables are quite common in relational database systems, because they temporarily store some intermediate results. In many applications they are useful, but Oracle requires additional resources to create them. Always consider whether the benefit they could bring is more than the cost to create them. Avoid staging tables when the information is not reused multiple times.

Some additional considerations:

- Storing intermediate results in staging tables could improve application performance. In general, whenever an intermediate result is usable by multiple following queries, it is worthwhile to store it in a staging table. The benefit of not retrieving data multiple times with a complex statement already at the second usage of the intermediate result is better than the cost to materialize it.
- Long and complex queries are hard to understand and optimize. Staging tables can break a complicated SQL statement into several smaller statements, and then store the result of each step.
- Consider using materialized views. These are precomputed tables comprising aggregated or joined data from fact and possibly dimension tables.

See Also: *Oracle9i Data Warehousing Guide* for detailed information on using materialized views

Restructuring the Indexes

Often, there is a beneficial impact on performance by restructuring indexes. This can involve the following:

- Remove nonselective indexes to speed the DML.
- Index performance-critical access paths.
- Consider reordering columns in existing concatenated indexes.
- Add columns to the index to improve selectivity.

Do not use indexes as a panacea. Application developers sometimes think that performance will improve if they create more indexes. If a single programmer creates an appropriate index, then this might indeed improve the application's performance. However, if 50 programmers each create an index, then application performance will probably be hampered.

Modifying or Disabling Triggers and Constraints

Using triggers consumes system resources. If you use too many triggers, then you can find that performance is adversely affected and you might need to modify or disable them.

Restructuring the Data

After restructuring the indexes and the statement, you can consider restructuring the data.

- Introduce derived values. Avoid `GROUP BY` in response-critical code.
- Review your data design. Change the design of your system if it can improve performance.
- Consider partitioning, if appropriate.

Maintaining Execution Plans Over Time

You can maintain the existing execution plan of SQL statements over time either using stored statistics or stored SQL execution plans. Storing optimizer statistics for tables will apply to all SQL statements that refer to those tables. Storing an execution plan (that is, plan stability) maintains the plan for a single SQL statement. If both statistics and a stored plan are available for a SQL statement, then the optimizer uses the stored plan.

See Also:

- [Chapter 3, "Gathering Optimizer Statistics"](#)
- [Chapter 7, "Using Plan Stability"](#)

Visiting Data as Few Times as Possible

Applications should try to access each row only once. This reduces network traffic and reduces database load. Consider doing the following:

- [Combine Multiples Scans with CASE Statements](#)

- [Use DML with RETURNING Clause](#)
- [Modify All the Data Needed in One Statement](#)

Combine Multiples Scans with CASE Statements

Often, it is necessary to calculate different aggregates on various sets of tables. Usually, this is done with multiple scans on the table, but it is easy to calculate all the aggregates with one single scan. Eliminating n-1 scans can greatly improve performance.

Combining multiple scans into one scan can be done by moving the `WHERE` condition of each scan into a `CASE` statement, which filters the data for the aggregation. For each aggregation, there could be another column that retrieves the data.

The following example asks for the count of all employees who earn less than 2000, between 2000 and 4000, and more than 4000 each month. This can be done with three separate queries:

```
SELECT COUNT (*)
FROM employees
WHERE salary < 2000;
```

```
SELECT COUNT (*)
FROM employees
WHERE salary BETWEEN 2000 AND 4000;
```

```
SELECT COUNT (*)
FROM employees
WHERE salary > 4000;
```

However, it is more efficient to run the entire query in a single statement. Each number is calculated as one column. The count uses a filter with the `CASE` statement to count only the rows where the condition is valid. For example:

```
SELECT COUNT (CASE WHEN salary < 2000
                  THEN 1 ELSE null END) count1,
       COUNT (CASE WHEN salary BETWEEN 2001 AND 4000
                  THEN 1 ELSE null END) count2,
       COUNT (CASE WHEN salary > 4000
                  THEN 1 ELSE null END) count3
FROM employees;
```

This is a very simple example. The ranges could be overlapping, the functions for the aggregates could be different, and so on.

Use DML with RETURNING Clause

When appropriate, use INSERT, UPDATE, or DELETE... RETURNING to select and modify data with a single call. This technique improves performance by reducing the number of calls to the database.

See Also: *Oracle9i SQL Reference* for syntax on the INSERT, UPDATE, and DELETE statements

Modify All the Data Needed in One Statement

When possible, use array processing. This means that an array of bind variable values is passed to Oracle for repeated execution. This is appropriate for iterative processes in which multiple rows of a set are subject to the same operation.

For example:

```
BEGIN
  FOR pos_rec IN (SELECT *
                 FROM order_positions
                 WHERE order_id = :id) LOOP
    DELETE FROM order_positions
           WHERE order_id = pos_rec.order_id AND
                 order_position = pos_rec.order_position;
  END LOOP;
DELETE FROM orders
WHERE order_id = :id;
END;
```

Alternatively, you could define a cascading constraint on orders. In the previous example, one SELECT and *n* DELETES are executed. When a user issues the DELETE on orders DELETE FROM orders WHERE order_id = :id, the database automatically deletes the positions with a single DELETE statement.

See Also: *Oracle9i Database Administrator's Guide* or *Oracle9i Heterogeneous Connectivity Administrator's Guide* for information on tuning distributed queries

Using Plan Stability

This chapter describes how to use plan stability to preserve performance characteristics.

This chapter contains the following sections:

- [Using Plan Stability to Preserve Execution Plans](#)
- [Using Plan Stability with the Cost-Based Optimizer](#)

Using Plan Stability to Preserve Execution Plans

Plan stability prevents certain database environment changes from affecting the performance characteristics of applications. Such changes include changes in optimizer statistics, changes to the optimizer mode settings, and changes to parameters affecting the sizes of memory structures, such as `SORT_AREA_SIZE` and `BITMAP_MERGE_AREA_SIZE`. Plan stability is most useful when you cannot risk any performance changes in an application.

Plan stability preserves execution plans in stored outlines. Oracle can create a public or private stored outline for one or all SQL statements. The optimizer then generates equivalent execution plans from the outlines when you enable the use of stored outlines. You can group outlines into categories and control which category of outlines Oracle uses to simplify outline administration and deployment.

The plans Oracle maintains in stored outlines remain consistent despite changes to a system's configuration or statistics. Using stored outlines also stabilizes the generated execution plan if the optimizer changes in subsequent Oracle releases. Plan stability also facilitates migration from the rule-based optimizer to the cost-based optimizer when you upgrade to a new Oracle release.

Note: If you develop applications for mass distribution, then you can use stored outlines to ensure that all customers access the same execution plans.

Using Hints with Plan Stability

The degree to which plan stability controls execution plans is dictated by how much Oracle's hint mechanism controls execution plans, because Oracle uses hints to record stored plans.

There is a one-to-one correspondence between SQL text and its stored outline. If you specify a different literal in a predicate, then a different outline applies. To avoid this, replace literals in applications with bind variables.

See Also: Oracle can force similar statements to share SQL by replacing literals with system-generated bind variables. This works with plan stability if the outline was generated using the `CREATE_STORED_OUTLINES` parameter, not the `CREATE OUTLINE` statement. Also, the outline must have been created with the `CURSOR_SHARING` parameter set to `SIMILAR` or `FORCE`, and the parameter must also set to `SIMILAR` or `FORCE` when attempting to use the outline. See [Chapter 14, "Memory Configuration and Use"](#) for more information.

Plan stability relies on preserving execution plans at a point in time when performance is satisfactory. In many environments, however, attributes for datatypes such as `dates` or `order numbers` can change rapidly. In these cases, permanent use of an execution plan can result in performance degradation over time as the data characteristics change.

This implies that techniques that rely on preserving plans in dynamic environments are somewhat contrary to the purpose of using cost-based optimization. Cost-based optimization attempts to produce execution plans based on statistics that accurately reflect the state of the data. Thus, you must balance the need to control plan stability with the benefit obtained from the optimizer's ability to adjust to changes in data characteristics.

How Outlines Use Hints

An outline consists primarily of a set of hints that is equivalent to the optimizer's results for the execution plan generation of a particular SQL statement. When Oracle creates an outline, plan stability examines the optimization results using the same data used to generate the execution plan. That is, Oracle uses the input to the execution plan to generate an outline, and not the execution plan itself.

Note: Oracle creates the `USER_OUTLINES` and `USER_OUTLINE_HINTS` views in the `SYS` tablespace based on data in the `OL$` and `OL$HINTS` tables, respectively. Direct manipulation of the `OL$`, `OL$HINTS`, and `OL$NODES` tables is prohibited.

You can embed hints in SQL statements, but this has no effect on how Oracle uses outlines. Oracle considers a SQL statement that you revised with hints to be different from the original SQL statement stored in the outline.

Storing Outlines

Oracle stores outline data in the `OL$`, `OL$HINTS`, and `OL$NODES` tables. Unless you remove them, Oracle retains outlines indefinitely.

The only effect outlines have on caching execution plans is that the outline's category name is used in addition to the SQL text to identify whether the plan is in cache. This ensures that Oracle does not use an execution plan compiled under one category to execute a SQL statement that Oracle should compile under a different category.

Enabling Plan Stability

Settings for several parameters, especially those ending with the suffix `_ENABLED`, must be consistent across execution environments for outlines to function properly. These parameters are:

- `QUERY_REWRITE_ENABLED`
- `STAR_TRANSFORMATION_ENABLED`
- `OPTIMIZER_FEATURES_ENABLE`

Using Supplied Packages to Manage Stored Outlines

The `DBMS_OUTLN` and `DBMS_OUTLN_EDIT` package provides procedures used for managing stored outlines and their outline categories.

Users need the `EXECUTE_CATALOG_ROLE` role to execute `DBMS_OUTLN`, but `public` has execute privileges on `DBMS_OUTLN_EDIT`. The `DBMS_OUTLN_EDIT` package is an invoker's rights package.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for detailed information on using `DBMS_OUTLN` and `DBMS_OUTLN_EDIT` procedures

Creating Outlines

Oracle can automatically create outlines for all SQL statements, or you can create them for specific SQL statements. In either case, the outlines derive their input from the optimizer.

Oracle creates stored outlines automatically when you set the parameter `CREATE_STORED_OUTLINES` to `true`. When activated, Oracle creates outlines for all

compiled SQL statements. You can create stored outlines for specific statements using the `CREATE OUTLINE` statement.

Note: You must ensure that schemas in which outlines are to be created have the `CREATE ANY OUTLINE` privilege. Otherwise, despite having turned on the `CREATE_STORED_OUTLINE` parameter, you will not find outlines in the database after you run the application.

Also, the default system tablespace can become exhausted if the `CREATE_STORED_OUTLINES` parameter is enabled and the running application has an abundance of literal SQL statements. If this happens, use the `DBMS_OUTLN.DROP_UNUSED` procedure to remove those literal SQL outlines.

The `CREATE_EDIT_TABLES` procedure in the `DBMS_OUTLN_EDIT` package creates tables in the invoker's schema. This is necessary for editing private outlines. This is callable by anyone with `EXECUTE` privilege on `DBMS_OUTLN_EDIT`.

See Also:

- *Oracle9i SQL Reference* for more information on the `CREATE OUTLINE` statement
- *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on the `DBMS_OUTLN` and `DBMS_OUTLN_EDIT` packages
- ["Using Outlines to Move to the Cost-Based Optimizer"](#) on page 7-13 for information on moving from the rule-based optimizer to the cost-based optimizer
- *Database Tuning with the Oracle Tuning Pack* for information on using the Outline Management and Outline Editor tools, which let you create, edit, delete, and manage stored outlines with an easy-to-use graphical interface

Using Category Names for Stored Outlines

Outlines can be categorized to simplify the management task. The `CREATE OUTLINE` statement allows for specification of a category. The `DEFAULT` category is chosen if unspecified. Likewise, the `CREATE_STORED_OUTLINES` parameter lets

you specify a category name, where specifying `true` produces outlines in the `DEFAULT` category.

If you specify a category name using the `CREATE_STORED_OUTLINES` parameter, then Oracle assigns all subsequently created outlines to that category until you reset the category name. Set the parameter to `false` to suspend outline generation.

If you set `CREATE_STORED_OUTLINES` to `true`, or if you use the `CREATE OUTLINE` statement without a category name, then Oracle assigns outlines to the category name of `DEFAULT`.

Note: The `CREATE_STORED_OUTLINES`, `USE_STORED_OUTLINES`, and `USE_PRIVATE_OUTLINES` parameters are system or session specific. They are not initialization parameters. For more information on these parameters, see the *Oracle9i SQL Reference*.

Using and Editing Stored Outlines

When you activate the use of stored outlines, Oracle always uses the cost-based optimizer. This is because outlines rely on hints, and to be effective, most hints require the cost-based optimizer.

To use stored outlines when Oracle compiles a SQL statement, set the system parameter `USE_STORED_OUTLINES` to `true` or to a category name. If you set `USE_STORED_OUTLINES` to `true`, then Oracle uses outlines in the default category. If you specify a category with the `USE_STORED_OUTLINES` parameter, then Oracle uses outlines in that category until you reset the parameter to another category name or until you suspend outline use by setting `USE_STORED_OUTLINES` to `false`. If you specify a category name and Oracle does not find an outline in that category that matches the SQL statement, then Oracle searches for an outline in the default category.

The designated outlines only control the compilation of SQL statements that have outlines. If you set `USE_STORED_OUTLINES` to `false`, then Oracle does not use outlines. When you set `USE_STORED_OUTLINES` to `false` and you set `CREATE_STORED_OUTLINES` to `true`, Oracle creates outlines but does not use them.

The `USE_PRIVATE_OUTLINES` parameter lets you control the use of private outlines. A private outline is an outline seen only in the current session and whose data resides in the current parsing schema. Any changes made to such an outline are not seen by any other session on the system, and applying a private outline to the compilation of a statement can only be done in the current session with the

`USE_PRIVATE_OUTLINES` parameter. Only when you explicitly choose to save your edits back to the public area are they seen by the rest of the users.

While the optimizer usually chooses optimal plans for queries, there are times when users know things about the execution environment that are inconsistent with the heuristics that the optimizer follows. By editing outlines directly, you can tune the SQL query without having to alter the application.

When a private outline is created, an error is returned if the prerequisite outline tables to hold the outline data do not exist in the local schema. These tables can be created using the `DBMS_OUTLN_EDIT.CREATE_EDIT_TABLES` procedure. You can also use the `UTLEDITOL.SQL` script.

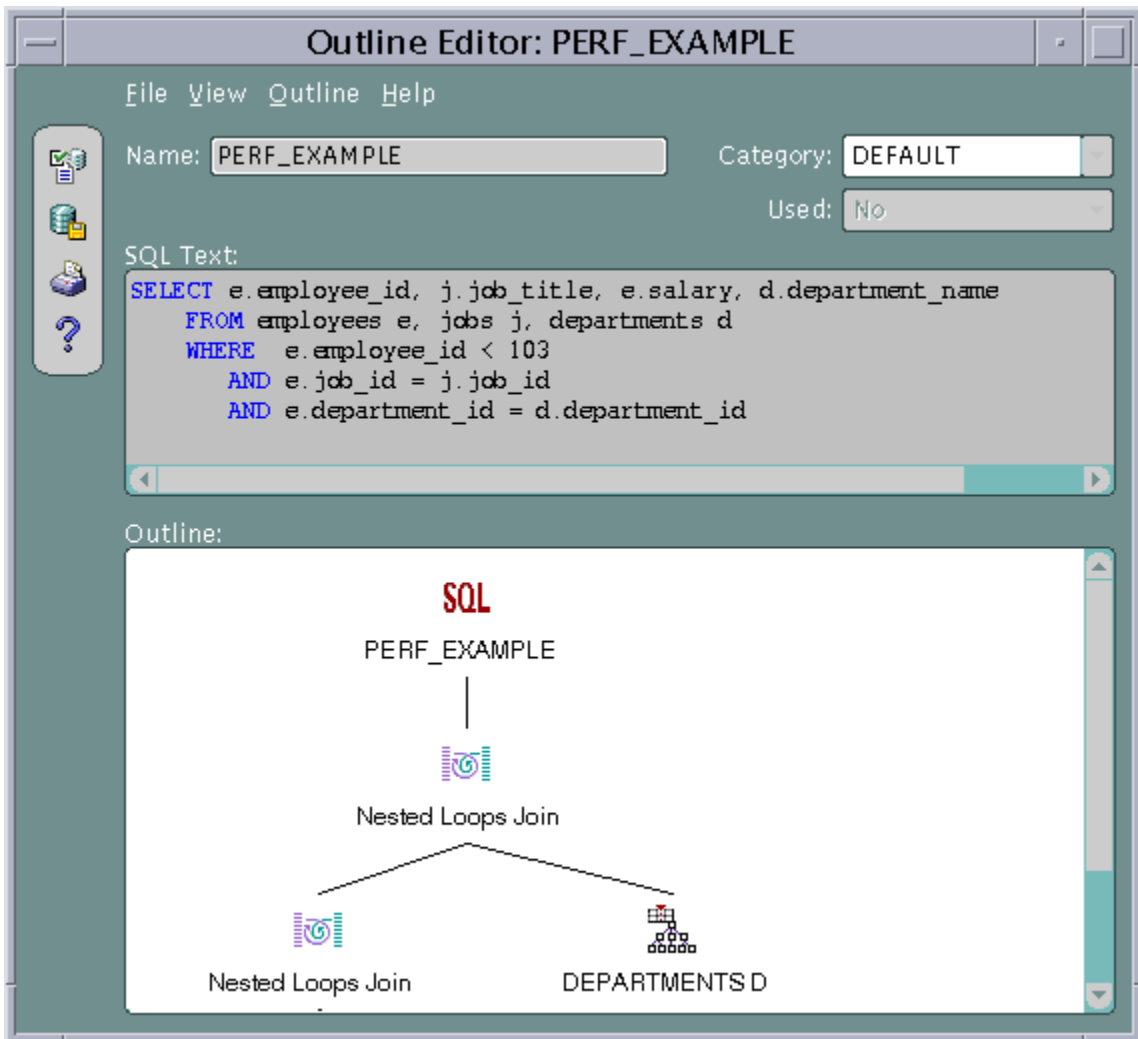
When the `USE_PRIVATE_OUTLINES` parameter is enabled and an outlined SQL statement is issued, the optimizer retrieves the outline from the session private area rather than the public area used when `USE_STORED_OUTLINES` is enabled. If no outline exists in the session private area, then the optimizer will not use an outline to compile the statement.

Any `CREATE OUTLINE` statement requires the `CREATE ANY OUTLINE` privilege. Specification of the `FROM` clause also requires the `SELECT` privilege. This privilege should be granted only to those users who would have the authority to view SQL text and hint text associated with the outlined statements. This role is required for the `CREATE OUTLINE FROM` command unless the issuer of the command is also the owner of the outline.

When you begin an editing session, `USE_PRIVATE_OUTLINES` should be set to the category to which the outline being edited belongs. When you are finished editing, this parameter should be set to `false` to restore the session to normal outline lookup according to the `USE_STORED_OUTLINES` parameter.

You can use the GUI Outline Editor of the Oracle Enterprise Manager Tuning Pack to update outlines. [Figure 7-1](#) is an illustration of the Outline Editor.

Figure 7-1 Outline Editor



See Also: *Database Tuning with the Oracle Tuning Pack* for more information on the GUI tool for editing outlines

Example of Editing Outlines

Assume that you want to edit the outline `o11`. The steps are as follows:

1. Connect to a schema from which the outlined statement can be executed, and ensure that the `CREATE ANY OUTLINE` and `SELECT` privileges have been granted.
2. Create outline editing tables locally with the `DBMS_OUTLN_EDIT.CREATE_EDIT_TABLES` procedure.
3. Clone the outline being edited to the private area using the following:

```
CREATE PRIVATE OUTLINE p_o11 FROM o11;
```

4. Edit the outline, either with the Outline Editor in Enterprise Manager or manually by querying the local `OL$HINTS` tables and performing DML against the appropriate hint tuples. `DBMS_OUTLN_EDIT.CHANGE_JOIN_POS` is available for changing join order.
5. If manually editing the outline, then resynchronize the stored outline definition using the following so-called identity statement:

```
CREATE PRIVATE OUTLINE p_o11 FROM PRIVATE p_o11;
```

You can also use `DBMS_OUTLN_EDIT.REFRESH_PRIVATE_OUTLINE` or `ALTER SYSTEM FLUSH SHARED_POOL` to accomplish this.

6. Test the edits. Set `USE_PRIVATE_OUTLINES=TRUE`, and issue the outline statement or run `EXPLAIN PLAN` on the statement.
7. If you want to preserve these edits for public use, then publicize the edits with the following statement.

```
CREATE OR REPLACE OUTLINE o11 FROM PRIVATE p_o11;
```

8. Disable private outline usage by setting the following:

```
USE_PRIVATE_OUTLINES=FALSE
```

See Also:

- *Oracle9i SQL Reference* for SQL syntax
- *Database Tuning with the Oracle Tuning Pack* for more information on the GUI tool for editing outlines
- *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on the DBMS_OUTLN and DBMS_OUTLN_EDIT packages

How to Tell If an Outline Is Being Used

You can test if an outline is being used with the V\$SQL view. Query the OUTLINE_CATEGORY column in conjunction with the SQL statement. If an outline was applied, then this column contains the category to which the outline belongs. Otherwise, it is NULL. The OUTLINE_SID column tells you if this particular cursor is using a public outline (value is 0) or a private outline (session's SID of the corresponding session using it).

For example:

```
SELECT OUTLINE_CATEGORY, OUTLINE_SID
FROM V$SQL
WHERE SQL_TEXT LIKE 'SELECT COUNT(*) FROM emp%';
```

Viewing Outline Data

You can access information about outlines and related hint data that Oracle stores in the data dictionary from the following views:

- USER_OUTLINES
- USER_OUTLINE_HINTS
- ALL_OUTLINES
- ALL_OUTLINE_HINTS
- DBA_OUTLINES
- DBA_OUTLINE_HINTS

Use the following syntax to obtain outline information from the USER_OUTLINES view, where the outline category is mycat:

```
SELECT NAME, SQL_TEXT
FROM USER_OUTLINES
WHERE CATEGORY='mycat';
```

Oracle responds by displaying the names and text of all outlines in category `mycat`.

To see all generated hints for the outline `name1`, use the following syntax:

```
SELECT HINT
FROM USER_OUTLINE_HINTS
WHERE NAME='name1';
```

Note: If necessary, you can use the procedure to move outline tables from one tablespace to another as described in ["Moving Outline Tables"](#) on page 7-11.

Moving Outline Tables

Oracle creates the `USER_OUTLINES` and `USER_OUTLINE_HINTS` views based on data in the `OL$` and `OL$HINTS` tables, respectively. Oracle creates these tables, and also the `OL$NODES` table, in the `SYS` tablespace using a schema called `OUTLN`. If outlines use too much space in the `SYS` tablespace, then you can move them. To do this, create a separate tablespace and move the outline tables into it using the following process.

1. The default system tablespace could become exhausted if the `CREATE_STORED_OUTLINES` parameter is on and if the running application has many literal SQL statements. If this happens, then use the `DBMS_OUTLN.DROP_UNUSED` procedure to remove those literal SQL outlines.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for detailed information on using the `DBMS_OUTLN` package

2. Export the `OL$`, `OL$HINTS`, and `OL$NODES` tables:

```
EXP OUTLN/OUTLN FILE = exp_file TABLES = 'OL$' 'OL$HINTS' 'OL$NODES'
```

3. Remove the previous `OL$`, `OL$HINTS`, and `OL$NODES` tables:

```
CONNECT OUTLN/outln_password;
DROP TABLE OL$;
CONNECT OUTLN/outln_password;
DROP TABLE OL$HINTS;
CONNECT OUTLN/outln_password;
DROP TABLE OL$NODES;
```

4. Create a new tablespace for the tables:

```
CREATE TABLESPACE outln_ts  
DATAFILE 'tspace.dat' SIZE 2MB  
DEFAULT STORAGE (INITIAL 10KB NEXT 20KB  
MINEXTENTS 1 MAXEXTENTS 999 PCTINCREASE 10) ONLINE;
```

5. Enter the following statement:

```
ALTER USER OUTLN DEFALUT TABLESPACE outln_ts;
```

6. Import the OL\$, OL\$HINTS, and OL\$NODES tables:

```
IMP OUTLN/outln_password  
FILE=exp_file TABLES = 'OL$' 'OL$HINTS' 'OL$NODES'
```

The **IMPORT** statement re-creates the OL\$, OL\$HINTS, and OL\$NODES tables in the schema named OUTLN, but the schema now resides in a new tablespace called OUTLN_TS.

Note: If Oracle8i outlines are imported into an Oracle9i database, then the DBMS_OUTLN.UPDATE_SIGNATURES procedure must be run. This updates the signatures of all outlines on the system so that they are compatible with Oracle9i semantics. If this step is not done, then no outlines from the Oracle8i database are used.

Using Plan Stability with the Cost-Based Optimizer

This section describes procedures you can use to significantly improve performance by taking advantage of cost-based optimizer functionality. Plan stability provides a way to preserve a system's targeted execution plans with satisfactory performance while also taking advantage of new cost-based optimizer features for the rest of the SQL statements.

Topics covered in this section are:

- [Using Outlines to Move to the Cost-Based Optimizer](#)
- [Upgrading and the Cost-Based Optimizer](#)

Using Outlines to Move to the Cost-Based Optimizer

If an application was developed using the rule-based optimizer, then a considerable amount of effort might have gone into manually tuning the SQL statements to optimize performance. You can use plan stability to leverage the effort that has already gone into performance tuning by preserving the behavior of the application when upgrading from rule-based to cost-based optimization.

By creating outlines for an application before switching to cost-based optimization, the plans generated by the rule-based optimizer can be used, while statements generated by newly written applications developed after the switch use cost-based plans. To create and use outlines for an application, use the following process.

Note: *Carefully read this procedure and consider its implications before executing it!*

1. Ensure that schemas in which outlines are to be created have the CREATE ANY OUTLINE privilege. For example, from SYS:

```
GRANT CREATE ANY OUTLINE TO user-name
```

2. Execute syntax similar to the following to designate; for example, the RBOCAT outline category.

```
ALTER SESSION SET CREATE_STORED_OUTLINES = rboecat;
```

3. Run the application long enough to capture stored outlines for all important SQL statements.

4. Suspend outline generation:

```
ALTER SESSION SET CREATE_STORED_OUTLINES = FALSE;
```

5. Gather statistics with the DBMS_STATS package.

6. Alter the parameter OPTIMIZER_MODE to CHOOSE.

7. Enter the following syntax to make Oracle use the outlines in category RBOCAT:

```
ALTER SESSION SET USE_STORED_OUTLINES = rboecat;
```

8. Run the application.

Subject to the limitations of plan stability, access paths for this application's SQL statements should be unchanged.

Note: If a query was not executed in step 2, then you can capture the old behavior of the query even after switching to cost-based optimization. To do this, change the optimizer mode to `RULE`, create an outline for the query, and then change the optimizer mode back to `CHOOSE`.

Upgrading and the Cost-Based Optimizer

When upgrading to a new Oracle release under cost-based optimization, there is always a possibility that some SQL statements will have their execution plans changed due to changes in the optimizer. While such changes benefit performance, you might have applications that perform so well that you would consider any changes in their behavior to be an unnecessary risk. For such applications, you can create outlines before the upgrade using the following procedure.

Note: *Carefully read this procedure and consider its implications before running it!*

1. Enter the following syntax to enable outline creation:

```
ALTER SESSION SET CREATE_STORED_OUTLINES = ALL_QUERIES;
```

2. Run the application long enough to capture stored outlines for all critical SQL statements.
3. Enter this syntax to suspend outline generation:

```
ALTER SESSION SET CREATE_STORED_OUTLINES = FALSE;
```

4. Upgrade the production system to the new version of the RDBMS.
5. Run the application.

After the upgrade, you can enable the use of stored outlines, or alternatively, you can use the outlines that were stored as a backup if you find that some statements exhibit performance degradation after the upgrade.

With the latter approach, you can selectively use the stored outlines for such problematic statements as follows:

1. For each problematic SQL statement, change the `CATEGORY` of the associated stored outline to a category name similar to this:

```
ALTER OUTLINE outline_name CHANGE CATEGORY TO problemcat;
```

2. Enter this syntax to make Oracle use outlines from the category `problemcat`.

```
ALTER SESSION SET USE_STORED_OUTLINES = problemcat;
```

Upgrading with a Test System

A test system, separate from the production system, can be useful for conducting experiments with optimizer behavior in conjunction with an upgrade. You can migrate statistics from the production system to the test system using `import/export`. This can alleviate the need to fill the tables in the test system with data.

You can move outlines between the systems by category. For example, after you create outlines in the `problemcat` category, export them by category using the query-based export option. This is a convenient and efficient way to export only selected outlines from one database to another without exporting all outlines in the source database. To do this, issue these statements:

```
EXP OUTLN/outln_password FILE=exp-file TABLES= 'OL$' 'OL$HINTS' 'OL$NODES'  
QUERY='WHERE CATEGORY="problemcat"'
```

Using the Rule-Based Optimizer

This chapter discusses Oracle's rule-based optimizer (RBO). In general, always use the cost-based approach. The rule-based approach is available for backward compatibility.

Note: Oracle Corporation strongly advises the use of cost-based optimization. Rule-based optimization will be deprecated in a future release.

See Also: [Chapter 1, "Introduction to the Optimizer"](#)

This chapter contains the following sections:

- [Overview of the Rule-Based Optimizer \(RBO\)](#)
- [Understanding Access Paths for the RBO](#)
- [Transforming and Optimizing Statements with the RBO](#)

Overview of the Rule-Based Optimizer (RBO)

Although Oracle supports the rule-based optimizer, you should design new applications to use the cost-based optimizer (CBO). You should also use the CBO for data warehousing applications, because the CBO supports enhanced features for DSS. Many new performance features, such as partitioned tables, improved star query processing, and materialized views, are only available with the CBO.

Note: If you have developed OLTP applications using Oracle version 6, and if you have tuned the SQL statements carefully based on the rules of the optimizer, then you might want to continue using the RBO when you upgrade these applications to a new Oracle release.

If you are using applications provided by third-party vendors, then check with the vendors to determine which type of optimizer is best suited to that application.

If `OPTIMIZER_MODE=CHOOSE`, if statistics do not exist, and if you do not add hints to SQL statements, then SQL statements use the RBO. You can use the RBO to access both relational data and object types. If `OPTIMIZER_MODE=FIRST_ROWS`, `FIRST_ROWS_n`, or `ALL_ROWS` and no statistics exist, then the CBO uses default statistics. Migrate existing applications to use the cost-based approach.

You can enable the CBO on a trial basis simply by collecting statistics. You can then return to the RBO by deleting the statistics or by setting either the value of the `OPTIMIZER_MODE` initialization parameter or the `OPTIMIZER_MODE` clause of the `ALTER SESSION` statement to `RULE`. You can also use this value if you want to collect and examine statistics for data without using the cost-based approach.

See Also: [Chapter 3, "Gathering Optimizer Statistics"](#) for an explanation of how to gather statistics

Understanding Access Paths for the RBO

Using the RBO, the optimizer chooses an execution plan based on the access paths available and the ranks of these access paths. Oracle's ranking of the access paths is heuristic. If there is more than one way to execute a SQL statement, then the RBO always uses the operation with the lower rank. Usually, operations of lower rank execute faster than those associated with constructs of higher rank.

The list shows access paths and their ranking:

RBO Path 1: Single Row by Rowid

RBO Path 2: Single Row by Cluster Join

RBO Path 3: Single Row by Hash Cluster Key with Unique or Primary Key

RBO Path 4: Single Row by Unique or Primary Key

RBO Path 5: Clustered Join

RBO Path 6: Hash Cluster Key

RBO Path 7: Indexed Cluster Key

RBO Path 8: Composite Index

RBO Path 9: Single-Column Indexes

RBO Path 10: Bounded Range Search on Indexed Columns

RBO Path 11: Unbounded Range Search on Indexed Columns

RBO Path 12: Sort Merge Join

RBO Path 13: MAX or MIN of Indexed Column

RBO Path 14: ORDER BY on Indexed Column

RBO Path 15: Full Table Scan

Details of the RBO Access Paths

Each of the following sections describes an access path, discusses when it is available, and shows the output generated for it by the `EXPLAIN PLAN` statement.

RBO Path 1: Single Row by Rowid

This access path is available only if the statement's `WHERE` clause identifies the selected rows by rowid or with the `CURRENT OF CURSOR` embedded SQL syntax supported by the Oracle precompilers. To execute the statement, Oracle accesses the table by rowid.

For example:

```
SELECT * FROM emp WHERE ROWID = 'AAAA7bAA5AAAA1UAAA';
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP

RBO Path 2: Single Row by Cluster Join

This access path is available for statements that join tables stored in the same cluster if both of the following conditions are true:

- The statement's `WHERE` clause contains conditions that equate each column of the cluster key in one table with the corresponding column in the other table.
- The statement's `WHERE` clause also contains a condition that guarantees the join returns only one row. Such a condition is likely to be an equality condition on the column(s) of a unique or primary key.

These conditions must be combined with `AND` operators. To execute the statement, Oracle performs a nested loops operation.

See Also: ["Nested Loop Outer Joins"](#) on page 1-51

For example, in the following statement, the `emp` and `dept` tables are clustered on the `deptno` column, and the `empno` column is the primary key of the `emp` table:

```
SELECT *
FROM emp, dept
WHERE emp.deptno = dept.deptno
AND emp.empno = 7900;
```


The EXPLAIN PLAN output for this statement might look like this:

```

OPERATION                OPTIONS                OBJECT_NAME
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS         BY ROWID              EMP
    INDEX                 UNIQUE SCAN           PK_EMP
    TABLE ACCESS         CLUSTER               DEPT

```

pk_emp is the name of an index that enforces the primary key.

RBO Path 3: Single Row by Hash Cluster Key with Unique or Primary Key

This access path is available if both of the following conditions are true:

- The statement's WHERE clause uses all columns of a hash cluster key in equality conditions. For composite cluster keys, the equality conditions must be combined with AND operators.
- The statement is guaranteed to return only one row, because the columns that make up the hash cluster key also make up a unique or primary key.

To execute the statement, Oracle applies the cluster's hash function to the hash cluster key value specified in the statement to obtain a hash value. Oracle then uses the hash value to perform a hash scan on the table.

For example:

In the following statement, the orders and line_items tables are stored in a hash cluster, and the orderno column is both the cluster key and the primary key of the orders table:

```

SELECT *
  FROM orders
 WHERE orderno = 65118968;

```

The EXPLAIN PLAN output for this statement might look like this:

```

OPERATION                OPTIONS                OBJECT_NAME
-----
SELECT STATEMENT
  TABLE ACCESS         HASH                  ORDERS

```

RBO Path 4: Single Row by Unique or Primary Key

This access path is available if the statement's `WHERE` clause uses all columns of a unique or primary key in equality conditions. For composite keys, the equality conditions must be combined with `AND` operators. To execute the statement, Oracle performs a unique scan on the index on the unique or primary key to retrieve a single rowid, and then accesses the table by that rowid.

For example:

In the following statement, the `empno` column is the primary key of the `emp` table:

```
SELECT *
  FROM emp
 WHERE empno = 7900;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP
INDEX	UNIQUE SCAN	PK_EMP

`pk_emp` is the name of the index that enforces the primary key.

RBO Path 5: Clustered Join

This access path is available for statements that join tables stored in the same cluster if the statement's `WHERE` clause contains conditions that equate each column of the cluster key in one table with the corresponding column in the other table. For a composite cluster key, the equality conditions must be combined with `AND` operators. To execute the statement, Oracle performs a nested loops operation.

See Also: ["Nested Loop Outer Joins"](#) on page 1-51

For example:

In the following statement, the `emp` and `dept` tables are clustered on the `deptno` column:

```
SELECT *
  FROM emp, dept
 WHERE emp.deptno = dept.deptno;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
NESTED LOOPS		
TABLE ACCESS	FULL	DEPT
TABLE ACCESS	CLUSTER	EMP

RBO Path 6: Hash Cluster Key

This access path is available if the statement's `WHERE` clause uses all the columns of a hash cluster key in equality conditions. For a composite cluster key, the equality conditions must be combined with `AND` operators. To execute the statement, Oracle applies the cluster's hash function to the hash cluster key value specified in the statement to obtain a hash value. Oracle then uses this hash value to perform a hash scan on the table.

For example: In the following statement, the `orders` and `line_items` tables are stored in a hash cluster, and the `orderno` column is the cluster key:

```
SELECT *
FROM line_items
WHERE orderno = 65118968;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	HASH	LINE_ITEMS

RBO Path 7: Indexed Cluster Key

This access path is available if the statement's `WHERE` clause uses all the columns of an indexed cluster key in equality conditions. For a composite cluster key, the equality conditions must be combined with `AND` operators.

To execute the statement, Oracle performs a unique scan on the cluster index to retrieve the rowid of one row with the specified cluster key value. Oracle then uses that rowid to access the table with a cluster scan. Because all rows with the same cluster key value are stored together, the cluster scan requires only a single rowid to find them all.

For example:

In the following statement, the `emp` table is stored in an indexed cluster, and the `deptno` column is the cluster key:

```
SELECT * FROM emp
WHERE deptno = 10;
```

The EXPLAIN PLAN output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	CLUSTER	EMP
INDEX	UNIQUE SCAN	PERS_INDEX

`pers_index` is the name of the cluster index.

RBO Path 8: Composite Index

This access path is available if the statement's WHERE clause uses all columns of a composite index in equality conditions combined with AND operators. To execute the statement, Oracle performs a range scan on the index to retrieve rowids of the selected rows, and then accesses the table by those rowids.

For example:

In the following statement, there is a composite index on the `job` and `deptno` columns:

```
SELECT *
FROM emp
WHERE job = 'CLERK'
AND deptno = 30;
```

The EXPLAIN PLAN output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP
INDEX	RANGE SCAN	JOB_DEPTNO_INDEX

`job_deptno_index` is the name of the composite index on the `job` and `deptno` columns.

RBO Path 9: Single-Column Indexes

This access path is available if the statement's WHERE clause uses the columns of one or more single-column indexes in equality conditions. For multiple single-column indexes, the conditions must be combined with AND operators.

If the `WHERE` clause uses the column of only one index, then Oracle executes the statement by performing a range scan on the index to retrieve the rowids of the selected rows, and then accesses the table by these rowids.

For example:

In the following statement, there is an index on the `job` column of the `emp` table:

```
SELECT *
  FROM emp
 WHERE job = 'ANALYST';
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP
INDEX	RANGE SCAN	JOB_INDEX

`job_index` is the index on `emp.job`.

If the `WHERE` clauses uses columns of many single-column indexes, then Oracle executes the statement by performing a range scan on each index to retrieve the rowids of the rows that satisfy each condition. Oracle then merges the sets of rowids to obtain a set of rowids of rows that satisfy all conditions. Oracle then accesses the table using these rowids.

Oracle can merge up to five indexes. If the `WHERE` clause uses columns of more than five single-column indexes, then Oracle merges five of them, accesses the table by rowid, and then tests the resulting rows to determine whether they satisfy the remaining conditions before returning them.

In the following statement, there are indexes on both the `job` and `deptno` columns of the `emp` table:

```
SELECT *
  FROM emp
 WHERE job = 'ANALYST'
    AND deptno = 20;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP

```

AND-EQUAL
INDEX          RANGE SCAN    JOB_INDEX
INDEX          RANGE SCAN    DEPTNO_INDEX
    
```

The AND-EQUAL operation merges the rowids obtained by the scans of the `job_index` and the `deptno_index`, resulting in a set of rowids of rows that satisfy the query.

RBO Path 10: Bounded Range Search on Indexed Columns

This access path is available if the statement's WHERE clause contains a condition that uses either the column of a single-column index or one or more columns that make up a leading portion of a composite index:

```

column = expr

column >[=] expr AND column <[=] expr

column BETWEEN expr AND expr

column LIKE 'c%'
    
```

Each of these conditions specifies a bounded range of indexed values that are accessed by the statement. The range is said to be bounded because the conditions specify both its least value and its greatest value. To execute such a statement, Oracle performs a range scan on the index, and then accesses the table by rowid.

This access path is not available if the expression `expr` references the indexed column.

For example:

In the following statement, there is an index on the `sal` column of the `emp` table:

```

SELECT *
  FROM emp
 WHERE sal BETWEEN 2000 AND 3000;
    
```

The EXPLAIN PLAN output for this statement might look like this:

```

OPERATION          OPTIONS          OBJECT_NAME
-----
SELECT STATEMENT
TABLE ACCESS              BY ROWID          EMP
INDEX                    RANGE SCAN        SAL_INDEX
    
```

`sal_index` is the name of the index on `emp.sal`.

In the following statement, there is an index on the `ename` column of the `emp` table:

```
SELECT *
  FROM emp
 WHERE ename LIKE 'S%';
```

RBO Path 11: Unbounded Range Search on Indexed Columns

This access path is available if the statement's `WHERE` clause contains one of the following conditions that use either the column of a single-column index or one or more columns of a leading portion of a composite index:

```
WHERE column >[=] expr
```

```
WHERE column <[=] expr
```

Each of these conditions specifies an unbounded range of index values accessed by the statement. The range is said to be unbounded, because the condition specifies either its least value or its greatest value, but not both. To execute such a statement, Oracle performs a range scan on the index, and then accesses the table by rowid.

For example:

In the following statement, there is an index on the `sal` column of the `emp` table:

```
SELECT *
  FROM emp
 WHERE sal > 2000;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP
INDEX	RANGE SCAN	SAL_INDEX

In the following statement, there is a composite index on the `order` and `line` columns of the `line_items` table:

```
SELECT *
  FROM line_items
 WHERE order > 65118968;
```

The access path is available, because the `WHERE` clause uses the `order` column, a leading portion of the index.

This access path is *not* available in the following statement, in which there is an index on the `order` and `line` columns:

```
SELECT *
  FROM line_items
 WHERE line < 4;
```

The access path is not available because the `WHERE` clause only uses the `line` column, which is not a leading portion of the index.

RBO Path 12: Sort Merge Join

This access path is available for statements that join tables that are not stored together in a cluster if the statement's `WHERE` clause uses columns from each table in equality conditions. To execute such a statement, Oracle uses a sort-merge operation. Oracle can also use a nested loops operation to execute a join statement.

See Also: ["Understanding Joins"](#) on page 1-40 for information on these operations

For example:

In the following statement, the `emp` and `dept` tables are not stored in the same cluster:

```
SELECT *
  FROM emp, dept
 WHERE emp.deptno = dept.deptno;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
MERGE JOIN		
SORT	JOIN	
TABLE ACCESS	FULL	EMP
SORT	JOIN	
TABLE ACCESS	FULL	DEPT

RBO Path 13: MAX or MIN of Indexed Column

This access path is available for a `SELECT` statement, and all of the following conditions are true:

- The query uses the `MAX` or `MIN` function to select the maximum or minimum value of either the column of a single-column index or the leading column of a composite index. The index cannot be a cluster index. The argument to the `MAX` or `MIN` function can be any expression involving the column, a constant, or the addition operator (+), the concatenation operation (||), or the `CONCAT` function.
- There are no other expressions in the select list.
- The statement has no `WHERE` clause or `GROUP BY` clause.

To execute the query, Oracle performs a full scan of the index to find the maximum or minimum indexed value. Because only this value is selected, Oracle need not access the table after scanning the index.

For example, in the following statement, there is an index on the `sal` column of the `emp` table:

```
SELECT MAX(sal) FROM emp;
```

The `EXPLAIN PLAN` output for this statement might look like this:

```
0      SELECT STATEMENT Optimizer=CHOOSE
1    0   SORT (AGGREGATE)
2    1   INDEX (FULL SCAN (MIN/MAX)) OF 'SAL_INDEX' (NON-UNIQUE)
```

RBO Path 14: ORDER BY on Indexed Column

This access path is available for a `SELECT` statement, and all of the following conditions are true:

- The query contains an `ORDER BY` clause that uses either the column of a single-column index or a leading portion of a composite index. The index cannot be a cluster index.
- There is a `PRIMARY KEY` or `NOT NULL` integrity constraint that guarantees that at least one of the indexed columns listed in the `ORDER BY` clause contains no nulls.
- The `NLS_SORT` initialization parameter is set to `BINARY`.

To execute the query, Oracle performs a range scan of the index to retrieve the rowids of the selected rows in sorted order. Oracle then accesses the table by these rowids.

For example:

In the following statement, there is a primary key on the `empno` column of the `emp` table:

```
SELECT *
  FROM emp
 ORDER BY empno;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	BY ROWID	EMP
INDEX	RANGE SCAN	PK_EMP

`pk_emp` is the name of the index that enforces the primary key. The primary key ensures that the column does not contain nulls.

RBO Path 15: Full Table Scan

This access path is available for any SQL statement, regardless of its `WHERE` clause conditions, except when its `FROM` clause contains `SAMPLE` or `SAMPLE BLOCK`.

Note that the full table scan is the lowest ranked access path on the list. This means that the RBO always chooses an access path that uses an index if one is available, even if a full table scan might execute faster.

The following conditions make index access paths unavailable:

- `column1 > column2`
- `column1 < column2`
- `column1 >= column2`
- `column1 <= column2`

where *column1* and *column2* are in the same table.

- `column IS NULL`
- `column IS NOT NULL`
- `column NOT IN`
- `column != expr`
- `column LIKE '%pattern'`

regardless of whether *column* is indexed.

- `expr = expr2`

where *expr* is an expression that operates on a column with an operator or function, regardless of whether the column is indexed.

- `NOT EXISTS` subquery
- `ROWNUM` pseudocolumn in a view
- Any condition involving a column that is not indexed

Any SQL statement that contains only these constructs and no others that make index access paths available must use full table scans.

For example: The following statement uses a full table scan to access the `emp` table:

```
SELECT *
  FROM emp;
```

The `EXPLAIN PLAN` output for this statement might look like this:

OPERATION	OPTIONS	OBJECT_NAME

SELECT STATEMENT		
TABLE ACCESS	FULL	EMP

Choosing Execution Plans for Joins with the RBO

Note: The following considerations apply to both the cost-based and rule-based approaches:

- The optimizer first determines whether joining two or more of the tables definitely results in a row source containing at most one row. The optimizer recognizes such situations based on `UNIQUE` and `PRIMARY KEY` constraints on the tables. If such a situation exists, then the optimizer places these tables first in the join order. The optimizer then optimizes the join of the remaining set of tables.
 - For join statements with outer join conditions, the table with the outer join operator must come after the other table in the condition in the join order. The optimizer does not consider join orders that violate this rule.
-
-

With the rule-based approach, the optimizer performs the following steps to choose an execution plan for a statement that joins R tables:

1. The optimizer generates a set of R join orders, each with a different table as the first table. The optimizer generates each potential join order using this algorithm:
 - To fill each position in the join order, the optimizer chooses the table with the most highly ranked available access path according to the ranks for access paths described in "[Understanding Access Paths for the RBO](#)" on page 8-3. The optimizer repeats this step to fill each subsequent position in the join order.
 - For each table in the join order, the optimizer also chooses the operation with which to join the table to the previous table or row source in the order. The optimizer does this by ranking the sort-merge operation as access path 12 and applying these rules:
 - * If the access path for the chosen table is ranked 11 or better, then the optimizer chooses a nested loops operation using the previous table or row source in the join order as the outer table.
 - * If the access path for the table is ranked lower than 12, and if there is an equijoin condition between the chosen table and the previous table or row source in join order, then the optimizer chooses a sort-merge operation.
 - * If the access path for the chosen table is ranked lower than 12, and if there is not an equijoin condition, then the optimizer chooses a nested loops operation with the previous table or row source in the join order as the outer table.
2. The optimizer then chooses among the resulting set of execution plans. The goal of the optimizer's choice is to maximize the number of nested loops join operations in which the inner table is accessed using an index scan. Because a nested loops join involves accessing the inner table many times, an index on the inner table can greatly improve the performance of a nested loops join.

Usually, the optimizer does not consider the order in which tables appear in the FROM clause when choosing an execution plan. The optimizer makes this choice by applying the following rules in order:

- The optimizer chooses the execution plan with the fewest nested-loops operations in which the inner table is accessed with a full table scan.

- If there is a tie, then the optimizer chooses the execution plan with the fewest sort-merge operations.
- If there is still a tie, then the optimizer chooses the execution plan for which the first table in the join order has the most highly ranked access path:
 - * If there is a tie among multiple plans whose first tables are accessed by the single-column indexes access path, then the optimizer chooses the plan whose first table is accessed with the most merged indexes.
 - * If there is a tie among multiple plans whose first tables are accessed by bounded range scans, then the optimizer chooses the plan whose first table is accessed with the greatest number of leading columns of the composite index.
- If there is still a tie, then the optimizer chooses the execution plan for which the first table appears later in the query's `FROM` clause.

Transforming and Optimizing Statements with the RBO

SQL is a very flexible query language; often, there are many statements you could use to achieve the same goal. Sometimes, the optimizer transforms one such statement into another that achieves the same goal if the second statement can be executed more efficiently.

This section discusses the following topics:

- [Transforming ORs into Compound Queries with the RBO](#)
- [Using Alternative SQL Syntax](#)

Transforming ORs into Compound Queries with the RBO

If a query contains a `WHERE` clause with multiple conditions combined with `OR` operators, then the optimizer transforms it into an equivalent compound query that uses the `UNION ALL` set operator if this makes it execute more efficiently:

- If each condition individually makes an index access path available, then the optimizer can make the transformation. The optimizer chooses an execution plan for the resulting statement that accesses the table multiple times using the different indexes, and then puts the results together.
- If any condition requires a full table scan because it does not make an index available, then the optimizer does not transform the statement. The optimizer

chooses a full table scan to execute the statement, and Oracle tests each row in the table to determine whether it satisfies any of the conditions.

See Also: ["Understanding Access Paths for the RBO"](#) on page 8-3 and ["How the CBO Transforms ORs into Compound Queries"](#) on page 2-30 for information on access paths and how indexes make them available

With the RBO, the optimizer makes this UNION ALL transformation, because each component query of the resulting compound query can be executed using an index. The RBO assumes that executing the compound query using two index scans is faster than executing the original query using a full table scan.

Using Alternative SQL Syntax

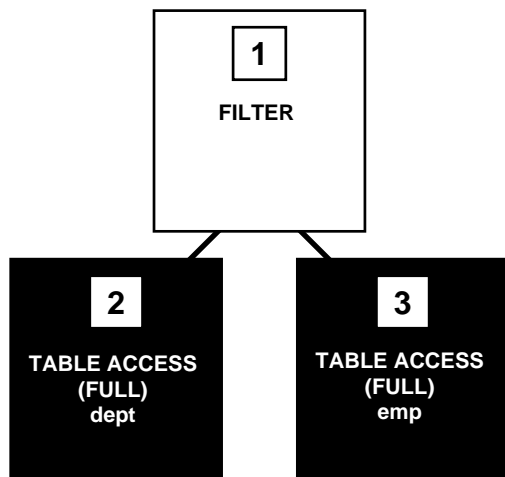
Because SQL is a flexible language, more than one SQL statement can meet the needs of an application. Although two SQL statements can produce the same result, Oracle might process one faster than the other. You can use the results of the EXPLAIN PLAN statement to compare the execution plans and costs of the two statements and determine which is more efficient.

This example shows the execution plans for two SQL statements that perform the same function. Both statements return all the departments in the dept table that have no employees in the emp table. Each statement searches the emp table with a subquery. Assume there is an index, deptno_index, on the deptno column of the emp table.

The first statement and its execution plan:

```
SELECT dname, deptno
   FROM dept
  WHERE deptno NOT IN
      (SELECT deptno FROM emp);
```

The execution plan for the transformed statement might look like the illustration in [Figure 8-1](#). The shaded boxes indicate steps that physically retrieve data and the clear boxes indicate steps that operate on data returned from the previous step.

Figure 8–1 Execution Plan with Two Full Table Scans

Step 3 of the output indicates that Oracle executes this statement by performing a full table scan of the `emp` table despite the index on the `deptno` column. This full table scan can be a time-consuming operation. Oracle does not use the index, because the subquery that searches the `emp` table does not have a `WHERE` clause that makes the index available.


However, this SQL statement selects the same rows by accessing the index:

```

SELECT dname, deptno
FROM dept
WHERE NOT EXISTS
  (SELECT deptno
   FROM emp
   WHERE dept.deptno = emp.deptno);
  
```

The execution plan for the transformed statement might look like the illustration in [Figure 8–2](#). The shaded boxes indicate steps that physically retrieve data and the clear boxes indicate steps that operate on data returned from the previous step.

Figure 8–2 Execution Plan with a Full Table Scan and an Index Scan



The `WHERE` clause of the subquery refers to the `deptno` column of the `emp` table, so the index `deptno_index` is used. The use of the index is reflected in step 3 of the execution plan. The index range scan of `deptno_index` takes less time than the full scan of the `emp` table in the first statement. Furthermore, the first query performs one full scan of the `emp` table for every `deptno` in the `dept` table. For these reasons, the second SQL statement is faster than the first.

If you have statements in an application that use the `NOT IN` operator, as the first query in this example does, then consider rewriting them so that they use the `NOT EXISTS` operator. This allows such statements to use an index if one exists.

Note: Alternative SQL syntax is effective only with the rule-based optimizer.

Part II

SQL-Related Performance Tools

Part II discusses Oracle's SQL-related performance tools. These tools examine the execution plan for a SQL statement, and determine whether the statement can be better optimized. You can get the execution plan from the `EXPLAIN PLAN SQL` statement, from querying `V$SQL_PLAN`, or from SQL trace.

In the development phase, use `EXPLAIN PLAN` to determine a good access plan, and then verify that it is the optimal plan through volume data testing. When evaluating a plan, examine the statement's actual resource consumption using `V$SQLAREA` with `V$SQL_PLAN`, Oracle Trace, or the SQL trace facility and `TKPROF`. The information in the `V$SQL_PLAN` view is very similar to the output of an `EXPLAIN PLAN` statement. However, `EXPLAIN PLAN` shows a theoretical plan that can be used if the statement were to be executed, whereas `V$SQL_PLAN` contains the actual plan used. Hence, querying `V$SQLAREA` in conjunction with `V$SQL_PLAN` provides similar results to using SQL Trace with `TKPROF`.

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#)

The autotrace tool lets you automatically get a report on the execution path used by the SQL optimizer and the statement execution statistics. It is useful for monitoring and tuning the performance of these statements. Oracle Trace is a GUI, event-driven data collection product, which the Oracle server uses to collect performance and resource utilization data.

The chapters in this part are:

- [Chapter 9, "Using EXPLAIN PLAN"](#)
- [Chapter 10, "Using SQL Trace and TKPROF"](#)
- [Chapter 11, "Using Autotrace in SQL*Plus"](#)
- [Chapter 12, "Using Oracle Trace"](#)

Using EXPLAIN PLAN

This chapter introduces execution plans, describes the SQL command `EXPLAIN PLAN`, and explains how to interpret its output. This chapter also provides procedures for managing outlines to control application performance characteristics.

This chapter contains the following sections:

- [Understanding EXPLAIN PLAN](#)
- [Creating the PLAN_TABLE Output Table](#)
- [Running EXPLAIN PLAN](#)
- [Displaying PLAN_TABLE Output](#)
- [Reading EXPLAIN PLAN Output](#)
- [Viewing Bitmap Indexes with EXPLAIN PLAN](#)
- [Viewing Partitioned Objects with EXPLAIN PLAN](#)
- [Viewing Parallel Execution with EXPLAIN PLAN](#)
- [CPU Costing Model](#)
- [EXPLAIN PLAN Restrictions](#)
- [PLAN_TABLE Columns](#)

See Also: *Oracle9i SQL Reference* for the syntax of `EXPLAIN PLAN`

Understanding EXPLAIN PLAN

The `EXPLAIN PLAN` statement displays execution plans chosen by the Oracle optimizer for `SELECT`, `UPDATE`, `INSERT`, and `DELETE` statements. A statement's execution plan is the sequence of operations Oracle performs to run the statement.

The row source tree is the core of the execution plan. It shows the following information:

- An ordering of the tables referenced by the statement
- An access method for each table mentioned in the statement
- A join method for tables affected by join operations in the statement
- Data operations like filter, sort, or aggregation

In addition to the row source tree, the plan table contains information about the following:

- Optimization, such as the cost and cardinality of each operation
- Partitioning, such as the set of accessed partitions
- Parallel execution, such as the distribution method of join inputs

The `EXPLAIN PLAN` results let you determine whether the optimizer selects a particular execution plan, such as, nested loops join. It also helps you to understand the optimizer decisions, such as why the optimizer chose a nested loops join instead of a hash join, and lets you understand the performance of a query.

Note: Oracle Performance Manager charts and Oracle SQL Analyze can automatically create and display explain plans for you. For more information on using explain plans, see *Database Tuning with the Oracle Tuning Pack*.

How Execution Plans Can Change

With the cost-based optimizer, execution plans can and do change as the underlying costs change. `EXPLAIN PLAN` output shows how Oracle runs the SQL statement when the statement was explained. This can differ from the plan during actual execution for a SQL statement, because of differences in the execution environment and explain plan environment.

Execution plans can differ due to the following:

- [Different Schemas](#)
- [Different Costs](#)

Different Schemas

- The execution and explain plan happen on different databases.
- The user explaining the statement is different from the user running the statement. Two users might be pointing to different objects in the same database, resulting in different execution plans.
- Schema changes (usually changes in indexes) between the two operations.

Different Costs

Even if the schemas are the same, the optimizer can choose different execution plans if the costs are different. Some factors that affect the costs include the following:

- Data volume and statistics
- Bind variable types
- Initialization parameters - set globally or at session level

Minimizing Throw-Away

Examining an explain plan lets you look for throw-away in cases such as the following:

- Full scans
- Unselective range scans
- Late predicate filters
- Wrong join order
- Late filter operations

For example, in the following explain plan, the last step is a very unselective range scan that is executed 76563 times, accesses 11432983 rows, throws away 99% of them, and retains 76563 rows. Why access 11432983 rows to realize that only 76563 rows are needed?

Example 9–1 Looking for Throw-Away in an Explain Plan

Rows	Execution Plan

12	SORT AGGREGATE
2	SORT GROUP BY
76563	NESTED LOOPS
76575	NESTED LOOPS
19	TABLE ACCESS FULL CN_PAYRUNS_ALL
76570	TABLE ACCESS BY INDEX ROWID CN_POSTING_DETAILS_ALL
76570	INDEX RANGE SCAN (object id 178321)
76563	TABLE ACCESS BY INDEX ROWID CN_PAYMENT_WORKSHEETS_ALL
11432983	INDEX RANGE SCAN (object id 186024)

Looking Beyond Execution Plans

The execution plan operation alone cannot differentiate between well-tuned statements and those that perform poorly. For example, an EXPLAIN PLAN output that shows that a statement uses an index does not necessarily mean that the statement runs efficiently. Sometimes indexes can be extremely inefficient. In this case, you should examine the following:

- The columns of the index being used
- Their selectivity (fraction of table being accessed)

It is best to use EXPLAIN PLAN to determine an access plan, and then later prove that it is the optimal plan through testing. When evaluating a plan, examine the statement’s actual resource consumption. Use Oracle Trace or the SQL trace facility and TKPROF to examine individual SQL statement performance.

See Also: [Chapter 10, "Using SQL Trace and TKPROF"](#) for information on TKPROF interpretation

Creating the PLAN_TABLE Output Table

Before issuing an EXPLAIN PLAN statement, you must have a table to hold its output. PLAN_TABLE is the default sample output table into which the EXPLAIN PLAN statement inserts rows describing execution plans. Use the SQL script UTLXPLAN.SQL to create the PLAN_TABLE in your schema. The exact name and location of this script depends on your operating system. On Unix, it is located in the \$ORACLE_HOME/rdbms/admin directory.

For example, run the commands in [Example 9–2](#) from a SQL*Plus session to create the PLAN_TABLE in the HR schema.

Example 9–2 Creating a PLAN_TABLE

```
CONNECT HR/your_password
@$ORACLE_HOME/RDBMS/ADMIN/UTLXPLAN.SQL
```

Table created.

Oracle Corporation recommends that you drop and rebuild the `PLAN_TABLE` table after upgrading the version of the database because the columns might change. This can cause scripts to fail or cause `TKPROF` to fail, if you are specifying the table.

If you want an output table with a different name, then create `PLAN_TABLE` and rename it with the `RENAME SQL` statement.

Running EXPLAIN PLAN

To explain a SQL statement, use the following:

```
EXPLAIN PLAN FOR
  SQL_Statement
```

For example:

```
EXPLAIN PLAN FOR
SELECT last_name FROM employees;
```

This explains the plan into the `PLAN_TABLE` table. You can then select the execution plan from `PLAN_TABLE`. This is useful if you do not have any other plans in `PLAN_TABLE`, or if you only want to look at the last statement.

Identifying Statements for EXPLAIN PLAN

With multiple statements, you can specify a statement identifier and use that to identify your specific execution plan. Before using `SET STATEMENT ID`, remove any existing rows for that statement ID.

In [Example 9–3](#), `bad1` is specified as the statement identifier:

Example 9–3 Using EXPLAIN PLAN with the STATEMENT ID Clause

```
EXPLAIN PLAN
  SET STATEMENT_ID = 'bad1' FOR
SELECT last_name FROM employees;
```

Specifying Different Tables for EXPLAIN PLAN

You can specify the INTO clause to specify a different table.

Example 9–4 Using EXPLAIN PLAN with the INTO Clause

```
EXPLAIN PLAN
  INTO my_plan_table
  FOR
SELECT last_name FROM employees;
```

You can specify a statement Id when using the INTO clause.

```
EXPLAIN PLAN
  INTO my_plan_table
  SET STATEMENT_ID = 'bad1' FOR
SELECT last_name FROM employees;
```

See Also: *Oracle9i SQL Reference* for a complete description of EXPLAIN PLAN syntax.

Displaying PLAN_TABLE Output

After you have explained the plan, use the two scripts provided by Oracle to display the most recent plan table output:

- UTLXPLS.SQL - Shows plan table output for serial processing
- UTLXPLP.SQL - Shows plan table output with parallel execution columns

[Example 1–4, "EXPLAIN PLAN Output"](#) on page 1-20 is an example of the plan table output when using the UTLXPLS.SQL script.

If you have specified a statement identifier, then you can write your own script to query the PLAN_TABLE. For example:

- Start with ID = 0 and given STATEMENT_ID.
- Use the CONNECT BY clause to walk the tree from parent to child, the join keys being STATEMENT_ID = PRIOR STATEMENT_ID and PARENT_ID = PRIOR ID.
- Use the pseudo-column LEVEL (associated with CONNECT BY) to indent the children.

```
SELECT cardinality "Rows",
       lpad(' ',level-1)||operation||' '||
       options||' '||object_name "Plan"
FROM PLAN_TABLE
```



```
CONNECT BY prior id = parent_id
          AND prior statement_id = statement_id
START WITH id = 0
          AND statement_id = 'bad1'
ORDER BY id;
```

Rows Plan

```
SELECT STATEMENT
TABLE ACCESS FULL EMPLOYEES
```

The NULL in the Rows column indicates that the optimizer does not have any statistics on the table. Analyzing the table shows the following:

Rows Plan

```
16957 SELECT STATEMENT
16957 TABLE ACCESS FULL EMPLOYEES
```

You can also select the COST. This is useful for comparing execution plans or for understanding why the optimizer chooses one execution plan over another.

Note: These simplified examples are not valid for recursive SQL.

Reading EXPLAIN PLAN Output

This section uses progressively complex examples to illustrate execution plans.

See Also: [Appendix A, "Schemas Used in Performance Examples"](#)

The statement in [Example 9-5](#) is used to display the execution plan.

Example 9-5 Statement to display the EXPLAIN PLAN

```
SELECT lpad(' ',level-1)||operation||' '||options||' '||
       object_name "Plan"
FROM plan_table
CONNECT BY prior id = parent_id
          AND prior statement_id = statement_id
START WITH id = 0 AND statement_id = '&1'
ORDER BY id;
```

EXPLAIN PLAN Examples

The following are EXPLAIN PLAN examples.

Example 9-6 *EXPLAIN PLAN example_plan1*

```
EXPLAIN PLAN SET statement_id = 'example_plan1' FOR
SELECT full_name FROM per_all_people_f
WHERE UPPER(full_name) LIKE 'Pe%' ;
```

Plan

```
-----
SELECT STATEMENT
TABLE ACCESS FULL PER_ALL_PEOPLE_F
```

This plan shows execution of a SELECT statement. The table `per_all_people_f` is accessed using a full table scan.

- Every row in the table `per_all_people_f` is accessed, and the WHERE clause criteria is evaluated for every row.
- The SELECT statement returns the rows meeting the WHERE clause criteria.

Example 9-7 *EXPLAIN PLAN example_plan2*

```
EXPLAIN PLAN SET statement_id = 'example_plan2' FOR
SELECT full_name FROM per_all_people_f
WHERE full_name LIKE 'Pe%' ;
```

Plan

```
-----
SELECT STATEMENT
TABLE ACCESS BY INDEX ROWID PER_ALL_PEOPLE_F
INDEX RANGE SCAN PER_PEOPLE_F_N54
```

This plan shows execution of a SELECT statement.

- Index `per_people_f_n54` is used in a range scan operation.
- The table `per_all_people_f` is accessed through ROWID. ROWIDs are obtained from the index in the previous step for keys that meet the WHERE clause criteria. When the table is accessed, any additional WHERE clause conditions that could not be evaluated during the range scan (because the column is present in the table and not in the index) are also evaluated.
- The SELECT statement returns rows satisfying the WHERE clause conditions (evaluated in previous steps).

Example 9–8 EXPLAIN PLAN example_plan3

```
EXPLAIN PLAN SET statement_id = 'example_plan3' FOR
SELECT segment1, segment2, description, inventory_item_id
   FROM mtl_system_items msi
   WHERE segment1 = :b1
        AND segment2 LIKE '%-BOM'
AND NVL(end_date_active,sysdate+1) > SYSDATE ;
```

Plan

```
-----
SELECT STATEMENT
  TABLE ACCESS BY INDEX ROWID MTL_SYSTEM_ITEMS
    INDEX RANGE SCAN MTL_SYSTEM_ITEMS_N8
```

This plan shows execution of a SELECT statement.

- Index `mtl_system_items_n8` is used in a range scan operation. This is an index on (`segment1`, `segment2`, `segment3`). The range scan happens using the following condition:

```
segment1 = :b1
```

The rows that come out of this step satisfy all the WHERE clause criteria that can be evaluated with the index columns. Therefore, the following condition is also evaluated at this stage:

```
segment2 LIKE '%-BOM'
```

- The table `per_all_people_f` is accessed through ROWIDs obtained from the index in the previous step. When the table is accessed, any additional WHERE clause conditions that could not be evaluated during the range scan (because the column is present in the table and not in the index) are also evaluated. Therefore, the following condition is evaluated at this stage:

```
NVL(end_date_active,sysdate+1) > SYSDATE
```

- The SELECT statement returns rows satisfying the WHERE clause conditions (evaluated in previous steps).

Example 9–9 EXPLAIN PLAN example_plan4

```
EXPLAIN PLAN SET statement_id = 'example_plan4' FOR
SELECT h.order_number, l.revenue_amount, l.ordered_quantity
   FROM so_headers_all h, so_lines_all l
   WHERE h.customer_id = :b1
```

```
AND h.date_ordered > SYSDATE-30
AND l.header_id = h.header_id ;
```

Plan

```
-----
SELECT STATEMENT
  NESTED LOOPS
    TABLE ACCESS BY INDEX ROWID SO_HEADERS_ALL
      INDEX RANGE SCAN SO_HEADERS_N1
    TABLE ACCESS BY INDEX ROWID SO_LINES_ALL
      INDEX RANGE SCAN SO_LINES_N1
```

This plan shows execution of a `SELECT` statement.

- Index `so_headers_n1` is used in a range scan operation. This is an index on `customer_id`. The range scan happens using the following condition:


```
customer_id = :b1
```
- The table `so_headers_all` is accessed through ROWIDs obtained from the index in the previous step. When the table is accessed, any additional `WHERE` clause conditions that could not be evaluated during the range scan (because the column is present in the table and not in the index) are also evaluated. Therefore, the following condition is evaluated at this stage:


```
h.date_ordered > sysdate-30
```
- For every row from `so_headers_all` satisfying the `WHERE` clause conditions, a range scan is run on `so_lines_n1` using the following condition:


```
l.header_id = h.header_id
```
- The table `so_lines_all` is accessed through ROWIDs obtained from the index in the previous step. When the table is accessed, any additional `WHERE` clause conditions that could not be evaluated during the range scan (because the column is present in the table and not in the index) are also evaluated. There are no additional conditions to evaluate here.
- The `SELECT` statement returns rows satisfying the `WHERE` clause conditions (evaluated in previous steps).

Viewing Bitmap Indexes with EXPLAIN PLAN

Index row sources using bitmap indexes appear in the `EXPLAIN PLAN` output with the word `BITMAP` indicating the type of the index. Consider the sample query and plan in [Example 9-10](#).

Example 9-10 *EXPLAIN PLAN with Bitmap Indexes*

```
EXPLAIN PLAN FOR
  SELECT * FROM t
  WHERE c1 = 2
  AND c2 <> 6
  OR c3 BETWEEN 10 AND 20;

SELECT STATEMENT
  TABLE ACCESS T BY INDEX ROWID
    BITMAP CONVERSION TO ROWID
      BITMAP OR
        BITMAP MINUS
          BITMAP MINUS
            BITMAP INDEX C1_IND SINGLE VALUE
            BITMAP INDEX C2_IND SINGLE VALUE
          BITMAP INDEX C2_IND SINGLE VALUE
        BITMAP MERGE
          BITMAP INDEX C3_IND RANGE SCAN
```

In this example, the predicate `c1=2` yields a bitmap from which a subtraction can take place. From this bitmap, the bits in the bitmap for `c2 = 6` are subtracted. Also, the bits in the bitmap for `c2 IS NULL` are subtracted, explaining why there are two `MINUS` row sources in the plan. The `NULL` subtraction is necessary for semantic correctness unless the column has a `NOT NULL` constraint. The `TO ROWIDS` option is used to generate the `ROWIDS` that are necessary for the table access.

Note: Queries using bitmap join index indicate the bitmap join index access path. The operation for bitmap join index is the same as bitmap index.

Viewing Partitioned Objects with EXPLAIN PLAN

Use `EXPLAIN PLAN` to see how Oracle accesses partitioned objects for specific queries.

Partitions accessed after pruning are shown in the `PARTITION START` and `PARTITION STOP` columns. The row source name for the range partition is `PARTITION RANGE`. For hash partitions, the row source name is `PARTITION HASH`.

A join is implemented using partial partition-wise join if the `DISTRIBUTION` column of the plan table of one of the joined tables contains `PARTITION(KEY)`. Partial partition-wise join is possible if one of the joined tables is partitioned on its join column and the table is parallelized.

A join is implemented using full partition-wise join if the partition row source appears before the join row source in the `EXPLAIN PLAN` output. Full partition-wise joins are possible only if both joined tables are equi-partitioned on their respective join columns. Examples of execution plans for several types of partitioning follow.

Examples of Displaying Range and Hash Partitioning with EXPLAIN PLAN

Consider the following table, `emp_range`, partitioned by range on `hire_date` to illustrate how pruning is displayed. Assume that the tables `emp` and `dept` from a standard Oracle schema exist.

```
CREATE TABLE emp_range
PARTITION BY RANGE(hire_date)
(
    PARTITION emp_p1 VALUES LESS THAN (TO_DATE('1-JAN-1991','DD-MON-YYYY')),
    PARTITION emp_p2 VALUES LESS THAN (TO_DATE('1-JAN-1993','DD-MON-YYYY')),
    PARTITION emp_p3 VALUES LESS THAN (TO_DATE('1-JAN-1995','DD-MON-YYYY')),
    PARTITION emp_p4 VALUES LESS THAN (TO_DATE('1-JAN-1997','DD-MON-YYYY')),
    PARTITION emp_p5 VALUES LESS THAN (TO_DATE('1-JAN-1999','DD-MON-YYYY'))
)
AS SELECT * FROM employees;
```

For the first example, consider the following statement:

```
EXPLAIN PLAN FOR SELECT * FROM emp_range;
```

Enter the following to display the `EXPLAIN PLAN` output:

```
@?/RDBMS/ADMIN/UTLXPLS
```

Oracle displays something similar to the following:

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		105	8K	1		
PARTITION RANGE ALL					1	5
TABLE ACCESS FULL	EMP_RANGE	105	8K	1	1	5

6 rows selected.

A partition row source is created on top of the table access row source. It iterates over the set of partitions to be accessed. In this example, the partition iterator covers all partitions (option ALL), because a predicate was not used for pruning. The PARTITION_START and PARTITION_STOP columns of the PLAN_TABLE show access to all partitions from 1 to 5.

For the next example, consider the following statement:

```
EXPLAIN PLAN FOR SELECT * FROM emp_range
WHERE hire_date >= TO_DATE('1-JAN-1995','DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		3	54	1		
PARTITION RANGE ITERATOR					4	5
TABLE ACCESS FULL	EMP_RANGE	3	54	1	4	5

6 rows selected.

In the previous example, the partition row source iterates from partition 4 to 5, because we prune the other partitions using a predicate on hire_date.

Finally, consider the following statement:

```
EXPLAIN PLAN FOR SELECT * FROM emp_range
WHERE hire_date < TO_DATE('1-JAN-1991','DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		2	36	1		
TABLE ACCESS FULL	EMP_RANGE	2	36	1	1	1

5 rows selected.

In the previous example, only partition 1 is accessed and known at compile time; thus, there is no need for a partition row source.

Plans for Hash Partitioning

Oracle displays the same information for hash partitioned objects, except the partition row source name is `PARTITION HASH` instead of `PARTITION RANGE`. Also, with hash partitioning, pruning is only possible using equality or `IN`-list predicates.

Examples of Pruning Information with Composite Partitioned Objects

To illustrate how Oracle displays pruning information for composite partitioned objects, consider the table `emp_comp` that is range partitioned on `hire_date` and subpartitioned by hash on `department_id`.

```
CREATE TABLE emp_comp PARTITION BY RANGE(hire_date) SUBPARTITION BY
HASH(department_id)
SUBPARTITIONS 3
(
  PARTITION emp_p1 VALUES LESS THAN (TO_DATE('1-JAN-1991','DD-MON-YYYY')),
  PARTITION emp_p2 VALUES LESS THAN (TO_DATE('1-JAN-1993','DD-MON-YYYY')),
  PARTITION emp_p3 VALUES LESS THAN (TO_DATE('1-JAN-1995','DD-MON-YYYY')),
  PARTITION emp_p4 VALUES LESS THAN (TO_DATE('1-JAN-1997','DD-MON-YYYY')),
  PARTITION emp_p5 VALUES LESS THAN (TO_DATE('1-JAN-1999','DD-MON-YYYY'))
)
AS SELECT * FROM employees;
```

For the first example, consider the following statement:

```
EXPLAIN PLAN FOR SELECT * FROM emp_comp;
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		105	8K	1		
PARTITION RANGE ALL					1	5
PARTITION HASH ALL					1	3
TABLE ACCESS FULL	EMP_COMP	105	8K	1	1	15

7 rows selected.

This example shows the plan when Oracle accesses all subpartitions of all partitions of a composite object. Two partition row sources are used for that purpose: a range partition row source to iterate over the partitions and a hash partition row source to iterate over the subpartitions of each accessed partition.

In the following example, the range partition row source iterates from partition 1 to 5, because no pruning is performed. Within each partition, the hash partition row source iterates over subpartitions 1 to 3 of the current partition. As a result, the table access row source accesses subpartitions 1 to 15. In other words, it accesses all subpartitions of the composite object.

```
EXPLAIN PLAN FOR SELECT * FROM emp_comp
WHERE hire_date = TO_DATE('15-FEB-1997', 'DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		1	96	1		
PARTITION HASH ALL					1	3
TABLE ACCESS FULL	EMP_COMP	1	96	1	13	15

6 rows selected.

In the previous example, only the last partition, partition 5, is accessed. This partition is known at compile time, so we do not need to show it in the plan. The hash partition row source shows accessing of all subpartitions within that partition; that is, subpartitions 1 to 3, which translates into subpartitions 13 to 15 of the `emp_comp` table.

Now consider the following statement:

```
EXPLAIN PLAN FOR SELECT * FROM emp_comp WHERE department_id = 20;
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		2	200	1		
PARTITION RANGE ALL					1	5
TABLE ACCESS FULL	EMP_COMP	2	200	1		

6 rows selected.

In the previous example, the predicate `deptno = 20` enables pruning on the hash dimension within each partition, so Oracle only needs to access a single subpartition. The number of that subpartition is known at compile time, so the hash partition row source is not needed.

Finally, consider the following statement:

```
VARIABLE dno NUMBER;
EXPLAIN PLAN FOR SELECT * FROM emp_comp WHERE department_id = :dno;
```

Plan Table

Operation	Name	Rows	Bytes	Cost	Pstart	Pstop
SELECT STATEMENT		2	200	1		
PARTITION RANGE ALL					1	5
PARTITION HASH SINGLE					KEY	KEY
TABLE ACCESS FULL	EMP_COMP	2	200	1		

7 rows selected.

The last two examples are the same, except that `deptno = 20` has been replaced by `department_id = :dno`. In this last case, the subpartition number is unknown at compile time, and a hash partition row source is allocated. The option is `SINGLE` for that row source, because Oracle accesses only one subpartition within each partition. The `PARTITION_START` and `PARTITION_STOP` is set to `KEY`. This means that Oracle determines the number of the subpartition at run time.

Examples of Partial Partition-wise Joins

In the following example, `emp_range` is joined on the partitioning column and is parallelized. This enables use of partial partition-wise join, because the `dept` table is not partitioned. Oracle dynamically partitions the `dept` table before the join.

```
ALTER TABLE emp PARALLEL 2;
Table altered.
ALTER TABLE dept PARALLEL 2;
Table altered.
```

To show the plan for the query, enter:

```
EXPLAIN PLAN FOR SELECT /*+ ORDERED USE_HASH(D) */ ename, dname
FROM emp_range e, dept d
WHERE e.deptno = d.deptno
      AND e.hire_date > TO_DATE('29-JUN-1996', 'DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	TQ	IN-OUT	PQ Distrib	Pstart	Pstop
SELECT STATEMENT		1	51	3					
HASH JOIN		1	51	3	2,02	P->S	QC (RANDOM)		
PARTITION RANGE ITERATOR					2,02	PCWP		4	5
TABLE ACCESS FULL	EMP_RANGE	3	87	1	2,00	PCWP		4	5
TABLE ACCESS FULL	DEPT	21	462	1	2,01	P->P	PART (KEY)		

8 rows selected.

The plan shows that the optimizer selects partition-wise join, because the PQ Distrib column contains the text PART (KEY), or partition key.

In the next example, emp_comp is joined on its hash partitioning column, deptno, and is parallelized. This enables use of partial partition-wise join, because the dept table is not partitioned. Again, Oracle dynamically partitions the dept table.

```
ALTER TABLE emp_comp PARALLEL 2;
Table altered.
EXPLAIN PLAN FOR SELECT /*+ ORDERED USE_HASH(D) */ ename, dname
FROM emp_comp e, dept d
WHERE e.deptno = d.deptno
AND e.hiredate > TO_DATE('13-MAR-1995', 'DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	TQ	IN-OUT	PQ Distrib	Pstart	Pstop
SELECT STATEMENT		1	51	3					
HASH JOIN		1	51	3	0,01	P->S	QC (RANDOM)		
PARTITION RANGE ITERATOR					0,01	PCWP		4	5
PARTITION HASH ALL					0,01	PCWP		1	3
TABLE ACCESS FULL	EMP_COMP	3	87	1	0,01	PCWP		10	15
TABLE ACCESS FULL	DEPT	21	462	1	0,00	P->P	PART (KEY)		

9 rows selected.

Examples of Full Partition-wise Joins

In the following example, emp_comp and dept_hash are joined on their hash partitioning columns. This enables use of full partition-wise join. The PARTITION HASH row source appears on top of the join row source in the plan table output.

To create the table dept_hash, enter:

```
CREATE TABLE dept_hash
PARTITION BY HASH(deptno)
PARTITIONS 3
```

```
PARALLEL
AS SELECT * FROM dept;
```

To show the plan for the query, enter:

```
EXPLAIN PLAN FOR SELECT /*+ ORDERED USE_HASH(D) */ ename, dname
FROM emp_comp e, dept_hash d
WHERE e.deptno = d.deptno
AND e.hiredate > TO_DATE('29-JUN-1996', 'DD-MON-YYYY');
```

Plan Table

Operation	Name	Rows	Bytes	Cost	TQ	IN-OUT	PQ Distrib	Pstart	Pstop
SELECT STATEMENT		2	102	2					
PARTITION HASH ALL					4,00	PCWP		1	3
HASH JOIN		2	102	2	4,00	P->S	QC (RANDOM)		
PARTITION RANGE ITERATOR					4,00	PCWP		4	5
TABLE ACCESS FULL	EMP_COMP	3	87	1	4,00	PCWP		10	15
TABLE ACCESS FULL	DEPT_HASH	63	1K	1	4,00	PCWP		1	3

9 rows selected.

Examples of INLIST ITERATOR and EXPLAIN PLAN

An INLIST ITERATOR operation appears in the EXPLAIN PLAN output if an index implements an IN-list predicate. For example:

```
SELECT * FROM emp WHERE empno IN (7876, 7900, 7902);
```

The EXPLAIN PLAN output appears as follows:

OPERATION	OPTIONS	OBJECT_NAME
SELECT STATEMENT		
INLIST ITERATOR		
TABLE ACCESS	BY ROWID	EMP
INDEX	RANGE SCAN	EMP_EMPNO

The INLIST ITERATOR operation iterates over the next operation in the plan for each value in the IN-list predicate. For partitioned tables and indexes, the three possible types of IN-list columns are described in the following sections.

When the IN-List Column is an Index Column

If the IN-list column `empno` is an index column but not a partition column, then the plan is as follows (the IN-list operator appears before the table operation but after the partition operation):

OPERATION	OPTIONS	OBJECT_NAME	PARTITION_START	PARTITION_STOP
SELECT STATEMENT				
PARTITION RANGE	ALL		KEY(INLIST)	KEY(INLIST)
INLIST ITERATOR				
TABLE ACCESS	BY LOCAL INDEX ROWID	EMP	KEY(INLIST)	KEY(INLIST)
INDEX	RANGE SCAN	EMP_EMPNO	KEY(INLIST)	KEY(INLIST)

The `KEY(INLIST)` designation for the partition start and stop keys specifies that an IN-list predicate appears on the index start/stop keys.

When the IN-List Column is an Index and a Partition Column

If `empno` is an indexed and a partition column, then the plan contains an INLIST ITERATOR operation before the partition operation:

OPERATION	OPTIONS	OBJECT_NAME	PARTITION_START	PARTITION_STOP
SELECT STATEMENT				
INLIST ITERATOR				
PARTITION RANGE	ITERATOR		KEY(INLIST)	KEY(INLIST)
TABLE ACCESS	BY LOCAL INDEX ROWID	EMP	KEY(INLIST)	KEY(INLIST)
INDEX	RANGE SCAN	EMP_EMPNO	KEY(INLIST)	KEY(INLIST)

When the IN-List Column is a Partition Column

If `empno` is a partition column and there are no indexes, then no INLIST ITERATOR operation is allocated:

OPERATION	OPTIONS	OBJECT_NAME	PARTITION_START	PARTITION_STOP
SELECT STATEMENT				
PARTITION RANGE	INLIST		KEY(INLIST)	KEY(INLIST)
TABLE ACCESS	FULL	EMP	KEY(INLIST)	KEY(INLIST)

If `emp_empno` is a bitmap index, then the plan is as follows:

OPERATION	OPTIONS	OBJECT_NAME
SELECT STATEMENT		
INLIST ITERATOR		

```
TABLE ACCESS          BY INDEX ROWID      EMP
BITMAP CONVERSION    TO ROWIDS
BITMAP INDEX          SINGLE VALUE          EMP_EMPNO
```

Example of Domain Indexes and EXPLAIN PLAN

You can also use EXPLAIN PLAN to derive user-defined CPU and I/O costs for domain indexes. EXPLAIN PLAN displays these statistics in the OTHER column of PLAN_TABLE.

For example, assume table emp has user-defined operator CONTAINS with a domain index emp_resume on the resume column, and the index type of emp_resume supports the operator CONTAINS. Then the query:

```
SELECT * FROM emp WHERE CONTAINS(resume, 'Oracle') = 1
```


might display the following plan:

```
OPERATION              OPTIONS              OBJECT_NAME          OTHER
-----              -
SELECT STATEMENT
TABLE ACCESS            BY ROWID             EMP
DOMAIN INDEX            EMP_RESUME           CPU: 300, I/O: 4
```

Viewing Parallel Execution with EXPLAIN PLAN

Tuning a parallel query begins much like a non-parallel query tuning exercise by choosing the driving table. However, the rules governing the choice are different. In the non-parallel case, the best driving table is typically the one that produces fewest number of rows after limiting conditions are applied. The small number of rows are joined to larger tables using non-unique indexes. For example, consider a table hierarchy consisting of CUSTOMER, ACCOUNT, and TRANSACTION.

Figure 9-1 A Table Hierarchy



`CUSTOMER` is the smallest table while `TRANSACTION` is the largest. A typical OLTP query might be to retrieve transaction information about a particular customer's account. The query would drive from the `CUSTOMER` table. The goal in this case is to minimize logical I/O, which typically minimizes other critical resources including physical I/O and CPU time.

For parallel queries, the choice of the driving table is usually the largest table since parallel query can be utilized. Obviously, it would not be efficient to use parallel query on the this query, because only a few rows from each table are ultimately accessed. However, what if it were necessary to identify all customers that had transactions of a certain type last month? It would be more efficient to drive from the `TRANSACTION` table since there are no limiting conditions on the customer table. The rows from the `TRANSACTION` table would be joined to the `ACCOUNT` table, and finally to the `CUSTOMER` table. In this case, the indexes utilized on the `ACCOUNT` and `CUSTOMER` table are likely to be highly selective primary key or unique indexes, rather than non-unique indexes used in the first query. Since the `TRANSACTION` table is large and the column is un-selective, it would be beneficial to utilize parallel query driving from the `TRANSACTION` table.

```
Parallel operation
-----
PARALLEL_TO_SERIAL
PARALLEL_TO_PARALLEL
PARALLEL_COMBINED_WITH_PARENT
PARALLEL_FROM_SERIAL
PARALLEL_COMBINED_WITH_PARENT
PARALLEL_TO_PARALLEL
```

`PARALLEL_TO_PARALLEL` operations generally produce the best performance as long as the workloads in each step are relatively equivalent.

A `PARALLEL_COMBINED_WITH_PARENT` operation occurs when the step is performed simultaneously with the parent step.

A `PARALLEL_TO_SERIAL` operation which is always the step that occurs when rows from a parallel operation are consumed by the query coordinator. Another type of operation that does not occur in this query is a `SERIAL` operation. If these types of operations occur, consider making them parallel operations to improve performance since they too are potential bottlenecks.

If a parallel step produces many rows, the query coordinator may not be able to consume them as fast as they are being produced. There is little that can be done to improve this.

CPU Costing Model

Every database operation uses the CPU. In most cases, CPU utilization is as important as I/O; often it is the only contribution to the cost (in cases of in-memory sort, hash, predicate evaluation, and cached I/O). In Oracle9i the optimizer introduces a new model, which includes the cost of CPU utilization. Including CPU utilization in the cost model helps generate better plans.

According to the CPU costing model:

$$\text{Cost} = \left(\begin{array}{l} \#SRds * \text{sreadtim} + \\ \#MRds * \text{mreadtim} + \\ \#CPUCycles / \text{cpuspeed} \end{array} \right) / \text{sreadtim}$$

where:

- #SRDs is the number of single block reads
- #MRDs is the number of multi block reads
- #CPUCycles is the number of CPU Cycles *)
- sreadtim is the single block read time
- mreadtim is the multi block read time
- cpuspeed is the CPU cycles per second

CPUCycles includes CPU cost of query processing (pure CPU cost) and CPU cost of data retrieval (CPU cost of the buffer cache get).

This model is straightforward for serial execution. For parallel execution, necessary adjustments are made while computing estimates for #SRD, #MRD, and #CPUCycles.

EXPLAIN PLAN Restrictions

Oracle does not support EXPLAIN PLAN for statements performing implicit type conversion of date bind variables. With bind variables in general, the EXPLAIN PLAN output might not represent the real execution plan.

From the text of a SQL statement, TKPROF cannot determine the types of the bind variables. It assumes that the type is CHARACTER, and gives an error message if this is not the case. You can avoid this limitation by putting appropriate type conversions in the SQL statement.

See Also: [Chapter 10, "Using SQL Trace and TKPROF"](#)

PLAN_TABLE Columns

The `PLAN_TABLE` used by the `EXPLAIN PLAN` statement contains the columns listed in [Table 9-1](#).

Table 9-1 *PLAN_TABLE Columns*

Column	Type	Description
STATEMENT_ID	VARCHAR2(30)	Value of the optional <code>STATEMENT_ID</code> parameter specified in the <code>EXPLAIN PLAN</code> statement.
TIMESTAMP	DATE	Date and time when the <code>EXPLAIN PLAN</code> statement was issued.
REMARKS	VARCHAR2(80)	Any comment (of up to 80 bytes) you want to associate with each step of the explained plan. If you need to add or change a remark on any row of the <code>PLAN_TABLE</code> , then use the <code>UPDATE</code> statement to modify the rows of the <code>PLAN_TABLE</code> .
OPERATION	VARCHAR2(30)	Name of the internal operation performed in this step. In the first row generated for a statement, the column contains one of the following values: DELETE STATEMENT INSERT STATEMENT SELECT STATEMENT UPDATE STATEMENT See Table 9-4 for more information on values for this column.
OPTIONS	VARCHAR2(225)	A variation on the operation described in the <code>OPERATION</code> column. See Table 9-4 for more information on values for this column.
OBJECT_NODE	VARCHAR2(128)	Name of the database link used to reference the object (a table name or view name). For local queries using parallel execution, this column describes the order in which output from operations is consumed.
OBJECT_OWNER	VARCHAR2(30)	Name of the user who owns the schema containing the table or index.
OBJECT_NAME	VARCHAR2(30)	Name of the table or index.
OBJECT_INSTANCE	NUMERIC	Number corresponding to the ordinal position of the object as it appears in the original statement. The numbering proceeds from left to right, outer to inner with respect to the original statement text. View expansion results in unpredictable numbers.
OBJECT_TYPE	VARCHAR2(30)	Modifier that provides descriptive information about the object; for example, <code>NON-UNIQUE</code> for indexes.

Table 9–1 (Cont.) PLAN_TABLE Columns

Column	Type	Description
OPTIMIZER	VARCHAR2 (255)	Current mode of the optimizer.
SEARCH_COLUMNS	NUMERIC	Not currently used.
ID	NUMERIC	A number assigned to each step in the execution plan.
PARENT_ID	NUMERIC	The ID of the next execution step that operates on the output of the ID step.
POSITION	NUMERIC	For the first row of output, this indicates the optimizer's estimated cost of executing the statement. For the other rows, it indicates the position relative to the other children of the same parent.
COST	NUMERIC	Cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null. Cost is not determined for table access operations. The value of this column does not have any particular unit of measurement; it is merely a weighted value used to compare costs of execution plans. The value of this column is a function of the CPU_COST and IO_COST columns.
CARDINALITY	NUMERIC	Estimate by the cost-based approach of the number of rows accessed by the operation.
BYTES	NUMERIC	Estimate by the cost-based approach of the number of bytes accessed by the operation.
OTHER_TAG	VARCHAR2 (255)	Contents of the OTHER column. See Table 9–2 for more information on the possible values for this column.
PARTITION_START	VARCHAR2 (255)	<p>Start partition of a range of accessed partitions. It can take one of the following values:</p> <p><i>n</i> indicates that the start partition has been identified by the SQL compiler, and its partition number is given by <i>n</i>.</p> <p>KEY indicates that the start partition will be identified at run time from partitioning key values.</p> <p>ROW REMOVE_LOCATION indicates that the start partition (same as the stop partition) will be computed at run time from the location of each record being retrieved. The record location is obtained by a user or from a global index.</p> <p>INVALID indicates that the range of accessed partitions is empty.</p>

Table 9–1 (Cont.) PLAN_TABLE Columns

Column	Type	Description
PARTITION_STOP	VARCHAR2 (255)	<p>Stop partition of a range of accessed partitions. It can take one of the following values:</p> <p><i>n</i> indicates that the stop partition has been identified by the SQL compiler, and its partition number is given by <i>n</i>.</p> <p>KEY indicates that the stop partition will be identified at run time from partitioning key values.</p> <p>ROW REMOVE_LOCATION indicates that the stop partition (same as the start partition) will be computed at run time from the location of each record being retrieved. The record location is obtained by a user or from a global index.</p> <p>INVALID indicates that the range of accessed partitions is empty.</p>
PARTITION_ID	NUMERIC	Step that has computed the pair of values of the PARTITION_START and PARTITION_STOP columns.
OTHER	LONG	Other information that is specific to the execution step that a user might find useful.
DISTRIBUTION	VARCHAR2 (30)	<p>Method used to distribute rows from producer query servers to consumer query servers.</p> <p>See Table 9–3 for more information on the possible values for this column. For more information about consumer and producer query servers, see <i>Oracle9i Data Warehousing Guide</i>.</p>
CPU_COST	NUMERIC	CPU cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null. The value of this column is proportional to the number of machine cycles required for the operation.
IO_COST	NUMERIC	I/O cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null. The value of this column is proportional to the number of data blocks read by the operation.
TEMP_SPACE	NUMERIC	Temporary space, in bytes, used by the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, or for operations that don't use any temporary space, this column is null.

[Table 9–2](#) describes the values that can appear in the OTHER_TAG column.

Table 9–2 Values of OTHER_TAG Column of the PLAN_TABLE

OTHER_TAG Text (examples)	Meaning	Interpretation
blank		Serial execution.
SERIAL_FROM_REMOTE (S -> R)	Serial from remote	Serial execution at a remote site.
SERIAL_TO_PARALLEL (S -> P)	Serial to parallel	Serial execution; output of step is partitioned or broadcast to parallel execution servers.
PARALLEL_TO_PARALLEL (P -> P)	Parallel to parallel	Parallel execution; output of step is repartitioned to second set of parallel execution servers.
PARALLEL_TO_SERIAL (P -> S)	Parallel to serial	Parallel execution; output of step is returned to serial "query coordinator" process.
PARALLEL_COMBINED_WITH_PARENT (PWP)	Parallel combined with parent	Parallel execution; output of step goes to next step in same parallel process. No interprocess communication to parent.
PARALLEL_COMBINED_WITH_CHILD (PWC)	Parallel combined with child	Parallel execution; input of step comes from prior step in same parallel process. No interprocess communication from child.

Table 9–3 describes the values that can appear in the DISTRIBUTION column:

Table 9–3 Values of DISTRIBUTION Column of the PLAN_TABLE

DISTRIBUTION Text	Interpretation
PARTITION (ROWID)	Maps rows to query servers based on the partitioning of a table or index using the rowid of the row to UPDATE/DELETE.
PARTITION (KEY)	Maps rows to query servers based on the partitioning of a table or index using a set of columns. Used for partial partition-wise join, PARALLEL INSERT, CREATE TABLE AS SELECT of a partitioned table, and CREATE PARTITIONED GLOBAL INDEX.
HASH	Maps rows to query servers using a hash function on the join key. Used for PARALLEL JOIN or PARALLEL GROUP BY.
RANGE	Maps rows to query servers using ranges of the sort key. Used when the statement contains an ORDER BY clause.
ROUND-ROBIN	Randomly maps rows to query servers.

Table 9–3 (Cont.) Values of DISTRIBUTION Column of the PLAN_TABLE

DISTRIBUTION Text	Interpretation
BROADCAST	Broadcasts the rows of the entire table to each query server. Used for a parallel join when one table is very small compared to the other.
QC (ORDER)	The query coordinator consumes the input in order, from the first to the last query server. Used when the statement contains an ORDER BY clause.
QC (RANDOM)	The query coordinator consumes the input randomly. Used when the statement does not have an ORDER BY clause.

Table 9–4 lists each combination of OPERATION and OPTION produced by the EXPLAIN PLAN statement and its meaning within an execution plan.

Table 9–4 OPERATION and OPTION Values Produced by EXPLAIN PLAN

Operation	Option	Description
AND-EQUAL	.	Operation accepting multiple sets of rowids, returning the intersection of the sets, eliminating duplicates. Used for the single-column indexes access path.
BITMAP	CONVERSION	TO ROWIDS converts bitmap representations to actual rowids that can be used to access the table. FROM ROWIDS converts the rowids to a bitmap representation. COUNT returns the number of rowids if the actual values are not needed.
BITMAP	INDEX	SINGLE VALUE looks up the bitmap for a single key value in the index. RANGE SCAN retrieves bitmaps for a key value range. FULL SCAN performs a full scan of a bitmap index if there is no start or stop key.
BITMAP	MERGE	Merges several bitmaps resulting from a range scan into one bitmap.
BITMAP	MINUS	Subtracts bits of one bitmap from another. Row source is used for negated predicates. Can be used only if there are nonnegated predicates yielding a bitmap from which the subtraction can take place. An example appears in " Viewing Bitmap Indexes with EXPLAIN PLAN " on page 9-11.
BITMAP	OR	Computes the bitwise OR of two bitmaps.
BITMAP	AND	Computes the bitwise AND of two bitmaps.

Table 9–4 (Cont.) OPERATION and OPTION Values Produced by EXPLAIN PLAN

Operation	Option	Description
BITMAP	KEY ITERATION	Takes each row from a table row source and finds the corresponding bitmap from a bitmap index. This set of bitmaps are then merged into one bitmap in a following BITMAP MERGE operation.
CONNECT BY	.	Retrieves rows in hierarchical order for a query containing a CONNECT BY clause.
CONCATENATION	.	Operation accepting multiple sets of rows returning the union-all of the sets.
COUNT	.	Operation counting the number of rows selected from a table.
	STOPKEY	Count operation where the number of rows returned is limited by the ROWNUM expression in the WHERE clause.
DOMAIN INDEX	.	Retrieval of one or more rowids from a domain index. The options column contain information supplied by a user-defined domain index cost function, if any.
FILTER	.	Operation accepting a set of rows, eliminates some of them, and returns the rest.
FIRST ROW	.	Retrieval of only the first row selected by a query.
FOR UPDATE	.	Operation retrieving and locking the rows selected by a query containing a FOR UPDATE clause.
HASH JOIN (These are join operations.)	.	Operation joining two sets of rows and returning the result. This join method is useful for joining large data sets of data (DSS, Batch). The join condition is an efficient way of accessing the second table. CBO uses the smaller of the two tables/data sources to build a hash table on the join key in memory. Then it scans the larger table, probing the hash table to find the joined rows.
HASH JOIN	ANTI	Hash anti-join.
HASH JOIN	SEMI	Hash semi-join.
INDEX (These are access methods.)	UNIQUE SCAN	Retrieval of a single rowid from an index.
INDEX	RANGE SCAN	Retrieval of one or more rowids from an index. Indexed values are scanned in ascending order.
INDEX	RANGE SCAN DESCENDING	Retrieval of one or more rowids from an index. Indexed values are scanned in descending order.

Table 9–4 (Cont.) OPERATION and OPTION Values Produced by EXPLAIN PLAN

Operation	Option	Description
INDEX	FULL SCAN	Retrieval of all rowids from an index when there is no start or stop key. Indexed values are scanned in ascending order.
INDEX	FULL SCAN DESCENDING	Retrieval of all rowids from an index when there is no start or stop key. Indexed values are scanned in descending order.
INDEX	FAST FULL SCAN	Retrieval of all rowids (and column values) using multiblock reads. No sorting order can be defined. Compares to a full table scan on only the indexed columns. Only available with the cost based optimizer.
INDEX	SKIP SCAN	Retrieval of rowids from a concatenated index without using the leading column(s) in the index. Introduced in Oracle9i. Only available with the cost based optimizer.
INLIST ITERATOR	.	Iterates over the next operation in the plan for each value in the IN-list predicate.
INTERSECTION	.	Operation accepting two sets of rows and returning the intersection of the sets, eliminating duplicates.
MERGE JOIN (These are join operations.)	.	Operation accepting two sets of rows, each sorted by a specific value, combining each row from one set with the matching rows from the other, and returning the result.
MERGE JOIN	OUTER	Merge join operation to perform an outer join statement.
MERGE JOIN	ANTI	Merge anti-join.
MERGE JOIN	SEMI	Merge semi-join.
MERGE JOIN	CARTESIAN	Can result from 1 or more of the tables not having any join conditions to any other tables in the statement. Can occur even with a join and it may not be flagged as CARTESIAN in the plan.
CONNECT BY	.	Retrieval of rows in hierarchical order for a query containing a CONNECT BY clause.
MINUS	.	Operation accepting two sets of rows and returning rows appearing in the first set but not in the second, eliminating duplicates.
NESTED LOOPS (These are join operations.)	.	Operation accepting two sets of rows, an outer set and an inner set. Oracle compares each row of the outer set with each row of the inner set, returning rows that satisfy a condition. This join method is useful for joining small subsets of data (OLTP). The join condition is an efficient way of accessing the second table.
NESTED LOOPS	OUTER	Nested loops operation to perform an outer join statement.

Table 9–4 (Cont.) OPERATION and OPTION Values Produced by EXPLAIN PLAN

Operation	Option	Description
PARTITION	SINGLE	Access one partition.
PARTITION	ITERATOR	Access many partitions (a subset).
PARTITION	ALL	Access all partitions.
PARTITION	INLIST	Similar to iterator, but based on an IN-list predicate.
PARTITION	INVALID	Indicates that the partition set to be accessed is empty. Iterates over the next operation in the plan for each partition in the range given by the PARTITION_START and PARTITION_STOP columns. PARTITION describes partition boundaries applicable to a single partitioned object (table or index) or to a set of equi-partitioned objects (a partitioned table and its local indexes). The partition boundaries are provided by the values of PARTITION_START and PARTITION_STOP of the PARTITION. Refer to Table 9–1 for valid values of partition start/stop.
REMOTE	.	Retrieval of data from a remote database.
SEQUENCE	.	Operation involving accessing values of a sequence.
SORT	AGGREGATE	Retrieval of a single row that is the result of applying a group function to a group of selected rows.
SORT	UNIQUE	Operation sorting a set of rows to eliminate duplicates.
SORT	GROUP BY	Operation sorting a set of rows into groups for a query with a GROUP BY clause.
SORT	JOIN	Operation sorting a set of rows before a merge-join.
SORT	ORDER BY	Operation sorting a set of rows for a query with an ORDER BY clause.
TABLE ACCESS (These are access methods.)	FULL	Retrieval of all rows from a table.
TABLE ACCESS	SAMPLE	Retrieval of sampled rows from a table.
TABLE ACCESS	CLUSTER	Retrieval of rows from a table based on a value of an indexed cluster key.
TABLE ACCESS	HASH	Retrieval of rows from table based on hash cluster key value.
TABLE ACCESS	BY ROWID RANGE	Retrieval of rows from a table based on a rowid range.
TABLE ACCESS	SAMPLE BY ROWID RANGE	Retrieval of sampled rows from a table based on a rowid range.

Table 9–4 (Cont.) OPERATION and OPTION Values Produced by EXPLAIN PLAN

Operation	Option	Description
TABLE ACCESS	BY USER ROWID	If the table rows are located using user-supplied rowids.
TABLE ACCESS	BY INDEX ROWID	If the table is nonpartitioned and rows are located using index(es).
TABLE ACCESS	BY GLOBAL INDEX ROWID	If the table is partitioned and rows are located using only global indexes.
TABLE ACCESS	BY LOCAL INDEX ROWID	If the table is partitioned and rows are located using one or more local indexes and possibly some global indexes. Partition Boundaries: The partition boundaries might have been computed by: A previous PARTITION step, in which case the PARTITION_START and PARTITION_STOP column values replicate the values present in the PARTITION step, and the PARTITION_ID contains the ID of the PARTITION step. Possible values for PARTITION_START and PARTITION_STOP are NUMBER(n), KEY, INVALID. The TABLE ACCESS or INDEX step itself, in which case the PARTITION_ID contains the ID of the step. Possible values for PARTITION_START and PARTITION_STOP are NUMBER(n), KEY, ROW REMOVE_LOCATION (TABLE ACCESS only), and INVALID.
UNION	.	Operation accepting two sets of rows and returns the union of the sets, eliminating duplicates.
VIEW	.	Operation performing a view's query and then returning the resulting rows to another operation.

See Also: [Chapter 1, "Introduction to the Optimizer"](#)

10

Using SQL Trace and TKPROF

The SQL Trace facility and `TKPROF` are two basic performance diagnostic tools that can help you monitor and tune applications running against the Oracle Server.

This chapter contains the following sections:

- [Understanding SQL Trace and TKPROF](#)
- [Using the SQL Trace Facility and TKPROF](#)
- [Avoiding Pitfalls in TKPROF Interpretation](#)
- [Sample TKPROF Output](#)

Understanding SQL Trace and TKPROF

The SQL Trace facility and TKPROF let you accurately assess the efficiency of the SQL statements an application runs. For best results, use these tools with EXPLAIN PLAN rather than using EXPLAIN PLAN alone.

Understanding the SQL Trace Facility

The SQL Trace facility provides performance information on individual SQL statements. It generates the following statistics for each statement:

- Parse, execute, and fetch counts
- CPU and elapsed times
- Physical reads and logical reads
- Number of rows processed
- Misses on the library cache
- Username under which each parse occurred
- Each commit and rollback

You can enable the SQL Trace facility for a session or for an instance. When the SQL Trace facility is enabled, performance statistics for all SQL statements executed in a user session or in the instance are placed into trace files.

The additional overhead of running the SQL Trace facility against an application with performance problems is normally insignificant compared with the inherent overhead caused by the application's inefficiency.

Note: Try to enable SQL Trace only for statistics collection and on specific sessions. If you must enable the facility on an entire production environment, then you can minimize performance impact with the following:

- Maintain at least 25% idle CPU capacity.
 - Maintain adequate disk space for the USER_DUMP_DEST location.
 - Stripe disk space over sufficient disks.
-
-

Understanding TKPROF

You can run the `TKPROF` program to format the contents of the trace file and place the output into a readable output file. Optionally, `TKPROF` can also:

- Determine the execution plans of SQL statements
- Create a SQL script that stores the statistics in the database

`TKPROF` reports each statement executed with the resources it has consumed, the number of times it was called, and the number of rows which it processed. This information lets you easily locate those statements that are using the greatest resource. With experience or with baselines available, you can assess whether the resources used are reasonable given the work done.

Using the SQL Trace Facility and TKPROF

Follow these steps to use the SQL Trace facility and `TKPROF`:

1. Set initialization parameters for trace file management.
See "[Step 1: Setting Initialization Parameters for Trace File Management](#)" on page 10-4.
2. Enable the SQL Trace facility for the desired session, and run the application. This step produces a trace file containing statistics for the SQL statements issued by the application.
See "[Step 2: Enabling the SQL Trace Facility](#)" on page 10-5.
3. Run `TKPROF` to translate the trace file created in Step 2 into a readable output file. This step can optionally create a SQL script that can be used to store the statistics in a database.
See "[Step 3: Formatting Trace Files with TKPROF](#)" on page 10-6.
4. Interpret the output file created in Step 3.
See "[Step 4: Interpreting TKPROF Output](#)" on page 10-12.
5. Optionally, run the SQL script produced in Step 3 to store the statistics in the database.
See "[Step 5: Storing SQL Trace Facility Statistics](#)" on page 10-17.

In the following sections, each of these steps is discussed in depth.

Step 1: Setting Initialization Parameters for Trace File Management

When the SQL Trace facility is enabled for a session, Oracle generates a trace file containing statistics for traced SQL statements for that session. When the SQL Trace facility is enabled for an instance, Oracle creates a separate trace file for each process. Before enabling the SQL Trace facility:

1. Check the settings of the `TIMED_STATISTICS`, `MAX_DUMP_FILE_SIZE`, and `USER_DUMP_DEST` initialization parameters. See [Table 10-1](#).

Table 10-1 Initialization Parameters to Check Before Enabling SQL Trace

Parameter	Description
<code>TIMED_STATISTICS</code>	This enables and disables the collection of timed statistics, such as CPU and elapsed times, by the SQL Trace facility, as well as the collection of various statistics in the dynamic performance tables. The default value of false disables timing. A value of true enables timing. Enabling timing causes extra timing calls for low-level operations. This is a dynamic parameter. It is also a session parameter.
<code>MAX_DUMP_FILE_SIZE</code>	When the SQL Trace facility is enabled at the instance level, every call to the server produces a text line in a file in the operating system's file format. The maximum size of these files (in operating system blocks) is limited by this initialization parameter. The default is 500. If you find that the trace output is truncated, then increase the value of this parameter before generating another trace file. This is a dynamic parameter. It is also a session parameter.
<code>USER_DUMP_DEST</code>	This must fully specify the destination for the trace file according to the conventions of the operating system. The default value is the default destination for system dumps on the operating system. This value can be modified with <code>ALTER SYSTEM SET USER_DUMP_DEST= newdir</code> . This is a dynamic parameter. It is also a session parameter.

Note: Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

2. Devise a way of recognizing the resulting trace file.

Be sure you know how to distinguish the trace files by name. Oracle writes them to the user dump destination specified by `USER_DUMP_DEST`. However, this directory can soon contain many hundreds of files, usually with generated names. It might be difficult to match trace files back to the session or process that created them. You can tag trace files by including in your programs a statement like `SELECT program_name FROM DUAL`. You can then trace each file back to the process that created it.

3. If the operating system retains multiple versions of files, then be sure that the version limit is high enough to accommodate the number of trace files you expect the SQL Trace facility to generate.
4. The generated trace files can be owned by an operating system user other than yourself. This user must make the trace files available to you before you can use `TKPROF` to format them.

See Also: ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Step 2: Enabling the SQL Trace Facility

Enable the SQL Trace facility for the session by using one of the following:

- `DBMS_SESSION.SET_SQL_TRACE` procedure
- `ALTER SESSION SET SQL_TRACE = TRUE;`

You can enable SQL Trace in another session by using the `DBMS_SYSTEM.SET_SQL_TRACE_IN_SESSION` procedure.

Caution: Because running the SQL Trace facility increases system overhead, enable it only when tuning SQL statements, and disable it when you are finished.

You might need to modify an application to contain the `ALTER SESSION` statement. For example, to issue the `ALTER SESSION` statement in Oracle Forms, invoke Oracle Forms using the `-s` option, or invoke Oracle Forms (Design) using the `statistics` option. For more information on Oracle Forms, see the *Oracle Forms Reference*.

To disable the SQL Trace facility for the session, enter:

```
ALTER SESSION SET SQL_TRACE = FALSE;
```

The SQL Trace facility is automatically disabled for the session when the application disconnects from Oracle.

You can enable the SQL Trace facility for an instance by setting the value of the `SQL_TRACE` initialization parameter to `TRUE` in the initialization file.

```
SQL_TRACE = TRUE
```

After the instance has been restarted with the updated initialization parameter file, SQL Trace is enabled for the instance and statistics are collected for all sessions. If the SQL Trace facility has been enabled for the instance, you can disable it for the instance by setting the value of the `SQL_TRACE` parameter to `FALSE`.

Note: Setting `SQL_TRACE` to `TRUE` can have a severe performance impact. For more information, see *Oracle9i Database Reference*.

Step 3: Formatting Trace Files with TKPROF

TKPROF accepts as input a trace file produced by the SQL Trace facility, and it produces a formatted output file. TKPROF can also be used to generate execution plans.

After the SQL Trace facility has generated a number of trace files, you can:

- Run TKPROF on each individual trace file, producing a number of formatted output files, one for each session.

- Concatenate the trace files, and then run TKPROF on the result to produce a formatted output file for the entire instance.

TKPROF does not report COMMITs and ROLLBACKs that are recorded in the trace file.

Sample TKPROF Output

Sample output from TKPROF is as follows:

```
SELECT * FROM emp, dept
WHERE emp.deptno = dept.deptno;
```

call	count	cpu	elapsed	disk	query current	rows
Parse	1	0.16	0.29	3	13	0
Execute	1	0.00	0.00	0	0	0
Fetch	1	0.03	0.26	2	2	4

Misses in library cache during parse: 1
Parsing user id: (8) SCOTT

Rows	Execution Plan
14	MERGE JOIN
4	SORT JOIN
4	TABLE ACCESS (FULL) OF 'DEPT'
14	SORT JOIN
14	TABLE ACCESS (FULL) OF 'EMP'

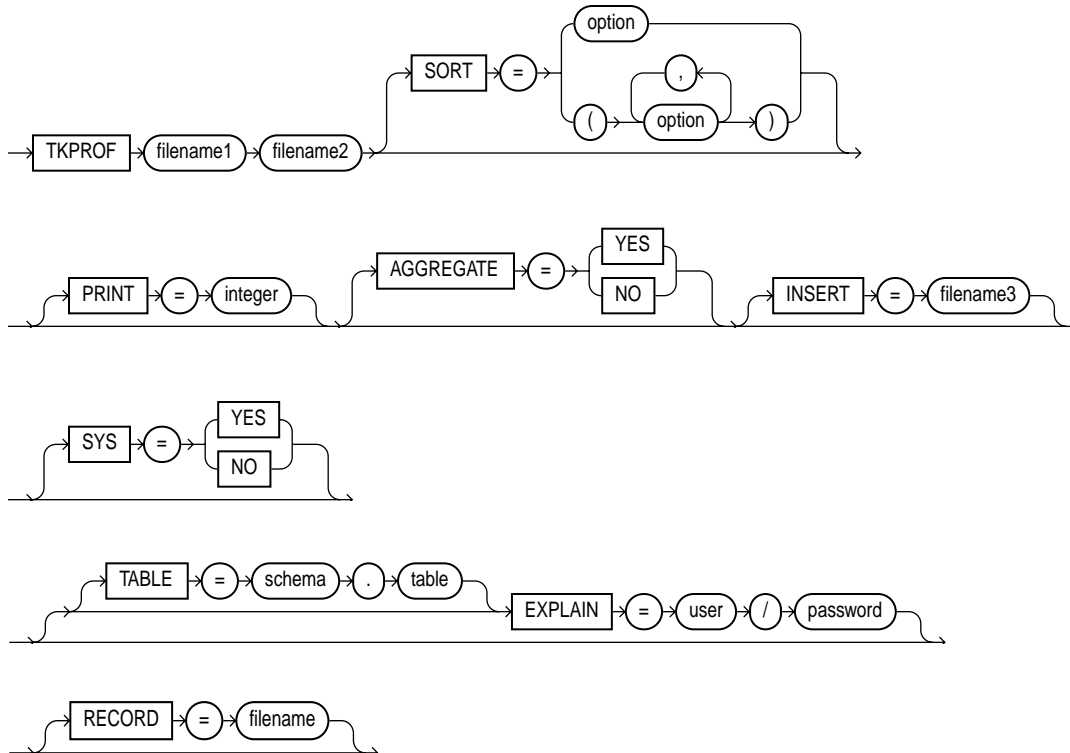
For this statement, TKPROF output includes the following information:

- The text of the SQL statement
- The SQL Trace statistics in tabular form
- The number of library cache misses for the parsing and execution of the statement.
- The user initially parsing the statement.
- The execution plan generated by EXPLAIN PLAN.

TKPROF also provides a summary of user level statements and recursive SQL calls for the trace file.

Syntax of TKPROF

TKPROF::=



If you invoke `TKPROF` without arguments, then online help is displayed. Use the arguments in [Table 10-2](#) with `TKPROF`.

Table 10-2 *TKPROF Arguments*

Argument	Description
<i>filename1</i>	Specifies the input file, a trace file containing statistics produced by the SQL Trace facility. This file can be either a trace file produced for a single session, or a file produced by concatenating individual trace files from multiple sessions.
<i>filename2</i>	Specifies the file to which <code>TKPROF</code> writes its formatted output.
WAITS	Specifies whether to record summary for any wait events found in the trace file. Values are <code>YES</code> or <code>NO</code> .

Table 10–2 (Cont.) TKPROF Arguments

Argument	Description
SORTS	Sorts traced SQL statements in descending order of specified sort option before listing them into the output file. If more than one option is specified, then the output is sorted in descending order by the sum of the values specified in the sort options. If you omit this parameter, then TKPROF lists statements into the output file in order of first use. Sort options are listed as follows:
PRSCNT	Number of times parsed.
PRSCPU	CPU time spent parsing.
PRSELA	Elapsed time spent parsing.
PRSDSK	Number of physical reads from disk during parse.
PRSQRY	Number of consistent mode block reads during parse.
PRSCU	Number of current mode block reads during parse.
PRSMIS	Number of library cache misses during parse.
EXECNT	Number of executes.
EXECPU	CPU time spent executing.
EXEELA	Elapsed time spent executing.
EXEDSK	Number of physical reads from disk during execute.
EXEDSK	Number of physical reads from disk during execute.
EXEQRY	Number of consistent mode block reads during execute.
EXECU	Number of current mode block reads during execute.
EXEROW	Number of rows processed during execute.
EXEMIS	Number of library cache misses during execute.
FCHCNT	Number of fetches.
FCHCPU	CPU time spent fetching.
FCHELA	Elapsed time spent fetching.
FCHDSK	Number of physical reads from disk during fetch.
FCHQRY	Number of consistent mode block reads during fetch.
FCHCU	Number of current mode block reads during fetch.
FCHROW	Number of rows fetched.

Table 10–2 (Cont.) TKPROF Arguments

Argument	Description
PRINT	Lists only the first integer sorted SQL statements from the output file. If you omit this parameter, then TKPROF lists all traced SQL statements. This parameter does not affect the optional SQL script. The SQL script always generates insert data for all traced SQL statements.
AGGREGATE	If you specify AGGREGATE = NO, then TKPROF does not aggregate multiple users of the same SQL text.
INSERT	Creates a SQL script that stores the trace file statistics in the database. TKPROF creates this script with the name <code>filename3</code> . This script creates a table and inserts a row of statistics for each traced SQL statement into the table.
SYS	Enables and disables the listing of SQL statements issued by the user SYS, or recursive SQL statements, into the output file. The default value of YES causes TKPROF to list these statements. The value of NO causes TKPROF to omit them. This parameter does not affect the optional SQL script. The SQL script always inserts statistics for all traced SQL statements, including recursive SQL statements.
TABLE	<p>Specifies the schema and name of the table into which TKPROF temporarily places execution plans before writing them to the output file. If the specified table already exists, then TKPROF deletes all rows in the table, uses it for the EXPLAIN PLAN statement (which writes more rows into the table), and then deletes those rows. If this table does not exist, then TKPROF creates it, uses it, and then drops it.</p> <p>The specified <i>user</i> must be able to issue INSERT, SELECT, and DELETE statements against the table. If the table does not already exist, then the user must also be able to issue CREATE TABLE and DROP TABLE statements. For the privileges to issue these statements, see the <i>Oracle9i SQL Reference</i>.</p> <p>This option allows multiple individuals to run TKPROF concurrently with the same user in the EXPLAIN value. These individuals can specify different TABLE values and avoid destructively interfering with each other's processing on the temporary plan table.</p> <p>If you use the EXPLAIN parameter without the TABLE parameter, then TKPROF uses the table PROF\$PLAN_TABLE in the schema of the user specified by the EXPLAIN parameter. If you use the TABLE parameter without the EXPLAIN parameter, then TKPROF ignores the TABLE parameter.</p> <p>If no plan table exists, TKPROF creates the table PROF\$PLAN_TABLE and then drops it at the end.</p>

Table 10–2 (Cont.) TKPROF Arguments

Argument	Description
EXPLAIN	Determines the execution plan for each SQL statement in the trace file and writes these execution plans to the output file. TKPROF determines execution plans by issuing the EXPLAIN PLAN statement after connecting to Oracle with the <i>user</i> and <i>password</i> specified in this parameter. The specified <i>user</i> must have CREATE SESSION system privileges. TKPROF takes longer to process a large trace file if the EXPLAIN option is used.
RECORD	Creates a SQL script with the specified filename with all of the nonrecursive SQL in the trace file. This can be used to replay the user events from the trace file.
WIDTH	An integer that controls the output line width of some TKPROF output, such as the explain plan. This parameter is useful for post-processing of TKPROF output.

Examples of TKPROF Statement

This section provides two brief examples of TKPROF usage. For an complete example of TKPROF output, see ["Sample TKPROF Output"](#) on page 10-23.

TKPROF Example 1 If you are processing a large trace file using a combination of SORT parameters and the PRINT parameter, then you can produce a TKPROF output file containing only the highest resource-intensive statements. For example, the following statement prints the 10 statements in the trace file that have generated the most physical I/O:

```
TKPROF ora53269.trc ora53269.prf SORT = (PRSDSK, EXEDSK, FCHDSK) PRINT = 10
```

TKPROF Example 2 This example runs TKPROF, accepts a trace file named dlsun12_jane_fg_sqlplus_007.trc, and writes a formatted output file named outputa.prf:

```
TKPROF dlsun12_jane_fg_sqlplus_007.trc OUTPUTA.PRF
EXPLAIN=scott/tiger TABLE=scott.temp_plan_table_a INSERT=STOREA.SQL SYS=NO
SORT=(EXECPU,FCHCPU)
```

This example is likely to be longer than a single line on the screen, and you might need to use continuation characters, depending on the operating system.

Note the other parameters in this example:

- The EXPLAIN value causes TKPROF to connect as the user scott and use the EXPLAIN PLAN statement to generate the execution plan for each traced SQL statement. You can use this to get access paths and row source counts.

- The `TABLE` value causes TKPROF to use the table `temp_plan_table_a` in the schema `scott` as a temporary plan table.
- The `INSERT` value causes TKPROF to generate a SQL script named `STOREA.SQL` that stores statistics for all traced SQL statements in the database.
- The `SYS` parameter with the value of `NO` causes TKPROF to omit recursive SQL statements from the output file. In this way, you can ignore internal Oracle statements such as temporary table operations.
- The `SORT` value causes TKPROF to sort the SQL statements in order of the sum of the CPU time spent executing and the CPU time spent fetching rows before writing them to the output file. For greatest efficiency, always use `SORT` parameters.

Step 4: Interpreting TKPROF Output

This section provides pointers for interpreting TKPROF output.

- [Tabular Statistics in TKPROF](#)
- [Library Cache Misses in TKPROF](#)
- [Statement Truncation in SQL Trace](#)
- [Identification of User Issuing the SQL Statement in TKPROF](#)
- [Execution Plan in TKPROF](#)
- [Deciding Which Statements to Tune](#)

While TKPROF provides a very useful analysis, the most accurate measure of efficiency is the actual performance of the application in question. At the end of the TKPROF output is a summary of the work done in the database engine by the process during the period that the trace was running.

Tabular Statistics in TKPROF

TKPROF lists the statistics for a SQL statement returned by the SQL Trace facility in rows and columns. Each row corresponds to one of three steps of SQL statement processing. Statistics are identified by the value of the `CALL` column. See [Table 10-3](#).

Table 10-3 CALL Column Values

CALL Value	Meaning
PARSE	Translates the SQL statement into an execution plan, including checks for proper security authorization and checks for the existence of tables, columns, and other referenced objects.
EXECUTE	Actual execution of the statement by Oracle. For INSERT, UPDATE, and DELETE statements, this modifies the data. For SELECT statements, this identifies the selected rows.
FETCH	Retrieves rows returned by a query. Fetches are only performed for SELECT statements.

The other columns of the SQL Trace facility output are combined statistics for all parses, all executes, and all fetches of a statement. The sum of `query` and `current` is the total number of buffers accessed, also called Logical I/Os (LIOs). See [Table 10-4](#).

Table 10-4 SQL Trace Statistics for Parses, Executes, and Fetches.

SQL Trace Statistic	Meaning
COUNT	Number of times a statement was parsed, executed, or fetched.
CPU	Total CPU time in seconds for all parse, execute, or fetch calls for the statement. This value is zero (0) if <code>TIMED_STATISTICS</code> is not turned on.
ELAPSED	Total elapsed time in seconds for all parse, execute, or fetch calls for the statement. This value is zero (0) if <code>TIMED_STATISTICS</code> is not turned on.
DISK	Total number of data blocks physically read from the datafiles on disk for all parse, execute, or fetch calls.
QUERY	Total number of buffers retrieved in consistent mode for all parse, execute, or fetch calls. Usually, buffers are retrieved in consistent mode for queries.
CURRENT	Total number of buffers retrieved in current mode. Buffers are retrieved in current mode for statements such as INSERT, UPDATE, and DELETE.

Statistics about the processed rows appear in the `ROWS` column. See [Table 10-5](#).

Table 10–5 SQL Trace Statistics for the ROWS Column

SQL Trace Statistic	Meaning
ROWS	Total number of rows processed by the SQL statement. This total does not include rows processed by subqueries of the SQL statement.

For `SELECT` statements, the number of rows returned appears for the fetch step. For `UPDATE`, `DELETE`, and `INSERT` statements, the number of rows processed appears for the execute step.

Note: The row source counts are displayed when a cursor is closed. In SQL*Plus, there is only one user cursor, so each statement executed causes the previous cursor to be closed; therefore, the row source counts are displayed. PL/SQL has its own cursor handling and does not close child cursors when the parent cursor is closed. Exiting (or reconnecting) causes the counts to be displayed.

Interpreting the Resolution of Statistics

Timing statistics have a resolution of one hundredth of a second; therefore, any operation on a cursor that takes a hundredth of a second or less might not be timed accurately. Keep this in mind when interpreting statistics. In particular, be careful when interpreting the results from simple queries that execute very quickly.

Understanding Recursive Calls

Sometimes, in order to execute a SQL statement issued by a user, Oracle must issue additional statements. Such statements are called recursive calls or recursive SQL statements. For example, if you insert a row into a table that does not have enough space to hold that row, then Oracle makes recursive calls to allocate the space dynamically. Recursive calls are also generated when data dictionary information is not available in the data dictionary cache and must be retrieved from disk.

If recursive calls occur while the SQL Trace facility is enabled, then `TKPROF` produces statistics for the recursive SQL statements and marks them clearly as recursive SQL statements in the output file. You can suppress the listing of Oracle internal recursive calls (for example, space management) in the output file by setting the `SYS` command-line parameter to `NO`. The statistics for a recursive SQL statement are included in the listing for that statement, not in the listing for the SQL statement that caused the recursive call. So, when you are calculating the total

resources required to process a SQL statement, consider the statistics for that statement as well as those for recursive calls caused by that statement.

Note: Recursive SQL statistics are not included for SQL-level operations. However, recursive SQL statistics *are* included for operations done under the SQL level, such as triggers. For more information, see "[Avoiding the Trigger Trap](#)" on page 10-23.

Library Cache Misses in TKPROF

TKPROF also lists the number of library cache misses resulting from parse and execute steps for each SQL statement. These statistics appear on separate lines following the tabular statistics. If the statement resulted in no library cache misses, then TKPROF does not list the statistic. In "[Sample TKPROF Output](#)" on page 10-7, the statement resulted in one library cache miss for the parse step and no misses for the execute step.

Statement Truncation in SQL Trace

The following SQL statements are truncated to 25 characters in the SQL Trace file:

```
SET ROLE
GRANT
ALTER USER
ALTER ROLE
CREATE USER
CREATE ROLE
```

Identification of User Issuing the SQL Statement in TKPROF

TKPROF also lists the user ID of the user issuing each SQL statement. If the SQL Trace input file contained statistics from multiple users and the statement was issued by more than one user, then TKPROF lists the ID of the last user to parse the statement. The user ID of all database users appears in the data dictionary in the column ALL_USERS.USER_ID.

Execution Plan in TKPROF

If you specify the EXPLAIN parameter on the TKPROF statement line, then TKPROF uses the EXPLAIN PLAN statement to generate the execution plan of each SQL statement traced. TKPROF also displays the number of rows processed by each step of the execution plan.

Note: Trace files generated immediately after instance startup contain data that reflects the activity of the startup process. In particular, they reflect a disproportionate amount of I/O activity as caches in the system global area (SGA) are filled. For the purposes of tuning, ignore such trace files.

See Also: [Chapter 9, "Using EXPLAIN PLAN"](#) for more information on interpreting execution plans

Deciding Which Statements to Tune

You need to find which SQL statements use the most CPU or disk resource. If the `TIMED_STATISTICS` parameter is on, then you can find high CPU activity in the `CPU` column. If `TIMED_STATISTICS` is not on, then check the `QUERY` and `CURRENT` columns.

See Also: ["Examples of TKPROF Statement"](#) on page 10-11 for examples of finding resource intensive statements

With the exception of locking problems and inefficient PL/SQL loops, neither the CPU time nor the elapsed time is necessary to find problem statements. The key is the number of block visits, both query (that is, subject to read consistency) and current (that is, *not* subject to read consistency). Segment headers and blocks that are going to be updated are acquired in current mode, but all query and subquery processing requests the data in query mode. These are precisely the same measures as the instance statistics `CONSISTENT GETS` and `DB BLOCK GETS`. You can find high disk activity in the `disk` column.

The following listing shows TKPROF output for one SQL statement as it appears in the output file:

```
SELECT *
FROM emp, dept
WHERE emp.deptno = dept.deptno;
```

call	count	cpu	elapsed	disk	query current	rows
Parse	11	0.08	0.18	0	0	0
Execute	11	0.23	0.66	0	3	6
Fetch	35	6.70	6.83	100	12326	2
total	57	7.01	7.67	100	12329	8

Misses in library cache during parse: 0

If it is acceptable to have 7.01 CPU seconds and to retrieve 824 rows, then you need not look any further at this trace output. In fact, a major use of TKPROF reports in a tuning exercise is to eliminate processes from the detailed tuning phase.

You can also see that 10 unnecessary parse call were made (because there were 11 parse calls for this one statement) and that array fetch operations were performed. You know this because more rows were fetched than there were fetches performed. A large gap between CPU and elapsed timings indicates Physical I/Os (PIOs).

Step 5: Storing SQL Trace Facility Statistics

You might want to keep a history of the statistics generated by the SQL Trace facility for an application, and compare them over time. TKPROF can generate a SQL script that creates a table and inserts rows of statistics into it. This script contains:

- A CREATE TABLE statement that creates an output table named TKPROF_TABLE.
- INSERT statements that add rows of statistics, one for each traced SQL statement, to the TKPROF_TABLE.

After running TKPROF, you can run this script to store the statistics in the database.

Generating the TKPROF Output SQL Script

When you run TKPROF, use the INSERT parameter to specify the name of the generated SQL script. If you omit this parameter, then TKPROF does not generate a script.

Editing the TKPROF Output SQL Script

After TKPROF has created the SQL script, you might want to edit the script before running it. If you have already created an output table for previously collected statistics and you want to add new statistics to this table, then remove the CREATE TABLE statement from the script. The script then inserts the new rows into the existing table.

If you have created multiple output tables, perhaps to store statistics from different databases in different tables, then edit the CREATE TABLE and INSERT statements to change the name of the output table.

Querying the Output Table

The following CREATE TABLE statement creates the TKPROF_TABLE:

```
CREATE TABLE TKPROF_TABLE (
  DATE_OF_INSERT    DATE,
  CURSOR_NUM        NUMBER,
  DEPTH             NUMBER,
  USER_ID           NUMBER,
  PARSE_CNT         NUMBER,
  PARSE_CPU         NUMBER,
  PARSE_ELAP        NUMBER,
  PARSE_DISK        NUMBER,
  PARSE_QUERY       NUMBER,
  PARSE_CURRENT     NUMBER,
  PARSE_MISS        NUMBER,
  EXE_COUNT         NUMBER,
  EXE_CPU           NUMBER,
  EXE_ELAP          NUMBER,
  EXE_DISK          NUMBER,
  EXE_QUERY         NUMBER,
  EXE_CURRENT       NUMBER,
  EXE_MISS          NUMBER,
  EXE_ROWS          NUMBER,
  FETCH_COUNT       NUMBER,
  FETCH_CPU         NUMBER,
  FETCH_ELAP        NUMBER,
  FETCH_DISK        NUMBER,
  FETCH_QUERY       NUMBER,
  FETCH_CURRENT     NUMBER,
  FETCH_ROWS        NUMBER,
  CLOCK_TICKS       NUMBER,
  SQL_STATEMENT     LONG);
```

Most output table columns correspond directly to the statistics that appear in the formatted output file. For example, the PARSE_CNT column value corresponds to the count statistic for the parse step in the output file.

The columns in [Table 10–6](#) help you identify a row of statistics.

Table 10–6 TKPROF_TABLE Columns for Identifying a Row of Statistics

Column	Description
SQL_STATEMENT	This is the SQL statement for which the SQL Trace facility collected the row of statistics. Because this column has datatype LONG, you cannot use it in expressions or WHERE clause conditions.

Table 10-6 (Cont.) TKPROF_TABLE Columns for Identifying a Row of Statistics

Column	Description
DATE_OF_INSERT	This is the date and time when the row was inserted into the table. This value is not exactly the same as the time the statistics were collected by the SQL Trace facility.
DEPTH	This indicates the level of recursion at which the SQL statement was issued. For example, a value of 0 indicates that a user issued the statement. A value of 1 indicates that Oracle generated the statement as a recursive call to process a statement with a value of 0 (a statement issued by a user). A value of n indicates that Oracle generated the statement as a recursive call to process a statement with a value of $n-1$.
USER_ID	This identifies the user issuing the statement. This value also appears in the formatted output file.
CURSOR_NUM	Oracle uses this column value to keep track of the cursor to which each SQL statement was assigned.

The output table does not store the statement's execution plan. The following query returns the statistics from the output table. These statistics correspond to the formatted output shown in the section "[Sample TKPROF Output](#)" on page 10-7.

```
SELECT * FROM TKPROF_TABLE;
```

Oracle responds with something similar to:

```
DATE_OF_INSERT CURSOR_NUM DEPTH USER_ID PARSE_CNT PARSE_CPU PARSE_ELAP
-----
21-DEC-1998          1      0      8          1          16          22

PARSE_DISK PARSE_QUERY PARSE_CURRENT PARSE_MISS EXE_COUNT EXE_CPU
-----
      3          11              0              1              1              0

EXE_ELAP EXE_DISK EXE_QUERY EXE_CURRENT EXE_MISS EXE_ROWS FETCH_COUNT
-----
      0          0          0              0              0              0              1

FETCH_CPU FETCH_ELAP FETCH_DISK FETCH_QUERY FETCH_CURRENT FETCH_ROWS
-----
      2          20           2           2              4              10

SQL_STATEMENT
-----
```

```
SELECT * FROM EMP, DEPT WHERE EMP.DEPTNO = DEPT.DEPTNO
```

Avoiding Pitfalls in TKPROF Interpretation

This section describes some fine points of TKPROF interpretation:

- [Avoiding the Argument Trap](#)
- [Avoiding the Read Consistency Trap](#)
- [Avoiding the Schema Trap](#)
- [Avoiding the Time Trap](#)
- [Avoiding the Trigger Trap](#)

Avoiding the Argument Trap

If you are not aware of the values being bound at run time, then it is possible to fall into the argument trap. `EXPLAIN PLAN` cannot determine the type of a bind variable from the text of SQL statements, and it always assumes that the type is `varchar`. If the bind variable is actually a number or a date, then TKPROF can cause implicit data conversions, which can cause inefficient plans to be executed. To avoid this, experiment with different data types in the query.

To avoid this problem, perform the conversion yourself.

See Also: ["EXPLAIN PLAN Restrictions"](#) on page 9-22 for information about TKPROF and bind variables

Avoiding the Read Consistency Trap

The next example illustrates the read consistency trap. Without knowing that an uncommitted transaction had made a series of updates to the `NAME` column, it is very difficult to see why so many block visits would be incurred.

Cases like this are not normally repeatable: if the process were run again, it is unlikely that another transaction would interact with it in the same way.

```
SELECT name_id
FROM cq_names
WHERE name = 'FLOOR';
```

call	count	cpu	elapsed	disk	query current	rows
----	-----	---	-----	----	-----	----
Parse	1	0.10	0.18	0	0 0	0

```
Execute      1    0.00      0.00      0      0      0      0
Fetch       1    0.11      0.21      2     101      0      1
```

```
Misses in library cache during parse: 1
Parsing user id: 01 (USER1)
```

```
Rows      Execution Plan
-----
0      SELECT STATEMENT
1      TABLE ACCESS (BY ROWID) OF 'CQ_NAMES'
2      INDEX (RANGE SCAN) OF 'CQ_NAMES_NAME' (NON_UNIQUE)
```

Avoiding the Schema Trap

This example shows an extreme (and thus easily detected) example of the schema trap. At first, it is difficult to see why such an apparently straightforward indexed query needs to look at so many database blocks, or why it should access any blocks at all in current mode.

```
SELECT name_id
FROM cq_names
WHERE name = 'FLOOR';
```

```
call      count      cpu      elapsed      disk  query current rows
-----
Parse      1      0.06      0.10      0      0      0      0
Execute    1      0.02      0.02      0      0      0      0
Fetch      1      0.23      0.30     31     31      3      1
```

```
Misses in library cache during parse: 0
Parsing user id: 02 (USER2)
```

```
Rows      Execution Plan
-----
0      SELECT STATEMENT
2340     TABLE ACCESS (BY ROWID) OF 'CQ_NAMES'
0      INDEX (RANGE SCAN) OF 'CQ_NAMES_NAME' (NON-UNIQUE)
```

Two statistics suggest that the query might have been executed with a full table scan. These statistics are the current mode block visits, plus the number of rows originating from the Table Access row source in the execution plan. The explanation is that the required index was built after the trace file had been produced, but before TKPROF had been run.

Generating a new trace file gives the following data:

```
SELECT name_id
FROM cq_names
WHERE name = 'FLOOR';
```

call	count	cpu	elapsed	disk	query current	rows
Parse	1	0.01	0.02	0	0	0
Execute	1	0.00	0.00	0	0	0
Fetch	1	0.00	0.00	0	2	1

```
Misses in library cache during parse: 0
Parsing user id: 02 (USER2)
```

Rows	Execution Plan
0	SELECT STATEMENT
1	TABLE ACCESS (BY ROWID) OF 'CQ_NAMES'
2	INDEX (RANGE SCAN) OF 'CQ_NAMES_NAME' (NON-UNIQUE)

One of the marked features of this correct version is that the parse call took 10 milliseconds of CPU time and 20 milliseconds of elapsed time, but the query apparently took no time at all to execute and perform the fetch. These anomalies arise because the clock tick of 10 milliseconds is too long relative to the time taken to execute and fetch the data. In such cases, it is important to get lots of executions of the statements, so that you have statistically valid numbers.

Avoiding the Time Trap

Sometimes, as in the following example, you might wonder why a particular query has taken so long.

```
UPDATE cq_names SET ATTRIBUTES = lower(ATTRIBUTES)
WHERE ATTRIBUTES = :att
```

call	count	cpu	elapsed	disk	query current	rows
Parse	1	0.06	0.24	0	0	0
Execute	1	0.62	19.62	22	526	12
Fetch	0	0.00	0.00	0	0	0

```
Misses in library cache during parse: 1
Parsing user id: 02 (USER2)
```

Rows	Execution Plan
------	----------------


```
-----
      0  UPDATE STATEMENT
2519  TABLE ACCESS (FULL) OF 'CQ_NAMES'
```

Again, the answer is interference from another transaction. In this case, another transaction held a shared lock on the table `cq_names` for several seconds before and after the update was issued. It takes a fair amount of experience to diagnose that interference effects are occurring. On the one hand, comparative data is essential when the interference is contributing only a short delay (or a small increase in block visits in the previous example). On the other hand, if the interference is contributing only a modest overhead, and the statement is essentially efficient, then its statistics might not need to be analyzed.

Avoiding the Trigger Trap

The resources reported for a statement include those for all of the SQL issued while the statement was being processed. Therefore, they include any resources used within a trigger, along with the resources used by any other recursive SQL (such as that used in space allocation). With the SQL Trace facility enabled, TKPROF reports these resources twice. Avoid trying to tune the DML statement if the resource is actually being consumed at a lower level of recursion.

If a DML statement appears to be consuming far more resources than you would expect, then check the tables involved in the statement for triggers and constraints that could be greatly increasing the resource usage.

Sample TKPROF Output

This section provides an extensive example of TKPROF output. Portions have been edited out for the sake of brevity.

Sample TKPROF Header

Copyright (c) Oracle Corporation 1979, 1999. All rights reserved.

Trace file: v80_ora_2758.trc

Sort options: default

```
count      = number of times OCI procedure was executed
cpu        = cpu time in seconds executing
elapsed    = elapsed time in seconds executing
disk       = number of physical reads of buffers from disk
query      = number of buffers gotten for consistent read
```

Sample TKPROF Output

current = number of buffers gotten in current mode (usually for update)
rows = number of rows processed by the fetch or execute call

The following statement encountered a error during parse:
select deptno, avg(sal) from emp e group by deptno
 having exists (select deptno from dept
 where dept.deptno = e.deptno
 and dept.budget > avg(e.sal)) order by 1
Error encountered: ORA-00904

Sample TKPROF Body

```
ALTER SESSION SET SQL_TRACE = true
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	0	0.00	0.00	0	0	0	0
Execute	1	0.00	0.10	0	0	0	0
Fetch	0	0.00	0.00	0	0	0	0
total	1	0.00	0.10	0	0	0	0

Misses in library cache during parse: 0

Misses in library cache during execute: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

```
SELECT emp.ename, dept.dname  
FROM emp, dept  
    WHERE emp.deptno = dept.deptno
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.11	0.13	2	0	1	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.00	0.00	2	2	4	14
total	3	0.11	0.13	4	2	5	14

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

```
-----  
    0 SELECT STATEMENT GOAL: CHOOSE  
   14 MERGE JOIN
```

```

4      SORT (JOIN)
4      TABLE ACCESS (FULL) OF 'DEPT'
14     SORT (JOIN)
14     TABLE ACCESS (FULL) OF 'EMP'

```

```
*****
```

```

SELECT a.ename name, b.ename manager
FROM emp a, emp b
  WHERE a.mgr = b.empno(+)

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.01	0.01	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.01	0.01	1	50	2	14
total	3	0.02	0.02	1	50	2	14

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 01 (USER01)

Rows Execution Plan

```

-----
0  SELECT STATEMENT  GOAL: CHOOSE
13  NESTED LOOPS (OUTER)
14  TABLE ACCESS (FULL) OF 'EMP'
13  TABLE ACCESS (BY ROWID) OF 'EMP'
26  INDEX (RANGE SCAN) OF 'EMP_IND' (NON-UNIQUE)

```

```
*****
```

```

SELECT ename, job, sal
FROM emp
WHERE sal =
  (SELECT max(sal)
   FROM emp)

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.00	0.00	0	12	4	1
total	3	0.00	0.00	0	12	4	1

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Sample TKPROF Output

Parsing user id: 01 (USER01)

Rows Execution Plan

```
-----  
      0 SELECT STATEMENT   GOAL: CHOOSE  
     14 FILTER  
     14 TABLE ACCESS (FULL) OF 'EMP'  
     14 SORT (AGGREGATE)  
     14 TABLE ACCESS (FULL) OF 'EMP'
```

```
SELECT deptno  
FROM emp  
WHERE job = 'clerk'  
GROUP BY deptno  
HAVING COUNT(*) >= 2
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.00	0.00	0	1	1	0
total	3	0.00	0.00	0	1	1	0

Misses in library cache during parse: 13

Optimizer goal: CHOOSE

Parsing user id: 01 (USER01)

Rows Execution Plan

```
-----  
      0 SELECT STATEMENT   GOAL: CHOOSE  
      0 FILTER  
      0 SORT (GROUP BY)  
     14 TABLE ACCESS (FULL) OF 'EMP'
```

```
SELECT dept.deptno, dname, job, ename  
FROM dept,emp  
WHERE dept.deptno = emp.deptno(+)  
ORDER BY dept.deptno
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.00	0.00	0	3	3	10
total	3	0.00	0.00	0	3	3	10

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 01 (USER01)

Rows Execution Plan

```
-----
      0 SELECT STATEMENT  GOAL: CHOOSE
     14 MERGE JOIN (OUTER)
        4 SORT (JOIN)
          4 TABLE ACCESS (FULL) OF 'DEPT'
         14 SORT (JOIN)
          14 TABLE ACCESS (FULL) OF 'EMP'
```

```
SELECT grade, job, ename, sal
FROM emp, salgrade
WHERE sal BETWEEN losal AND hisal
ORDER BY grade, job
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.04	0.06	2	16	1	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.01	0.01	1	10	12	10
total	3	0.05	0.07	3	26	13	10

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

```
-----
      0 SELECT STATEMENT  GOAL: CHOOSE
     14 SORT (ORDER BY)
     14 NESTED LOOPS
        5 TABLE ACCESS (FULL) OF 'SALGRADE'
       70 TABLE ACCESS (FULL) OF 'EMP'
```

```
SELECT LPAD(' ',level*2)||ename org_chart, level, empno, mgr, job, deptno
FROM emp
CONNECT BY prior empno = mgr
START WITH ename = 'clark'
       OR ename = 'blake'
ORDER BY deptno
```

Sample TKPROF Output

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.01	0.01	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.01	0.01	0	1	2	0
total	3	0.02	0.02	0	1	2	0

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

```

-----
0  SELECT STATEMENT  GOAL: CHOOSE
0  SORT (ORDER BY)
0  CONNECT BY
14 TABLE ACCESS (FULL) OF 'EMP'
0  TABLE ACCESS (BY ROWID) OF 'EMP'
0  TABLE ACCESS (FULL) OF 'EMP'

```

```

*****
CREATE TABLE TKOPTKP (a number, b number)

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.01	0.01	1	0	1	0
Fetch	0	0.00	0.00	0	0	0	0
total	2	0.01	0.01	1	0	1	0

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

```

-----
0  CREATE TABLE STATEMENT  GOAL: CHOOSE

```

```

*****
INSERT INTO TKOPTKP
VALUES (1,1)

```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.07	0.09	0	0	0	0

Execute	1	0.01	0.20	2	2	3	1
Fetch	0	0.00	0.00	0	0	0	0
-----	-----	-----	-----	-----	-----	-----	-----
total	2	0.08	0.29	2	2	3	1

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

 0 INSERT STATEMENT GOAL: CHOOSE

 INSERT INTO TKOPTKP SELECT * FROM TKOPTKP

call	count	cpu	elapsed	disk	query	current	rows
-----	-----	-----	-----	-----	-----	-----	-----
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.02	0.02	0	2	3	11
Fetch	0	0.00	0.00	0	0	0	0
-----	-----	-----	-----	-----	-----	-----	-----
total	2	0.02	0.02	0	2	3	11

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Rows Execution Plan

 0 INSERT STATEMENT GOAL: CHOOSE
 12 TABLE ACCESS (FULL) OF 'TKOPTKP'

 SELECT *
 FROM TKOPTKP
 WHERE a > 2

call	count	cpu	elapsed	disk	query	current	rows
-----	-----	-----	-----	-----	-----	-----	-----
Parse	1	0.01	0.01	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	1	0.00	0.00	0	1	2	10
-----	-----	-----	-----	-----	-----	-----	-----
total	3	0.01	0.01	0	1	2	10

Misses in library cache during parse: 1

Optimizer goal: CHOOSE

Parsing user id: 02 (USER02)

Sample TKPROF Output

```
Rows      Execution Plan
-----
          0  SELECT STATEMENT   GOAL: CHOOSE
          24  TABLE ACCESS (FULL) OF 'TKOPTKP'
*****
```

Sample TKPROF Summary

OVERALL TOTALS FOR ALL NON-RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
Parse	18	0.40	0.53	30	182	3	0
Execute	19	0.05	0.41	3	7	10	16
Fetch	12	0.05	0.06	4	105	66	78
total	49	0.50	1.00	37	294	79	94

Misses in library cache during parse: 18

Misses in library cache during execute: 1

OVERALL TOTALS FOR ALL RECURSIVE STATEMENTS

call	count	cpu	elapsed	disk	query	current	rows
Parse	69	0.49	0.60	9	12	8	0
Execute	103	0.13	0.54	0	0	0	0
Fetch	213	0.12	0.27	40	435	0	162
total	385	0.74	1.41	49	447	8	162

Misses in library cache during parse: 13

19 user SQL statements in session.

69 internal SQL statements in session.

88 SQL statements in session.

17 statements EXPLAINED in this session.

Trace file: v80_ora_2758.trc

Trace file compatibility: 7.03.02

Sort options: default

1 session in tracefile.

19 user SQL statements in trace file.

69 internal SQL statements in trace file.

88 SQL statements in trace file.

41 unique SQL statements in trace file.

17 SQL statements EXPLAINED using schema:


```
SCOTT.prof$plan_table  
  Default table was used.  
  Table was created.  
  Table was dropped.  
1017 lines in trace file.
```

Using Autotrace in SQL*Plus

This chapter provides information on using the Autotrace feature in SQL*Plus and the iSQL*Plus statistics report.

This chapter contains the following sections:

- [Overview of the Autotrace Report](#)
- [Collecting Timing Statistics](#)
- [Tracing Parallel and Distributed Queries](#)
- [SYSTEM Variables Influencing SQL*Plus Performance](#)
- [iSQL*Plus Server Statistics Report](#)

Overview of the Autotrace Report

In SQL*Plus you can automatically get a report on the execution path used by the SQL optimizer and the statement execution statistics. The report is generated after a successful SQL DML statement, such as `SELECT`, `DELETE`, `UPDATE` or `INSERT`. It is useful for monitoring and tuning the performance of these DML statements.

Configuring the Autotrace Report

You can control the report by setting the `AUTOTRACE` system variable. See [Table 11-1](#).

Table 11-1 Autotrace Settings

Autotrace Setting	Result
<code>SET AUTOTRACE OFF</code>	No <code>AUTOTRACE</code> report is generated. This is the default.
<code>SET AUTOTRACE ON EXPLAIN</code>	The <code>AUTOTRACE</code> report shows only the optimizer execution path.
<code>SET AUTOTRACE ON STATISTICS</code>	The <code>AUTOTRACE</code> report shows only the SQL statement execution statistics.
<code>SET AUTOTRACE ON</code>	The <code>AUTOTRACE</code> report includes both the optimizer execution path and the SQL statement execution statistics.
<code>SET AUTOTRACE TRACEONLY</code>	Similar to <code>SET AUTOTRACE ON</code> , but suppresses the printing of the user's query output, if any. If <code>STATISTICS</code> is enabled, query data is still fetched, but not printed.

Setups Required for the Autotrace Report

To use this feature, the `PLUSTRACE` role must be granted to the user, such as `HR`. DBA privileges are required to grant the `PLUSTRACE` role.

Additionally, a `PLAN_TABLE` table must be created in the user's schema, such as the `HR` schema. For information on creating the `PLAN_TABLE`, see "[Creating the PLAN_TABLE Output Table](#)" on page 9-4.

To create the `PLUSTRACE` role and grant it to the DBA, run the commands in [Example 11-1](#) from a SQL*Plus session.

Example 11-1 Creating the PLUSTRACE Role

```
CONNECT / AS SYSDBA
@$ORACLE_HOME/SQLPLUS/ADMIN/PLUSTRCE.SQL
```

```

drop role plustrace;
Role dropped.
create role plustrace;
Role created.
.
grant plustrace to dba with admin option;
Grant succeeded.

```

To grant the `PLUSTRACE` role to the HR user, run the commands in [Example 11–2](#) from a SQL*Plus session.

Example 11–2 Granting the PLUSTRACE Role

```

CONNECT / AS SYSDBA
GRANT PLUSTRACE TO HR;
Grant succeeded.

```

See Also:

- *Oracle9i SQL Reference* for more information about granting roles and creating the `PLAN_TABLE` table
- *SQL*Plus User's Guide and Reference* for more information about the `PLUSTRACE` role and about the `AUTOTRACE` variable of the `SET` command
- *Oracle9i Database Administrator's Guide* for information about roles

Execution Plans for SQL Statements

An execution plan shows the SQL optimizer's query execution path. Each line of the execution plan has a sequential line number. SQL*Plus also displays the line number of the parent operation. For a discussion and an example of an execution plan, see "[Understanding Execution Plans](#)" on page 1-18.

The execution plan output is generated using the `EXPLAIN PLAN` command. The format of the output can vary depending on your setups. The format of the columns of the `PLAN_TABLE` may be altered with the `COLUMN` command. For example, to stop the `PARENT_ID` column being displayed, enter the following:

```
COLUMN PARENT_ID NOPRINT
```

The default formats can be found in the site profile, such as, `GLOGIN.SQL`.

See Also: [Chapter 9, "Using EXPLAIN PLAN"](#) for more information about generating and interpreting the output of EXPLAIN PLAN

Database Statistics for SQL Statements

The statistics are recorded by the server when your statement executes and indicate the system resources required to execute your statement. The results include the statistics listed in [Table 11-2](#).

Table 11-2 Database Statistics

Database Statistic Name	Description
recursive calls	Number of recursive calls generated at both the user and system level. Oracle maintains tables used for internal processing. When Oracle needs to make a change to these tables, it internally generates an internal SQL statement, which in turn generates a recursive call.
db block gets	Number of times a CURRENT block was requested.
consistent gets	Number of times a consistent read was requested for a block.
physical reads	Total number of data blocks read from disk. This number equals the value of "physical reads direct" plus all reads into buffer cache.
redo size	Total amount of redo generated in bytes.
bytes sent via SQL*Net to client	Total number of bytes sent to the client from the foreground processes.
bytes received via SQL*Net from client	Total number of bytes received from the client over Oracle Net.
SQL*Net roundtrips to/from client	Total number of Oracle Net messages sent to and received from the client.
sorts (memory)	Number of sort operations that were performed completely in memory and did not require any disk writes.
sorts (disk)	Number of sort operations that required at least one disk write.
rows processed	Number of rows processed during the operation.

The client referred to in the statistics is SQL*Plus. Oracle Net refers to the generic process communication between SQL*Plus and the server, regardless of whether Oracle Net is installed. You cannot change the default format of the statistics report.

See Also:

- *Oracle9i Database Reference* for a more complete list of database statistics
- [Chapter 3, "Gathering Optimizer Statistics"](#) for more information about the statistics and how to interpret them

Tracing Statements Examples

This section shows examples of the use of the AUTOTRACE feature.

Tracing Statements for Performance Statistics and Query Execution Path

If the SQL buffer contains the following statement:

```
SELECT E.LAST_NAME, E.SALARY, J.JOB_TITLE
       FROM EMPLOYEES E, JOBS J
       WHERE E.JOB_ID=J.JOB_ID AND E.SALARY>12000;
```

The statement can be automatically traced when it is run with the following:

```
SET AUTOTRACE ON
/
```

The output is similar to the following:

LAST_NAME	SALARY	JOB_TITLE
King	24000	President
Kochhar	17000	Administration Vice President
De Haan	17000	Administration Vice President
Russell	14000	Sales Manager
Partners	13500	Sales Manager
Hartstein	13000	Marketing Manager

6 rows selected.

Execution Plan

```
-----
0      SELECT STATEMENT Optimizer=CHOOSE
1      0      TABLE ACCESS (BY INDEX ROWID) OF 'EMPLOYEES'
2      1      NESTED LOOPS
3      2      TABLE ACCESS (FULL) OF 'JOBS'
4      2      INDEX (RANGE SCAN) OF 'EMP_JOB_IX' (NON-UNIQUE)
```

Statistics

```
-----  
0 recursive calls  
2 db block gets  
34 consistent gets  
0 physical reads  
0 redo size  
848 bytes sent via SQL*Net to client  
503 bytes received via SQL*Net from client  
4 SQL*Net roundtrips to/from client  
0 sorts (memory)  
0 sorts (disk)  
6 rows processed
```

Note: Your output may vary depending on the version of the server to which you are connected and the configuration of the server.

Tracing Statements Without Displaying Query Data

To trace the same statement without displaying the query data, enter the following:

```
SET AUTOTRACE TRACEONLY
```

This option is useful when you are tuning a large query, but do not want to see the query output.

Tracing Statements Using a Database Link

To trace a statement using a database link, enter:

```
SET AUTOTRACE TRACEONLY EXPLAIN  
SELECT * FROM EMPLOYEES@MY_LINK;
```

Execution Plan

```
-----  
0          SELECT STATEMENT (REMOTE) Optimizer=CHOOSE  
1    0    TABLE ACCESS (FULL) OF 'EMPLOYEES'  MY_LINK.DB_DOMAIN
```

The execution plan shows that the table being accessed on line 1 of the execution plan is through the database link MY_LINK.DB_DOMAIN.

Collecting Timing Statistics

Use the SQL*Plus `TIMING` command to collect and display data on the amount of computer resources used to run one or more commands or blocks. `TIMING` collects data for an elapsed period of time, saving the data on commands run during the period in a timer.

To delete all timers, enter `CLEAR TIMING` at the command prompt.

See Also: *SQL*Plus User's Guide and Reference* for information about the `TIMING` command

Tracing Parallel and Distributed Queries

When you trace a statement in a parallel or distributed query, the execution plan shows the cost based optimizer estimates of the number of rows, called the cardinality. In general, the cost, cardinality and bytes at each node represent cumulative results. For example, the cost of a join node accounts for not only the cost of completing the join operations, but also the entire costs of accessing the relations in that join.

[Example 11-3](#) is an example of tracing statements with the parallel query option. The output varies depending on the configuration of a system. In the execution plan output, items marked with an asterisk (*) denote a parallel or remote operation. Each operation is explained in the second part of the report. The execution plan in [Example 11-3](#) consists of columns that contain information such as:

- Line number of each execution step
- Function of the SQL statement
- Text of the query for the Oracle Real Application Clusters database or remote database

For more information about the columns in the `PLAN_TABLE`, see "[PLAN_TABLE Columns](#)" on page 9-23.

Example 11-3 *Tracing Statements With Parallel Query Option*

To trace a parallel query running the parallel query option:

```
SQL> CREATE TABLE D2_t1 (unique1 NUMBER) PARALLEL -(DEGREE 6);
Table created.
```

```
SQL> CREATE TABLE D2_t2 (unique1 NUMBER) PARALLEL -(DEGREE 6);
Table created.
```

```
SQL> CREATE UNIQUE INDEX d2_i_unique1 ON d2_t1(unique1);
Index created.
```

```
SQL> SET LONG 500 LONGCHUNKSIZE 500
```

```
SQL> SET AUTOTRACE ON EXPLAIN
```

```
SQL> SELECT /*+ INDEX(B,D2_I_UNIQUE1) USE_NL(B) ORDERED -*/ COUNT (A.UNIQUE1)
FROM D2_T2 A, D2_T1 B
WHERE A.UNIQUE1 = B.UNIQUE1;
```

Execution Plan

```
-----
0      SELECT STATEMENT Optimizer=CHOOSE (Cost=1 Card=1 Bytes=26)
1    0  SORT (AGGREGATE)
2    1  SORT* (AGGREGATE)                                :Q2000
3    2  NESTED LOOPS* (Cost=1 Card=41 Bytes=1066)       :Q2000
4    3    TABLE ACCESS* (FULL) OF 'D2_T2' (Cost=1 Card=41 Byte :Q2000 s=533)
5    3      INDEX* (UNIQUE SCAN) OF 'D2_I_UNIQUE1' (UNIQUE) :Q2000

2 PARALLEL_TO_SERIAL                                SELECT /*+ PIV_SSF */ SYS_OP_MSR(COUNT(A1.C0
)) FROM (SELECT /*+ ORDERED NO_EXPAND USE_NL
(A3) INDEX(A3 "D2_I_UNIQUE1") */ A2.C0 C0,A3
.ROWID C1,A3."UNIQUE1" C2 FROM (SELECT /*+ N
O_EXPAND ROWID(A4) */ A4."UNIQUE1" C0 FROM "
D2_T2" PX_GRANULE(0, BLOCK_RANGE, DYNAMIC)
A4) A2,"D2_T1" A3 WHERE A2.C0=A3."UNIQUE1")
A1

3 PARALLEL_COMBINED_WITH_PARENT
4 PARALLEL_COMBINED_WITH_PARENT
5 PARALLEL_COMBINED_WITH_PARENT
```

Line 0 of the execution plan shows the cost based optimizer estimates the number of rows at 1, taking 26 bytes. The total cost of the statement is 1. Lines 2, 3, 4 and 5 are marked with asterisks, denoting parallel operations. For example, the NESTED LOOPS step on line 3 is a PARALLEL_TO_SERIAL operation. PARALLEL_TO_SERIAL operations execute a SQL statement to produce output serially. Line 2 also shows that the parallel query server had the identifier Q2000.

Numbers identifying parallel report lines cross reference the line of the parent report. For example, in the last line of the example:

```
4 PARALLEL_COMBINED_WITH_PARENT
```

The 4 refers to the 4 3 TABLE ACCESS*... line in the execution plan.

Monitoring Disk Reads and Buffer Gets

To monitor disk reads and buffer gets, execute the following command:

```
SET AUTOTRACE ON TRACEONLY STATISTICS
```

[Example 11-4](#) shows typical results.

Example 11-4 Monitoring Disk Reads and Buffer Gets

Statistics

```
-----
70 recursive calls
0 db block gets
591 consistent gets
404 physical reads
0 redo size
315 bytes sent via SQL*Net to client
850 bytes received via SQL*Net from client
3 SQL*Net roundtrips to/from client
3 sorts (memory)
0 sorts (disk)
0 rows processed
```

If consistent gets or physical reads is high relative to the amount of data returned, it indicates that the query is expensive and needs to be reviewed for optimization. For example, if you are expecting less than 1,000 rows back and consistent gets is 1,000,000 and physical reads is 10,000, further optimization is needed.

Note: You can also monitor disk reads and buffer gets using V\$SQL or TKPROF.

SYSTEM Variables Influencing SQL*Plus Performance

The following variables can influence SQL*Plus performance.

SET APPINFO OFF

Sets automatic registering of scripts through the DBMS_APPLICATION_INFO package. Setting APPINFO OFF disables the registering and monitoring of performance and resource usage of scripts. This reduction in overheads may improve performance.

SET ARRAYSIZE

Sets the number of rows, called a batch, that SQL*Plus will fetch from the database at one time. Valid values are 1 to 5000. A large value increases the efficiency of queries and subqueries that fetch many rows, but requires more memory. Values over approximately 100 provide little added performance. `ARRAYSIZE` has no effect on the results of SQL*Plus operations other than increasing efficiency.

SET DEFINE OFF

Controls whether SQL*Plus parses scripts for substitution variables. If `DEFINE` is `OFF`, SQL*Plus does not parse scripts for substitution variables. If your script does not use substitution variables, setting `DEFINE OFF` may result in some performance gains.

SET FLUSH OFF

Controls when output is sent to the user's display device. `OFF` allows the host operating system to buffer output which may improve performance by reducing the amount of program input and output.

Use `OFF` only when you run a script that does not require user interaction and whose output you do not need to see until the script finishes running.

`SET FLUSH` is not supported in *i*SQL*Plus.

SET SERVEROUTPUT

Controls whether SQL*Plus checks for and displays DBMS output. If `SERVEROUTPUT` is `OFF`, SQL*Plus does not check for DBMS output and does not display output after applicable SQL or PL/SQL statements. Suppressing this output checking and display may result in performance gains.

SET TRIMOUT ON

Determines whether SQL*Plus allows trailing blanks at the end of each displayed line. `ON` removes blanks at the end of each line, which may improve performance especially when you access SQL*Plus from a slow communications device.

`TRIMOUT ON` does not affect spooled output.

`SET TRIMOUT` is not supported in *i*SQL*Plus.

SET TRIMSPPOOL ON

Determines whether SQL*Plus allows trailing blanks at the end of each spooled line. ON removes blanks at the end of each line, which may improve performance especially when you access SQL*Plus from a slow communications device. TRIMSPPOOL ON does not affect terminal output.

SET TRIMSPPOOL is not supported in iSQL*Plus.

iSQL*Plus Server Statistics Report

The iSQL*Plus Server statistics report provides static environment information as well as dynamic information about iSQL*Plus sessions. The active statistics in the report are useful for monitoring and tuning iSQL*Plus. You can display the iSQL*Plus Server statistics report with:

```
http://machine_name.domain:port/isqlplusdba?statistics={active|full}
[&refresh=number]
```

where

- *machine_name.domain* is the URL of the Oracle HTTP Server for which you want to generate iSQL*Plus Server statistics.
- *port* is the port number used by the iSQL*Plus Server. The default is 7777.
- *?statistics={active|full}* specifies the level of detail to report.
 - *full* gives all possible statistics and is the default.
 - *active* gives dynamically changing session statistics for the iSQL*Plus Server. These statistics are also included at the end of the full report.
- [*&refresh=number*] optionally specifies the time in seconds before the statistics report is automatically refreshed. The minimum value is 10 seconds.

To run the report, you must have Oracle HTTP Server authentication to access the iSQL*Plus DBA URL. However, an Oracle9i login is not required because there is no connection to a database. To maximize resource availability, each user of iSQL*Plus should have a database schema profile with appropriately defined limits.

See Also: For more information about the iSQL*Plus Server statistics report, including the full set of statistics, see *SQL*Plus User's Guide and Reference*

Active Statistics

The active statistics report shows the current values for the statistics listed in [Table 11-3](#). These statistics provide useful feedback for tuning the iSQL*Plus Server.

Table 11-3 Active Statistics

Statistic	Description
Sessions active	The number of concurrent active sessions, or the number of people currently logged in to iSQL*Plus.
Sessions since startup	The cumulative count of sessions established since the iSQL*Plus Server started.
Maximum concurrent sessions	The maximum or peak number of concurrent sessions since the iSQL*Plus Server started.
Sessions expired since startup	The cumulative count of the number of sessions timed-out due to inactivity since the iSQL*Plus Server started.
Requests active	The number of concurrent active HTTP requests. Each request corresponds to a user action such as clicking a button, and the processing of that request by iSQL*Plus. Requests active has a maximum value set by <code>iSQLPlusNumberOfThreads</code> . If it reaches this limit and user response time is poor, then response time may be improved by increasing <code>iSQLPlusNumberOfThreads</code> .
Requests since startup	The cumulative count of active HTTP requests since the iSQL*Plus server was started.
Next expiry operation (minutes)	The number of minutes (rounded down) until the next expiry process.
Expiry operations since startup	The number of times the expiry process has run since the iSQL*Plus Server started.
Hash table collisions	The number of active sessions that currently have a hash table collision. Compare this with Sessions active to see if there is a current problem with collisions.
Hash table collisions since startup	The cumulative count of the sessions that have had a hash table collision since the iSQL*Plus Server started. Compare this with Sessions since startup to see if there is an ongoing problem with collisions.

Interpreting Active Statistics

The following notes provide some interpretation of the active statistics.

Increasing Number of Threads

If users have more idle time compared to active time, then a higher value of `iSQLPlusNumberOfThreads` may be needed.

- Active time is when a user request is in progress, and a thread to process it is consumed.
- Idle time is when a user request has been processed, and the associated processing thread is available for use by another iSQL*Plus session.

Each thread can handle one user request. A request begins when a user clicks a button or follows a command link in iSQL*Plus, and finishes when all results have been returned to the user.

When setting the value of `iSQLPlusNumberOfThreads`, note the following:

- If `iSQLPlusHashTableSize` is not specified in the `isqlplus.conf` file, its value increases when `iSQLPlusNumberOfThreads` is increased.
- If the `iSQLPlusNumberOfThreads` value is set too small, iSQL*Plus may not be able to handle the request load. Indications of this problem are:
 - Frequent Connection refused errors in `ORACLE_HOME/Apache/Apache/logs/error_log`
 - An Oracle HTTP Server internal error:
The server encountered an internal error or misconfiguration and was unable to complete your request.

Increasing Hash Table Size

If users typically have more idle time compared to active time, then a higher value of `iSQLPlusHashTableSize` is needed for a given value of `iSQLPlusNumberOfThreads`. Each user session consumes one entry in the hash table even if the session is idle.

Reducing Timeout Interval

If large numbers of sessions are being timed out, it is an indication that users are not logging out cleanly, and sessions may be remaining idle. In this case, and if the iSQL*Plus Server load is high, you may want to consider reducing the `iSQLPlusTimeOutInterval` to more aggressively time out sessions.

Idle Timeout

The idle timeout is the time the Oracle HTTP Server waits for results from iSQL*Plus. The parameter value for the FastCGI timeout parameter, `-idle-timeout`, is set to 3600 seconds. This value is likely to prevent iSQL*Plus timing out before the Web browser and is sufficient for many long queries to return results before iSQL*Plus times out.

The idle timeout should not be confused with the `iSQLPlusTimeOutInterval` which manages the lifetime of a user's session.

Using Oracle Trace

This chapter describes how to use Oracle Trace to collect Oracle server event data.

The topics in this chapter include:

- [Overview of Oracle Trace](#)
- [Collecting Oracle Trace Data](#)
- [Accessing Oracle Trace Collection Results](#)
- [Oracle Server Events](#)
- [Troubleshooting Oracle Trace](#)

Note: Oracle Trace will be deprecated in a future release. Oracle Corporation strongly advises the use of SQL Trace and TKPROF instead.

See Also: [Chapter 10, "Using SQL Trace and TKPROF"](#)

Overview of Oracle Trace

Oracle Trace is a general-purpose event-driven data collection product, which the Oracle server uses to collect performance and resource utilization data, such as SQL parse, execute, and fetch statistics, and wait statistics.

Event Data

An **event** is the occurrence of some activity within a product. Oracle Trace collects data for predefined events occurring within a software product instrumented with the Oracle Trace API. That is, the product has embedded calls to the Oracle Trace API. An example of an event is a parse or fetch.

There are two types of events:

- Point events
Point events represent an instantaneous occurrence of something in the product. An example of a point event is an error occurrence.
- Duration events
Duration events have a beginning and an ending. An example of a duration event is a transaction. Duration events can have other events occur within them; for example, an error can occur *within* a transaction.

The Oracle server has more than a dozen events. The following are three of these events:

- Database Connection: A point event that records data, such as the server login user name.
- SQL Parse: One of the series of SQL processing duration events. This event records a large set of data, such as sorts, resource use, and cursor numbers.
- RowSource: Data about the execution plan, such as SQL operation, position, object identification, and number of rows processed by a single row source within an execution plan.

Event Sets

Oracle Trace events can be organized into **event sets** that restrict the data collection to specific events. You can establish event sets for performance monitoring, auditing, diagnostics, or any logical event grouping.

Each event set is described by its own **product definition file** (.pdf). The product definition file is a set of events and their associated data items. The complete set of

events defined for an instrumented product is referred to as the ALL event set. Other event sets are then derived from the ALL set. For example, the Oracle Server includes an event set known as the EXPERT set. This set includes SQL event data used by the Oracle Expert tuning application, but excludes other events, such as wait events.

Accessing Collected Data

During a collection, Oracle Trace stores event data in memory and periodically writes it to a collection binary file. This method ensures low resource overhead associated with the collection process. You can access event data collected in the binary file by formatting the data to database tables, which makes the data available for fast, flexible access. These database tables are called Oracle Trace formatter tables.

Collecting Oracle Trace Data

You can collect Oracle Trace data using one of the following mechanisms:

- Oracle Trace Command-Line Interface (CLI)
- Database `ORACLE_TRACE_*` initialization parameters
- Oracle Trace stored procedures run from PL/SQL

Using the Oracle Trace Command-Line Interface

You can control Oracle Trace server collections with the Oracle Trace CLI (command-line interface). The CLI is invoked by the `OTRCCOL` executable for the following functions:

- `OTRCCOL START job_id input_parameter_file`
- `OTRCCOL STOP job_id input_parameter_file`
- `OTRCCOL CHECK collection_name`
- `OTRCCOL FORMAT input_parameter_file`
- `OTRCCOL DCF col_name cdf_file`
- `OTRCCOL DFD col_name username password service [col_id]`

The `job_id` parameter should be set to a value of 1.

The input parameter file contains specific parameter values required for each function, as shown in the following examples. `col_name` (collection name) and `cdf_file` (collection definition file) are initially defined in the `START` function input parameter file.

Note: The server parameter `ORACLE_TRACE_ENABLE` must be set to `true` to allow Oracle Trace to collect any server event data. This is required for Oracle8 and newer servers.

Starting Collections

The `OTRCCOL START` function invokes a collection based on parameter values contained in the input parameter file. For example:

```
OTRCCOL START 1 my_start_input_file
```

where file `my_start_input_file` contains at least the following input parameters:

```
col_name= collection name
cdf_file= collection name.cdf
dat_file= collection name.dat
fdf_file= facility definition file.fdf
```

The server event sets that can be used as values for the `fdf_file` parameter are `ORACLE`, `ORACLEC`, `ORACLEED`, `ORACLEE`, and `ORACLESM`, plus `CONNECT`, `SQL_ONLY`, `SQL_PLAN`, `SQL_TXN`, `SQLSTATS`, `SQLWAITS`, and `WAITS`.

See Also: [Table 12-2](#) for a description of the server event sets

Collection `.cdf` and `.dat` files are created in the directory `$ORACLE_HOME/otrace/admin/cdf` by default for collections started using the CLI (or PL/SQL procedures), unless overridden by `EPC_COLLECTION_PATH` environment variable.

Note: This chapter refers to file path names on UNIX-based systems. For the exact path on other operating systems, see your Oracle platform-specific documentation. A complete discussion of these parameters is provided in *Oracle9i Database Reference*.

The `collections name` parameter can be any valid unique filename.

For Oracle database collections, one additional parameter is required:

```
regid= 1 192216243 0 0 5 database_SID
```

Note: Older CLI versions required this syntax exactly, with no whitespace before '=' but at least some whitespace after '='. This is no longer true: CLI is not whitespace sensitive.

The `regid` parameter record identifies a database by SID where Oracle Trace collection is to be performed. The six elements making up the `regid` parameter record are as follows, in this order:

- flag (this should be 1 for this usage)
- vendor number (Oracle's vendor number is 192216243)
- `cf_num`
- `cf_val`
- facility number (the Oracle server is facility number 5)
- database SID

The `cf_num` and `cf_val` elements should set to zero in this basic Oracle database collection `regid`.

Limiting Data Collected

There are several ways of limiting the amount of data collected. For example, additional `regid` records can be specified to reduce the amount of collected data, and nonzero `cf_num` and `cf_val` can be specified in those situations. In the Oracle server, Oracle Trace cross facility item 6 (`cf_num = 6`) is reserved to record database userID values.

Restrict By User

In the Oracle server, Oracle Trace cross facility item 6 (`cf_num = 6`) is reserved to record database userID values. So, for example, if you provide an additional `regid` record with `cf_num = 6` and `cf_val = some_DB_userID`, then the collection of database event data is limited to only those events performed by that database user.

If you are interested only in collecting database activity for users 23 and 45, then you would provide the following 3 `regid` records:

```
regid= 1 192216243 0 0 5 ORCL
```

```
regid= 1 192216243 6 23 5 ORCL  
regid= 1 192216243 6 45 5 ORCL
```

Restrict By Process

The input parameter file used by the CLI when starting a collection can also contain the following optional parameters, for both database and nondatabase Oracle Trace collections:

```
prores= process restriction  
max_cdf= maximum collection file size
```

If no process restriction records are specified, then there are no restrictions on which processes can take part in the collection. If process restrictions are used, then one or more process ID (PID) values can be specified, as well as the operating system username for the owner of each process of interest.

Setting a Limit on File Size

The `max_cdf` parameter is often useful, in several different modes of use. This parameter specifies the maximum amount of Oracle Trace data that should be collected, in bytes (in other words, size of the `collection.dat` file).

A zero value indicates that no limit should be imposed; otherwise, a positive value up to 2 GB can be specified to stop the data collection when that size limit is reached. In addition, a negative value can be specified (but not less than -2 GB), which instructs Oracle Trace to collect data in its "circular data file" mode: when `collection.dat` reaches magnitude(`max_cdf`), then save that data (and delete any previously saved dat file), and then start collecting to a new `collection.dat` file. This limits the total amount of disk space used, but allows Oracle Trace data collection to continue until you manually stop collection.

Checking the Status of a Collection

Verify that the collection was started.

```
otrccol check collection_name
```

The collection should show as active, not active, or not found.

Stopping Collections

The `OTRCCOL STOP` function halts a running collection as follows:

```
OTRCCOL STOP 1 my_stop_input_file
```

where `my_stop_input_file` contains the collection name and `cdf_file` name.

Formatting Collections

The `OTRCCOL FORMAT` function formats the binary collection file to Oracle tables. An example of the `FORMAT` statement is:

```
OTRCCOL FORMAT my_format_input_file
```

where `my_format_input_file` contains the following input parameters:

```
username= database username
password= database password
service= database service name
cdf_file= usually same as collection_name.cdf
full_format= 0/1
```

A `full_format` value of 1 produces a full format. A `full_format` value of 0 produces a partial format, which only formats new data; in other words, data collected since any previous format.

See Also: ["Formatting Oracle Trace Data to Oracle Tables"](#) on page 12-13 for more information on formatting an Oracle Trace collection using the `otrcfmt` utility program.

Deleting Collections

The `OTRCCOL DCF` (delete collection files) function deletes collection `.cdf` and `.dat` files for a specific collection. The `OTRCCOL DFD` (delete formatted data) function deletes formatted data from the Oracle Trace formatter tables for a specific collection. You can specify an optional `col_id` parameter for a selective DFD, where more than one `col_id` has been created for a collection by multiple (full) formats.

Using Initialization Parameters to Control Oracle Trace

Six Oracle database initialization parameters are set up by default to control Oracle Trace. By logging into a privileged account in the database and executing the `SHOW PARAMETER ORACLE_TRACE` statement, you see the following parameters:

Table 12-1 Oracle Trace Initialization Parameters

Name	Type	Default Value
<code>ORACLE_TRACE_COLLECTION_NAME</code>	string	[null]

Table 12–1 Oracle Trace Initialization Parameters

Name	Type	Default Value
ORACLE_TRACE_COLLECTION_PATH	string	\$ORACLE_HOME/otrace/admin/cdf
ORACLE_TRACE_COLLECTION_SIZE	integer	5242880
ORACLE_TRACE_ENABLE	boolean	false
ORACLE_TRACE_FACILITY_NAME	string	oracled
ORACLE_TRACE_FACILITY_PATH	string	\$ORACLE_HOME/otrace/admin/cdf

You can modify these Oracle Trace server parameters to allow Oracle Trace collection of server event data and use them by adding them to the initialization file.

However, this method for controlling the Oracle Trace collection is rather inflexible: the collection name cannot be changed without performing a database shutdown. (For Oracle releases prior to 8.1.7, the collection can only be stopped by doing a shutdown, then setting `ORACLE_TRACE_ENABLE = FALSE` before restarting.) However, with `ORACLE_TRACE_ENABLE = TRUE` but `ORACLE_TRACE_COLLECTION_NAME = ""` [that is, empty name string], Oracle Trace collections of database event data can be performed using one of the other collection control mechanisms; for example, the Oracle Trace CLI. These other mechanisms are more flexible than the database initialization parameters. In general, they are preferred over using parameters for collection control.

Enabling Oracle Trace Collections

The `ORACLE_TRACE_ENABLE` database initialization parameter is `false` by default. This disables any collection of Oracle Trace data for that server, regardless of the mechanism used to control the collection.

Setting `ORACLE_TRACE_ENABLE` to `true` in `DBinit.ora` enables Oracle Trace collections for the server, but it does not necessarily start a collection when the database instance is started. If the database parameters alone are to be used to start an Oracle Trace collection of database event data, then all 6 `ORACLE_TRACE_*` parameters must be specified, or have nonnull values by default. Typically, this means that both `ORACLE_TRACE_ENABLE` must be set to `true` and a nonnull `ORACLE_TRACE_COLLECTION_NAME` must be provided (up to 16 characters in length).

Note: The collection name is also used to form the *collection name.cdf* and *.dat* binary file names, so 8.3 file naming conventions may apply on some platforms. 8.3 file naming means systems where filenames are restricted to 8 or fewer characters, plus a file extension of 3 or fewer characters.

ORACLE_TRACE_ENABLE is now a dynamic parameter (as of Oracle8i, Release 3), so it can be set to `true` or `false` while the database is running. This can be done for the current database session or for all sessions (including future ones), using `ALTER SESSION` or `ALTER SYSTEM` statements. When the database is subsequently shut down and then restarted, the `DBinit.ora` setting for ORACLE_TRACE_ENABLE is again used to initially enable or disable Oracle Trace collection of database event data.

Determining the Event Set that Oracle Trace Collects

The ORACLE_TRACE_FACILITY_NAME database initialization parameter specifies the event set that Oracle Trace collects, if the database parameters are used to control data collection. The default for this parameter is ORACLEED (in other words, Oracle "default" event set).

Note: The ORACLE_TRACE_FACILITY_NAME parameter does not use a file extension. So, the `.fdf` extension should not be specified as part of this parameter.

With database parameters set to start an Oracle Trace Collection, if the database does not begin collecting data, then check the following:

- The event set file, identified by ORACLE_TRACE_FACILITY_NAME, with an `.fdf` extension, should be in the directory specified by the ORACLE_TRACE_FACILITY_PATH initialization parameter. The exact directory that this parameter specifies is platform-specific.
- The following files should exist in the Oracle Trace admin directory: `COLLECT.DAT`, `FACILITY.DAT` (or `PROCESS.DAT` for Oracle 7.3), and `REGID.DAT`. If they do not, then run the `OTRCCREF` executable to create or re-create them.

- The Oracle Trace parameters should be set to the values that you changed in the initialization file. Use Instance Manager to identify Oracle Trace parameter settings.
- Look for an `EPC_ERROR.LOG` file to see more information about why a collection failed. Oracle Trace creates the `EPC_ERROR.LOG` file in the current default directory if the Oracle Trace Collection Services `OTRCCOL` image must log an error.
- Look for `*.trc` files in the directory specified by the server `USER_DUMP_DEST` initialization parameter. Searching for "epc" in the `*.trc` files might give errors. These errors and their descriptions may be found in `$ORACLE_HOME/otrace/msgs/epcus.msg` (assuming US installation), depending on availability for a given platform.

Controlling Oracle Trace Collections from PL/SQL

Oracle provides an additional Oracle Trace library that allows control of both database and nondatabase Oracle Trace collections from PL/SQL.

Note: Older versions of the Oracle server, back to release 7.3.3, provided other stored procedures to start and stop Oracle Trace database collections, but with significant limitations. For example, only database sessions already active when the collection was started participated in the collection: no database event data was collected for sessions that began after the Oracle Trace collection started.

As of Oracle8, these limitations were eliminated for database collections started through the Oracle Trace CLI. However, these limitations still applied to database collections controlled through the older Oracle7 stored procedures. The new procedures provided with Oracle Trace 8.1.7 remove these limitations, permitting the same level of collection control as the Oracle Trace CLI, and for both database and nondatabase collections.

Both the name and the location of this new library are platform-dependent. On Unix platforms, the library is `$ORACLE_HOME/lib/libtracepls9.so`.

On Win32 platforms (for example, Windows NT), the library is `%ORACLE_HOME%\bin\oratracepls9.dll`.

The `otrace/admin` directory contains two new SQL scripts that can be used to define a database `LIBRARY` object for this library, and to define the procedures that can be used to call out to the library from PL/SQL:

- `OTRCPLSLIB.SQL`
- `OTRCPLSCMD.SQL`

Note: By default, `OTRCPLSLIB.SQL` grants access to the `LIBRARY` to all database users. You can edit this SQL script to restrict access as appropriate.

In addition, the `otrace/demo` directory contains several SQL scripts showing PL/SQL examples that start, stop, and then format an Oracle Trace collection. These are respectively:

- `OTRCPLSSC.SQL`
- `OTRCPLSCC.SQL`
- `OTRCPLSFC.SQL`

In the "start collection" example script `OTRCPLSSC.SQL`, the `regid_list` contains only a single element: "1 192216243 0 0 5 ORCL". The inner double quotes are required to form a single `regid` string from its six components. These components are the following, in the order shown:

- flag (this should be 1 for this usage)
- vendor number (Oracle's vendor number is 192216243)
- `cf_num`
- `cf_val`
- facility number (Oracle server is facility number 5)
- database SID

For an Oracle Trace database collection, a `regid` string like this example is required, basically to identify the database SID and to specify that you are collecting for an Oracle server. The `cf_num` and `cf_val` should be zero in this basic `regid` record.

Additional `regid` records can be specified in order to reduce the amount of collected data. This is when the `cf_num` and `cf_val` items are used. In the Oracle server, Oracle Trace cross facility item 6 (`cf_num = 6`) is reserved to record database

userID values. So, if you provide an additional `regid` record with `cf_num = 6` and a `cf_val = some_DB_userID`, then the collection of database event data is limited to only those events performed by that database user. For example, if you are only interested in collecting database activity for users 23 and 45, then the `regid_list` consists of three records:

```
regid_list  VARCHAR2(256) := '1 192216243 0 0 5 ORCL",
                                     "1 192216243 6 23 5 ORCL",
                                     "1 192216243 6 45 5 ORCL";
```

Note: For better readability, the layout of this `regid` string has been simplified.

Similarly, the `fdf_list` argument could specify the name of a single `.fdf` file (facility definition file). Typically, this is the case. However, more than one `.fdf` could be specified in `fdf_list` if multiple facilities are involved in the collection. Of course, only one `.fdf` can be specified for any given facility; for example, the database.

On the other hand, the process restriction list `prores_list` can be empty. This indicates that there are to be no restrictions on which processes can take part in the collection. If process restrictions are used, then one or more process ID (PID) values can be specified, as well as the operating system username for the owner of each process.

Other arguments in the "start collection" example in `OTRCP_LSSC.SQL` are single numeric or string values, as shown. For example, the collection name and maximum collection data file size specified by the `col_name` and `maxsize` variables, respectively.

Accessing Oracle Trace Collection Results

Running an Oracle Trace collection produces the following collection files:

- `collection_name.cdf` is the binary Oracle Trace collection definition file for the collection. It describes what is to be collected.
- `collection_name.dat` is the output file containing the collected Oracle Trace event data in binary form. When newly created, the empty `.dat` file contains only 35 bytes of file header information.

You can access the Oracle Trace data in the collection files in the following ways:

- The data can be formatted to Oracle tables for SQL access and reporting.
- You can create Oracle Trace reports from the binary file.

Formatting Oracle Trace Data to Oracle Tables

You can format Oracle Trace binary collection data to Oracle database tables, and you can then access this formatted data using SQL or other tools. The Oracle Trace format produces a separate table for each event type collected; for example, a parse event table is created to store data for all database parse events that were recorded during a server collection.

Note: For Oracle server releases 7.3.4 and later, the Oracle Trace formatter automatically creates the formatter tables as needed.

Use the following syntax to format an Oracle Trace collection with the `OTRCFMT` formatter utility:

```
OTRCFMT [options] collection_name.cdf [user/password@database]
```

If collection `.cdf` and `.dat` are not located in the current default directory, then specify the full file path for the `.cdf` file.

If you omit `user/password@database` (or any part of it, such as password or database), then `OTRCFMT` prompts you for this information.

Oracle Trace allows data to be formatted while a collection is occurring. By default, Oracle Trace formats only the portion of the collection that has not been formatted previously. If you want to reformat the entire collection file, then use the optional parameter `-f` (which generate a new collection ID in the formatter tables).

Oracle Trace provides several sample SQL scripts that you can use to access the formatted server event data tables. These are located in `OTRCRPT*.SQL` in the `otrace` directory tree.

Note: Because there are incompatibilities between at least some versions of the formatter tables, use a separate database schema for each version of the Oracle Trace formatter.

Running the Oracle Trace Reporting Utility

The Oracle Trace reporting utility displays data for items associated with each occurrence of a server event. These reports can be quite large. You can control the report output by using optional statement qualifiers. Use the following report utility statement syntax:

```
OTRCREP [options] collection name.cdf
```

If collection `.cdf` and `.dat` are not located in the current default directory, then specify the full file path for the `.cdf` file.

First, you might want to run a report called `PROCESS.txt`. You can produce this report to provide a listing of specific process identifiers for which you want to run another report.

You can manipulate the output of the Oracle Trace reporting utility by using the following optional report qualifiers:

Qualifier	Description
<code>output_path</code>	Specifies a full output path for the report files. If this path is not specified, then the files are placed in the current directory.
<code>-p[pid]</code>	Organizes event data by process. If you specify a process ID (PID), then you have one file with all the events generated by that process in chronological order. If you omit the PID, then you have one file for each process that participated in the collection. The output files are named <code>collection_name_ppid.txt</code> .
<code>-P</code>	Produces a report file name <code>PROCESS.txt</code> that lists all processes that participated in the collection. It does not include event data. You can produce this report first to determine the specific processes for which you want to produce more detailed reports.
<code>-w#</code>	Sets report width, such as <code>-w132</code> . The default is 80 characters.
<code>-l#</code>	Sets the number of report lines for each page. The default is 63 lines for each page.
<code>-h</code>	Suppresses all event and item report headers, producing a shorter report.
<code>-a</code>	Creates a report containing all the events for all products, in the order they occur in the data collection (<code>.dat</code>) file.

Default `OTRCREP` report output, with no optional qualifiers specified, consists of one text file for each event type collected. Data from all participating processes are combined in each of these text files.

Oracle Server Events

The following sections describe events that have been instrumented in Oracle Server. Most of the events are useful for performance analysis and tuning and workload analysis by Oracle Expert.

There are two types of events: **point** events and **duration** events. Point events represent an instantaneous occurrence of something in the instrumented product. An example of a point event is an error occurrence. Duration events have a beginning and ending. An example of a duration event is a transaction. Duration events can have other events occur within them; for example, the occurrence of an error within a transaction.

[Table 12-2](#) lists the Oracle Server events instrumented for Oracle Trace. For more detailed descriptions, refer to the section for the event in which you are interested.

Table 12-2 Oracle Server Events

Event	Description	Type of Event
1 Connection	Connection to a database.	Point
2 Disconnect	Disconnection from a database.	Point
3 ErrorStack	Code stack for core dump.	Point
4 Migration	Session migration between shared server processes.	Point
5 ApplReg	Application context information.	Point
6 RowSource	Row information. For Oracle Server release 8.0.2 and higher, this also includes data about the execution plan.	Point
7 SQLSegment	Text of SQL statement.	Point
8 Parse	SQL parsing information.	Duration
9 Execute	Information for execution of SQL.	Duration
10 Fetch	Actual row retrieval information.	Duration
11 LogicalTX	The first time a database command is performed that may change the database status.	Duration
12 PhysicalTX	Event marking a definite change in database status.	Duration
13 Wait	Generic WAIT event. Context is provided in the event strings.	Point

Data Items Collected for Events

Specific kinds of information, known as items, are associated with each event. There are three types of items:

- [Resource Utilization Items](#)
- [Cross-Product Items](#)
- [Items Specific to Oracle Server Events](#)

Resource Utilization Items

Oracle Trace has a standard set of items, called resource utilization items, that it can collect for any instrumented application, including the Oracle Server. In addition, all duration events in the Oracle Server include items for database statistics specific to the Oracle Server.

The standard resource utilization items are described in [Table 12-3](#).

An Oracle Trace collection can be formatted into Oracle tables for access, analysis, and reporting. The last column contains the data type for data items formatted to the Oracle database.

Table 12-3 *Standard Resource Utilization Items*

Item Name	Description	Item ID	Datatype of Formatted Data
UCPU	Amount of CPU time in user mode	129	number
SCPU	Amount of CPU time in system mode	130	number
INPUT_IO	Number of times file system performed input	131	number
OUTPUT_IO	Number of times file system performed output	132	number
PAGEFAULTS	Number of hard and soft page faults	133	number
PAGEFAULT_IO	Number of hard page faults	134	number
MAXRS_SIZE	Maximum resident set size used	135	number

The implementation of the item is platform specific; if the item is not implemented, the value is 0. For example, currently only CPU times are recorded on Windows NT.

Cross-Product Items

Oracle Trace provides a set of 14 items called cross-product items (also known as cross-facility items for historical reasons). These data items allow programmers to relate events for different products. For example, a business transaction may generate events in two products: an application and the database. The cross-product data items allow these disparate events to be joined for analysis of the entire business transaction.

Cross-product items are reserved for specific products or product types as described in [Table 12-4](#). If you do not use the products for which items are reserved, then you can use those items for your own purposes.

Table 12-4 Cross-Product Items

Item Name	Layer	Description	Item ID	Datatype of Formatted Data
CROSS_FAC_1	Application	Application ID. For use by high-level applications such as Oracle Financials, third-party or customer applications	136	number
CROSS_FAC_2	Oracle Forms	Oracle Forms ID	137	number
CROSS_FAC_3	Oracle Net	Remote node connection ID	138	number
CROSS_FAC_4	Oracle Server	Transaction ID	139	number
CROSS_FAC_5	Oracle Server	Hash_ID of SQL statement	140	number
CROSS_FAC_6	Oracle Server release 8.x	User ID	141	number
CROSS_FAC_7	Oracle Server release 8.x	Wait type	142	number
CROSS_FAC_8	n/a	Not reserved	143	number
CROSS_FAC_9	n/a	Not reserved	144	number
CROSS_FAC_10	n/a	Not reserved	145	number
CROSS_FAC_11	n/a	Not reserved	146	number
CROSS_FAC_12	n/a	Not reserved	147	number
CROSS_FAC_13	n/a	Not reserved	148	number
CROSS_FAC_14	n/a	Not reserved	149	number

Note: In this version of Oracle Trace, the term “facility” has been changed to “product”. Therefore, the items named `CROSS_FAC_x` are cross-product items.

Cross-product item 1 (referred to as `CROSS_FAC_1`) contains data only if data is supplied by an instrumented application.

Cross-product item 2 (`CROSS_FAC_2`) is reserved for use by a future release of Oracle Forms. Instrumented applications and Oracle Forms pass identification data to the Oracle Server collection through these cross-product items.

Cross-product item 3 (`CROSS_FAC_3`) is reserved for use by Oracle Net. Oracle Net supplies the connection ID to Oracle Trace through `CROSS_FAC_3`. `CROSS_FAC_3` is the key element in coordinating client/server or multitier Oracle Trace collections. Oracle Trace uses the Oracle Net global connection ID as the common element to match in the merger, for example the client and server collection files. The global connection ID is the same for the client and the server connection.

Most Oracle Server events record cross-product items 1 through 6. (Cache I/O does not.)

Items Specific to Oracle Server Events

The Oracle Server product (or facility) definition files (that is, *.`fdf`) defines items specific to the Oracle Server. Use the item’s number to locate it within the list. The formatted datatype describes how the Oracle Trace formatter defines the item when it formats data into an Oracle database.

The Oracle Server items are listed in [Table 12-5](#).

Table 12-5 Oracle Server Items

Item Name	Description	Item Number	Formatted Datatype
<code>App_Action</code>	Action name set by using the <code>DBMS_APPLICATION_INFO.SET_MODULE</code> procedure	23	<code>varchar2</code> (1020)
<code>App_Module</code>	Module name set using the <code>DBMS_APPLICATION_INFO.SET_MODULE</code> procedure	22	<code>varchar2</code> (1020)
<code>Commit_Abort</code>	Indicates if a transaction committed or aborted	24	<code>number</code>

Table 12–5 Oracle Server Items (Cont.)

Item Name	Description	Item Number	Formatted Datatype
Consistent_Gets	Number of blocks retrieved in consistent mode (did not change the data and therefore did not create any locks or conflicts with other users)	104	number
CPU_Session	CPU session	112	number
Current_UID	Current user ID	36	number
Cursor_Number	Number of cursor associated with SQL statement	25	number
DB_Block_Change	Number of blocks changed	102	number
DB_Block_Gets	Number of blocks retrieved in current mode. For large queries, this item tells how many sections of the database (logical pages) were fetched to retrieve all needed records.	103	number
Deferred_Logging	Value used by Oracle Trace internally	14	number
Depth	Recursive level at which SQL statement is processed	32	number
Description	Depends upon event in which it occurs (for example, wait event description)	43	varchar2 (1020)
Elapsed_Session	Elapsed time for the session	113	number
End_of_Fetch	Flag set if data retrieved is last data from query	38	number
Lib_Cache_Addr	Address of SQL statement in library cache	27	varchar2(64)
Login_UID	Internal ID within the Oracle database that identifies the user ID for the session	15	number
Login_UserName	Internal ID within the Oracle database that identifies the system account name for the session	16	varchar2 (1020)
Missed	Flag set if SQL statement was missing in library cache	33	number

Table 12–5 Oracle Server Items (Cont.)

Item Name	Description	Item Number	Formatted Datatype
Object_ID ¹	Object ID of the row source	46	number
Operation ¹	Text of the operation	47	varchar2 (1020)
Operation_ID ¹	Position of the operation within the execution plan for a statement	28	number
Optimizer_Mode	Oracle optimizer mode	35	varchar2(128)
Oracle_Cmd_Type	Oracle command number	34	number
Oracle_PID	Oracle process ID	11	number
OS_Image	Operating system image (program name)	42	long
OS_Mach	Operating system host machine	20	varchar2 (1020)
OS_Term	Operating system terminal	19	varchar2 (1020)
OS_UName	Operating system username	18	varchar2 (1020)
P1	The definition of P1 depends upon the event in which it occurs.	1	number
P2	The definition of P2 depends upon the event in which it occurs.	2	number
P3	The definition of P3 depends upon the event in which it occurs.	3	number
P4	The definition of P4 depends upon the event in which it occurs.	4	number
P5	The definition of P5 depends upon the event in which it occurs.	5	number
P6	The definition of P6 depends upon the event in which it occurs.	6	number
P7	The definition of P7 depends upon the event in which it occurs.	7	number

Table 12–5 Oracle Server Items (Cont.)

Item Name	Description	Item Number	Formatted Datatype
P8	The definition of P8 depends upon the event in which it occurs.	8	number
P9	The definition of P9 depends upon the event in which it occurs.	9	number
P10	The definition of P10 depends upon the event in which it occurs.	10	number
Parent_Op_ID ¹	Parent operation	44	number
PGA_Memory	Process Global Area memory	101	number
Physical Reads	Number of blocks read from disk	105	number
Position ¹	Position within events having same parent operation	45	number
Position_ID ²	Position of the operation within the execution plan for a statement	28	number
Redo_Entries	Number of redo entries made by process	106	number
Redo_Size	Size of redo entries	107	number
Row_Count	Number of rows processed	29	number
Schema_UID	Schema user ID	37	number
Session_Index	Oracle session ID	12	number
Session_Serial	Session serial number	13	number
SID	Text version of session ID	17	varchar2 (1020)
Sort_Disk	Number of disk sorts performed	110	number
Sort_Memory	Number of memory sorts performed	109	number
Sort_Rows	Total number of rows sorted	111	number
SQL_Text	Text of SQL statement	31	long
SQL_Text_Hash	Pointer to SQL statement	26	number
SQL_Text_Segment	Address of SQL text	30	number

Table 12–5 Oracle Server Items (Cont.)

Item Name	Description	Item Number	Formatted Datatype
T_Scan_Rows_Got	Rows processed during full table scans	108	number
TX_ID	Unique identifier for a transaction that consists of rollback segment number, slot number, and wrap number	41	varchar2(72)
TX_SO_Addr	The address of the transaction state object	40	varchar2(64)
TX_Type	Type of the transaction. Value is a bitmap (for example, 2 active transaction, 0X10 space transaction, 0X20 recursive transaction).	39	number
UGA_Memory	User Global Area session memory	100	number
Wait_Time	Elapsed time, in hundredths of seconds, for the wait event	21	number

¹ Item specific to Oracle Server release 8.0.2 and higher

² Replaced by `Operation_ID` for Oracle Server release 8.0.2 and higher

Items Associated with Each Event

The following sections describe each event in more detail and provide tables that list the items associated with each event. For item descriptions, refer to [Table 12–5](#).

Note: Prior to Oracle Server release 8.0.5, cross-product items 1-5 were set by the server code. Starting with Oracle Server release 8.0.5, cross-product item 6 was added (and cross-product item 7 for wait events.)

When you format data, Oracle Trace creates a table for each event type collected. The name of the event data table is **V_vendor#_F_product#_E_event#_version**, where **version** is the number of the Oracle Server release. Any periods in the product version are replaced with underscores. You can use the `otrcsyn.sql` script to create synonyms for these tables.

Note: The following tables use Oracle7 Server names for example purposes.

The Oracle Trace formatter creates a column for each event item. For point events, the column name is the same as the item name. For duration events, the items for the start event have `_START` appended to the item name and the items for the end event have `_END` appended to the item name.

The formatter automatically includes additional columns for collection number, process identifier, and timestamp information as described in [Table 12-6](#).

Table 12-6 Additional Columns Included by Oracle Trace Formatter

Column Name	Description	Datatype
<code>collection_number</code>	collection ID, automatically assigned by the formatter	number(4)
<code>epid</code>	process ID	number(8)
<code>timestamp</code>	logged time for point events	date
<code>timestamp_nano</code>	fraction of seconds of logged time for point events	number
<code>timestamp_start</code>	duration event start time	date
<code>timestamp_nano_start</code>	fraction of seconds of duration event start time	number
<code>timestamp_end</code>	duration event end time	date
<code>timestamp_nano_end</code>	fraction of seconds of duration event end time	number

Event Statistics Block

Items relating to database performance appear in several events. For convenience, these items are referenced as the Event Statistics Block. The items in the Event Statistics block are shown in the following list:

```

UGA_Memory
PGA_Memory
DB_Block_Change
DB_Block_Gets
Consistent_Gets
Physical_Reads
Redo_Entries
Redo_Size
T_Scan_Rows_Got
Sort_Memory
Sort_Disk
Sort_Rows
CPU_Session
Elapsed_Session

```

Connection Event

The Connection event (`event=1`) records every time a connection is made to a database. Items associated with the Connection event are shown in the following list:

```
Session_Index
Session_Serial
Oracle_PID
Login_UID
Login_UName
SID
OS_UName
OS_Term
OS_Mach
OS_Image
Cross-Product Items 1-6
```

The Oracle Server uses the combination of `Session_Index` and `Session_Serial` to uniquely identify a connection. Oracle Net uses the connection ID, stored in `CROSS_FAC_3`, to uniquely identify a connection.

Disconnect Event

The Disconnect event records every time a database disconnection is made. Items associated with the Disconnect event are shown in the following list.

```
Session_Index
Session_Serial
Event Statistics Block
Oracle_PID
Cross-Product Items 1-6
```

A Disconnect event corresponds to at most one Connection event. Therefore, the same fields uniquely identify a disconnect: either the combination of `Session_Index` and `Session_Serial`, or `CROSS_FAC_3`.

ErrorStack Event

The ErrorStack event identifies the process that has the error. Items associated with the ErrorStack event are shown in the following list.

```
Session_Index
Session_Serial
Oracle_PID
P1
```


P2
P3
P4
P5
P6
P7
P8

Cross-Product Items 1-6

The `ErrorStack` event does not have an explicit identifier. The combination of `Session_Index`, `Session_Serial`, `Timestamp`, and `Timestamp_Nano` should uniquely identify a specific `ErrorStack` event.

Migration Event

The `Migration` event is logged each time a session migrates to a shared server process. This event is currently disabled in the Oracle server code.

Items associated with the `Migration` event are shown in the following list:

`Session_Index`
`Session_Serial`
`Oracle_PID`
Cross-Product Items 1-6

The `Migration` event does not have an explicit identifier. The combination of `Session_Index`, `Session_Serial`, `Timestamp`, and `Timestamp_Nano` should uniquely identify a specific `Migration` event.

ApplReg Event

The `ApplReg` event (event=5) registers with Oracle Trace where the application is at a certain point in time. Items associated with the `ApplReg` event are shown in the following list:

`Session_Index`
`Session_Serial`
`App_Module`
`App_Action`
Cross-Product Items 1-6

The `ApplReg` event does not have an explicit identifier. The combination of `Session_Index`, `Session_Serial`, `Timestamp`, and `Timestamp_Nano` should uniquely identify a specific `ApplReg` event.

RowSource Event

The RowSource event logs the number of rows processed by a single row source within an execution plan. Items associated with the RowSource event are shown in the following list:

Session_Index
Session_Serial
Cursor_Number
Position_ID
Row_Count
Cross-Product Items 1-5

The combination of Session_Index, Session_Serial, Cursor_Number, and Position_ID uniquely identifies a RowSource event.

RowSource Event

The RowSource event logs the number of rows processed by a single row source within an execution plan. Items associated with the RowSource event for Oracle Server release 8.0.2 or higher are shown in the following list:

Session_Index
Session_Serial
Cursor_Number
Operation_ID
Row_Count
Parent_Op_ID
Position
Object_ID
Operation
Cross-Product Items 1-6

The combination of Session_Index, Session_Serial, Cursor_Number, and Operation_ID uniquely identifies a RowSource event.

Note: The text in the Operation item is equivalent to information about the execution plan, which is similar to data that can be obtained by running explain plan.

SQLSegment Event

The SQLSegment event is a description of a SQL statement. Items associated with the SQLSegment event are shown in the following list:

Session_Index
Session_Serial
Cursor_Number
SQL_Text_Hash
Lib_Cache_Addr
SQL_Text_Segment
SQL_Text
Cross-Product Items 1-6

A SQL segment does not have an explicit identifier. The `SQL_Text_Hash` field is always the same for each occurrence of a SQL statement but multiple statements can have the same hash value. If a statement is forced out of the library cache and then swapped back in, the same statement can have multiple values for `Lib_Cache_Addr`. The combination of `Session_Index`, `Session_Serial`, `SQL_Text_Hash`, and `Lib_Cache_Addr` usually should identify a particular SQL statement for a session. If you add `Cursor_Number`, you identify a particular occurrence of a SQL statement within the session.

Wait Event

The wait event shows the total waiting time in hundredths of seconds for all responses. Items associated with the wait event are shown in the following list:

Session_Index
Session_Serial
P1
P2
P3
Description
Cross-Product Items 1-7

The wait event does not have an explicit identifier. The combination of `Session_Index`, `Session_Serial`, `Description`, `Timestamp`, and `Timestamp_Nano` should uniquely identify a specific wait event.

Parse Event

The Parse event records the start and end of the parsing phase during the processing of a SQL statement. The parsing phase occurs when the SQL text is read in and broken down (parsed) into its various components. Tables and fields are identified, as well as which fields are sort criteria and which information needs to be returned. Items associated with the parse event are shown in the following lists:

Items for Start of Parse Event	Items for End of Parse Event
Session_Index	Session_Index
Session_Serial	Session_Serial
Event Statistics Block	Event Statistics Block
Cursor_Number	Cursor_Number
Resource Items	Depth
Cross-Product Items 1-6	Missed
	Oracle_Cmd_Type
	Optimizer_Mode
	Current_UID
	Schema_UID
	SQL_Text_Hash
	Lib_Cache_Addr
	Resource Items

The combination of `Session_Index`, `Session_Serial`, `Cursor_Number`, and `SQL_Text_Hash` uniquely identifies a specific parse event.

Execute Event

The Execute event is where the query plan is executed. That is, the parsed input is analyzed to determine exact access methods for retrieving the data, and the data is prepared for fetch if necessary. Items associated with the Execute event are shown in the following lists:

Items for Start of Execute Event	Items for End of Execute Event
Session_Index	Session_Index
Session_Serial	Session_Serial
Event Statistics Block	Event Statistics Block
Cursor_Number	Cursor_Number
Resource Items	Depth
	Missed

Items for Start of Execute Event	Items for End of Execute Event
	Row_Count
	SQL_Text_Hash
	Lib_Cache_Addr
	Resource Items

The combination of `Session_Index`, `Session_Serial`, `Cursor_Number`, and `SQL_Text_Hash` uniquely identifies a specific Execute event.

Fetch Event

The Fetch event is the actual return of the data. Multiple fetches can be performed for the same statement to retrieve all the data. Items associated with the Fetch event are shown in the following lists:

Items for Start of Fetch Event	Items for End of Fetch Event
<code>Session_Index</code>	<code>Session_Index</code>
<code>Session_Serial</code>	<code>Session_Serial</code>
Event Statistics Block	Event Statistics Block
<code>Cursor_Number</code>	<code>Cursor_Number</code>
Resource Items	Depth
Cross-Product Items 1-6	Row_Count
	End_of_Fetch
	SQL_Text_Hash
	Lib_Cache_Addr
	Resource Items

The combination of `Session_Index`, `Session_Serial`, `Cursor_Number`, `SQL_Text_Hash`, `Timestamp`, and `Timestamp_Nano` uniquely identifies a specific Fetch event.

LogicalTX Event

The LogicalTX event logs the start and end of a logical transaction (that is, statements issued that may cause a change to the database status). Items associated with the LogicalTX event are shown in the following lists:

Items for Start of LogicalTX Event	Items for End of LogicalTX Event
Session_Index	Session_Index
Session_Serial	Session_Serial
Event Statistics Block	Event Statistics Block
TX_Type	TX_Type
TX_SO_Addr	TX_SO_Addr
Resource Items	Resource Items
Cross-Product Items 1-6	

The transaction identifier stored in `CROSS_FAC_4` should uniquely identify a specific transaction. Or, use `Session_Index`, `Session_Serial`, and `TX_SO_Addr`.

PhysicalTX Event

The PhysicalTX event logs the start and end of a physical transaction (that is, statements issued that caused a change in database status). Items associated with the PhysicalTX event are shown in the following lists:

Items for Start of PhysicalTX Event	Items for End of PhysicalTX Event
Session_Index	Session_Index
Session_Serial	Session_Serial
Event Statistics Block	Event Statistics Block
TX_Type	TX_Type
TX_ID	TX_ID
Resource Items	Commit_Abort
Cross-Product Items 1-6	Resource Items

The transaction identifier stored in `CROSS_FAC_4` should uniquely identify a specific transaction.

Event Set File Names

Oracle Trace events can be organized into **event sets** that restrict the data collection to specific events. You can establish event sets for performance monitoring, auditing, diagnostics, or any logical event grouping.

Table 12–7 Server Event Set File Names

Event Set File Name (.fdf)	Description
CONNECT	CONNECT_DISCONNECT event set. Collects statistics about connects to the database and disconnects from the database.
ORACLE	ALL event set. Collects all statistics for the Oracle Server including wait events.
ORACLEC	CACHEIO event set. Collects caching statistics for buffer cache I/O.
ORACLEDD	Oracle Server DEFAULT event set. Collects statistics for the Oracle Server.
ORACLEEE	EXPERT event set. Collects statistics for the Oracle Expert application.
ORACLESM	SUMMARY event set. Collects workload statistics for the Summary Advisor application.
SQL_ONLY	SQL_TEXT_ONLY event set. Collects statistics about connects to the database, disconnects from the database, and SQL text.
SQL_PLAN	SQL_STATS_AND_PLAN event set. Collect statistics about connects to the database, disconnects from the database, SQL statistics, SQL text, and row source (EXPLAIN PLAN).
SQLSTATS	SQL_AND_STATS event set. Collects SQL text and statistics only.

Table 12–7 Server Event Set File Names (Cont.)

Event Set File Name (.fdf)	Description
SQL_TXN	SQL_TXNS_AND_STATS event set. Collects statistics about connects to the database, disconnects from the database, transactions, SQL text and statistics, and row source (EXPLAIN PLAN).
SQLWAITS	SQL_AND_WAIT_STATS event set. Collects statistics about connects to the database, disconnects from the database, row source (EXPLAIN PLAN), SQL text and statistics, and wait events.
WAITS	WAIT_EVENTS event set. Collects statistics about connects to the database, disconnects from the database, and wait events.

See Also: ["Event Sets"](#) on page 12-2

Troubleshooting Oracle Trace

Use the following sections to troubleshoot problems while using Oracle Trace.

Oracle Trace Configuration

If you suspect an Oracle Trace configuration problem:

- Examine the `EPC_ERROR.LOG` file for details of any logged Oracle Trace errors.
- Look for the administration files on `$ORACLE_HOME/otrace/admin` (*.dat files) and run `otrccref` to re-create the Oracle Trace *.dat files if the files do not exist.
- Verify that the .fdf files are in the `$ORACLE_HOME/otrace/admin/fdf` directory.
- Verify the correct version of Oracle Trace Collection Services match the appropriate Oracle Server version
 - ⌘ `$ORACLE_HOME/bin/otrccol version`

If the Returned Value Is: Then the Command-Line Interface Release Is:

1	733
2	803
3	734
4	804
5	805
6	813
7	814
8	815
9	806
10	816
11	817
12	901

- To verify that a collection is currently running, use the command-line interface to check the status:

```
% $ORACLE_HOME/bin/otrccol check collection_name
```

To test the CLI:

1. CLI needs to run from a privileged account, for example, the Oracle operating system user account.
2. The Oracle home and SID environmental variables must be set properly.

To check settings on UNIX:

```
printenv ORACLE_HOME
printenv ORACLE_SID
```

To set settings on UNIX:

```
setenv ORACLE_HOME path
setenv ORACLE_SID sid
```

There should be one CLI for each ORACLE_HOME. For example, if you have two Oracle Server release 7.3.3 instances sharing the same ORACLE_HOME, there should be only one CLI.

3. Verify that the collection name has not already been used before you start the collection.

Look for *collection name.cdf* and *.dat* files in:

- `$ORACLE_HOME/otrace/admin/cdf` directory
 - The directory specified in database parameter `ORACLE_TRACE_COLLECTION_PATH`
 - The directory specified by `EPC_COLLECTION_PATH` environment variable
4. If you want to generate database activity for this collection, connect to the database.
 - For Oracle Server release 7.3.x, connect to the service before you create your collection.
 - For Oracle Server release 8.0, you can connect to the database anytime and the processes are registered.

Server Environment

If you suspect a server environment problem, verify the following:

- The server node has sufficient disk space for the collection output files. If there is not sufficient disk space, the collection stops. To limit the size of collections, use the Trace option for limiting collection size. For a description on how to limit collection size, see the Using the Oracle Trace CLI section in this chapter.

To solve the immediate problem, stop the collection, and free up space so Oracle Trace can end the collection.

Initially limiting the collection to specific users or wait events also helps to limit the amount of data collected. Limiting users and wait events is available for Oracle Server releases 8.0.4 and higher.
- Your session does not participate in more than five collections. Sessions log data for the five most recent collections. Thus, if you have more than five collections, data is missing for the oldest collection.

Missing Data

Wait Times Were Not Collected Wait times are collected only if the `INITsid.ORA` parameter, `TIMED_STATISTICS`, is set to `true`.

Note: Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER_SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

See Also: ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Missing SQL Statement from Collection If an expected SQL statement does not appear to be in your collection, it may be because a small amount of data in the Oracle Trace data collection buffers may not have been flushed out to the collection data file, even though the collection has been stopped. Additional database activity should flush these buffers to disk, and shutting down the database also forces a flush of these buffers.

Collection Is Too Large There may be times when a collection is too large. Starting with Oracle Server release 8.0.4, you can collect data for specific users and specific wait event types to minimize the size of the collection. Because, almost always, the server is waiting for a latch, lock, or resource, wait event data for a brief collection can be quite extensive.

Collection Is Empty In Oracle8 databases (prior to Oracle 8.1.7), the `ORACLE_TRACE_ENABLE` parameter in the `INITsid.ORA` file on the server must be set to `true` before the database is started. Starting with Oracle 8.1.7 it is dynamic and may be modified through `ALTER_SESSION` or `ALTER_SYSTEM`. (For Oracle7 the `ORACLE_TRACE_ENABLE` parameter should be left as `false`, unless you are using the `init.ora` parameter method to start or stop collections.) You can also see this problem if there are too many collections running concurrently.

Oracle Trace Could Not Access Memory

On Windows NT systems, if you are running Oracle Trace collections and an error occurs indicating Oracle Trace could not access memory, the `collect.dat` file has become full. You must create a new `.dat` file by running the `otrccref.exe` image

located in the `$ORACLE_HOME/bin` directory. However, database services must be shutdown to release the `collect.dat` file for the `otrccref` script to be able to create the new `collect.dat` file. You can also increase the number of `collect.dat` records to more than the default of 36 records (for example, `otrccref -c50` to create a new `otrace/admin/collect.dat` file containing 50 records).

Oracle7 Stored Procedures

If the attempt to collect Oracle Trace data for an Oracle7 database results in the message "Error starting/stopping Oracle7 database collection," this may be due to missing database stored procedures that Oracle Trace uses to start and stop Oracle7 collections.

- Check for stored procedures (for Oracle Server releases 7.3.x)

To check for stored procedures using the Oracle Enterprise Manager console, use the Navigator and the following path:

```
Networks > Databases > your_database > Schema Objects > Packages > SYS
```

- Look for stored procedures starting with `DBMS_ORACLE_TRACE_XXX`.

To check for stored procedures using Oracle Server Manager or Oracle SQL*Plus Worksheet:

```
select object_name from dba_objects where object_name like '%TRACE%'
and object_type = 'PACKAGE';
OBJECT_NAME
DBMS_ORACLE_TRACE_AGENT
DBMS_ORACLE_TRACE_USER
2 rows selected.
```

For Oracle7, Oracle Trace required that these stored procedures be installed on the database. These SQL scripts may be automatically run during database installation depending on the platform-specific installation procedures. If they are not executed during database installation, you must run these scripts manually. You can add these stored procedures to the database by running the `otrcsvr.sql` script from `$ORACLE_HOME/otrace/admin`) from a privileged database account (`SYS` or `INTERNAL`). To run the script, set the default to the path where the script is located. This script runs other scripts that do not have the path specified. These other scripts fail if you are not in the directory where these scripts will run.

EPC_ERROR.LOG File

The EPC_ERROR.LOG file provides information about the collection processing, specifically the Oracle Trace Collection Services errors.

The EPC_ERROR.LOG file is created in the current default directory.

See Also: *Oracle Enterprise Manager Messages Manual* for general information about causes and actions for most Oracle Trace messages

Formatter Tables

Oracle Server releases 7.3.4 and 8.0.4 and later automatically create the formatter tables. Prior to Oracle Server releases 7.3.4 and 8.0.4, you must run the `otrcfmtc.sql` script from Oracle Server Manager or Oracle SQL*Plus Worksheet as the user who will be formatting the data. If you must manually execute `otrcfmtc.sql` to create the formatter tables, use the SQL script from the same Oracle home as your collections to be formatted.

The `otrcfmtc.sql` script is located in the `$ORACLE_HOME/otrace/admin` directory.

Formatting error might be due to one of the following causes:

1. The user did not run the script to create the formatter tables (valid for releases of Oracle Server prior to 7.3.4 and 8.0.4).
2. The formatter tables were not created by the same user ID that was used when the collection was created (valid for releases of Oracle Server prior to 7.3.4 and 8.0.4).

Look for EPC_COLLECTION.

To check for formatter tables using SQL Worksheet:

```
CONNECT username/password@service name
DESCRIBE epc_collection
```


Part III

Creating a Database for Good Performance

Part III describes how to configure a database for good performance.

The chapters in this part are:

- [Chapter 13, "Building a Database for Performance"](#)
- [Chapter 14, "Memory Configuration and Use"](#)
- [Chapter 15, "I/O Configuration and Design"](#)
- [Chapter 16, "Understanding Operating System Resources"](#)
- [Chapter 17, "Configuring Instance Recovery Performance"](#)
- [Chapter 18, "Configuring Undo and Temporary Segments"](#)
- [Chapter 19, "Configuring Shared Servers"](#)

Building a Database for Performance

Although performance modifications can be made to both the database and to the Oracle instance at a later time, much can be gained by carefully designing the database for the intended needs.

This chapter contains the following sections:

- [Initial Database Creation](#)
- [Creating Tables for Good Performance](#)
- [Loading and Indexing Data](#)
- [Initial Instance Configuration](#)
- [Setting up Operating System, Database, and Network Monitoring](#)

Note: This chapter is an overview of Oracle's new methodology for designing a database for performance. Before reading this chapter, it is important to read the information in the *Oracle9i Database Performance Planning* manual. For detailed information on memory and I/O configuration, see the other chapters in Part III.

Initial Database Creation

One of the first stages in managing a database is the initial database creation. This section describes important steps in the creation of an Oracle database.

Database Creation Using the Installer

The Oracle Installer lets you create a database during software installation or at a later time using the Database Creation Assistant. This is an efficient way of creating databases for small to medium size environments, and it provides a straightforward graphical user interface. However, this procedure sets some limits on the possibilities for various options, and it is therefore not recommended for database creation in larger environments.

Manual Database Creation

A manual approach provides full freedom for different configuration options. This is especially important for larger environments. A manual approach typically involves designing multiple parameter files for use during the various stages, running SQL scripts for the initial `CREATE DATABASE` and subsequent `CREATE TABLESPACE` statements, running necessary data dictionary scripts, and so on.

Parameters Necessary for Initial Database Creation

The initialization parameter file is read whenever an Oracle instance is started, including the very first start before the database is created. There are very few parameters that cannot be modified at a later time. The most important parameters to set correctly at database creation time are listed in [Table 13-1](#).

Table 13-1 Important Initialization Parameters for Database Creation

Parameter	Description
<code>DB_BLOCK_SIZE</code>	Sets the size of the Oracle database blocks stored in the database files and cached in the SGA. The range of values depends on the operating system, but it is typically powers of two in the range 2048 to 16384. Common values are 4096 or 8192 for transaction processing systems and higher values for database warehouse systems.

Table 13–1 (Cont.) Important Initialization Parameters for Database Creation

Parameter	Description
DB_NAME and DB_DOMAIN	Set the name of the database and the domain name of the database, respectively. Although they can be changed at a later time, it is highly advisable to set these correctly before the initial creation. The names chosen must be reflected in the SQL*Net configuration as well.
COMPATIBLE	Specifies the release with which the Oracle server must maintain compatibility. It lets you take advantage of the maintenance improvements of a new release immediately in your production systems without testing the new functionality in your environment. If your application was designed for a specific release of Oracle, and you are actually installing a later release, then you might want to set this parameter to the version of the previous release.

Note: Compatibility must be set to Release 2 (9.2) or higher to have a locally managed SYSTEM tablespace. If EXTENT MANAGEMENT LOCAL is specified at CREATE DATABASE time, the database cannot be migrated back to a pre-9.2 release and the SYSTEM tablespace cannot be migrated to dictionary-managed. All tablespaces subsequently created must be locally managed.

The CREATE DATABASE Statement

The first SQL statement that is executed after startup of the initial instance is the CREATE DATABASE statement. This creates the initial system tablespace, creates the initial redo logfiles, and sets certain database options. The options listed in [Table 13–2](#) cannot be changed or can only be changed with difficulty at a later time.

Table 13–2 Database Options for Initial Creation

Database Options	Description
Character set	The character set specified by this option identifies the character set used for SQL text, for the internal data dictionary, and most importantly for text stored as datatypes CHAR, VARCHAR, or VARCHAR2. After data including any type of national characters has been loaded, the character set cannot be changed.

Table 13–2 (Cont.) Database Options for Initial Creation

Database Options	Description
National character set	This character set is used for the datatypes NCHAR, NVARCHAR, and NVARCHAR2. In general, as with the regular character set, it cannot be changed.
SQL .BSQ file	This creates the internal data dictionary. For information on modifying this file, see Chapter 15, "I/O Configuration and Design" .
Location of initial datafile	The initial datafile(s) that will make up the system tablespace should be set with care. They can be modified later, but the procedure involves a complete shutdown of the instance.
EXTENT MANAGEMENT LOCAL	Use the EXTENT MANAGEMENT LOCAL clause to create a locally managed SYSTEM tablespace. AUTOALLOCATE is the default; extent sizes are chosen by the system. EXTENT SIZE UNIFORM is not available for the SYSTEM tablespace.
Default temporary tablespace	If you specify EXTENT MANAGEMENT LOCAL, then you must also specify the default temporary tablespace.
MAXDATAFILES	The maximum number of data files.
MAXLOGFILES	The maximum number of log files.

Note: Compatibility must be set to Release 2 (9.2) or higher to have a locally managed SYSTEM tablespace. If EXTENT MANAGEMENT LOCAL is specified at CREATE DATABASE time, the database cannot be migrated back to a pre-9.2 release and the SYSTEM tablespace cannot be migrated to dictionary-managed. All tablespaces subsequently created must be locally managed.

Example 13–1 Sample CREATE DATABASE Script

```
CONNECT SYS/ORACLE AS SYSDBA
STARTUP NOMOUNT pfile=/u01/admin/init_create.ora
CREATE DATABASE "dbname"
DATAFILE '/u01/oradata/system01.dbf' size 200M
LOGFILE '/u02/oradata/redo01.dbf' size 100M,
        '/u02/oradata/redo02.dbf' size 100M
CHARACTER SET "WE8ISO8859P1"
NATIONAL CHARACTER SET "UTF8"
EXTENT MANAGEMENT LOCAL
DEFAULT TEMPORARY TABLESPACE mytemp TEMPFILE 'temp.f' SIZE 1000M
```

```
MAXDATAFILES = 50  
MAXLOGFILES = 5;
```

See Also: *Oracle9i SQL Reference* for detailed information about the CREATE DATABASE statement

Running Data Dictionary Scripts

After running the CREATE DATABASE statement, certain catalog scripts must be executed. They are found in the rdbms/admin directory on UNIX or the rdbms\admin directory on Windows, under the ORACLE_HOME directory. The following scripts must be executed:

- CATALOG.SQL - Needed for all normal data dictionary views
- CATPROC.SQL - Needed to load the initial PL/SQL environment

Note: When specific options or features are in use (for example, Java or replication), more scripts are necessary. These are documented with each individual option.

Example 13–2 Executing Required Data Dictionary Scripts

```
CONNECT SYS/ORACLE AS SYSDBA  
@@CATALOG  
@@CATPROC
```

The use of the double at-sign forces execution of these scripts from the proper directory.

Sizing Redo Log Files

The size of the redo log files can influence performance, because the behavior of the database writer and archiver processes depend on the redo log sizes. Generally, larger redo log files provide better performance. Small log files can increase checkpoint activity and reduce performance. Because the recommendation on I/O distribution for high performance is to use separate disks for the redo log files, there is no reason not to make them large. A potential problem with large redo log files is that these are a single point of failure if redo log mirroring is not in effect.

It is not possible to provide a specific size recommendation for redo log files, but redo log files in the range of a hundred megabytes to a few gigabytes are considered reasonable. Size your online redo log files according to the amount of redo your

system generates. A rough guide is to switch logs at most once every twenty minutes.

The complete set of required redo log files can be created during database creation. After they are created, the size of a redo log size cannot be changed. However, new, larger files can be added later, and the original (smaller) ones can subsequently be dropped.

Not much can be done to speed up the creation of the initial database and the loading of necessary data dictionary views from catalog SQL files. These steps must be run serially after each other.

Note: Although the size of the redo log files does not affect LGWR performance, it can affect DBWR and checkpoint behavior.

Creating Subsequent Tablespaces

After creating the initial database, several extra tablespaces must be created. All databases should have at least three tablespaces in addition to the system tablespace: a temporary tablespace, which is used for things like sorting; a rollback tablespace, which is used to store rollback segments or is designated as the automatic undo management segment; and at least one tablespace for actual application use. In most cases, applications require several tablespaces. For extremely large tablespaces with many datafiles, multiple `ALTER TABLESPACE X ADD DATAFILE Y` statements can also be run in parallel.

During tablespace creation, the datafiles that make up the tablespace are initialized with special "empty" block images. `TEMPFILES` are not initialized.

Oracle does this to ensure that all datafiles can be written in their entirety, but this can obviously be a lengthy process if done serially. Therefore, run multiple `CREATE TABLESPACE` statements concurrently to speed up the tablespace creation process. See the SQL statement in [Example 13-3](#) on page 13-7. The most efficient way to do this is to run one SQL script for each set of available disks.

For permanent tables, the choice between local and global extent management on tablespace creation can have a large effect on performance. For any permanent tablespace that has moderate to large insert, modify, or delete operations compared to reads, local extent management should be chosen.

Creating Permanent Tablespaces - Automatic Segment-Space Management

For permanent tablespaces, Oracle Corporation recommends using automatic segment-space management. Such tablespaces (often referred to as bitmapped tablespaces) are locally managed tablespaces with bitmap segment space management. They are available with Oracle9i Release 1 (9.0.1) and later.

See Also: *Oracle9i Database Administrator's Guide* for more information on creating and using automatic segment-space management for tablespaces

Creating Temporary Tablespaces

Temporary tablespaces can be dictionary-managed or locally managed. With Oracle9i Release 1 (9.0.1) and later, Oracle Corporation recommends use of locally managed temporary tablespaces. [Example 13–3](#) shows how you can create a temporary tablespace with local extent management:

Example 13–3 Creating a Temporary Tablespace

```
CONNECT SYSTEM/MANAGER
CREATE TABLESPACE appdata DATAFILE
  '/u02/oradata/appdata01.dbf' size 1000M;
CREATE TEMPORARY TABLESPACE mytemp TEMPFILE 'temp.f' SIZE 1000M;
```

In another session:

```
CONNECT SYSTEM/MANAGER
CREATE TABLESPACE appindex DATAFILE
  '/u03/oradata/appindex01.dbf' size 1000M;
```

Creating Tables for Good Performance

When installing applications, an initial step is to create all necessary tables and indexes. When you create a segment, such as a table, Oracle allocates space in the database for the data. If subsequent database operations cause the data volume to increase and exceed the space allocated, then Oracle extends the segment.

When creating tables and indexes, note the following:

- Specify automatic segment-space management for tablespaces
 - This allows Oracle to automatically manage segment space for best performance.
- Set storage options carefully

Applications should carefully set storage options for the intended use of the table or index. This includes setting the value for `PCTFREE`. Note that using automatic segment-space management eliminates the need to specify `PCTUSED`.

Note: Use of free lists is no longer encouraged. To use automatic segment-space management, create locally managed tablespaces, with the segment space management clause set to `AUTO`.

- Set the `INITTRANS` value higher if necessary

Each datablock has a number of transaction entries that are used for row locking purposes. Initially, this number is specified by the `INITTRANS` parameter, and the default value (1 for tables, 2 for indexes) is generally sufficient. However, if a table or index is known to have many rows for each block with a high possibility of many concurrent updates, then it is beneficial to set a higher value. This must be done at the `CREATE TABLE/CREATE INDEX` time to ensure that it is set for all blocks of the object.

See Also:

- *Oracle9i Real Application Clusters Concepts* for more information about Real Application Clusters storage considerations
- *Oracle9i Real Application Clusters Administration* for more information about associating instances and users with free lists and free list groups
- *Oracle9i Real Application Clusters Deployment and Performance* for more information about configuring multiblock lock assignments
- *Oracle9i Real Application Clusters Deployment and Performance* for more information about using free lists and free list groups in Real Application Clusters

Note: Note that the operation of creating tables is relatively fast, and not much is gained by doing it in parallel.

Data Segment Compression

Data segment compression reduces disk use and memory use (specifically, the buffer cache), often leading to a better scaleup for read-only operations. Data segment compression can also speed up query execution.

Tuning to Achieve a Better Compression Ratio

Oracle9i Release 2 (9.2) achieves a good compression ratio in many cases with no special tuning. However, if you need a better compression ratio, tuning can improve it slightly in some cases and very substantially in other cases.

Heap-organized block-level compression works by eliminating column value repetitions within individual blocks. This is done by moving such repeated column values into a shared block-level symbol table and replacing occurrences with references into the symbol table. Therefore, the compression ratio is higher in blocks that have more repeated values. As a database administrator or application developer, you can take advantage of this fact by reorganizing rows in the segments that you want to compress, to increase the likelihood of such repetitions.

For a single-column table, you can order table rows by the column value, using a `CREATE TABLE AS SELECT` with an `ORDER BY` clause.

You can also apply this method to a table in which one column has low cardinality and other columns have high cardinalities: Order the table rows by the low cardinality column.

The following views contain information about column cardinalities within segments:

- `ALL_TAB_COL_STATISTICS`
- `ALL_PART_COL_STATISTICS`
- `ALL_SUBPART_COL_STATISTICS`

Example 13–4 Estimating Data Segment Compression/Decompression Ratio

You can estimate the compression or decompression ratio for a table `table_t` by using the following procedure, which allows automatic cleanup:

1. Make sampling repeatable:

```
ALTER SESSION SET EVENTS '10193 trace name context forever, level 1';
```

2. Spawn a `DBMS_JOB` to execute once after 1 minute with the following body (cleanup procedure):

```
LOCK TABLE table_t$a1 IN EXCLUSIVE MODE;  
DROP TABLE table_t$a1;  
DROP TABLE table_t$a2;
```

3. Create an empty compressed table:

```
CREATE TABLE table_t$a1 COMPRESS AS SELECT * FROM table_t WHERE ROWNUM < 1;  
LOCK TABLE table_t$a1 IN SHARE MODE;
```

4. Create an empty uncompressed table:

```
CREATE TABLE table_t$a2 NOCOMPRESS AS SELECT * FROM table_t  
WHERE ROWNUM < 1;  
INSERT /*+ APPEND */ INTO table_t$a1 SELECT * FROM table_t  
SAMPLE BLOCK(x,y);  
INSERT /*+ APPEND */ INTO table_t$a2 SELECT * FROM table_t  
SAMPLE BLOCK(x,y);
```

The data segment compression ratio is the number of blocks in table `table_t$a1` divided by number of blocks in table `table_t$a2`.

Note: The first time you add a compressed partition to a partitioned table that currently contains only uncompressed partitions, you must either drop all existing bitmap indexes and bitmap index partitions or mark them `UNUSABLE` prior to adding the compressed partition, even if the new partition contains no data. Operations for which this condition is relevant are adding, splitting, merging, and moving partitions.

See Also: *Oracle9i SQL Reference* for block group sampling syntax
`SAMPLE BLOCK(x,y)`

Loading and Indexing Data

Many applications need to load data as part of the initial application installation process. This can be fixed data, such as postal code or other type of lookup data, or it can be actual data originating in older systems. Oracle's SQL*Loader tool is the most efficient way to load a substantial amount of data.

Using SQL*Loader for Good Performance

When running SQL*Loader, you specify to use either the conventional path or the direct path. The conventional path uses ordinary SQL `INSERT` statements to load data into the tables, which means that the loading can be done concurrently with other database operations. However, the loading then is also limited by the normal `INSERT` performance. For quick loading of large amounts of data, choose the direct path. With the direct path, the loading process bypasses SQL and loads directly into the database blocks. During this type of load, normal operation on the table (or partition for partitioned tables) cannot be performed.

The following tips could help speed up the data loading process using SQL*Loader:

- Use fixed length fields rather than delimited or quoted fields. This reduces the time needed for SQL*Loader to read and interpret the input file.
- Run SQL*Loader locally rather than through a network connection.
- Load numbers read from the input file as text rather than in binary format. This might be faster, especially if the numbers are loaded without any computation. This is true for integers and for floating point numbers.
- For very large load operations, use SQL*Loader to do parallel data load.

See Also: *Oracle9i Database Utilities* for detailed information on SQL*Loader

Efficient Index Creation

The most efficient way to create indexes is to create them after data has been loaded. By doing this, space management becomes much simpler, and no index maintenance takes place for each row inserted. SQL*Loader automatically does this, but if you are using other methods to do initial data load, you might need to do this manually. Additionally, index creation can be done in parallel using the `PARALLEL` clause of the `CREATE INDEX` statement. However, SQL*Loader is not able to do this, so you must manually create indexes in parallel after loading data.

Specifying Memory for Sorting Data

During index creation on tables that contain data, the data must be sorted. This sorting is done in the fastest possible way, if all available memory is used for sorting. Oracle recommends that you enable automatic sizing of SQL working areas by setting the `PGA_AGGREGATE_TARGET` initialization parameter.

See Also:

- ["Configuring the PGA Working Memory"](#) on page 14-48 for information on PGA memory management
- *Oracle9i Database Reference* for information on the PGA_AGGREGATE_TARGET initialization parameter

Specifying Memory for SQL Work Area with SORT_AREA_SIZE

The memory for the SQL work area can also be controlled with the SORT_AREA_SIZE initialization parameter.

Note: Oracle does not recommend using the SORT_AREA_SIZE parameter unless the instance is configured with the shared server option. Oracle recommends that you enable automatic sizing of SQL working areas by setting PGA_AGGREGATE_TARGET instead. SORT_AREA_SIZE is retained for backward compatibility.

The value of the SORT_AREA_SIZE parameter should be set using the following rules:

1. Find the amount of available memory by subtracting the size of the SGA and the size of the operating system from the total system memory.
2. Divide this amount by the number of parallel slaves that you will use; this is typically the same as the number of CPUs.
3. Subtract a process overhead, typically a five to ten megabytes, to get the value for SORT_AREA_SIZE.

Note: You can also save time on index creating operations, or fast rebuilds, with on the fly statistics generation.

[Example 13-5](#) is an example of setting the SORT_AREA_SIZE parameter.

Example 13-5 Example of Creating Indexes Efficiently

A system with 512 Mb memory runs an Oracle instance with a 100 Mb SGA, and the operating system uses 50 Mb. The memory available for sorting is 362 Mb, which equals 512 minus 50 minus 100. If the system has four CPUs running with four parallel slaves, then each of these will have 90 Mb available. 10 Mb is set aside for

process overhead, so `SORT_AREA_SIZE` should be set to 80 Mb. This can be done either in the initialization file or for each session with the following statement:

```
ALTER SESSION SET SORT_AREA_SIZE = 80000000;
```

Initial Instance Configuration

A running Oracle instance is configured using startup parameters, which are set in the initialization parameter file. These parameters influence the behavior of the running instance, including influencing performance. In general, a very simple initialization file with few relevant settings covers most situations, and the initialization file should not be the first place you expect to do performance tuning, except for the few parameters shown in [Table 13-4](#).

See Also: [Chapter 14, "Memory Configuration and Use"](#)

[Table 13-3](#) describes the parameters necessary in a minimal initialization file. Although these parameters are necessary they have no performance impact:

Table 13-3 *Necessary Initialization Parameters Without Performance Impact*

Parameter	Description
DB_NAME	Name of the database. This should match the ORACLE_SID environment variable.
DB_DOMAIN	Location of the database in Internet dot notation.
OPEN_CURSORS	Limit on the maximum number of cursors (active SQL statements) for each session. The setting is application-dependent, and the default, in many cases, is sufficient.
CONTROL_FILES	Set to contain at least two files on different disk drives to prevent failures from control file loss.
DB_FILES	Set to the maximum number of files that can assigned to the database.

[Table 13-4](#) includes the most important parameters to set with performance implications:

Table 13–4 Important Initialization Parameters With Performance Impact

Parameter	Description
DB_BLOCK_SIZE	Sets the database block size.
DB_CACHE_SIZE	Size of the buffer cache in the SGA. There are no good and simple rules to set a value, which is very application dependent, but typical values are in the range of twenty to fifty for each user session. More often, this value is set too high than too low. DB_BLOCK_BUFFERS has been deprecated.
SHARED_POOL_SIZE	Sets the size of the shared pool in the SGA. The setting is application-dependent, but it is typically in the range of a few to a few tens of megabytes for each user session.
PROCESSES	Sets the maximum number of processes that can be started by that instance. This is the most important primary parameter to set, because many other parameter values are deduced from this.
SESSIONS	This is set by default from the value of processes. However, if you are using the shared server, then the deduced value is likely to be insufficient.
JAVA_POOL_SIZE	If you are using Java stored procedures, then this parameter should be set depending on the actual requirements of memory for the Java environment.
LOG_ARCHIVE_XXX	Enables redo log archiving. See <i>Oracle9i User-Managed Backup and Recovery Guide</i> .
ROLLBACK_SEGMENTS	Allocates one or more rollback segments by name to this instance. If you set this parameter, the instance acquires all of the rollback segments named in this parameter, even if the number of rollback segments exceeds the minimum number required by the instance (calculated as <code>TRANSACTIONS / TRANSACTIONS_PER_ROLLBACK_SEGMENT</code>).

Example of a Minimal Initialization File

In many cases, only the parameters mentioned in the following example need to be set to appropriate values for the Oracle instance to be reasonable well-tuned. Here is an example of such an initialization file:

```
DB_NAME = finance
DB_DOMAIN = hq.company.com
CONTROL_FILES = ('/u01/database/control1.dbf', '/u02/database/control2.dbf')
DB_BLOCK_SIZE = 8192
DB_BLOCK_BUFFERS = 12000 # this is approximately 100 Mb
DB_FILES = 200 # Maximum 200 files in the database
```

```
SHARED_POOL_SIZE = 100000000 # 100 Mb
PROCESSES = 80 # Would be good for approximately 70
# directly connected users
# log_archive_XXX
# Set various archiving parameters
```

Configuring Undo Space

Oracle needs undo space to keep information for read consistency, for recovery, and for actual rollback statements. This is kept either in rollback segments or in one or more automatic undo management tablespaces.

The `V$UNDOSTAT` view contains statistics for monitoring and tuning undo space. Using this view, you can better estimate the amount of undo space required for the current workload. Oracle also uses this information to help tune undo usage in the system. The `V$ROLLSTAT` view contains information about the behavior of the undo segments in the undo tablespaces.

See Also:

- *Oracle9i Database Administrator's Guide* for detailed information on managing undo space using rollback segments or using automatic undo management
- [Chapter 24, "Dynamic Performance Views for Tuning"](#) for more information on the `V$UNDOSTAT` and `V$ROLLSTAT` views

Setting up Operating System, Database, and Network Monitoring

To effectively diagnose performance problems, it is vital to have an established performance baseline for later comparison when the system is running poorly. Without a baseline data point, it can be very difficult to identify new problems. For example, perhaps the volume of transactions on the system has increased, or the transaction profile or application has changed, or the number of users has increased.

After the database is created, tables are created, data is loaded and indexed, and the instance is configured, it is time to set up monitoring procedures.

See Also: [Chapter 20, "Oracle Tools to Gather Database Statistics"](#)

Memory Configuration and Use

Proper sizing and effective use of the Oracle memory caches greatly improves database performance. This chapter explains how to allocate memory to Oracle memory caches, and how to use those caches.

This chapter contains the following sections:

- [Understanding Memory Allocation Issues](#)
- [Configuring and Using the Buffer Cache](#)
- [Configuring and Using the Shared Pool and Large Pool](#)
- [Configuring and Using the Java Pool](#)
- [Configuring and Using the Redo Log Buffer](#)
- [Configuring the PGA Working Memory](#)

Understanding Memory Allocation Issues

Oracle stores information in memory caches and on disk. Memory access is much faster than disk access. Disk scans (physical I/O) take a significant amount of time, compared with memory access, typically in the order of 10 milliseconds. Physical I/O also increases the CPU resources required, because of the path length in device drivers and operating system event schedulers. For this reason, it is more efficient for data requests for frequently accessed objects to be satisfied solely by memory, rather than also requiring disk access.

A performance goal is to reduce the physical I/O overhead as much as possible, either by making it more likely that the required data is in memory or by making the process of retrieving the required data more efficient.

Oracle Memory Caches

The main Oracle memory caches that affect performance are:

- Shared pool
- Large pool
- Java pool
- Buffer cache
- Log buffer
- Process-private memory (for example, used for sorting, hashing, and so on)

The size of these memory caches is configurable using initialization configuration parameters. The values for these parameters are also dynamically configurable using the `ALTER SYSTEM` statement (except for the log buffer and java pool, which are static after startup).

Dynamically Changing Cache Sizes

You can dynamically reconfigure the sizes of the shared pool, the large pool, the buffer cache, and the process-private memory.

Memory for the shared pool, large pool, java pool, and buffer cache is allocated in units of granules. Generally speaking, on most platforms, the size of a granule is 4 MB if the total SGA size is less than 128 MB, and it is 16 MB for larger SGAs. There may be some platform dependency; for example, on 32-bit Windows NT, the granule size is 8 MB for SGAs larger than 128 MB. Please consult your operating system's documentation for more details.

The granule size that is currently being used for SGA can be viewed in the view `V$SGA_DYNAMIC_COMPONENTS`. The same granule size is used for all dynamic components in the SGA.

You can, if necessary, decrease the size of one cache and reallocate that memory to another cache. You can expand the total SGA size to a value equal to the `SGA_MAX_SIZE` parameter.

Note: `SGA_MAX_SIZE` cannot be dynamically resized.

The maximum amount of memory usable by the instance is determined at instance startup by the initialization parameter `SGA_MAX_SIZE`. You can specify `SGA_MAX_SIZE` to be larger than the sum of all of the memory components (such as buffer cache and shared pool); otherwise, `SGA_MAX_SIZE` defaults to the actual size used by those components. Setting `SGA_MAX_SIZE` larger than the sum of memory used by all of the components lets you dynamically increase a cache size without needing to decrease the size of another cache.

Viewing Information About Dynamic Resize Operations

With Oracle9i Release 2 (9.2) and higher, you have the following views that provide information about dynamic SGA resize operations:

- `V$SGA_CURRENT_RESIZE_OPS`: Information about SGA resize operations that are currently in progress. An operation can be a grow or a shrink of a dynamic SGA component.
- `V$SGA_RESIZE_OPS`: Information about the last 100 completed SGA resize operations. This does not include any operations currently in progress.
- `V$SGA_DYNAMIC_COMPONENTS`: Information about the dynamic components in SGA. This view summarizes information based on all completed SGA resize operations since startup.
- `V$SGA_DYNAMIC_FREE_MEMORY`: Information about the amount of SGA memory available for future dynamic SGA resize operations.

See Also:

- *Oracle9i Database Concepts* for more information about dynamic SGA
- *Oracle9i Database Reference* for detailed column information for these views

Application Considerations

With memory configuration, it is important to size the cache appropriately for the application's needs. Conversely, tuning the application's use of the caches can greatly reduce resource requirements. Efficient use of the Oracle memory caches also reduces the load on related resources, such as the latches that protect the caches, the CPU, and the I/O system.

For best performance, you should consider the following:

- The application should be designed and coded to interact with Oracle efficiently.
- Memory allocations to Oracle memory structures should best reflect the needs of the application.

Making changes or additions to an existing application might require resizing Oracle memory structures to meet the needs of your modified application.

Operating System Memory Use

For most operating systems, it is important to consider the following:

Reduce paging

Paging occurs when an operating system transfers memory-resident pages to disk solely to allow new pages to be loaded into memory. Many operating systems page to accommodate large amounts of information that do not fit into real memory. On most operating systems, paging reduces performance.

Use the operating system utilities to examine the operating system, to identify whether there is a lot of paging on your system. If there is, then the total memory on the system might not be large enough to hold everything for which you have allocated memory. Either increase the total memory on your system, or decrease the amount of memory allocated.

Fit the SGA into main memory

Because the purpose of the SGA is to store data in memory for fast access, the SGA should be within main memory. If pages of the SGA are swapped to disk, then the data is no longer quickly accessible. On most operating systems, the disadvantage of paging significantly outweighs the advantage of a large SGA.

Note: The `LOCK_SGA` parameter can be used to lock the SGA into physical memory and prevent it from being paged out.

To see how much memory is allocated to the SGA and each of its internal structures, enter the following SQL*Plus statement:

```
SHOW SGA
```

The output of this statement will look similar to the following:

```
Total System Global Area 840205000 bytes
Fixed Size                 279240 bytes
Variable Size             520093696 bytes
Database Buffers         318767104 bytes
Redo Buffers              1064960 bytes
```

Allow adequate memory to individual users

When sizing the SGA, ensure that you allow enough memory for the individual server processes and any other programs running on the system.

See Also: Your operating system hardware and software documentation, as well as the Oracle documentation specific to your operating system, for more information on tuning operating system memory usage

Iteration During Configuration

Configuring memory allocation involves distributing available memory to Oracle memory structures, depending on the needs of the application. The distribution of memory to Oracle structures can affect the amount of physical I/O necessary for Oracle to operate. Having a good first initial memory configuration also provides an indication of whether the I/O system is effectively configured.

It might be necessary to repeat the steps of memory allocation after the initial pass through the process. Subsequent passes let you make adjustments in earlier steps, based on changes in later steps. For example, decreasing the size of the buffer cache lets you increase the size of another memory structure, such as the shared pool.

Configuring and Using the Buffer Cache

For many types of operations, Oracle uses the buffer cache to store blocks read from disk. Oracle bypasses the buffer cache for particular operations, such as sorting and parallel reads. For operations that use the buffer cache, this section explains the following:

- [Using the Buffer Cache Effectively](#)
- [Sizing the Buffer Cache](#)
- [Interpreting and Using the Buffer Cache Advisory Statistics](#)
- [Considering Multiple Buffer Pools](#)

Using the Buffer Cache Effectively

To use the buffer cache effectively, SQL statements for the application should be tuned to avoid unnecessary resource consumption. To ensure this, verify that frequently executed SQL statements and SQL statements that perform many buffer gets have been tuned.

See Also: [Chapter 6, "Optimizing SQL Statements"](#) for information on how to do this

Sizing the Buffer Cache

When configuring a new instance, it is impossible to know the correct size for the buffer cache. Typically, a database administrator makes a first estimate for the cache size, then runs a representative workload on the instance and examines the relevant statistics to see whether the cache is under or over configured.

Buffer Cache Advisory Statistics

A number of statistics can be used to examine buffer cache activity. These include the following:

- `V$DB_CACHE_ADVICE`
- Buffer cache hit ratio

Using `V$DB_CACHE_ADVICE` This view is populated when the `DB_CACHE_ADVICE` parameter is set to `ON`. This view shows the simulated miss rates for a range of potential buffer cache sizes.

Each cache size simulated has its own row in this view, with the predicted physical I/O activity that would take place for that size. The `DB_CACHE_ADVICE` parameter is dynamic, so the advisory can be enabled and disabled dynamically to allow you to collect advisory data for a specific workload.

Two minor overheads are associated with this advisory:

- CPU: When the advisory is enabled, there is a small increase in CPU usage, because additional bookkeeping is required.
- Memory: The advisory requires memory to be allocated from the shared pool (about 100 bytes for each buffer).

Oracle9i Release 2 (9.2) uses DBA-based sampling to gather cache advisory statistics. Sampling substantially reduces both CPU and memory overhead associated with bookkeeping. Sampling is not used for a buffer pool if the number of buffers in that buffer pool is small to begin with.

To use `V$DB_CACHE_ADVICE`, the parameter `DB_CACHE_ADVICE` should be set to `ON`, and a representative workload should be running on the instance. Allow the workload to stabilize before querying the `V$DB_CACHE_ADVICE` view.

The following SQL statement returns the predicted I/O requirement for the default buffer pool for various cache sizes:

```

COLUMN size_for_estimate          FORMAT 999,999,999,999 heading 'Cache Size (MB)'
COLUMN buffers_for_estimate       FORMAT 999,999,999 heading 'Buffers'
COLUMN estd_physical_read_factor  FORMAT 999.90 heading 'Estd Phys|Read Factor'
COLUMN estd_physical_reads       FORMAT 999,999,999 heading 'Estd Phys| Reads'

SELECT size_for_estimate, buffers_for_estimate, estd_physical_read_factor, estd_physical_reads
FROM V$DB_CACHE_ADVICE
WHERE name = 'DEFAULT'
      AND block_size = (SELECT value FROM V$PARAMETER WHERE name = 'db_block_size')
      AND advice_status = 'ON';

```

The following output shows that if the cache was 212 MB, rather than the current size of 304 MB, the estimated additional number of physical reads would be over 17 million (17,850,847). Increasing the cache size beyond its current size would not provide a significant benefit.

Cache Size (MB)	Buffers	Estd Phys Read Factor	Estd Phys Reads	
30	3,802	18.70	192,317,943	10% of Current Size
60	7,604	12.83	131,949,536	


91	11,406	7.38	75,865,861	
121	15,208	4.97	51,111,658	
152	19,010	3.64	37,460,786	
182	22,812	2.50	25,668,196	
212	26,614	1.74	17,850,847	
243	30,416	1.33	13,720,149	
273	34,218	1.13	11,583,180	
304	38,020	1.00	10,282,475	Current Size
334	41,822	.93	9,515,878	
364	45,624	.87	8,909,026	
395	49,426	.83	8,495,039	
424	53,228	.79	8,116,496	
456	57,030	.76	7,824,764	
486	60,832	.74	7,563,180	
517	64,634	.71	7,311,729	
547	68,436	.69	7,104,280	
577	72,238	.67	6,895,122	
608	76,040	.66	6,739,731	200% of Current Size

This view assists in cache sizing by providing information that predicts the number of physical reads for each potential cache size. The data also includes a physical read factor, which is a factor by which the current number of physical reads is estimated to change if the buffer cache is resized to a given value.

Note: With Oracle, physical reads do not necessarily indicate disk reads; physical reads may well be satisfied from the file system cache.

The relationship between successfully finding a block in the cache and the size of the cache is not always a smooth distribution. When sizing the buffer pool, avoid the use of additional buffers that contribute little or nothing to the cache hit ratio. In the example illustrated in [Figure 14-1](#) on page 14-9, only narrow bands of increments to the cache size may be worthy of consideration.

Figure 14–1 Physical I/O and Buffer Cache Size



Examining [Figure 14–1](#) leads to the following observations:

- The benefit from increasing buffers from point A to point B is considerably higher than from point B to point C.
- The decrease in the physical I/O between points A and B and points B and C is not smooth, as indicated by the dotted line in the graph.

Calculating the Buffer Cache Hit Ratio The buffer cache hit ratio calculates how often a requested block has been found in the buffer cache without requiring disk access. This ratio is computed using data selected from the dynamic performance view `V$SYSSTAT`. The buffer cache hit ratio can be used to verify the physical I/O as predicted by `V$DB_CACHE_ADVICE`.

The statistics in [Table 14–1](#) are used to calculate the hit ratio.

Table 14–1 Statistics for Calculating the Hit Ratio

Statistic	Description
session logical reads	The total number of requests to access a block, whether in memory or on disk.

Table 14–1 (Cont.) Statistics for Calculating the Hit Ratio

Statistic	Description
physical reads	The total number of requests to access a data block that resulted in access to datafiles on disk. The block could have been read into the cache or read into local memory by a direct read.
physical reads direct	The number of blocks read, bypassing the buffer cache, excluding direct reads for large objects (LOBs).
physical reads direct (lob)	The number of blocks read while reading LOBs, bypassing the buffer cache.

Example 14–1 has been simplified by using values selected directly from the V\$SYSSTAT table, rather than over an interval. It is best to calculate the delta of these statistics over an interval while your application is running, then use them to determine the hit ratio.

See Also: [Chapter 20, "Oracle Tools to Gather Database Statistics"](#) for more information on collecting statistics over an interval

Example 14–1 Calculating the Buffer Cache Hit Ratio

```
SELECT NAME, VALUE
   FROM V$SYSSTAT
 WHERE NAME IN ('session logical reads','physical reads',
              'physical reads direct','physical reads direct (lob)');
```

The output of this query will look similar to the following:

```
NAME                                                    VALUE
-----
session logical reads                                  464905358
physical reads                                        10380487
physical reads direct                                 86850
physical reads direct (lob)                          0
```

Calculate the hit ratio for the buffer cache with the following formula:

$$\text{Hit Ratio} = 1 - ((\text{physical reads} - \text{physical reads direct} - \text{physical reads direct (lob)}) / (\text{db block gets} + \text{consistent gets} - \text{physical reads direct} - \text{physical reads direct (lob)}))$$

Based on the sample statistics in the example, the buffer cache hit ratio is equal to **.978 or 97.8%**.

Interpreting and Using the Buffer Cache Advisory Statistics

There are many factors to examine before considering whether to increase or decrease the buffer cache size. For example, you should examine `V$DB_CACHE_ADVICE` data and the buffer cache hit ratio.

A low cache hit ratio does not imply that increasing the size of the cache would be beneficial for performance. A good cache hit ratio could wrongly indicate that the cache is adequately sized for the workload.

To interpret the buffer cache hit ratio, you should consider the following:

- Repeated scanning of the same large table or index can artificially inflate a poor cache hit ratio. Examine frequently executed SQL statements with a large number of buffer gets, to ensure that the execution plan for such SQL statements is optimal. If possible, avoid repeated scanning of frequently accessed data by performing all of the processing in a single pass or by optimizing the SQL statement.
- If possible, avoid requerying the same data, by caching frequently accessed data in the client program or middle tier.
- Blocks encountered during a long full table scan are not put at the head of the list of last recently used (LRU) blocks. Therefore, the blocks are aged out faster than blocks read when performing indexed lookups or small table scans. Thus, poor hit ratios when valid large full table scans are occurring should also be considered when interpreting the buffer cache data.

Note: Short table scans are scans performed on tables under a certain size threshold. The definition of a small table is the maximum of 2% of the buffer cache and 20, whichever is bigger.

- In any large database running OLTP applications in any given unit of time, most rows are accessed either one or zero times. On this basis, there might be little purpose in keeping the block in memory for very long following its use.
- A common mistake is to continue increasing the buffer cache size. Such increases have no effect if you are doing full table scans or operations that do not use the buffer cache.

Increasing Memory Allocated to the Buffer Cache

As a general rule, investigate increasing the size of the cache if the cache hit ratio is low and your application has been tuned to avoid performing full table scans.

To increase cache size, first set the `DB_CACHE_ADVICE` parameter to `ON`, and let the cache statistics stabilize. Examine the advisory data in the `V$DB_CACHE_ADVICE` view to determine the next increment required to significantly decrease the amount of physical I/O performed. If it is possible to allocate the required extra memory to the buffer cache without causing the host operating system to page, then allocate this memory. To increase the amount of memory allocated to the buffer cache, increase the value of the parameter `DB_CACHE_SIZE`.

If required, resize the buffer pools dynamically, rather than shutting down the instance to perform this change.

Note: When the cache is resized, `DB_CACHE_ADVICE` is set to `OFF`. Also, `V$DB_CACHE_ADVICE` shows the advisory for the old value of the cache. This value remains until `DB_CACHE_ADVICE` is explicitly set back to `READY` or `ON`.

The `DB_CACHE_SIZE` parameter specifies the size of the default cache for the database's standard block size. To create and use tablespaces with block sizes different than the database's standard block sizes (such as to support transportable tablespaces), you must configure a separate cache for each block size used. The `DB_nK_CACHE_SIZE` parameter can be used to configure the nonstandard block size needed (where n is 2, 4, 8, 16 or 32 and n is not the standard block size).

Note: The process of choosing a cache size is the same, regardless of whether the cache is the default standard block size cache, the `KEEP` or `RECYCLE` cache, or a nonstandard block size cache.

See Also: *Oracle9i Database Reference* and *Oracle9i Database Administrator's Guide* for more information on using the `DB_nK_CACHE_SIZE` parameters

Reducing Memory Allocated to the Buffer Cache

If the cache hit ratio is high, then the cache is probably large enough to hold the most frequently accessed data. Check `V$DB_CACHE_ADVICE` data to see whether decreasing the cache size significantly causes the number of physical I/Os to increase. If not, and if you require memory for another memory structure, then you might be able to reduce the cache size and still maintain good performance. To make the buffer cache smaller, reduce the size of the cache by changing the value for the parameter `DB_CACHE_SIZE`.

Considering Multiple Buffer Pools

A single default buffer pool is generally adequate for most systems. However, users with detailed knowledge of an application's buffer pool might benefit from configuring multiple buffer pools.

With segments that have atypical access patterns, store blocks from those segments in two different buffer pools: the `KEEP` pool and the `RECYCLE` pool. A segment's access pattern may be atypical if it is constantly accessed (that is, hot) or infrequently accessed (for example, a large segment accessed by a batch job only once a day).

Multiple buffer pools let you address these differences. You can use a `KEEP` buffer pool to maintain frequently accessed segments in the buffer cache, and a `RECYCLE` buffer pool to prevent objects from consuming unnecessary space in the cache. When an object is associated with a cache, all blocks from that object are placed in that cache. Oracle maintains a `DEFAULT` buffer pool for objects that have not been assigned to a specific buffer pool. The default buffer pool is of size `DB_CACHE_SIZE`. Each buffer pool uses the same LRU replacement policy (for example, if the `KEEP` pool is not large enough to store all of the segments allocated to it, then the oldest blocks age out of the cache).

By allocating objects to appropriate buffer pools, you can:

- Reduce or eliminate I/Os
- Isolate or limit an object to a separate cache

Random Access to Large Segments

A problem can occur with an LRU aging method when a very large segment is accessed with a large or unbounded index range scan. Here, very large means large compared to the size of the cache. Any single segment that accounts for a substantial portion (more than 10%) of nonsequential physical reads can be considered very large. Random reads to a large segment can cause buffers that contain data for other segments to be aged out of the cache. The large segment ends up consuming a large percentage of the cache, but it does not benefit from the cache.

Very frequently accessed segments are not affected by large segment reads because their buffers are warmed frequently enough that they do not age out of the cache. However, the problem affects warm segments that are not accessed frequently enough to survive the buffer aging caused by the large segment reads. There are three options for solving this problem:

1. If the object accessed is an index, find out whether the index is selective. If not, tune the SQL statement to use a more selective index.
2. If the SQL statement is tuned, you can move the large segment into a separate `RECYCLE` cache so that it does not affect the other segments. The `RECYCLE` cache should be smaller than the `DEFAULT` buffer pool, and it should reuse buffers more quickly than the `DEFAULT` buffer pool.
3. Alternatively, you can move the small warm segments into a separate `KEEP` cache that is not used at all for large segments. The `KEEP` cache can be sized to minimize misses in the cache. You can make the response times for specific queries more predictable by putting the segments accessed by the queries in the `KEEP` cache to ensure that they do not age out.

Oracle Real Application Cluster Instances

You can create multiple buffer pools for each database instance. The same set of buffer pools need not be defined for each instance of the database. Among instances, the buffer pools can be different sizes or not defined at all. Tune each instance according to the application requirements for that instance.

Using Multiple Buffer Pools

To define a default buffer pool for an object, use the `BUFFER_POOL` keyword of the `STORAGE` clause. This clause is valid for `CREATE` and `ALTER TABLE`, `CLUSTER`, and `INDEX` SQL statements. After a buffer pool has been specified, all subsequent blocks read for the object are placed in that pool.

If a buffer pool is defined for a partitioned table or index, then each partition of the object inherits the buffer pool from the table or index definition, unless you override it with a specific buffer pool.

When the buffer pool of an object is changed using the `ALTER` statement, all buffers currently containing blocks of the altered segment remain in the buffer pool they were in before the `ALTER` statement. Newly loaded blocks and any blocks that have aged out and are reloaded go into the new buffer pool.

See Also: *Oracle9i SQL Reference* for information on specifying `BUFFER_POOL` in the `STORAGE` clause

Buffer Pool Data in V\$DB_CACHE_ADVICE

V\$DB_CACHE_ADVICE can be used to size all pools configured on an instance. Make the initial cache size estimate, run the representative workload, then simply query the V\$DB_CACHE_ADVICE view for the pool you want to use.

For example, to query data from the KEEP pool:

```
SELECT SIZE_FOR_ESTIMATE, BUFFERS_FOR_ESTIMATE, ESTD_PHYSICAL_READ_FACTOR, ESTD_PHYSICAL_READS
FROM V$DB_CACHE_ADVICE
WHERE NAME          = 'KEEP'
AND BLOCK_SIZE     = (SELECT VALUE FROM V$PARAMETER WHERE NAME = 'db_block_size')
AND ADVICE_STATUS = 'ON';
```

Buffer Pool Hit Ratios

The data in V\$SYSSTAT reflects the logical and physical reads for all buffer pools within one set of statistics. To determine the hit ratio for the buffer pools individually, you must query the V\$BUFFER_POOL_STATISTICS view. This view maintains statistics for each pool on the number of logical reads and writes.

The buffer pool hit ratio can be determined using the following formula:

$$\text{hit ratio} = 1 - [\text{physical reads} / (\text{block gets} + \text{consistent gets})]$$

The values *physical reads*, *block gets*, and *consistent gets* can be obtained from the following query:

```
SELECT NAME, PHYSICAL_READS, DB_BLOCK_GETS, CONSISTENT_GETS,
       1 - (PHYSICAL_READS / (DB_BLOCK_GETS + CONSISTENT_GETS)) "Hit Ratio"
FROM V$BUFFER_POOL_STATISTICS;
```

Determining Which Segments Have Many Buffers in the Pool

The V\$BH view shows the data object ID of all blocks that currently reside in the SGA. To determine which segments have many buffers in the pool, you can use one of the two methods described in this section. You can either look at the buffer cache usage pattern for all segments ([Method 1](#)) or examine the usage pattern of a specific segment, ([Method 2](#)).

Method 1 The following query counts the number of blocks for all segments that reside in the buffer cache at that point in time. Depending on buffer cache size, this might require a lot of sort space.

```
COLUMN object_name FORMAT a40
COLUMN number_of_blocks FORMAT 999,999,999,999
```

```
SELECT o.object_name, COUNT(1) number_of_blocks
   FROM DBA_OBJECTS o, V$BH bh
  WHERE o.object_id = bh.objd
        AND o.owner   != 'SYS'
  GROUP BY o.object_name
  ORDER BY count(1);
```

OBJECT_NAME	NUMBER_OF_BLOCKS
OA_PREF_UNIQ_KEY	1
SYS_C002651	1
..	
DS_PERSON	78
OM_EXT_HEADER	701
OM_SHELL	1,765
OM_HEADER	5,826
OM_INSTANCE	12,644

Method 2 Use the following steps to determine the percentage of the cache used by an individual object at a given point in time:

1. Find the Oracle internal object number of the segment by entering the following query:

```
SELECT DATA_OBJECT_ID, OBJECT_TYPE
   FROM DBA_OBJECTS
  WHERE OBJECT_NAME = UPPER('SEGMENT_NAME');
```

Because two objects can have the same name (if they are different types of objects), use the OBJECT_TYPE column to identify the object of interest.

2. Find the number of buffers in the buffer cache for SEGMENT_NAME:

```
SELECT COUNT(*) BUFFERS
   FROM V$BH
  WHERE objd = data_object_id_value;
```

where *data_object_id_value* is from step 1.

3. Find the number of buffers in the instance:

```
SELECT NAME, BLOCK_SIZE, SUM(BUFFERS)
   FROM V$BUFFER_POOL
  GROUP BY NAME, BLOCK_SIZE
  HAVING SUM(BUFFERS) > 0;
```


4. Calculate the ratio of buffers to total buffers to obtain the percentage of the cache currently used by `SEGMENT_NAME`:

```
% cache used by segment_name = [buffers(Step2)/total buffers(Step3)]
```

Note: This technique works only for a single segment. You must run the query for each partition for a partitioned object.

KEEP Pool

If there are certain segments in your application that are referenced frequently, then store the blocks from those segments in a separate cache called the `KEEP` buffer pool. Memory is allocated to the `KEEP` buffer pool by setting the parameter `DB_KEEPCACHE_SIZE` to the required size. The memory for the `KEEP` pool is not a subset of the default pool. Typical segments that can be kept are small reference tables that are used frequently. Application developers and DBAs can determine which tables are candidates.

You can check the number of blocks from candidate tables by querying `V$BH`, as described in "[Determining Which Segments Have Many Buffers in the Pool](#)" on page 14-15.

Note: The `NOCACHE` clause has no effect on a table in the `KEEP` cache.

The goal of the `KEEP` buffer pool is to retain objects in memory, thus avoiding I/O operations. The size of the `KEEP` buffer pool, therefore, depends on the objects that you want to keep in the buffer cache. You can compute an approximate size for the `KEEP` buffer pool by adding together the blocks used by all objects assigned to this pool. If you gather statistics on the segments, you can query `DBA_TABLES.BLOCKS` and `DBA_TABLES.EMPTY_BLOCKS` to determine the number of blocks used.

Calculate the hit ratio by taking two snapshots of system performance at different times, using the previous query. Subtract the more recent values for `physical reads`, `block gets`, and `consistent gets` from the older values, and use the results to compute the hit ratio.

A buffer pool hit ratio of 100% might not be optimal. Often, you can decrease the size of your `KEEP` buffer pool and still maintain a sufficiently high hit ratio. Allocate blocks removed from the `KEEP` buffer pool to other buffer pools.

Note: If an object grows in size, then it might no longer fit in the `KEEP` buffer pool. You will begin to lose blocks out of the cache.

Each object kept in memory results in a trade-off. It is beneficial to keep frequently-accessed blocks in the cache, but retaining infrequently-used blocks results in less space for other, more active blocks.

RECYCLE Pool

It is possible to configure a `RECYCLE` buffer pool for blocks belonging to those segments that you do not want to remain in memory. The `RECYCLE` pool is good for segments that are scanned rarely or are not referenced frequently. If an application accesses the blocks of a very large object in a random fashion, then there is little chance of reusing a block stored in the buffer pool before it is aged out. This is true regardless of the size of the buffer pool (given the constraint of the amount of available physical memory). Consequently, the object's blocks need not be cached; those cache buffers can be allocated to other objects.

Memory is allocated to the `RECYCLE` buffer pool by setting the parameter `DB_RECYCLE_CACHE_SIZE` to the required size. This memory for the `RECYCLE` buffer pool is not a subset of the default pool.

Do not discard blocks from memory too quickly. If the buffer pool is too small, then blocks can age out of the cache before the transaction or SQL statement has completed execution. For example, an application might select a value from a table, use the value to process some data, and then update the record. If the block is removed from the cache after the `SELECT` statement, then it must be read from disk again to perform the update. The block should be retained for the duration of the user transaction.

Configuring and Using the Shared Pool and Large Pool

Oracle uses the shared pool to cache many different types of data. Cached data includes the textual and executable forms of PL/SQL blocks and SQL statements, dictionary cache data, and other data.

Proper use and sizing of the shared pool can reduce resource consumption in at least four ways:

1. Parse overhead is avoided if the SQL statement is already in the shared pool. This saves CPU resources on the host and elapsed time for the end user.

2. Latching resource usage is significantly reduced, which results in greater scalability.
3. Shared pool memory requirements are reduced, because all applications use the same pool of SQL statements and dictionary resources.
4. I/O resources are saved, because dictionary elements that are in the shared pool do not require disk access.

This section covers the following:

- [Shared Pool Concepts](#)
- [Using the Shared Pool Effectively](#)
- [Sizing the Shared Pool](#)
- [Interpreting Shared Pool Statistics](#)
- [Using the Large Pool](#)
- [Using CURSOR_SPACE_FOR_TIME](#)
- [Caching Session Cursors](#)
- [Configuring the Reserved Pool](#)
- [Keeping Large Objects to Prevent Aging](#)
- [CURSOR_SHARING for Existing Applications](#)

Shared Pool Concepts

The main components of the shared pool are the library cache and the dictionary cache. The library cache stores the executable (parsed or compiled) form of recently referenced SQL and PL/SQL code. The dictionary cache stores data referenced from the data dictionary. Many of the caches in the shared pool automatically increase or decrease in size, as needed, including the library cache and the dictionary cache. Old entries are aged out of these caches to accommodate new entries when the shared pool does not have free space.

A cache miss on the data dictionary cache or library cache is more expensive than a miss on the buffer cache. For this reason, the shared pool should be sized to ensure that frequently used data is cached.

A number of features make large memory allocations in the shared pool: for example, the shared server, parallel query, or Recovery Manager. Oracle recommends segregating the SGA memory used by these features by configuring a distinct memory area, called the large pool.

See Also: ["Using the Large Pool"](#) on page 14-35 for more information on configuring the large pool

Allocation of memory from the shared pool is performed in chunks. This allows large objects (over 5k) to be loaded into the cache without requiring a single contiguous area, hence reducing the possibility of running out of enough contiguous memory due to fragmentation.

Infrequently, Java, PL/SQL, or SQL cursors may make allocations out of the shared pool that are larger than 5k. To allow these allocations to occur most efficiently, Oracle segregates a small amount of the shared pool. This memory is used if the shared pool does not have enough space. The segregated area of the shared pool is called the reserved pool.

See Also: ["Configuring the Reserved Pool"](#) on page 14-40 for more information on the reserved area of the shared pool

Dictionary Cache Concepts

Information stored in the data dictionary cache includes usernames, segment information, profile data, tablespace information, and sequence numbers. The dictionary cache also stores descriptive information, or metadata, about schema objects. Oracle uses this metadata when parsing SQL cursors or during the compilation of PL/SQL programs.

Library Cache Concepts

The library cache holds executable forms of SQL cursors, PL/SQL programs, and Java classes. This section focuses on tuning as it relates to cursors, PL/SQL programs, and Java classes. These are collectively referred to as application code.

When application code is run, Oracle attempts to reuse existing code if it has been executed previously and can be shared. If the parsed representation of the statement does exist in the library cache and it can be shared, then Oracle reuses the existing code. This is known as a soft parse, or a library cache hit.

If Oracle is unable to use existing code, then a new executable version of the application code must be built. This is known as a hard parse, or a library cache miss.

See Also: ["SQL Sharing Criteria"](#) on page 14-21 for details on when a SQL and PL/SQL statements can be shared

Library cache misses can occur on either the parse step or the execute step when processing a SQL statement.

When an application makes a parse call for a SQL statement, if the parsed representation of the statement does not already exist in the library cache, then Oracle parses the statement and stores the parsed form in the shared pool. This is a hard parse. You might be able to reduce library cache misses on parse calls by ensuring that all shareable SQL statements are in the shared pool whenever possible.

If an application makes an execute call for a SQL statement, and if the executable portion of the previously built SQL statement has been aged out (that is, deallocated) from the library cache to make room for another statement, then Oracle implicitly reparses the statement, creating a new shared SQL area for it, and executes it. This also results in a hard parse. Usually, you can reduce library cache misses on execution calls by allocating more memory to the library cache.

In order to perform a hard parse, Oracle uses more resources than during a soft parse. Resources used for a soft parse include CPU and library cache latch gets. Resources required for a hard parse include additional CPU, library cache latch gets, and shared pool latch gets.

SQL Sharing Criteria

Oracle automatically determines whether a SQL statement or PL/SQL block being issued is identical to another statement currently in the shared pool.

Oracle performs the following steps for the comparison:

1. The text of the statement issued is compared to existing statements in the shared pool.
2. The text of the statement is hashed. If there is no matching hash value, then the SQL statement does not currently exist in the shared pool, and a hard parse is performed.
3. If there is a matching hash value for an existing SQL statement in the shared pool, then Oracle compares the text of the matched statement to the text of the statement hashed to see if they are identical. The text of the SQL statements or PL/SQL blocks must be identical, character for character, including spaces, case, and comments. For example, the following statements cannot use the same shared SQL area:

```
SELECT * FROM employees;  
SELECT * FROM Employees;  
SELECT * FROM employees;
```

Usually, SQL statements that differ only in literals cannot use the same shared SQL area. For example, the following SQL statements do not resolve to the same SQL area:

```
SELECT count(1) FROM employees WHERE manager_id = 121;  
SELECT count(1) FROM employees WHERE manager_id = 247;
```

The only exception to this rule is when the parameter `CURSOR_SHARING` has been set to `SIMILAR` or `FORCE`. Similar statements can share SQL areas when the `CURSOR_SHARING` parameter is set to `SIMILAR` or `FORCE`. The costs and benefits involved in using `CURSOR_SHARING` are explained later in this section.

See Also: *Oracle9i Database Reference* for more information on the `CURSOR_SHARING` parameter

4. The objects referenced in the issued statement are compared to the referenced objects of all existing statements in the shared pool to ensure that they are identical.

References to schema objects in the SQL statements or PL/SQL blocks must resolve to the same object in the same schema.

For example, if two users each issue the following SQL statement:

```
SELECT * FROM employees;
```

and they each have their own `employees` table, then this statement is not considered identical, because the statement references different tables for each user.

5. Bind variables in the SQL statements must match in name, datatype, and length.

For example, the following statements cannot use the same shared SQL area, because the bind variable names differ:

```
SELECT * FROM employees WHERE department_id = :department_id;  
SELECT * FROM employees WHERE department_id = :dept_id;
```

Many Oracle products (such as Oracle Forms and the precompilers) convert the SQL before passing statements to the database. Characters are uniformly changed to uppercase, white space is compressed, and bind variables are renamed so that a consistent set of SQL statements is produced.

6. The session's environment must be identical. Items compared include the following:
 - Optimization approach and goal. SQL statements must be optimized using the same optimization approach and, in the case of the cost-based approach, the same optimization goal.
 - Session-configurable parameters such as `SORT_AREA_SIZE`.

Using the Shared Pool Effectively

An important purpose of the shared pool is to cache the executable versions of SQL and PL/SQL statements. This allows multiple executions of the same SQL or PL/SQL code to be performed without the resources required for a hard parse, which results in significant reductions in CPU, memory, and latch usage.

The shared pool is also able to support unshared SQL in data warehousing applications, which execute low-concurrency, high-resource SQL statements. In this situation, using unshared SQL with literal values is recommended. Using literal values rather than bind variables allows the optimizer to make good column selectivity estimates, thus providing an optimal data access plan.

See Also: *Oracle9i Data Warehousing Guide*

In an OLTP system, there are a number of ways to ensure efficient use of the shared pool and related resources. Discuss the following items with application developers and agree on strategies to ensure that the shared pool is used effectively:

- [Shared Cursors](#)
- [Maintaining Connections](#)
- [Single-User Logon and Qualified Table Reference](#)
- [Use of PL/SQL](#)
- [Avoid Performing DDL](#)
- [Cache Sequence Numbers](#)
- [Cursor Access and Management](#)

Efficient use of the shared pool in high-concurrency OLTP systems significantly reduces the probability of parse-related application scalability issues.

Shared Cursors

Reuse of shared SQL for multiple users running the same application, avoids hard parsing. Soft parses provide a significant reduction in the use of resources such as the shared pool and library cache latches. To share cursors, do the following:

- Use bind variables rather than literals in SQL statements whenever possible. For example, the following two statements cannot use the same shared area because they do not match character for character:

```
SELECT employee_id FROM employees WHERE department_id = 10;  
SELECT employee_id FROM employees WHERE department_id = 20;
```

By replacing the literals with a bind variable, only one SQL statement would result, which could be executed twice:

```
SELECT employee_id FROM employees WHERE department_id = :dept_id;
```

Note: For existing applications where rewriting the code to use bind variables is impractical, it is possible to use the `CURSOR_SHARING` initialization parameter to avoid some of the hard parse overhead. For more information see section "[CURSOR_SHARING for Existing Applications](#)" on page 14-43.

- Avoid application designs that result in large numbers of users issuing dynamic, unshared SQL statements. Typically, the majority of data required by most users can be satisfied using preset queries. Use dynamic SQL where such functionality is required.
- Be sure that users of the application do not change the optimization approach and goal for their individual sessions.
- Establish the following policies for application developers:
 - Standardize naming conventions for bind variables and spacing conventions for SQL statements and PL/SQL blocks.
 - Consider using stored procedures whenever possible. Multiple users issuing the same stored procedure use the same shared PL/SQL area automatically. Because stored procedures are stored in a parsed form, their use reduces runtime parsing.

Maintaining Connections

Large OLTP applications with middle tiers should maintain connections, rather than connecting and disconnecting for each database request. Maintaining connections saves CPU resources and database resources, such as latches.

Single-User Logon and Qualified Table Reference

Large OLTP systems where users log in to the database as their own user ID can benefit from explicitly qualifying the segment owner, rather than using public synonyms. This significantly reduces the number of entries in the dictionary cache. For example:

```
SELECT employee_id FROM hr.employees WHERE department_id = :dept_id;
```

An alternative to qualifying table names is to connect to the database through a single user ID, rather than individual user IDs. User-level validation can take place locally on the middle tier. Reducing the number of distinct userIDs also reduces the load on the dictionary cache.

Use of PL/SQL

Using stored PL/SQL packages can overcome many of the scalability issues for systems with thousands of users, each with individual user sign-on and public synonyms. This is because a package is executed as the owner, rather than the caller, which reduces the dictionary cache load considerably.

Note: Oracle Corporation encourages the use of definer-rights packages to overcome scalability issues. The benefits of reduced dictionary cache load are not as obvious with invoker-rights packages.

Avoid Performing DDL

Avoid performing DDL operations on high-usage segments during peak hours. Performing DDL on such segments often results in the dependent SQL being invalidated and hence reparsed on a later execution.

Cache Sequence Numbers

Allocating sufficient cache space for frequently updated sequence numbers significantly reduces the frequency of dictionary cache locks, which improves

scalability. The `CACHE` keyword on the `CREATE SEQUENCE` or `ALTER SEQUENCE` statement lets you configure the number of cached entries for each sequence.

See Also: *Oracle9i SQL Reference* for details on the `CREATE SEQUENCE` and `ALTER SEQUENCE` statements

Cursor Access and Management

Depending on the Oracle application tool you are using, it is possible to control how frequently your application performs parse calls.

The frequency with which your application either closes cursors or reuses existing cursors for new SQL statements affects the amount of memory used by a session and often the amount of parsing performed by that session.

An application that closes cursors or reuses cursors (for a different SQL statement), does not need as much session memory as an application that keeps cursors open. Conversely, that same application may need to perform more parse calls, using extra CPU and Oracle resources.

Cursors associated with SQL statements that are not executed frequently can be closed or reused for other statements, because the likelihood of reexecuting (and reparsing) that statement is low.

Extra parse calls are required when a cursor containing a SQL statement that will be reexecuted is closed or reused for another statement. Had the cursor remained open, it could have been reused without the overhead of issuing a parse call.

The ways in which you control cursor management depends on your application development tool. The following sections introduce the methods used for some Oracle tools.

See Also:

- The tool-specific documentation for more information about each tool
- *Oracle9i Database Concepts* for more information on cursor sharing and management

Reducing Parse Calls with OCI When using Oracle Call Interface (OCI), do not close and reopen cursors that you will be reexecuting. Instead, leave the cursors open, and change the literal values in the bind variables before execution.

Avoid reusing statement handles for new SQL statements when the existing SQL statement will be reexecuted in the future.

Reducing Parse Calls with the Oracle Precompilers When using the Oracle precompilers, you can control when cursors are closed by setting precompiler clauses. In Oracle mode, the clauses are as follows:

- `HOLD_CURSOR = YES`
- `RELEASE_CURSOR = NO`
- `MAXOPENCURSORS = desired_value`

Oracle Corporation recommends that you not use ANSI mode, in which the values of `HOLD_CURSOR` and `RELEASE_CURSOR` are switched.

The precompiler clauses can be specified on the precompiler command line or within the precompiler program. With these clauses, you can employ different strategies for managing cursors during execution of the program.

See Also: Your language's precompiler manual for more information on these clauses

Reducing Parse Calls with SQLJ Prepare the statement, then reexecute the statement with the new values for the bind variables. The cursor stays open for the duration of the session.

Reducing Parse Calls with JDBC Avoid closing cursors if they will be reexecuted, because the new literal values can be bound to the cursor for reexecution. Alternatively, JDBC provides a SQL statement cache within the JDBC client using the `setStmtCacheSize()` method. Using this method, JDBC creates a SQL statement cache that is local to the JDBC program.

See Also: *Oracle9i JDBC Developer's Guide and Reference* for more information on using the JDBC SQL statement cache

Reducing Parse Calls with Oracle Forms With Oracle Forms, it is possible to control some aspects of cursor management. You can exercise this control either at the trigger level, at the form level, or at run time.

Sizing the Shared Pool

When configuring a brand new instance, it is impossible to know the correct size to make the shared pool cache. Typically, a DBA makes a first estimate for the cache size, then runs a representative workload on the instance, and examines the relevant statistics to see whether the cache is under-configured or over-configured.

For most OLTP applications, shared pool size is an important factor in application performance. Shared pool size is less important for applications that issue a very limited number of discrete SQL statements, such as data support systems (DSS).

If the shared pool is too small, then extra resources are used to manage the limited amount of available space. This consumes CPU and latching resources, and causes contention.

Optimally, the shared pool should be just large enough to cache frequently accessed objects. Having a significant amount of free memory in the shared pool is a waste of memory.

Shared Pool: Library Cache Statistics

When sizing the shared pool, the goal is to ensure that SQL statements that will be executed multiple times are cached in the library cache, without allocating too much memory.

The statistic that shows the amount of reloading (that is, reparsing) of a previously cached SQL statement that was aged out of the cache is the `RELOADS` column in the `V$LIBRARYCACHE` view. In an application that reuses SQL effectively, on a system with an optimal shared pool size, the `RELOADS` statistic will have a value near zero.

The `INVALIDATIONS` column in `V$LIBRARYCACHE` view shows the number of times library cache data was invalidated and had to be reparsed. `INVALIDATIONS` should be near zero. This means SQL statements that could have been shared were invalidated by some operation (for example, a DDL). This statistic should be near zero on OLTP systems during peak loads.

Another key statistic is the amount of free memory in the shared pool at peak times. The amount of free memory can be queried from `V$SGASTAT`, looking at the free memory for the shared pool. Optimally, free memory should be as low as possible, without causing any reloads on the system.

Lastly, a broad indicator of library cache health is the library cache hit ratio. This value should be considered along with the other statistics discussed in this section and other data, such as the rate of hard parsing and whether there is any shared pool or library cache latch contention.

These statistics are discussed in more detail in the following section.

V\$LIBRARYCACHE

You can monitor statistics reflecting library cache activity by examining the dynamic performance view `V$LIBRARYCACHE`. These statistics reflect all library cache activity since the most recent instance startup.

Each row in this view contains statistics for one type of item kept in the library cache. The item described by each row is identified by the value of the `NAMESPACE` column. Rows with the following `NAMESPACE` values reflect library cache activity for SQL statements and PL/SQL blocks:

- `SQL AREA`
- `TABLE/PROCEDURE`
- `BODY`
- `TRIGGER`

Rows with other `NAMESPACE` values reflect library cache activity for object definitions that Oracle uses for dependency maintenance.

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#) for information on columns in the `V$LIBRARYCACHE` view

To examine each namespace individually, use the following query:

```
SELECT namespace
       , pins
       , pinhits
       , reloads
       , invalidations
FROM V$LIBRARYCACHE
ORDER BY namespace;
```

The output of this query could look like the following:

NAMESPACE	PINS	PINHITS	RELOADS	INVALIDATIONS
BODY	8870	8819	0	0
CLUSTER	393	380	0	0
INDEX	29	0	0	0
OBJECT	0	0	0	0
PIPE	55265	55263	0	0
SQL AREA	21536413	21520516	11204	2
TABLE/PROCEDURE	10775684	10774401	0	0
TRIGGER	1852	1844	0	0

To calculate the library cache hit ratio, use the following formula:

$$\text{Library Cache Hit Ratio} = \text{sum}(\text{pinhits}) / \text{sum}(\text{pins})$$

Using the library cache hit ratio formula, the cache hit ratio is the following:

```
SUM(PINHITS) / SUM(PINS)
-----
.999466248
```

Note: These queries return data from instance startup, rather than statistics gathered during an interval; interval statistics can better pinpoint the problem.

See Also: [Chapter 20, "Oracle Tools to Gather Database Statistics"](#) for information on how gather information over an interval

Examining the returned data leads to the following observations:

- For the `SQL AREA` namespace, there were 21,536,413 executions.
- 11,204 of the executions resulted in a library cache miss, requiring Oracle to implicitly reparse a statement or block or reload an object definition because it aged out of the library cache (that is, a `RELOAD`).
- SQL statements were invalidated two times, again causing library cache misses.
- The hit percentage is about 99.94%. This means that only .06% of executions resulted in reparsing.

The amount of free memory in the shared pool is reported in `V$SGASTAT`. Report the current value from this view using the following query:

```
SELECT * FROM V$SGASTAT
WHERE NAME = 'free memory'
AND POOL = 'shared pool';
```

The output will be similar to the following:

```
POOL          NAME                               BYTES
-----
shared pool free memory                4928280
```

If free memory is always available in the shared pool, then increasing the size of the pool offers little or no benefit. However, just because the shared pool is full does not necessarily mean there is a problem. It may be indicative of a well-configured system.

Shared Pool Advisory Statistics

The amount of memory available for the library cache can drastically affect the parse rate of an Oracle instance. With Oracle9i, Release 2 (9.2) or higher, the shared pool advisory statistics provide a database administrator with information about library cache memory and predict how changes in the size of the shared pool can affect the parse rate.

The shared pool advisory statistics track the library cache's use of shared pool memory and predict how the library cache will behave in shared pools of different sizes. Two fixed views provide the information to determine how much memory the library cache is using, how much is currently pinned, how much is on the shared pool's LRU list, as well as how much time might be lost or gained by changing the size of the shared pool.

The following views of the shared pool advisory statistics are available. These views display any data when shared pool advisory is on. These statistics reset when the advisory is turned off.

V\$SHARED_POOL_ADVICE This view displays information about estimated parse time savings in different sizes of shared pool. The sizes range from 50% to 200% of current shared pool size, in equal intervals. The value of the interval depends on current shared pool size.

Parse time saved refers to the amount of time saved by keeping library cache memory objects in the shared pool, as opposed to having to reload these object.

V\$LIBRARY_CACHE_MEMORY This view displays information about memory allocated to library cache memory objects in different namespaces. A memory object is an internal grouping of memory for efficient management. A library cache object may consist of one or more memory objects.

See Also:

- ["V\\$SHARED_POOL_ADVICE"](#) on page 24-45
- ["V\\$LIBRARY_CACHE_MEMORY"](#) on page 24-16

Shared Pool: Dictionary Cache Statistics

Typically, if the shared pool is adequately sized for the library cache, it will also be adequate for the dictionary cache data.

Misses on the data dictionary cache are to be expected in some cases. On instance startup, the data dictionary cache contains no data. Therefore, any SQL statement

issued is likely to result in cache misses. As more data is read into the cache, the likelihood of cache misses decreases. Eventually, the database reaches a steady state, in which the most frequently used dictionary data is in the cache. At this point, very few cache misses occur.

Each row in the `V$ROWCACHE` view contains statistics for a single type of data dictionary item. These statistics reflect all data dictionary activity since the most recent instance startup. The columns in the `V$ROWCACHE` view that reflect the use and effectiveness of the data dictionary cache are listed in [Table 14-2](#).

Table 14-2 *V\$ROWCACHE Columns*

Column	Description
PARAMETER	Identifies a particular data dictionary item. For each row, the value in this column is the item prefixed by <code>dc_</code> . For example, in the row that contains statistics for file descriptions, this column has the value <code>dc_files</code> .
GETS	Shows the total number of requests for information on the corresponding item. For example, in the row that contains statistics for file descriptions, this column has the total number of requests for file description data.
GETMISSES	Shows the number of data requests which were not satisfied by the cache, requiring an I/O.
MODIFICATIONS	Shows the number of times data in the dictionary cache was updated.

Use the following query to monitor the statistics in the `V$ROWCACHE` view over a period of time while your application is running. The derived column `PCT_SUCC_GETS` can be considered the item-specific hit ratio:

```
column parameter format a21
column pct_succ_gets format 999.9
column updates format 999,999,999

SELECT parameter
       , sum(gets)
       , sum(getmisses)
       , 100*sum(gets - getmisses) / sum(gets) pct_succ_gets
       , sum(modifications) updates
FROM V$ROWCACHE
WHERE gets > 0
GROUP BY parameter;
```


The output of this query will be similar to the following:

PARAMETER	SUM(GETS)	SUM(GETMISSES)	PCT_SUCC_GETS	UPDATES
dc_database_links	81	1	98.8	0
dc_free_extents	44876	20301	54.8	40,453
dc_global_oids	42	9	78.6	0
dc_histogram_defs	9419	651	93.1	0
dc_object_ids	29854	239	99.2	52
dc_objects	33600	590	98.2	53
dc_profiles	19001	1	100.0	0
dc_rollback_segments	47244	16	100.0	19
dc_segments	100467	19042	81.0	40,272
dc_sequence_grants	119	16	86.6	0
dc_sequences	26973	16	99.9	26,811
dc_synonyms	6617	168	97.5	0
dc_tablespace_quotas	120	7	94.2	51
dc tablespaces	581248	10	100.0	0
dc_used_extents	51418	20249	60.6	42,811
dc_user_grants	76082	18	100.0	0
dc_usernames	216860	12	100.0	0
dc_users	376895	22	100.0	0

Examining the data returned by the sample query leads to these observations:

- There are large numbers of misses and updates for used extents, free extents, and segments. This implies that the instance had a significant amount of dynamic space extension.
- Based on the percentage of successful gets, and comparing that statistic with the actual number of gets, the shared pool is large enough to store dictionary cache data adequately.

It is also possible to calculate an overall dictionary cache hit ratio using the following formula; however, summing up the data over all the caches will lose the finer granularity of data:

```
SELECT (SUM(GETS - GETMISSES - FIXED)) / SUM(GETS) "ROW CACHE" FROM V$ROWCACHE;
```

Interpreting Shared Pool Statistics

Shared pool statistics indicate adjustments that can be made. The following sections describe some of your choices.

Increasing Memory Allocation

Increasing the amount of memory for the shared pool increases the amount of memory available to both the library cache and the dictionary cache.

Allocating Additional Memory for the Library Cache To ensure that shared SQL areas remain in the cache after their SQL statements are parsed, increase the amount of memory available to the library cache until the `V$LIBRARYCACHE.RELOADS` value is near zero. To increase the amount of memory available to the library cache, increase the value of the initialization parameter `SHARED_POOL_SIZE`. The maximum value for this parameter depends on your operating system. This measure reduces implicit reparsing of SQL statements and PL/SQL blocks on execution.

To take advantage of additional memory available for shared SQL areas, you might also need to increase the number of cursors permitted for a session. You can do this by increasing the value of the initialization parameter `OPEN_CURSORS`.

Allocating Additional Memory to the Data Dictionary Cache Examine cache activity by monitoring the `GETS` and `GETMISSES` columns. For frequently accessed dictionary caches, the ratio of total `GETMISSES` to total `GETS` should be less than 10% or 15%, depending on the application.

Consider increasing the amount of memory available to the cache if all of the following are true:

- Your application is using the shared pool effectively. See "[Using the Shared Pool Effectively](#)" on page 14-23.
- Your system has reached a steady state, any of the item-specific hit ratios are low, and there are a large numbers of gets for the caches with low hit ratios.

Increase the amount of memory available to the data dictionary cache by increasing the value of the initialization parameter `SHARED_POOL_SIZE`.

Reducing Memory Allocation

If your `RELOADS` are near zero, and if you have a small amount of free memory in the shared pool, then the shared pool is probably large enough to hold the most frequently accessed data.

If you always have significant amounts of memory free in the shared pool, and if you would like to allocate this memory elsewhere, then you might be able to reduce the shared pool size and still maintain good performance.

To make the shared pool smaller, reduce the size of the cache by changing the value for the parameter `SHARED_POOL_SIZE`.

Using the Large Pool

Unlike the shared pool, the large pool does not have an LRU list. Oracle does not attempt to age objects out of the large pool.

You should consider configuring a large pool if your instance uses any of the following:

- Parallel query

Parallel query uses shared pool memory to cache parallel execution message buffers.

See Also: *Oracle9i Data Warehousing Guide* for more information on sizing the large pool with parallel query

- Recovery Manager

Recovery Manager uses the shared pool to cache I/O buffers during backup and restore operations. For I/O server processes and backup and restore operations, Oracle allocates buffers that are a few hundred kilobytes in size.

See Also: *Oracle9i Recovery Manager User's Guide* for more information on sizing the large pool when using Recovery Manager

- Shared server

In a shared server architecture, the session memory for each client process is included in the shared pool.

Tuning the Large Pool and Shared Pool for the Shared Server Architecture

As Oracle allocates shared pool memory for shared server session memory, the amount of shared pool memory available for the library cache and dictionary cache decreases. If you allocate this session memory from a different pool, then Oracle can use the shared pool primarily for caching shared SQL and not incur the performance overhead from shrinking the shared SQL cache.

Oracle recommends using the large pool to allocate the shared server-related User Global Area (UGA), rather than using the shared pool. This is because Oracle uses the shared pool to allocate Shared Global Area (SGA) memory for other purposes,

such as shared SQL and PL/SQL procedures. Using the large pool instead of the shared pool decreases fragmentation of the shared pool.

To store shared server-related UGA in the large pool, specify a value for the initialization parameter `LARGE_POOL_SIZE`. To see which pool (shared pool or large pool) the memory for an object resides in, check the column `POOL` in `V$SGASTAT`. The large pool is not configured by default; its minimum value is 300K. If you do not configure the large pool, then Oracle uses the shared pool for shared server user session memory.

Configure the size of the large pool based on the number of simultaneously active sessions. Each application requires a different amount of memory for session information, and your configuration of the large pool or SGA should reflect the memory requirement. For example, assuming that the shared server requires 200K to 300K to store session information for each active session. If you anticipate 100 active sessions simultaneously, then configure the large pool to be 30M, or increase the shared pool accordingly if the large pool is not configured.

Note: If a shared server architecture is used, then Oracle allocates some fixed amount of memory (about 10K) for each configured session from the shared pool, even if you have configured the large pool. The `CIRCUITS` initialization parameter specifies the maximum number of concurrent shared server connections that the database allows.

See Also:

- *Oracle9i Database Concepts* for more information about the large pool
- *Oracle9i Database Reference* for complete information about initialization parameters

Determining an Effective Setting for Shared Server UGA Storage The exact amount of UGA Oracle uses depends on each application. To determine an effective setting for the large or shared pools, observe UGA use for a typical user and multiply this amount by the estimated number of user sessions.

Even though use of shared memory increases with shared servers, the total amount of memory use decreases. This is because there are fewer processes; therefore, Oracle uses less PGA memory with shared servers when compared to dedicated server environments.

Note: For best performance with sorts using shared servers, set `SORT_AREA_SIZE` and `SORT_AREA_RETAINED_SIZE` to the same value. This keeps the sort result in the large pool instead of having it written to disk.

Checking System Statistics in the V\$SESSTAT View Oracle collects statistics on total memory used by a session and stores them in the dynamic performance view `V$SESSTAT`. [Table 14-3](#) lists these statistics.

Table 14-3 V\$SESSTAT Statistics Reflecting Memory

Statistic	Description
session UGA memory	The value of this statistic is the amount of memory in bytes allocated to the session.
Session UGA memory max	The value of this statistic is the maximum amount of memory in bytes ever allocated to the session.

To find the value, query `V$STATNAME`. If you are using a shared server, you can use the following query to decide how much larger to make the shared pool. Issue the following queries while your application is running:

```
SELECT SUM(VALUE) || ' BYTES' "TOTAL MEMORY FOR ALL SESSIONS"
  FROM V$SESSTAT, V$STATNAME
 WHERE NAME = 'session uga memory'
 AND V$SESSTAT.STATISTIC# = V$STATNAME.STATISTIC#;
```

```
SELECT SUM(VALUE) || ' BYTES' "TOTAL MAX MEM FOR ALL SESSIONS"
  FROM V$SESSTAT, V$STATNAME
 WHERE NAME = 'session uga memory max'
 AND V$SESSTAT.STATISTIC# = V$STATNAME.STATISTIC#;
```

These queries also select from the dynamic performance view `V$STATNAME` to obtain internal identifiers for `session memory` and `max session memory`. The results of these queries could look like the following:

```
TOTAL MEMORY FOR ALL SESSIONS
-----
157125 BYTES

TOTAL MAX MEM FOR ALL SESSIONS
-----
417381 BYTES
```

The result of the first query indicates that the memory currently allocated to all sessions is 157,125 bytes. This value is the total memory with a location that depends on how the sessions are connected to Oracle. If the sessions are connected to dedicated server processes, then this memory is part of the memories of the user processes. If the sessions are connected to shared server processes, then this memory is part of the shared pool.

The result of the second query indicates that the sum of the maximum sizes of the memories for all sessions is 417,381 bytes. The second result is greater than the first because some sessions have deallocated memory since allocating their maximum amounts.

If you use a shared server architecture, you can use the result of either of these queries to determine how much larger to make the shared pool. The first value is likely to be a better estimate than the second unless nearly all sessions are likely to reach their maximum allocations at the same time.

Limiting Memory Use for Each User Session by Setting PRIVATE_SGA You can set the `PRIVATE_SGA` resource limit to restrict the memory used by each client session from the SGA. `PRIVATE_SGA` defines the number of bytes of memory used from the SGA by a session. However, this parameter is used rarely, because most DBAs do not limit SGA consumption on a user-by-user basis.

See Also: *Oracle9i SQL Reference*, `ALTER RESOURCE COST` statement, for more information about setting the `PRIVATE_SGA` resource limit

Reducing Memory Use with Three-Tier Connections If you have a high number of connected users, then you can reduce memory usage by implementing three-tier connections. This by-product of using a transaction process (TP) monitor is feasible only with pure transactional models, because locks and uncommitted DMLs cannot be held between calls. A shared server environment offers the following advantages:

- It is much less restrictive of the application design than a TP monitor.
- It dramatically reduces operating system process count and context switches by enabling users to share a pool of servers.
- It substantially reduces overall memory usage, even though more SGA is used in shared server mode.

Using `CURSOR_SPACE_FOR_TIME`

If you have no library cache misses, then you might be able to accelerate execution calls by setting the value of the initialization parameter `CURSOR_SPACE_FOR_TIME` to `true`. This parameter specifies whether a cursor can be deallocated from the library cache to make room for a new SQL statement. `CURSOR_SPACE_FOR_TIME` has the following values meanings:

- If `CURSOR_SPACE_FOR_TIME` is set to `false` (the default), then a cursor can be deallocated from the library cache regardless of whether application cursors associated with its SQL statement are open. In this case, Oracle must verify that the cursor containing the SQL statement is in the library cache.
- If `CURSOR_SPACE_FOR_TIME` is set to `true`, then a cursor can be deallocated only when all application cursors associated with its statement are closed. In this case, Oracle need not verify that a cursor is in the cache, because it cannot be deallocated while an application cursor associated with it is open.

Setting the value of the parameter to `true` saves Oracle a small amount of time and can slightly improve the performance of execution calls. This value also prevents the deallocation of cursors until associated application cursors are closed.

Do not set the value of `CURSOR_SPACE_FOR_TIME` to `true` if you have found library cache misses on execution calls. Such library cache misses indicate that the shared pool is not large enough to hold the shared SQL areas of all concurrently open cursors. If the value is `true`, and if the shared pool has no space for a new SQL statement, then the statement cannot be parsed, and Oracle returns an error saying that there is no more shared memory. If the value is `false`, and if there is no space for a new statement, then Oracle deallocates an existing cursor. Although deallocating a cursor could result in a library cache miss later (only if the cursor is reexecuted), it is preferable to an error halting your application because a SQL statement cannot be parsed.

Do not set the value of `CURSOR_SPACE_FOR_TIME` to `true` if the amount of memory available to each user for private SQL areas is scarce. This value also prevents the deallocation of private SQL areas associated with open cursors. If the private SQL areas for all concurrently open cursors fills your available memory so that there is no space for a new SQL statement, then the statement cannot be parsed. Oracle returns an error indicating that there is not enough memory.

Caching Session Cursors

If an application repeatedly issues parse calls on the same set of SQL statements, then the reopening of the session cursors can affect system performance. Session

cursors can be stored in a session cursor cache. This feature can be particularly useful for applications that use Oracle Forms, because switching from one form to another closes all session cursors associated with the first form.

Oracle checks the library cache to determine whether more than three parse requests have been issued on a given statement. If so, then Oracle assumes that the session cursor associated with the statement should be cached and moves the cursor into the session cursor cache. Subsequent requests to parse that SQL statement by the same session then find the cursor in the session cursor cache.

To enable caching of session cursors, you must set the initialization parameter `SESSION_CACHED_CURSORS`. The value of this parameter is a positive integer specifying the maximum number of session cursors kept in the cache. An LRU algorithm removes entries in the session cursor cache to make room for new entries when needed.

You can also enable the session cursor cache dynamically with the statement:

```
ALTER SESSION SET SESSION_CACHED_CURSORS = value;
```

To determine whether the session cursor cache is sufficiently large for your instance, you can examine the session statistic `session cursor cache hits` in the `V$SYSSTAT` view. This statistic counts the number of times a parse call found a cursor in the session cursor cache. If this statistic is a relatively low percentage of the total parse call count for the session, then consider setting `SESSION_CACHED_CURSORS` to a larger value.

Configuring the Reserved Pool

Although Oracle breaks down very large requests for memory into smaller chunks, on some systems there might still be a requirement to find a contiguous chunk (for example, over 5 KB) of memory. (The default minimum reserved pool allocation is 4,400 bytes.)

If there is not enough free space in the shared pool, then Oracle must search for and free enough memory to satisfy this request. This operation could conceivably hold the latch resource for detectable periods of time, causing minor disruption to other concurrent attempts at memory allocation.

Hence, Oracle internally reserves a small memory area in the shared pool that can be used if the shared pool does not have enough space. This reserved pool makes allocation of large chunks more efficient.

By default, Oracle configures a small reserved pool. This memory can be used for operations such as PL/SQL and trigger compilation or for temporary space while

loading Java objects. After the memory allocated from the reserved pool is freed, it returns to the reserved pool.

You probably will not need to change the default amount of space Oracle reserves. However, if necessary, the reserved pool size can be changed by setting the `SHARED_POOL_RESERVED_SIZE` initialization parameter. This parameter sets aside space in the shared pool for unusually large allocations.

For large allocations, Oracle attempts to allocate space in the shared pool in the following order:

1. From the unreserved part of the shared pool.
2. From the reserved pool. If there is not enough space in the unreserved part of the shared pool, then Oracle checks whether the reserved pool has enough space.
3. From memory. If there is not enough space in the unreserved and reserved parts of the shared pool, then Oracle attempts to free enough memory for the allocation. It then retries the unreserved and reserved parts of the shared pool.

Using `SHARED_POOL_RESERVED_SIZE` The default value for `SHARED_POOL_RESERVED_SIZE` is 5% of the `SHARED_POOL_SIZE`. This means that, by default, the reserved list is configured.

If you set `SHARED_POOL_RESERVED_SIZE` to more than half of `SHARED_POOL_SIZE`, then Oracle signals an error. Oracle does not let you reserve too much memory for the reserved pool. The amount of operating system memory, however, might constrain the size of the shared pool. In general, set `SHARED_POOL_RESERVED_SIZE` to 10% of `SHARED_POOL_SIZE`. For most systems, this value is sufficient if you have already tuned the shared pool. If you increase this value, then the database takes memory from the shared pool. (This reduces the amount of unreserved shared pool memory available for smaller allocations.)

Statistics from the `V$SHARED_POOL_RESERVED` view help you tune these parameters. On a system with ample free memory to increase the size of the SGA, the goal is to have the value of `REQUEST_MISSES` equal zero. If the system is constrained for operating system memory, then the goal is to not have `REQUEST_FAILURES` or at least prevent this value from increasing.

If you cannot achieve these target values, then increase the value for `SHARED_POOL_RESERVED_SIZE`. Also, increase the value for `SHARED_POOL_SIZE` by the same amount, because the reserved list is taken from the shared pool.

See Also: *Oracle9i Database Reference* for details on setting the `LARGE_POOL_SIZE` parameter

When `SHARED_POOL_RESERVED_SIZE` Is Too Small The reserved pool is too small when the value for `REQUEST_FAILURES` is more than zero and increasing. To resolve this, increase the value for the `SHARED_POOL_RESERVED_SIZE` and `SHARED_POOL_SIZE` accordingly. The settings you select for these parameters depend on your system's SGA size constraints.

Increasing the value of `SHARED_POOL_RESERVED_SIZE` increases the amount of memory available on the reserved list without having an effect on users who do not allocate memory from the reserved list.

When `SHARED_POOL_RESERVED_SIZE` Is Too Large Too much memory might have been allocated to the reserved list if:

- `REQUEST_MISSES` is zero or not increasing
- `FREE_MEMORY` is greater than or equal to 50% of `SHARED_POOL_RESERVED_SIZE` minimum

If either of these conditions is true, then decrease the value for `SHARED_POOL_RESERVED_SIZE`.

When `SHARED_POOL_SIZE` is Too Small The `V$SHARED_POOL_RESERVED` fixed view can also indicate when the value for `SHARED_POOL_SIZE` is too small. This can be the case if `REQUEST_FAILURES` is greater than zero and increasing.

If you have enabled the reserved list, then decrease the value for `SHARED_POOL_RESERVED_SIZE`. If you have not enabled the reserved list, then you could increase `SHARED_POOL_SIZE`.

Keeping Large Objects to Prevent Aging

After an entry has been loaded into the shared pool, it cannot be moved. Sometimes, as entries are loaded and aged, the free memory can become fragmented.

Use the PL/SQL package `DBMS_SHARED_POOL` to manage the shared pool. Shared SQL and PL/SQL areas age out of the shared pool according to a least recently used (LRU) algorithm, similar to database buffers. To improve performance and prevent reparsing, you might want to prevent large SQL or PL/SQL areas from aging out of the shared pool.

The `DBMS_SHARED_POOL` package lets you keep objects in shared memory, so that they do not age out with the normal LRU mechanism. By using the `DBMS_SHARED_POOL` package and by loading the SQL and PL/SQL areas before memory fragmentation occurs, the objects can be kept in memory. This ensures that memory is available, and it prevents the sudden, inexplicable slowdowns in user response time that occur when SQL and PL/SQL areas are accessed after aging out.

The `DBMS_SHARED_POOL` package is useful for the following:

- When loading large PL/SQL objects, such as the `STANDARD` and `DIUTIL` packages. When large PL/SQL objects are loaded, user response time may be affected if smaller objects that need to age out of the shared pool to make room. In some cases, there might be insufficient memory to load the large objects.
- Frequently executed triggers. You might want to keep compiled triggers on frequently used tables in the shared pool.
- `DBMS_SHARED_POOL` supports sequences. Sequence numbers are lost when a sequence ages out of the shared pool. `DBMS_SHARED_POOL` keeps sequences in the shared pool, thus preventing the loss of sequence numbers.

To use the `DBMS_SHARED_POOL` package to pin a SQL or PL/SQL area, complete the following steps:

1. Decide which packages or cursors to pin in memory.
2. Start up the database.
3. Make the call to `DBMS_SHARED_POOL.KEEP` to pin your objects.

This procedure ensures that your system does not run out of shared memory before the kept objects are loaded. By pinning the objects early in the life of the instance, you prevent memory fragmentation that could result from pinning a large portion of memory in the middle of the shared pool.

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for specific information on using `DBMS_SHARED_POOL` procedures

CURSOR_SHARING for Existing Applications

One of the first stages of parsing is to compare the text of the statement with existing statements in the shared pool to see if the statement can be shared. If the statement differs textually in any way, then Oracle does not share the statement.

Exceptions to this are possible when the parameter `CURSOR_SHARING` has been set to `SIMILAR` or `FORCE`. When this parameter is used, Oracle first checks the shared

pool to see if there is an identical statement in the shared pool. If an identical statement is not found, then Oracle searches for a similar statement in the shared pool. If the similar statement is there, then the parse checks continue to verify the executable form of the cursor can be used. If the statement is not there, then a hard parse is necessary to generate the executable form of the statement.

Similar SQL Statements

Statements that are identical, except for the values of some literals, are called similar statements. Similar statements pass the textual check in the parse phase when the `CURSOR_SHARING` parameter is set to `SIMILAR` or `FORCE`. Textual similarity does not guarantee sharing. The new form of the SQL statement still needs to go through the remaining steps of the parse phase to ensure that the execution plan of the preexisting statement is equally applicable to the new statement.

See Also: ["SQL Sharing Criteria"](#) on page 14-21 for more details on the various checks performed

CURSOR_SHARING

Setting `CURSOR_SHARING` to `EXACT` allows SQL statements to share the SQL area only when their texts match exactly. This is the default behavior. Using this setting, similar statements cannot be shared; only textually exact statements can be shared.

Setting `CURSOR_SHARING` to either `SIMILAR` or `FORCE` allows similar statements to share SQL. The difference between `SIMILAR` and `FORCE` is that `SIMILAR` forces similar statements to share the SQL area without deteriorating execution plans. Setting `CURSOR_SHARING` to `FORCE` forces similar statements to share the executable SQL area, potentially deteriorating execution plans. Hence, `FORCE` should be used as a last resort, when the risk of suboptimal plans is outweighed by the improvements in cursor sharing.

When to use CURSOR_SHARING

The `CURSOR_SHARING` initialization parameter can solve some performance problems. It has the following values: `FORCE`, `SIMILAR`, and `EXACT` (default). Using this parameter provides benefit to existing applications that have many similar SQL statements.

Note: Oracle does not recommend setting `CURSOR_SHARING` to `FORCE` in a DSS environment or if you are using complex queries. Also, star transformation is not supported with `CURSOR_SHARING` set to either `SIMILAR` or `FORCE`. For more information, see the "[OPTIMIZER_FEATURES_ENABLE Parameter](#)" on page 1-56.

The optimal solution is to write sharable SQL, rather than rely on the `CURSOR_SHARING` parameter. This is because although `CURSOR_SHARING` does significantly reduce the amount of resources used by eliminating hard parses, it requires some extra work as a part of the soft parse to find a similar statement in the shared pool.

Note: Setting `CURSOR_SHARING` to `SIMILAR` or `FORCE` causes an increase in the maximum lengths (as returned by `DESCRIBE`) of any selected expressions that contain literals (in a `SELECT` statement). However, the actual length of the data returned does not change.

Consider setting `CURSOR_SHARING` to `SIMILAR` or `FORCE` if you can answer 'yes' to both of the following questions:

1. Are there statements in the shared pool that differ only in the values of literals?
2. Is the response time low due to a very high number of library cache misses?

Caution: Setting `CURSOR_SHARING` to `FORCE` or `SIMILAR` prevents any outlines generated with literals from being used if they were generated with `CURSOR_SHARING` set to `EXACT`.

To use stored outlines with `CURSOR_SHARING=FORCE` or `SIMILAR`, the outlines must be generated with `CURSOR_SHARING` set to `FORCE` or `SIMILAR` and with the `CREATE_STORED_OUTLINES` parameter.

Using `CURSOR_SHARING = SIMILAR` (or `FORCE`) can significantly improve cursor sharing on some applications that have many similar statements, resulting in reduced memory usage, faster parses, and reduced latch contention.

Configuring and Using the Java Pool

If your application uses Java, you should investigate whether you need to modify the default configuration for the Java pool.

See Also: *Oracle9i Java Developer's Guide*

Configuring and Using the Redo Log Buffer

Server processes making changes to data blocks in the buffer cache generate redo data into the log buffer. LGWR begins writing to copy entries from the redo log buffer to the online redo log if any of the following are true:

- The log buffer becomes one third full.
- LGWR is posted by a server process performing a COMMIT or ROLLBACK.
- DBWR posts LGWR to do so.

When LGWR writes redo entries from the redo log buffer to a redo log file or disk, user processes can then copy new entries over the entries in memory that have been written to disk. LGWR usually writes fast enough to ensure that space is available in the buffer for new entries, even when access to the redo log is heavy.

A larger buffer makes it more likely that there is space for new entries, and also gives LGWR the opportunity to efficiently write out redo records (too small a log buffer on a system with large updates means that LGWR is continuously flushing redo to disk so that the log buffer remains 2/3 empty).

On machines with fast processors and relatively slow disks, the processors might be filling the rest of the buffer in the time it takes the redo log writer to move a portion of the buffer to disk. A larger log buffer can temporarily mask the effect of slower disks in this situation. Alternatively, you can do one of the following:


- Improve the checkpointing or archiving process
- Improve the performance of log writer (perhaps by moving all online logs to fast raw devices)

Good usage of the redo log buffer is a simple matter of:

- Batching commit operations for batch jobs, so that log writer is able to write redo log entries efficiently
- Using NOLOGGING operations when you are loading large quantities of data

The size of the redo log buffer is determined by the initialization parameter `LOG_BUFFER`. The log buffer size cannot be modified after instance startup.

Figure 14–2 Redo Log Buffer



Sizing the Log Buffer

Applications that insert, modify, or delete large volumes of data usually need to change the default log buffer size. The log buffer is small compared with the total SGA size, and a modestly sized log buffer can significantly enhance throughput on systems that perform many updates.

A reasonable first estimate for such systems is to make the log buffer 1 MB. On most systems, sizing the log buffer larger than 1m does not provide any performance benefit. Increasing the log buffer size does not have any negative implications on performance or recoverability. It merely uses extra memory.

Log Buffer Statistics

The statistic `REDO BUFFER ALLOCATION RETRIES` reflects the number of times a user process waits for space in the redo log buffer. This statistic can be queried through the dynamic performance view `V$SYSSTAT`.

Use the following query to monitor these statistics over a period of time while your application is running:

```
SELECT NAME, VALUE
   FROM V$SYSSTAT
  WHERE NAME = 'redo buffer allocation retries';
```

The value of `redo buffer allocation retries` should be near zero over an interval. If this value increments consistently, then processes have had to wait for space in the redo log buffer. The wait can be caused by the log buffer being too small or by checkpointing. Increase the size of the redo log buffer, if necessary, by changing the value of the initialization parameter `LOG_BUFFER`. The value of this parameter is expressed in bytes. Alternatively, improve the checkpointing or archiving process.

Another data source is to check whether the `log buffer space wait` event is not a significant factor in the wait time for the instance; if not, the log buffer size is most likely adequate.

Configuring the PGA Working Memory

The Program Global Area (PGA) is a private memory region containing data and control information for a server process. Access to it is exclusive to that server process and is read and written only by the Oracle code acting on behalf of it. An example of such information is the runtime area of a cursor. Each time a cursor is executed, a new runtime area is created for that cursor in the PGA memory region of the server process executing that cursor.

Note: Part of the runtime area can be located in the SGA when using shared servers.

For complex queries (for example, decision support queries), a big portion of the runtime area is dedicated to work areas allocated by memory intensive operators, such as the following:

- Sort-based operators (for example, `ORDER BY`, `GROUP BY`, `ROLLUP`, window functions)
- Hash-join
- Bitmap merge
- Bitmap create
- Write buffers used by bulk load operations

A sort operator uses a work area (the sort area) to perform the in-memory sort of a set of rows. Similarly, a hash-join operator uses a work area (the hash area) to build a hash table from its left input.

The size of a work area can be controlled and tuned. Generally, bigger work areas can significantly improve the performance of a particular operator at the cost of higher memory consumption. Ideally, the size of a work area is big enough that it can accommodate the input data and auxiliary memory structures allocated by its associated SQL operator. This is known as the optimal size of a work area. When the size of the work area is smaller than optimal, the response time increases, because an extra pass is performed over part of the input data. This is known as the one-pass size of the work area. Under the one-pass threshold, when the size of a work area is far too small compared to the input data size, multiple passes over the input data are needed. This could dramatically increase the response time of the operator. This is known as the multi-pass size of the work area. For example, a serial sort operation that needs to sort 10GB of data needs a little more than 10GB to run optimal and at least 40MB to run one-pass. If this sort gets less than 40MB, then it must perform several passes over the input data.

The goal is to have most work areas running with an optimal size (for example, more than 90% or even 100% for pure OLTP systems), while a smaller fraction of them are running with a one-pass size (for example, less than 10%). Multi-pass execution should be avoided. Even for DSS systems running large sorts and hash-joins, the memory requirement for the one-pass executions is relatively small. A system configured with a reasonable amount of PGA memory should not need to perform multiple passes over the input data.

Prior to Oracle9i, the maximum size of these working areas was controlled using the `SORT_AREA_SIZE`, `HASH_AREA_SIZE`, `BITMAP_MERGE_AREA_SIZE` and `CREATE_BITMAP_AREA_SIZE` parameters. Setting these parameters is difficult, because the maximum work area size is ideally selected based on the data input size and the total number of work areas active in the system. These two factors vary a lot from one work area to another and from one point in time to another. Thus, the various *_AREA_SIZE parameters are hard to tune under the best of circumstances.

With Oracle9i, you can simplify and improve the way PGA memory is allocated, by enabling automatic PGA memory management. In this mode, Oracle dynamically adjusts the size of the portion of the PGA memory dedicated to work areas, based on an overall PGA memory target explicitly set by the DBA. To enable automatic PGA memory management, you have to set the initialization parameter `PGA_AGGREGATE_TARGET`, as described in the following section.

Note: This mechanism cannot be used for shared server connections.

Automatic PGA Memory Management

When running under the automatic PGA memory management mode, sizing of work areas for all dedicated sessions becomes automatic. Thus, the `*_AREA_SIZE` parameters are ignored by all sessions running in that mode. At any given time, the total amount of PGA memory available to active work areas in the instance is automatically derived from the `PGA_AGGREGATE_TARGET` initialization parameter. This amount is set to the value of `PGA_AGGREGATE_TARGET` minus the amount of PGA memory allocated by other components of the system (for example, PGA memory allocated by sessions). The resulting PGA memory is then assigned to individual active work areas, based on their specific memory requirements.

Under automatic PGA memory management mode, the main goal of Oracle is to honor the `PGA_AGGREGATE_TARGET` limit set by the DBA, by controlling dynamically the amount of PGA memory allotted to SQL work areas. At the same time, Oracle tries to maximize the performance of all the memory-intensive SQL operators, by maximizing the number of work areas that are using an optimal amount of PGA memory (cache memory). The rest of the work areas are executed in one-pass mode, unless the PGA memory limit set by the DBA with the parameter `PGA_AGGREGATE_TARGET` is so low that multi-pass execution is required to reduce even more the consumption of PGA memory and honor the PGA target limit.

When configuring a brand new instance, it is hard to know precisely the appropriate setting for `PGA_AGGREGATE_TARGET`. You can determine this setting in three stages:

1. Make a first estimate for `PGA_AGGREGATE_TARGET`, based on a rule of thumb.
2. Run a representative workload on the instance and monitor performance, using PGA statistics collected by Oracle, to see whether the maximum PGA size is under-configured or over-configured.
3. Tune `PGA_AGGREGATE_TARGET`, using Oracle's PGA advice statistics.

The following sections explain this in detail:

- [Setting `PGA_AGGREGATE_TARGET` Initially](#)
- [Monitoring the Performance of the Automatic PGA Memory Management](#)
- [Tuning `PGA_AGGREGATE_TARGET`](#)

Setting `PGA_AGGREGATE_TARGET` Initially

The value of the `PGA_AGGREGATE_TARGET` initialization parameter (for example 100000 KB, 2500 MB, or 50 GB) should be set based on the total amount of memory

available for the Oracle instance. This value can then be tuned and dynamically modified at the instance level. [Example 14–2](#) illustrates a typical situation.

Example 14–2 Initial Setting of PGA_AGGREGATE_TARGET

Assume that an Oracle instance is configured to run on a system with 4 GB of physical memory. Part of that memory should be left for the operating system and other non-Oracle applications running on the same hardware system. You might decide to dedicate only 80% of the available memory to the Oracle instance, or 3.2 GB.

You must then divide the resulting memory between the SGA and the PGA.

- For OLTP systems, the PGA memory typically accounts for a small fraction of the total memory available (for example, 20%), leaving 80% for the SGA.
- For DSS systems running large, memory-intensive queries, PGA memory can typically use up to 70% of that total (up to 2.2 GB in this example).

Good initial values for the parameter `PGA_AGGREGATE_TARGET` might be:

- For OLTP: $PGA_AGGREGATE_TARGET = (total_mem * 80\%) * 20\%$
- For DSS: $PGA_AGGREGATE_TARGET = (total_mem * 80\%) * 50\%$

where *total_mem* is the total amount of physical memory available on the system.

In this example, with a value of *total_mem* equal to 4 GB, you can initially set `PGA_AGGREGATE_TARGET` to 1600 MB for the DSS system and to 655 MB for the OLTP system.

Monitoring the Performance of the Automatic PGA Memory Management

Before starting the tuning process, you need to know how to monitor and interpret the key statistics collected by Oracle to help in assessing the performance of the automatic PGA memory management component. Several dynamic performance views are available for this purpose:

- [V\\$PGASTAT](#)
- [V\\$PROCESS](#)
- [V\\$SQL_WORKAREA_HISTOGRAM](#)
- [V\\$SQL_WORKAREA_ACTIVE](#)
- [V\\$SQL_WORKAREA](#)

V\$PGASTAT This view gives instance-level statistics on the PGA memory usage and the automatic PGA memory manager. For example:

```
SELECT * FROM V$PGASTAT;
```

The output of this query might look like the following:

NAME	VALUE	UNIT
aggregate PGA target parameter	524288000	bytes
aggregate PGA auto target	463435776	bytes
global memory bound	25600	bytes
total PGA inuse	9353216	bytes
total PGA allocated	73516032	bytes
maximum PGA allocated	698371072	bytes
total PGA used for auto workareas	0	bytes
maximum PGA used for auto workareas	560744448	bytes
total PGA used for manual workareas	0	bytes
maximum PGA used for manual workareas	0	bytes
over allocation count	0	
total bytes processed	4.0072E+10	bytes
total extra bytes read/written	3.1517E+10	bytes
cache hit percentage	55.97	percent

The main statistics displayed in V\$PGASTAT are as follows:

- aggregate PGA target parameter: This is the current value of the initialization parameter `PGA_AGGREGATE_TARGET`, here set to 500 MB. If you do not set this parameter, its value is 0 and automatic management of the PGA memory is disabled.
- aggregate PGA auto target: This gives the amount of PGA memory Oracle can use for work areas running in automatic mode. This amount is dynamically derived from the value of the parameter `PGA_AGGREGATE_TARGET` and the current work area workload. Hence, it is continuously adjusted by Oracle. If this value is small compared to the value of `PGA_AGGREGATE_TARGET`, then a lot of PGA memory is used by other components of the system (for example, PL/SQL or Java memory) and little is left for sort work areas. You must ensure that enough PGA memory is left for work areas running in automatic mode.
- global memory bound: This gives the maximum size of a work area executed in `AUTO` mode. This value is continuously adjusted by Oracle to reflect the current state of the work area workload. The global memory bound generally decreases when the number of active work areas is increasing in the

system. As a rule of thumb, the value of the global bound should not decrease to less than one megabyte. If it does, then the value of `PGA_AGGREGATE_TARGET` should probably be increased.

- `total PGA allocated`: This gives the current amount of PGA memory allocated by the instance. Oracle tries to keep this number less than the value of `PGA_AGGREGATE_TARGET`. However, it is possible for the PGA allocated to exceed that value by a small percentage and for a short period of time, when the work area workload is increasing very rapidly or when the initialization parameter `PGA_AGGREGATE_TARGET` is set to a too small value.
- `total PGA used for auto workareas`: This indicates how much PGA memory is currently consumed by work areas running under automatic memory management mode. This number can be used to determine how much memory is consumed by other consumers of the PGA memory (for example, PL/SQL or Java):

`PGA other = total PGA allocated - total PGA used for auto workareas`

- `over allocation count`: This statistic is cumulative from instance start-up. Over-allocating PGA memory can happen if the value of `PGA_AGGREGATE_TARGET` is too small to accommodate the `PGA other` component in the previous equation plus the minimum memory required to execute the work area workload. When this happens, Oracle cannot honor the initialization parameter `PGA_AGGREGATE_TARGET`, and extra PGA memory needs to be allocated. If over-allocation occurs, you should increase the value of `PGA_AGGREGATE_TARGET` using the information provided by the advice view `V$PGA_TARGET_ADVICE`.
- `total bytes processed`: This is the number of bytes processed by memory-intensive SQL operators since instance start-up. For example, the number of byte processed is the input size for a sort operation. This number is used to compute the cache hit percentage metric.
- `extra bytes read/written`: When a work area cannot run optimally, one or more extra passes is performed over the input data. `extra bytes read/written` represents the number of bytes processed during these extra passes since instance start-up. This number is also used to compute the cache hit percentage.
- `cache hit percentage`: This metric is computed by Oracle to reflect the performance of the PGA memory component. It is cumulative from instance start-up. A value of 100% means that all work areas executed by the system since instance start-up have used an optimal amount of PGA memory. This is,

of course, ideal but rarely happens except maybe for pure OLTP systems. In reality, some work areas run one-pass or even multi-pass, depending on the overall size of the PGA memory. When a work area cannot run optimally, one or more extra passes is performed over the input data. This reduces the cache hit percentage in proportion to the size of the input data and the number of extra passes performed. [Example 14-3](#) shows how cache hit percentage is affected by extra passes.

Example 14-3 Calculating Cache Hit Percentage

Consider a simple example: Four sort operations have been executed, three were small (1 MB of input data) and one was bigger (100 MB of input data). The total number of bytes processed (BP) by the four operations is 103 MB. If one of the small sorts runs one-pass, an extra pass over 1 MB of input data is performed. This 1 MB value is the number of extra bytes read/written, or EBP. The cache hit percentage is calculated by the following formula:

$$BP \times 100 / (BP + EBP)$$

The cache hit percentage in this case is 99.03%, almost 100%. This value reflects the fact that only one of the small sorts had to perform an extra pass while all other sorts were able to run optimally. Hence, the cache hit percentage is almost 100%, because this extra pass over 1 MB represents a tiny overhead. On the other hand, if the big sort is the one to run one-pass, then EBP is 100 MB instead of 1 MB, and the cache hit percentage falls to 50.73%, because the extra pass has a much bigger impact.

V\$PROCESS This view has one row for each Oracle process connected to the instance. The columns `PGA_USED_MEM`, `PGA_ALLOC_MEM` and `PGA_MAX_MEM` can be used to monitor the PGA memory usage of these processes. For example:

```
SELECT PROGRAM, PGA_USED_MEM, PGA_ALLOC_MEM, PGA_MAX_MEM
FROM V$PROCESS;
```

The output of this query might look like the following:

PROGRAM	PGA_USED_MEM	PGA_ALLOC_MEM	PGA_MAX_MEM
PSEUDO	0	0	0
oracle@miflo (PMON)	120463	234291	234291
oracle@miflo (DBW0)	1307179	1817295	1817295
oracle@miflo (LGWR)	4343655	4849203	4849203
oracle@miflo (CKPT)	194999	332583	332583
oracle@miflo (SMON)	179923	775311	775323

oracle@miflo (RECO)	129719	242803	242803
oracle@miflo (TNS V1-V3)	1400543	1540627	1540915
oracle@miflo (P000)	299599	373791	635959
oracle@miflo (P001)	299599	373791	636007
oracle@miflo (P002)	299599	373791	570471
oracle@miflo (P003)	303899	373791	636007
oracle@miflo (P004)	299599	373791	635959

V\$SQL_WORKAREA_HISTOGRAM This view shows the number of work areas executed with optimal memory size, one-pass memory size, and multi-pass memory size since instance start-up. Statistics in this view are subdivided into buckets that are defined by the optimal memory requirement of the work area. Each bucket is identified by a range of optimal memory requirements specified by the values of the columns `LOW_OPTIMAL_SIZE` and `HIGH_OPTIMAL_SIZE`.

[Examples 14-4](#) and [14-5](#) show two ways of using `V$SQL_WORKAREA_HISTOGRAM`.

Example 14-4 Querying V\$SQL_WORKAREA_HISTOGRAM: Nonempty Buckets

Consider a sort operation that requires 3 MB of memory to run optimally (cached). Statistics about the work area used by this sort are placed in the bucket defined by `LOW_OPTIMAL_SIZE = 2097152` (2 MB) and `HIGH_OPTIMAL_SIZE = 4194303` (4 MB minus 1 byte), because 3 MB falls within that range of optimal sizes. Statistics are segmented by work area size, because the performance impact of running a work area in optimal, one-pass or multi-pass mode depends mainly on the size of that work area.

The following query shows statistics for all nonempty buckets. Empty buckets are removed with the predicate `where total_execution != 0`.

```
SELECT LOW_OPTIMAL_SIZE/1024 low_kb,
       (HIGH_OPTIMAL_SIZE+1)/1024 high_kb,
       OPTIMAL_EXECUTIONS, ONEPASS_EXECUTIONS, MULTIPASSES_EXECUTIONS
FROM V$SQL_WORKAREA_HISTOGRAM
WHERE TOTAL_EXECUTIONS != 0;
```

The result of the query might look like the following:

LOW_KB	HIGH_KB	OPTIMAL_EXECUTIONS	ONEPASS_EXECUTIONS	MULTIPASSES_EXECUTIONS
8	16	156255	0	0
16	32	150	0	0
32	64	89	0	0
64	128	13	0	0

128	256	60	0	0
256	512	8	0	0
512	1024	657	0	0
1024	2048	551	16	0
2048	4096	538	26	0
4096	8192	243	28	0
8192	16384	137	35	0
16384	32768	45	107	0
32768	65536	0	153	0
65536	131072	0	73	0
131072	262144	0	44	0
262144	524288	0	22	0

The query result shows that, in the 1024 KB to 2048 KB bucket, 551 work areas used an optimal amount of memory, while 16 ran in one-pass mode and none ran in multi-pass mode. It also shows that all work areas under 1 MB were able to run in optimal mode.

Example 14–5 Querying V\$SQL_WORKAREA_HISTOGRAM: Percent Optimal

You can also use V\$SQL_WORKAREA_HISTOGRAM to find the percentage of times work areas were executed in optimal, one-pass, or multi-pass mode since start-up. This query only considers work areas of a certain size, with an optimal memory requirement of at least 64 KB.

```
SELECT optimal_count, round(optimal_count*100/total, 2) optimal_perc,
       onepass_count, round(onepass_count*100/total, 2) onepass_perc,
       multipass_count, round(multipass_count*100/total, 2) multipass_perc
FROM
  (SELECT decode(sum(total_executions), 0, 1, sum(total_executions)) total,
         sum(OPTIMAL_EXECUTIONS) optimal_count,
         sum(ONEPASS_EXECUTIONS) onepass_count,
         sum(MULTIPASSES_EXECUTIONS) multipass_count
   FROM v$sql_workarea_histogram
   WHERE low_optimal_size > 64*1024);
```

The output of this query might look like the following:

OPTIMAL_COUNT	OPTIMAL_PERC	ONEPASS_COUNT	ONEPASS_PERC	MULTIPASS_COUNT	MULTIPASS_PERC
2239	81.63	504	18.37	0	0

This result shows that 81.63% of these work areas have been able to run using an optimal amount of memory. The rest (18.37%) ran one-pass. None of them ran multi-pass. Such behavior is preferable, for the following reasons:

- Multi-pass mode can severely degrade performance. A high number of multi-pass work areas has an exponentially adverse effect on the response time of its associated SQL operator.
- Running one-pass does not require a large amount of memory; only 22 MB is required to sort 1 GB of data in one-pass mode.

V\$SQL_WORKAREA_ACTIVE This view can be used to display the work areas that are active (or executing) in the instance. Small active sorts (under 64 KB) are excluded from the view. Use this view to precisely monitor the size of all active work areas and to determine if these active work areas spill to a temporary segment.

[Example 14-6](#) shows a typical query of this view:

Example 14-6 Querying V\$SQL_WORKAREA_ACTIVE

```
SELECT to_number(decode(SID, 65535, NULL, SID)) sid,
       operation_type OPERATION,
       trunc(EXPECTED_SIZE/1024) ESIZE,
       trunc(ACTUAL_MEM_USED/1024) MEM,
       trunc(MAX_MEM_USED/1024) "MAX MEM",
       NUMBER_PASSES PASS,
       trunc(TEMPSEG_SIZE/1024) TSIZE
FROM V$SQL_WORKAREA_ACTIVE
ORDER BY 1,2;
```

The output of this query might look like the following:

SID	OPERATION	ESIZE	MEM	MAX MEM	PASS	TSIZE
8	GROUP BY (SORT)	315	280	904	0	
8	HASH-JOIN	2995	2377	2430	1	20000
9	GROUP BY (SORT)	34300	22688	22688	0	
11	HASH-JOIN	18044	54482	54482	0	
12	HASH-JOIN	18044	11406	21406	1	120000

This output shows that session 12 (column `SID`) is running a hash-join having its work area running in one-pass mode (`PASS` column). This work area is currently using 11406 KB of memory (`MEM` column) and has used, in the past, up to 21406 KB of PGA memory (`MAX MEM` column). It has also spilled to a temporary segment of size 120000 KB. Finally, the column `ESIZE` indicates the maximum amount of memory that the PGA memory manager expects this hash-join to use. This maximum is dynamically computed by the PGA memory manager according to workload.

When a work area is deallocated—that is, when the execution of its associated SQL operator is complete—the work area is automatically removed from the `V$SQL_WORKAREA_ACTIVE` view.

V\$SQL_WORKAREA Oracle maintains cumulative work area statistics for each loaded cursor whose execution plan uses one or more work areas. Every time a work area is deallocated, the `V$SQL_WORKAREA` table is updated with execution statistics for that work area.

`V$SQL_WORKAREA` can be joined with `V$SQL` to relate a work area to a cursor. It can even be joined to `V$SQL_PLAN` to precisely determine which operator in the plan uses a work area.

[Example 14-7](#) shows three typical queries on the `V$SQL_WORKAREA` dynamic view:

Example 14-7 Querying V\$SQL_WORKAREA

The following query finds the top 10 work areas requiring most cache memory:

```
SELECT *
FROM
  ( SELECT workarea_address, operation_type, policy, estimated_optimal_size
    FROM V$SQL_WORKAREA
    ORDER BY estimated_optimal_size )
WHERE ROWNUM <= 10;
```

The following query finds the cursors with one or more work areas that have been executed in one or even multiple passes:

```
col sql_text format A80 wrap
SELECT sql_text, sum(ONEPASS_EXECUTIONS) onepass_cnt,
       sum(MULTIPASSES_EXECUTIONS) mpass_cnt
FROM V$SQL s, V$SQL_WORKAREA wa
WHERE s.address = wa.address
GROUP BY sql_text
HAVING sum(ONEPASS_EXECUTIONS+MULTIPASSES_EXECUTIONS)>0;
```

Using the hash value and address of a particular cursor, the following query displays the cursor execution plan, including information about the associated work areas.

```
col "O/I/M" format a10
col name format a20
SELECT operation, options, object_name name,
       trunc(bytes/1024/1024) "input(MB)",
       trunc(last_memory_used/1024) last_mem,
       trunc(estimated_optimal_size/1024) optimal_mem,
```

```

trunc(estimated_onepass_size/1024) onepass_mem,
decode(optimal_executions, null, null,
       optimal_executions||'/'||onepass_executions||'/'||
       multipasses_executions) "O/I/M"
FROM V$SQL_PLAN p, V$SQL_WORKAREA w
WHERE p.address=w.address(+)
      AND p.hash_value=w.hash_value(+)
      AND p.id=w.operation_id(+)
      AND p.address='88BB460C'
      AND p.hash_value=3738161960;

```

OPERATION	OPTIONS	NAME	input(MB)	LAST_MEM	OPTIMAL_ME	ONEPASS_ME	O/I/M
SELECT STATE							
SORT	GROUP BY		4582	8	16	16	16/0/0
HASH JOIN	SEMI		4582	5976	5194	2187	16/0/0
TABLE ACCESS FULL		ORDERS	51				
TABLE ACCESS FUL		LINEITEM	1000				

You can get the address and hash value from the V\$SQL view by specifying a pattern in the query. For example:

```

SELECT address, hash_value
FROM V$SQL
WHERE sql_text LIKE '%my_pattern%';

```

Tuning PGA_AGGREGATE_TARGET

To help you tune the initialization parameter PGA_AGGREGATE_TARGET, Oracle provides two PGA advice performance views:

- [VSPGA_TARGET_ADVICE](#)
- [VSPGA_TARGET_ADVICE_HISTOGRAM](#)

By examining these two views, you no longer need to use an empirical approach to tune the value of PGA_AGGREGATE_TARGET. Instead, you can use the content of these views to determine how key PGA statistics will be impacted if you change the value of PGA_AGGREGATE_TARGET.

In both views, values of PGA_AGGREGATE_TARGET used for the prediction are derived from fractions and multiples of the current value of that parameter, to assess possible higher and lower values. Values used for the prediction range from 10 MB to a maximum of 256 GB.

Oracle generates PGA advice performance views by recording the workload history and then simulating this history for different values of `PGA_AGGREGATE_TARGET`. The simulation process happens in the background and continuously updates the workload history to produce the simulation result. You can view the result at any time by querying `V$PGA_TARGET_ADVICE` or `V$PGA_TARGET_ADVICE_HISTOGRAM`.

To enable automatic generation of PGA advice performance views, make sure the following parameters are set:

- `PGA_AGGREGATE_TARGET`, to enable automatic PGA memory management. Set the initial value as described in ["Setting PGA_AGGREGATE_TARGET Initially"](#) on page 14-50.
- `STATISTICS_LEVEL`. Set this to `TYPICAL` (the default) or `ALL`; setting this parameter to `BASIC` turns off generation of PGA performance advice views.

The content of these PGA advice performance views is reset at instance start-up or when `PGA_AGGREGATE_TARGET` is altered.

Note: Simulation cannot include all factors of real execution, so derived statistics might not exactly match up with real performance statistics. You should always monitor the system after changing `PGA_AGGREGATE_TARGET`, to verify that the new performance is what you expect.

V\$PGA_TARGET_ADVICE This view predicts how the statistics cache hit percentage and over allocation count in `V$PGASTAT` will be impacted if you change the value of the initialization parameter `PGA_AGGREGATE_TARGET`.

[Example 14-8](#) shows a typical query of this view:

Example 14-8 Querying V\$PGA_TARGET_ADVICE

```
SELECT round(PGA_TARGET_FOR_ESTIMATE/1024/1024) target_mb,
       ESTD_PGA_CACHE_HIT_PERCENTAGE cache_hit_perc,
       ESTD_OVERALLOC_COUNT
FROM V$PGA_TARGET_ADVICE;
```


The output of this query might look like the following:

TARGET_MB	CACHE_HIT_PERC	ESTD_OVERALLOC_COUNT
63	23	367
125	24	30

250	30	3
375	39	0
500	58	0
600	59	0
700	59	0
800	60	0
900	60	0
1000	61	0
1500	67	0
2000	76	0
3000	83	0
4000	85	0

The result of the this query can be plotted as shown in [Example 14-3](#):

Figure 14-3 Graphical Representation of V\$PGA_TARGET_ADVICE



The curve shows how the PGA cache hit percentage improves as the value of `PGA_AGGREGATE_TARGET` increases. The shaded zone in the graph is the over allocation zone, where the value of the column `ESTD_OVERALLOCATION_COUNT` is nonzero. It indicates that `PGA_AGGREGATE_TARGET` is too small to even meet the minimum PGA memory needs. If `PGA_AGGREGATE_TARGET` is set within the over allocation zone, the memory manager will over-allocate memory and actual PGA memory consumed will be more than the limit you set. It is therefore meaningless to set a value of `PGA_AGGREGATE_TARGET` in that zone. In this particular example `PGA_AGGREGATE_TARGET` should be set to at least 375 MB.

Note: Even though the theoretical maximum for the PGA cache hit percentage is 100%, there is a practical limit on the maximum size of a work area, which may prevent this theoretical maximum from being reached, even if you further increase `PGA_AGGREGATE_TARGET`. This should happen only in large DSS systems where the optimal memory requirement is large and might cause the value of the cache hit percentage to taper off at a lower percentage, like 90%.

Beyond the over allocation zone, the value of the PGA cache hit percentage increases rapidly. This is due to an increase in the number of work areas which run optimally or one-pass and a decrease in the number of multi-pass executions. At some point, somewhere around 500 MB in this example, there is an inflection in the curve that corresponds to the point where most (probably all) work areas can run optimally or at least one-pass. After this inflection, the cache hit percentage keeps increasing, though at a lower pace, up to the point where it starts to taper off and shows only slight improvement with increase in `PGA_AGGREGATE_TARGET`. In [Figure 14-3](#), this happens when `PGA_AGGREGATE_TARGET` reaches 3 GB. At that point, the cache hit percentage is 83% and only improves marginally (by 2%) with one extra gigabyte of PGA memory. In this particular example, 3 GB is probably the optimal value for `PGA_AGGREGATE_TARGET`.

Ideally, `PGA_AGGREGATE_TARGET` should be set at the optimal value, or at least to the maximum value possible in the region beyond the over allocation zone. As a rule of thumb, the PGA cache hit percentage should be higher than 60%, since at 60% the system is almost processing double the number of bytes it actually needs to process in an ideal situation. Using this particular example, it makes sense to set `PGA_AGGREGATE_TARGET` to at least 500 MB and as close as possible to 3 GB. But the right setting for the parameter `PGA_AGGREGATE_TARGET` really depends on

how much memory can be dedicated to the PGA component. Generally, adding PGA memory requires reducing memory for some of the SGA components, like the shared pool or the buffer cache. This is because the overall memory dedicated to the Oracle instance is often bound by the amount of physical memory available on the system. As a result, any decisions to increase PGA memory must be taken in the larger context of the available memory in the system and the performance of the various SGA components (which you monitor with shared pool advisory and buffer cache advisory statistics). If memory cannot be taken away from the SGA, you might consider adding more physical memory to the system.

See Also:

- ["Shared Pool Advisory Statistics"](#) on page 14-31
- ["Buffer Cache Advisory Statistics"](#) on page 14-6

How to Tune PGA_AGGREGATE_TARGET You can use the following steps as a tuning guideline in tuning PGA_AGGREGATE_TARGET:

1. Set PGA_AGGREGATE_TARGET so there is no memory over-allocation; avoid setting it in the over-allocation zone. In [Example 14-8](#), PGA_AGGREGATE_TARGET should be set to at least 375 MB.
2. After eliminating over-allocations, aim at maximizing the PGA cache hit percentage, based on your response-time requirement and memory constraints. In [Example 14-8](#), assume you have a limit *X* on memory you can allocate to PGA.
 - If this limit *X* is beyond the optimal value, then you would set PGA_AGGREGATE_TARGET to the optimal value. After this point, the incremental benefit with higher memory allocation to PGA_AGGREGATE_TARGET is very small. In [Example 14-8](#), if you have 10 GB to dedicate to PGA, set PGA_AGGREGATE_TARGET to 3 GB, the optimal value. The remaining 7 GB is dedicated to the SGA.
 - If the limit *X* is less than the optimal value, then you would set PGA_AGGREGATE_TARGET to *X*. In [Example 14-8](#), if you have only 2 GB to dedicate to PGA, set PGA_AGGREGATE_TARGET to 2 GB and accept a cache hit percentage of 75%.

Finally, like most statistics collected by Oracle that are cumulative since instance start-up, you can take a snapshot of the view at the beginning and at the end of a time interval. You can then derive the predicted statistics for that time interval as follows:


```

estd_overalloc_count = (difference in estd_overalloc_count between the two snapshots)
                        (difference in bytes_processed between the two snapshots)
estd_pga_cache_hit_percentage = -----
                                (difference in bytes_processed + extra_bytes_rw between the two snapshots )

```

V\$PGA_TARGET_ADVICE_HISTOGRAM This view predicts how the statistics displayed by the performance view V\$SQL_WORKAREA_HISTOGRAM will be impacted if you change the value of the initialization parameter PGA_AGGREGATE_TARGET. You can use the dynamic view V\$PGA_TARGET_ADVICE_HISTOGRAM to view detailed information on the predicted number of optimal, one-pass and multi-pass work area executions for the set of PGA_AGGREGATE_TARGET values you use for the prediction.

The V\$PGA_TARGET_ADVICE_HISTOGRAM view is identical to the V\$SQL_WORKAREA_HISTOGRAM view, with two additional columns to represent the PGA_AGGREGATE_TARGET values used for the prediction. Therefore, any query executed against the V\$SQL_WORKAREA_HISTOGRAM view can be used on this view, with an additional predicate to select the desired value of PGA_AGGREGATE_TARGET.

Example 14–9 Querying V\$PGA_TARGET_ADVICE_HISTOGRAM

The following query displays the predicted content of V\$SQL_WORKAREA_HISTOGRAM for a value of the initialization parameter PGA_AGGREGATE_TARGET set to twice its current value.

```

SELECT LOW_OPTIMAL_SIZE/1024 low_kb, (HIGH_OPTIMAL_SIZE+1)/1024 high_kb,
       estd_optimal_executions estd_opt_cnt,
       estd_onepass_executions estd_onepass_cnt,
       estd_multipasses_executions estd_mpass_cnt
FROM v$pga_target_advice_histogram
WHERE pga_target_factor = 2
      AND estd_total_executions != 0
ORDER BY 1;

```

The output of this query might look like the following.

LOW_KB	HIGH_KB	ESTD_OPTIMAL_CNT	ESTD_ONEPASS_CNT	ESTD_MPASS_CNT
8	16	156107	0	0
16	32	148	0	0
32	64	89	0	0
64	128	13	0	0
128	256	58	0	0

256	512	10	0	0
512	1024	653	0	0
1024	2048	530	0	0
2048	4096	509	0	0
4096	8192	227	0	0
8192	16384	176	0	0
16384	32768	133	16	0
32768	65536	66	103	0
65536	131072	15	47	0
131072	262144	0	48	0
262144	524288	0	23	0

The output shows that increasing `PGA_AGGREGATE_TARGET` by a factor of 2 will allow all work areas under 16 MB to execute in optimal mode.

See Also: *Oracle9i Database Reference*

V\$SYSSTAT and V\$SESSTAT

Statistics in the `V$SYSSTAT` and `V$SESSTAT` views show the total number of work areas executed with optimal memory size, one-pass memory size, and multi-pass memory size. These statistics are cumulative since the instance or the session was started.

The following query gives the total number and the percentage of times work areas were executed in these three modes since the instance was started:

```
SELECT name profile, cnt, decode(total, 0, 0, round(cnt*100/total)) percentage
FROM (SELECT name, value cnt, (sum(value) over ()) total
FROM V$SYSSTAT
WHERE name like 'workarea exec%');
```

The output of this query might look like the following:

PROFILE	CNT	PERCENTAGE
workarea executions - optimal	5395	95
workarea executions - onepass	284	5
workarea executions - multipass	0	0

Configuring SORT_AREA_SIZE

Tuning sort operations using `SORT_AREA_SIZE` is only relevant for configurations running the Oracle shared server option or for configurations not running under the

automatic memory management mode. In the later case, Oracle Corporation strongly recommends switching to the automatic memory management mode, because it is easier to manage and often outperforms a manually-tuned system.

This section describes the following:

- [Fundamentals of Sorts](#)
- [Recognizing Memory and Disk Sorts](#)
- [Application Characteristics](#)
- [Considerations with SORT_AREA_SIZE](#)
- [Considerations with SORT_AREA_RETAINED_SIZE](#)
- [Using NOSORT to Create Indexes Without Sorting](#)
- [Using GROUP BY NOSORT](#)

Fundamentals of Sorts

A sort is an operation that orders data according to certain criteria before the data is returned to the requestor.

Operations that perform sorts include the following:

- CREATE INDEX
- SELECT ORDER BY
- SELECT DISTINCT
- SELECT GROUP BY
- SELECT ... CONNECT BY
- SELECT ... CONNECT BY ROLLUP
- Sort merge joins

See Also: *Oracle9i Database Concepts* for a list of SQL statements that perform sorts

When the `WORKAREA_SIZE_POLICY` parameter is set to `MANUAL`, the maximum amount of memory allocated for a sort is defined by the parameter `SORT_AREA_SIZE`. If the sort operation is not able to completely fit into `SORT_AREA_SIZE` memory, then the sort is separated into phases. The temporary output of each phase is stored in temporary segments on disk. The tablespace in which these sort segments are created is the user's temporary tablespace.

When Oracle writes sort operations to disk, it writes out partially sorted data in sorted runs. After all the data has been received by the sort, Oracle merges the runs to produce the final sorted output. If the sort area is not large enough to merge all the runs at once, then subsets of the runs are merged in several merge passes. If the sort area is larger, then there are fewer, longer runs produced. A larger sort area also means that the sort can merge more runs in one merge pass.

Note: Reads and writes performed as a part of disk sort operations bypass the buffer cache.

Recognizing Memory and Disk Sorts

Oracle collects statistics that reflect sort activity and stores them in dynamic performance views, such as V\$SQLAREA and V\$SYSSTAT.

Table 14–4 lists the statistics from V\$SYSSTAT that reflect sort behavior.

Table 14–4 V\$SYSSTAT Statistics Reflecting Sort Behavior

Statistic	Description
sorts (memory)	The number of sorts small enough to be performed entirely in memory without I/O to temporary sort segments on disk.
sorts (disk)	The number of sorts too large to be performed entirely in memory, requiring I/O to temporary sort segments on disk.

For example, the following query monitors these statistics:

```
SELECT NAME, VALUE
   FROM V$SYSSTAT
  WHERE NAME IN ('sorts (memory)', 'sorts (disk)');
```

The output of this query might look like the following:

```
NAME                                                    VALUE
-----
sorts (memory)                                         430418566
sorts (disk)                                           33255
```

In addition, to find individual SQL statements that are performing sorts, query the V\$SQLAREA view. Order the rows by SORTS to identify the SQL statements performing the most sorts. For example:

```
SELECT HASH_VALUE, SQL_TEXT, SORTS, EXECUTIONS
```

```
FROM V$SQLAREA  
ORDER BY SORTS;
```

In an OLTP environment, the best solution is to investigate whether the SQL statements can be tuned to avoid the sort activity.

Application Characteristics

For best performance in OLTP systems, most sorts should occur solely within memory. Sorts written to disk can adversely affect performance. If your OLTP application frequently performs sorts that do not fit into sort area size, and if the application has been tuned to avoid unnecessary sorting, then consider increasing the `SORT_AREA_SIZE` parameter for the whole instance.

If there are only a few programs that perform larger than average sorts that sort to disk, then it is possible to modify `SORT_AREA_SIZE` at the session level only for that workload or application (for example, building an index).

DSS applications typically access large volumes of data. These types of applications are expected to perform sorts to disk, purely because of the nature of the application and the data volumes involved. In DSS applications, it is important to identify the optimal `SORT_AREA_SIZE` to allow the disk sorts to perform most efficiently. Allocating more memory to sorts does not necessarily mean that the sort will be faster.

Considerations with `SORT_AREA_SIZE`

The main consideration when choosing `SORT_AREA_SIZE` is balancing memory usage with sort performance.

Since Oracle8 release 8.0, sorts do not allocate the whole of `SORT_AREA_SIZE` in one memory allocation at the beginning of the sort. The memory is allocated in `DB_BLOCK_SIZE` chunks when required, up to `SORT_AREA_SIZE`.

This means that increasing `SORT_AREA_SIZE` memory is a concern when the majority of processes on the system perform sorts and use the maximum allocation. In this situation, increasing `SORT_AREA_SIZE` for the instance as a whole results in more memory being allocated from the operating system (for dedicated connections; that is, if a shared server environment is not used). This is not necessarily a problem if the system has free memory available. However, if there is not enough free memory, then this causes paging or swapping.

If a shared server environment is used, then the additional memory is allocated out of the shared pool or the large pool when it is configured (that is, when the `LARGE_POOL_SIZE` initialization parameter is specified). The actual amount of memory

used in the shared pool is the lesser of `SORT_AREA_SIZE`, `SORT_AREA_RETAINED_SIZE`, and the actual allocation used by the sort.

If the `SORT_AREA_SIZE` is too small, then the sort is not performed as efficiently as possible. This means that sorts that could have been memory-only sorts will be disk sorts, or, alternatively, that the number of sort runs required to process the sort could be larger than necessary. Both of these situations can severely degrade performance.

Remember that there is a point after which increasing the `SORT_AREA_SIZE` no longer provides a performance benefit.

Increasing `SORT_AREA_SIZE` `SORT_AREA_SIZE` is a dynamically modifiable initialization parameter that specifies the maximum amount of memory to use for each sort. If a significant number of sorts require disk I/O to temporary segments, then your application's performance might benefit from increasing the value of `SORT_AREA_SIZE`. Alternatively in a DSS environment, increasing `SORT_AREA_SIZE` is not likely to make the sort a memory-only sort; however, depending on the current value and the new value chosen, it could make the sort faster.

The maximum value of this parameter depends on your operating system. You need to determine what size `SORT_AREA_SIZE` makes sense for your system.

Considerations with `SORT_AREA_RETAINED_SIZE`

The `SORT_AREA_RETAINED_SIZE` parameter determines the lower memory limit to which Oracle reduces the size of the sort area after the sort has started sending the sorted data to the user or to the next part of the query.

With dedicated connections, the freed memory is not released to the operating system, rather the freed memory is made available to the session for reuse.

However, if the connection is through shared server, then there could be a memory benefit to setting `SORT_AREA_RETAINED_SIZE`. If this parameter is set after the sort has completed, then the sorted data is stored in the SGA. The amount of memory used in the SGA is the lesser of the actual usage or `SORT_AREA_RETAINED_SIZE` if it is set; otherwise, it is `SORT_AREA_SIZE`. This is why setting `SORT_AREA_RETAINED_SIZE` could be of use with a shared server environment.

Note: Connections made to the database through shared servers usually should not perform large sorts.

Although there might be a memory saving with shared server, setting `SORT_AREA_RETAINED_SIZE` causes additional I/O to write and read data to and from temporary segments on disk (if the sort requires more than `SORT_AREA_RETAINED_SIZE` bytes).

Using NOSORT to Create Indexes Without Sorting

One cause of sorting is the creation of indexes. Creating an index for a table involves sorting all rows in the table based on the values of the indexed columns. However, Oracle lets you create indexes without sorting. If the rows in the table are loaded in ascending order, then you can create the index faster without sorting.

The NOSORT Clause To create an index without sorting, load the rows into the table in ascending order of the indexed column values. Your operating system might provide a sorting utility to sort the rows before you load them. When you create the index, use the `NOSORT` clause on the `CREATE INDEX` statement. For example, the following `CREATE INDEX` statement creates the index `my_emp_name_ix` on the `last_name` column of the `employees` table without sorting the rows in the `employees` table:

```
CREATE INDEX my_emp_name_ix
  ON employees(last_name)
  NOSORT;
```

In this SQL example, it is assumed that the rows in the table are loaded in ascending order of the indexed column values.

Note: Specifying `NOSORT` in a `CREATE INDEX` statement negates the use of `PARALLEL INDEX CREATE`, even if `PARALLEL (DEGREE n)` is specified.

When to Use the NOSORT Clause Presorting your data and loading it in order might not be the fastest way to load a table.

- If you have a multiple-CPU computer, then you might be able to load data faster using multiple processors in parallel, each processor loading a different portion of the data. To take advantage of parallel processing, load the data without sorting it first. Then, create the index without the `NOSORT` clause.
- If you have a single-CPU computer, then sort your data before loading, if possible. Then, create the index with the `NOSORT` clause.

Using GROUP BY NOSORT

Sorting can be avoided when performing a `GROUP BY` operation when you know that the input data is already ordered, so that all rows in each group are clumped together. This can be the case if the rows are being retrieved from an index that matches the grouped columns, or if a sort merge join produces the rows in the right order. `ORDER BY` sorts can be avoided in the same circumstances. When no sort takes place, the `EXPLAIN PLAN` output indicates `GROUP BY NOSORT`.

I/O Configuration and Design

The I/O subsystem is a vital component of an Oracle database. This chapter introduces fundamental I/O concepts, discusses the I/O requirements of different parts of the database, and provides sample configurations for I/O subsystem design.

This chapter includes the following topics:

- [Understanding I/O](#)
- [Basic I/O Configuration](#)

Understanding I/O

The performance of many software applications is inherently limited by disk I/O. Applications that spend the majority of CPU time waiting for I/O activity to complete are said to be I/O-bound.

Oracle is designed so that if an application is well written, its performance should not be limited by I/O. Tuning I/O can enhance the performance of the application if the I/O system is operating at or near capacity and is not able to service the I/O requests within an acceptable time. However, tuning I/O cannot help performance if the application is not I/O-bound (for example, when CPU is the limiting factor).

Designing I/O Layouts

Consider the following database requirements when designing an I/O system:

1. Storage, such as minimum bytes of disk
2. Availability, such as continuous (24 x 7) or business hours only
3. Performance, such as I/O throughput and application response times

Many I/O designs plan for storage and availability requirements with the assumption that performance will not be an issue. This is not always the case. Optimally, the number of disks and controllers to be configured should be determined by I/O throughput and redundancy requirements. Then, the size of disks can be determined by the storage requirements.

Disk Performance and Reliability

For any database, the I/O subsystem is critical for system availability, performance, and data integrity. A weakness in any of these areas can render the database system unstable, unscalable, or untrustworthy.

All I/O subsystems use magnetic disk drives. Conventional magnetic disk drives contain moving parts. Because these moving parts have a design life, they are subject to tolerances in manufacturing that make their reliability and performance inconsistent. Not all drives that are theoretically identical perform the same, and they can break down over time. When assembling a large disk configuration, you need only to look at the mean time between failures (MTBF) of disk drives and the number of disks to see that disk failures are a common occurrence. This is unfortunate, because the core assets of any system (the data) reside on the disks.

Disk Technology

The main component of any I/O subsystem, the disk drive, has barely changed over the last few years. The only changes are the increase in drive capacity from under 1 gigabyte to over 50 gigabytes, and small improvements in disk access time (rpm) and, hence, throughput. Improvements made in CPUs, on the other hand, have doubled their performance every 18 months.

There is a disk drive paradox that says that if you size the number of disks required by disk capacity, then you need fewer and fewer disks over time, because the disks increase in size. However, if you size the number of disks by performance, then you must double the number of disks for each CPU every 18 months. Dealing with this paradox has proved to be a problem in many system configurations, especially those designed to be cheaper by using fewer disks.

In addition to the size of the actual disks manufactured, the way disk subsystems are connected has changed. On smaller systems, in general, disk drives are connected individually to the host machine by small computer system interfaces (SCSI) through one of a number of disk controllers. For high-end systems, disk mirroring and striping are essential for performance and availability. These requirements lead to hundreds of disks connected through complex wiring configurations.

However, although disk technology itself has not greatly changed, I/O subsystems have evolved into disk arrays that overcome many of the described problems. The systems perform disk mirroring, provide hot swapping of disks, and in many cases, provide simpler connections to the host by fiber interfaces. The more sophisticated disk arrays are, in fact, small computers themselves with their own CPU, battery-backed memory cache for high-performance resilient writes, dial-home diagnostics, and proxy backup to allow backup without going through the host operating system.

What Is Disk Contention?

Disk contention occurs when multiple processes try to access the same disk simultaneously. Most disks have limits on both the number of accesses (I/O operations each second) and the amount of data they can transfer each second (I/O data rate, or **throughput**). When these limits are reached, processes must wait to access the disk.

Load Balancing and Striping

A performance engineer's goal is to distribute the I/O load evenly across the available devices. This is known as load balancing or distributing I/O. Historically, load balancing had to be performed manually. Database administrators would determine the I/O requirements and characteristics of each datafile and design the I/O layout based on which files could be placed together on a single disk to distribute the activity evenly over all disks.

If a particular table or index was I/O-intensive, then to further distribute the I/O load, DBAs also had the option of manually distributing, or striping, the data. This was achieved by splitting the object's extents over separate datafiles, then distributing the datafiles over devices.

Fundamentally, striping divides data into small portions and stores these portions in separate files on separate disks. This allows multiple processes to access different portions of the data concurrently without disk contention. Operating systems, hardware vendors, and third-party software vendors provide the tools for striping a heavily used file across many physical devices simply, thus making the job of balancing I/O significantly easier for DBAs.

Striping is helpful both in OLTP environments, to optimize random access to tables with many rows, and also in DSS environments, to allow parallel operations to scan large volumes of data quickly. Striping techniques are productive when the load redistribution eliminates or reduces some form of queue. Striping files can also improve RMAN backup and restore performance.

If the concurrent load is too heavy for the available hardware, then striping does not alleviate the problem.

Striping and RAID

You must also consider the recoverability requirements of each particular system. Configurations with redundant arrays of inexpensive disks (RAID) provide improved data reliability, while offering the option of striping. The RAID level chosen should depend on performance and cost. Different RAID levels are suited to different types of applications, depending on their I/O characteristics.

The following list provides a brief overview of the most popular RAID configurations used for database files, along with the applications that best suit them. These descriptions are very general. Consult with your hardware and software vendors for details about implementing a specific configuration:

RAID 0: Striping

Files are striped across many physical disks. This configuration provides read and write performance, but not reliability.

Note: Although RAID 0 provides the best read and write performance, it is not a true RAID system because it does not allow for redundancy. Oracle recommends that you do not place production database files on RAID 0 systems.

RAID 1: Mirroring

Physical disks are used to store a specified number n of concurrently maintained copies of a file. The number of disks required is n times m , where m is the number of disks required to store the original files. RAID 1 provides good reliability and good read rates. Sometimes, writes can be costly, because n writes are required to maintain n copies.

RAID 0+1: Striping and Mirroring

This level combines the technologies of RAID 0 and RAID 1. It is widely used because it provides good reliability and better read and write performance than RAID 1.

Note: Mirroring can cause I/O bottlenecks. Generally, the process of writing to each mirror is done in parallel and does not cause a bottleneck. However, if each mirror is striped differently, then the I/O does not complete until the slowest member of the mirror is complete. To avoid I/O problems, stripe the copies using the same number of disks as used for the primary database.

RAID 5: Striping and Redundancy

RAID 5 striping is similar to striping in RAID 0. Recoverability of lost data due to disk failure is achieved by storing parity data regarding all disks in a stripe, with storage uniform throughout disks in the group. Compared to RAID 1, the benefit is the saving in disk cost. RAID 5 provides good reliability. Sequential reads benefit the most, while write performance can suffer. This configuration might not be ideal for write-intensive applications.

Recent RAID 5 implementations avoid many of the traditional RAID 5 limits by installing large amounts of battery-backed memory (NVRAM). The cache offsets the performance penalty in two ways:

- It helps defer the physical I/O operations, because the small write penalty does not bottleneck the individual I/O call that initiated it.
- It helps a disk array to group small writes into larger batches that don't incur the small-write penalty.

The drawbacks are:

- The cache is expensive, which can defeat the RAID 5 cost advantage (for each byte of storage) over RAID 1.
- I/O throughput requirement might exceed the cache capacity. If I/O calls keep coming fast enough, the cache can fill up, triggering all I/O to the array to cease until the array resynchronizes with the physical disk.

Balancing Budget, Performance, and Availability

The choice of I/O subsystem and how to design the layout is an exercise in balancing budget, performance, and availability. It is essential to be able to expand the I/O subsystem as throughput requirements grow and more I/O is performed. To achieve a scalable I/O system, the chosen system must be able to evolve over time, with a minimum of downtime. In many cases, this involves configuring the system to achieve 95% optimal performance, forsaking 5% performance for ease of configuration and eventual upgrades.

Basic I/O Configuration

This section describes the basic information to be gathered and decisions to be made when defining a system's I/O configuration. You want to keep the configuration as simple as possible, while maintaining the required availability, recoverability, and performance. The more complex a configuration becomes, the more difficult it is to administer, maintain, and tune.

Determining Application I/O Characteristics

This section describes the basic I/O characteristics and requirements that must be determined. These I/O characteristics influence decisions on the type of technology required and how to configure that technology. I/O requirements include the performance, space, and recoverability needs specific to the site.

To design an efficient I/O subsystem, you need the following information:

- [I/O Rate: Read Rate and Write Rate](#)
- [I/O Concurrency](#)
- [I/O Size](#)
- [Availability](#)
- [Storage Size](#)

When defining a system's I/O configuration, you need to take into account the maximum load that will be generated by parallel operations such as parallel index creation and parallel scans. You only get one chance to configure your disks, and if enough load distribution is not built in, then parallel operations will not scale. This is especially important for log and archive disks.

I/O Rate: Read Rate and Write Rate

The read rate is the number of reads each second. The write rate is the number of writes each second. The sum of the read rate and the write rate is the I/O rate (the number of I/O operations each second). For a system that performs well, your application's I/O rate should be a significant factor when determining the absolute minimum number of disks and controllers required.

A system that has a high write rate might also benefit from the following configuration options:

- Using raw devices or third-party software that enables writes directly to disk, avoiding reading from and writing to the UNIX buffer cache
- Using an I/O disk array that has a memory cache large enough to sustain destaging dirty blocks without filling up the cache
- Avoiding RAID 5 configuration for write-intensive applications

See Also: Your vendor-specific documentation to see whether high write rates can be sustained without a performance hit

I/O Concurrency

I/O concurrency measures the number of distinct processes that simultaneously have I/O requests to the I/O subsystem. From an Oracle perspective, I/O concurrency is considered the number of processes concurrently issuing I/O requests. A high degree of I/O concurrency implies that there are many distinct

processes simultaneously issuing I/O requests. A low degree of concurrency implies that few processes are simultaneously issuing I/O requests.

Note: High I/O rate systems are typified by high concurrency and small block requests. High data rate systems are typified by low concurrency and large block requests.

I/O Size

I/O size is the size of the I/O request from Oracle's perspective. The minimum I/O size is the operating system block size, while the maximum is typically a factor of the Oracle block size multiplied by the multiblock read count.

Although the I/O size can vary depending on the type of operation, there are some reasonable, general estimates that can be made, depending on the nature of the application.

- With a DSS system, the majority of the I/Os are typically going to be large, approximately $n * DB_BLOCK_SIZE$.
- With an OLTP system, the I/Os are predominantly going to be the size of `DB_BLOCK_SIZE`.

Factors Affecting I/O Size and Concurrency In an Oracle system, the various files have different requirements of the disk subsystem. The requirements are distinguished by the rate of reading and writing of data and by the concurrency of these operations. High rates are relatively easily sustained by the disk subsystem when the concurrency is low. It is important that design of the disk subsystem take the following factors into consideration:

- Operational parameters. These are the configurable and nonmodifiable operational parameters for Oracle, the operating system, the hardware, and any third party software such as a logical volume master (LVM).

See Also: [Table 15-2](#) on page 15-13

- Typical I/O size and degree of concurrency for each file. These vary considerably from one site to another, even for the same type of file. For example, a site that performs many concurrent disk sorts has different characteristics for its `TEMP` tablespace than a site that performs predominantly small, in-memory sorts.

Table 15–1 shows the read and write rate and concurrency for a theoretical installation. The example describes a high-query, high-update OLTP application that performs predominantly in-memory sorts, with indexes that remain cached in the buffer cache.

Table 15–1 *Read and Write Concurrency Rates*

Component	Read Rate	Write Rate	Concurrency
Archive logs	High	High	Low
Redo logs	High	High	Low
Undo segment tablespaces	Low	High	High
TEMP tablespaces	Low	Low	High
Index tablespaces	Low	Medium	High
Data tablespaces	High	Medium	High
Application log and output files	Low	Medium	High
Binaries (shared)	Low	Low	High

Availability

Site-specific availability requirements impose additional specifications on the disk storage technology and layout. Typical considerations include the appropriate RAID technology to meet recoverability requirements and any Oracle-specific safety measures, such as mirroring of redo logs, archive logs, and control files.

Dedicating separate disks to mirroring redo log files is an important safety precaution. This ensures that the datafiles and the redo log files cannot both be lost in a single disk failure.

See Also: *Oracle9i User-Managed Backup and Recovery Guide* for more information on recoverability

Storage Size

With the large disks available today, if the performance and availability requirements are satisfied, then in most cases the storage needs have already been met. If, however, the system stores voluminous data online, and the data does not have high concurrency or throughput requirements, then you can fall short of storage requirements. In such cases, you need to consider the following methods of maximizing storage space:

- Use different RAID configurations.
- Use more disks.
- Use larger disks.

I/O Configuration Decisions

Depending on the hardware and software available, you should determine the following system characteristics during the I/O layout phase:

- Number of disks available for database files
- RAID level
- Stripe depth and stripe width
- Whether to use raw devices or file system
- Whether the system can perform asynchronous I/O

To maximize performance, the goal of system configuration should be to distribute the I/O load for the database as evenly as possible over the available disks.

In addition to distributing I/O, other configuration concerns should include the expected growth on the system and the required recoverability.

Know Your I/O System

You should be aware of the capabilities of the I/O system. This information is available by reviewing hardware and software documentation (if applicable), and also by performing tests at your site. Tests can include varying any of the following factors:

- I/O sizes
- Raw devices compared to file systems
- Read and write performance
- Sequential and random access
- Asynchronous I/O compared to synchronous I/O
- Configuration of features your system has for enhancing I/O

Performing this type of research can provide insight into potential configurations that can be used for the final design.

Match I/O Requirements with the I/O System

If you have a benchmark that simulates the load placed on the I/O system by the application, or if there is an existing production system, then look at the Oracle statistics to determine the I/O rates for each file and for the database as a whole.

See Also: *Oracle9i Database Performance Planning* for more information on operating system I/O statistics

To determine the Oracle file I/O statistics, look at the following:

- Number of physical reads (V\$FILESTAT.PHYRDS)
- Number of physical writes (V\$FILESTAT.PHYWRTS)
- Average read time
- Number of physical I/Os. This number is the sum of the physical reads and physical writes.

Assuming that the Oracle buffer cache is adequately sized, the physical reads and physical writes statistics are useful.

Another important statistic is the average number of I/Os each second. To determine this number, sample the V\$FILESTAT data over an interval, add the physical reads and writes, and then divide this total by the elapsed time in seconds. The general formula is as follows:

Average I/O per second = (physical reads + physical writes) / elapsed seconds

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#) for details on how to calculate the delta information in V\$FILESTAT

To estimate the I/O requirements, scale the I/O statistics with the expected workload on the new system. Comparing the scaled data with the disk capabilities can potentially identify whether there will be a mismatch between the I/O requirements of the new application and the capabilities of the I/O system.

Also identify any I/O-intensive operations that are not part of the typical load and have I/O rates that greatly peak over the average. Ensure that your system is able to sustain these rates. Operations such as index builds, data loads, and batch processing fall into this category.

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#) for details on how to calculate the delta information in `V$FILESTAT`

FILESYSTEMIO_OPTIONS

With Oracle9i Release 2 (9.2), you can use the `FILESYSTEMIO_OPTIONS` initialization parameter to enable or disable asynchronous I/O or direct I/O on file system files. This parameter is platform-specific and has a default value that is best for a particular platform. It can be dynamically changed to update the default setting.

See Also: ["FILESYSTEMIO_OPTIONS Initialization Parameter"](#) on page 16-4

Lay Out the Files Using Operating System or Hardware Striping

If your operating system has LVM software or hardware-based striping, then it is possible to distribute I/O using these tools. Decisions to be made when using an LVM or hardware striping include **stripe depth** and **stripe width**.

- Stripe depth is the size of the stripe, sometimes called stripe unit.
- Stripe width is the product of the stripe depth and the number of drives in the striped set.

Choose these values wisely so that the system is capable of sustaining the required throughput. For an Oracle database, reasonable stripe depths range from 256 KB to 4 MB. Different types of applications benefit from different stripe depths. The optimal stripe depth and stripe width depend on the following:

- [Requested I/O Size](#)
- [Concurrency of I/O Requests](#)
- [Alignment of Physical Stripe Boundaries with Block Size Boundaries](#)
- [Variable Transfer Rate](#)
- [Manageability of the Proposed System](#)

Requested I/O Size

[Table 15–2](#) lists the Oracle and operating system parameters that you can use to set I/O size:

Table 15–2 Oracle and Operating System Operational Parameters

Parameter	Description
DB_BLOCK_SIZE	The size of single-block I/O requests. This parameter is also used in combination with multiblock parameters to determine multiblock I/O request size.
OS block size	Determines I/O size for redo log and archive log operations.
Maximum OS I/O size	Places an upper bound on the size of a single I/O request.
DB_FILE_MULTIBLOCK_READ_COUNT	The maximum I/O size for full table scans is computed by multiplying this parameter with DB_BLOCK_SIZE. (the upper value is subject to operating system limits).
SORT_AREA_SIZE	Determines I/O sizes and concurrency for sort operations.
HASH_AREA_SIZE	Determines the I/O size for hash operations.

In addition to I/O size, the degree of concurrency also helps in determining the ideal stripe depth. Consider the following when choosing stripe width and stripe depth:

- On low-concurrency (sequential) systems, ensure that no single I/O visits the same disk twice. For example, assume that the stripe width is four disks, and the stripe depth is 32k. If a single 1MB I/O request (for example, for a full table scan) is issued by an Oracle server process, then each disk in the stripe must perform four I/Os to return the requested data. To avoid this situation, the size of the average I/O should be smaller than the stripe width multiplied by the stripe depth. If this is not the case, then a single I/O request made by Oracle to the operating system results in multiple physical I/O requests to the same disk.
- On high-concurrency (random) systems, ensure that no single I/O request is broken up into more than one physical I/O call. Failing to do this multiplies the number of physical I/O requests performed in your system, which in turn can severely degrade the I/O response times.

Concurrency of I/O Requests

In a system with a high degree of concurrent small I/O requests, such as in a traditional OLTP environment, it is beneficial to keep the stripe depth large. Using stripe depths larger than the I/O size is called **coarse grain striping**. In high-concurrency systems, the stripe depth can be

$$n * DB_BLOCK_SIZE$$

where n is greater than 1.

Coarse grain striping allows a disk in the array to service several I/O requests. In this way, a large number of concurrent I/O requests can be serviced by a set of striped disks with minimal I/O setup costs. Coarse grain striping strives to maximize overall I/O throughput. Multiblock reads, as in full table scans, will benefit when stripe depths are large and can be serviced from one drive. Parallel query in a DSS environment is also a candidate for coarse grain striping. This is because there are many individual processes, each issuing separate I/Os. If coarse grain striping is used in systems that do not have high concurrent requests, then hot spots could result.

In a system with a few large I/O requests, such as in a traditional DSS environment or a low-concurrency OLTP system, then it is beneficial to keep the stripe depth small. This is called **fine grain striping**. In such systems, the stripe depth is

$n * DB_BLOCK_SIZE$

where n is smaller than the multiblock read parameters (such as `DB_FILE_MULTIBLOCK_READ_COUNT`).

Fine grain striping allows a single I/O request to be serviced by multiple disks. Fine grain striping strives to maximize performance for individual I/O requests or response time.

Alignment of Physical Stripe Boundaries with Block Size Boundaries

On some Oracle ports, an Oracle block boundary may not align with the stripe. If your stripe depth is the same size as the Oracle block, then a single I/O issued by Oracle might result in two physical I/O operations.

This is not optimal in an OLTP environment. To ensure a higher probability of one logical I/O resulting in no more than one physical I/O, the minimum stripe depth should be at least twice the Oracle block size. [Table 15-3](#) shows recommended minimum stripe depth for random access and for sequential reads.

Table 15-3 Minimum Stripe Depth

Disk Access	Minimum Stripe Depth
Random reads and writes	The minimum stripe depth is twice the Oracle block size.
Sequential reads	The minimum stripe depth is twice the value of <code>DB_FILE_MULTIBLOCK_READ_COUNT</code> , multiplied by the Oracle block size.

See Also: The specific documentation for your platform

Variable Transfer Rate

The transfer rate for a disk drive is not the same for all portions of a disk. The outer sectors of a disk drive move the disk head faster than the inner sectors, leading to a faster transfer rate for the outer sectors.

The outside portions of a disk drive hold more data than the inside portions.

The transfer rate of a disk drive can vary by a factor of two from the inner edge of the drive to the outer edge. It is therefore beneficial to place frequently accessed data toward the outer edge of a disk drive.

Manageability of the Proposed System

With an LVM, the simplest configuration to manage is one with a single striped volume over all available disks. In this case, the stripe width encompasses all available disks. All database files reside within that volume, effectively distributing the load evenly. This single-volume layout provides adequate performance in most situations.

A single-volume configuration is viable only when used in conjunction with RAID technology that allows easy recoverability, such as RAID 1. Otherwise, losing a single disk means losing all files concurrently and, hence, performing a full database restore and recovery.

In addition to performance, there is a manageability concern: the design of the system must allow disks to be added simply, to allow for database growth. The challenge is to do so while keeping the load balanced evenly.

For example, an initial configuration can involve the creation of a single striped volume over 64 disks, each disk being 16 GB. This is total disk space of 1 terabyte (TB) for the primary data. Sometime after the system is operational, an additional 80 GB (that is, five disks) must be added to account for future database growth.

The options for making this space available to the database include creating a second volume that includes the five new disks. However, an I/O bottleneck might develop, if these new disks are unable to sustain the I/O throughput required for the files placed on them.

Another option is to increase the size of the original volume. LVMs are becoming sophisticated enough to allow dynamic reconfiguration of the stripe width, which allows disks to be added while the system is online. This begins to make the placement of all files on a single striped volume feasible in a production environment.

If your LVM is unable to support dynamically adding disks to the stripe, then it is likely that you need to choose a smaller, more manageable stripe width. Then, when new disks are added, the system can grow by a stripe width.

In the preceding example, eight disks might be a more manageable stripe width. This is only feasible if eight disks are capable of sustaining the required number of I/Os each second. Thus, when extra disk space is required, another eight-disk stripe can be added, keeping the I/O balanced across the volumes.

Note: The smaller the stripe width becomes, the more likely it is that you will need to spend time distributing the files on the volumes, and the closer the procedure becomes to manually distributing I/O.

Manually Distributing I/O

If your system does not have an LVM or hardware striping, then I/O must be manually balanced across the available disks by distributing the files according to each file's I/O requirements. In order to make decisions on file placement, you should be familiar with the I/O requirements of the database files and the capabilities of the I/O system. If you are not familiar with this data and do not have a representative workload to analyze, you can make a first guess and then tune the layout as the usage becomes known.

To stripe disks manually, you need to relate a file's storage requirements to its I/O requirements.

1. Evaluate database disk-storage requirements by checking the size of the files and the disks.
2. Identify the expected I/O throughput for each file. Determine which files have the highest I/O rate and which do not have many I/Os. Lay out the files on all the available disks so as to even out the I/O rate.

One popular approach to manual I/O distribution suggests separating a frequently used table from its index. This is not correct. During the course of a transaction, the index is read first, and then the table is read. Because these I/Os occur sequentially, the table and index can be stored on the same disk without contention. It is not sufficient to separate a datafile simply because the datafile contains indexes or table data. The decision to segregate a file should be made only when the I/O rate for that file affects database performance.

When to Separate Files

Regardless of whether you use operating system striping or manual I/O distribution, if the I/O system or I/O layout is not able to support the I/O rate required, then you need to separate files with high I/O rates from the remaining files. You can identify such files either at the planning stage or after the system is live.

The decision to segregate files should only be driven by I/O rates, recoverability concerns, or manageability issues. (For example, if your LVM does not support dynamic reconfiguration of stripe width, then you might need to create smaller stripe widths to be able to add n disks at a time to create a new stripe of identical configuration.)

Before segregating files, verify that the bottleneck is truly an I/O issue. The data produced from investigating the bottleneck identifies which files have the highest I/O rates.

See Also: ["Identifying and Gathering Data on Resource-Intensive SQL"](#) on page 6-3

Tables, Indexes, and TEMP Tablespaces

If the files with high I/O are datafiles belonging to tablespaces that contain tables and indexes, then identify whether the I/O for those files can be reduced by tuning SQL or application code.

If the files with high-I/O are datafiles that belong to the `TEMP` tablespace, then investigate whether to tune the SQL statements performing disk sorts to avoid this activity, or to tune the sorting.

After the application has been tuned to avoid unnecessary I/O, if the I/O layout is still not able to sustain the required throughput, then consider segregating the high-I/O files.

See Also: ["Identifying and Gathering Data on Resource-Intensive SQL"](#) on page 6-3

Redo Log Files

If the high-I/O files are redo log files, then consider splitting the redo log files from the other files. Possible configurations can include the following:

- Placing all redo logs on one disk without any other files. Also consider availability; members of the same group should be on different physical disks and controllers for recoverability purposes.
- Placing each redo log group on a separate disk that does not store any other files.
- Striping the redo log files across several disks, using an operating system striping tool. (Manual striping is not possible in this situation.)
- Avoiding the use of RAID 5 for redo logs.


Redo log files are written sequentially by the Log Writer (LGWR) process. This operation can be made faster if there is no concurrent activity on the same disk. Dedicating a separate disk to redo log files usually ensures that LGWR runs smoothly with no further tuning necessary. If your system supports asynchronous I/O but this feature is not currently configured, then test to see if using this feature is beneficial. Performance bottlenecks related to LGWR are rare.

Archived Redo Logs

If the archiver is slow, then it might be prudent to prevent I/O contention between the archiver process and LGWR by ensuring that archiver reads and LGWR writes are separated. This is achieved by placing logs on alternating drives.

For example, suppose a system has four redo log groups, each group with two members. To create separate-disk access, the eight log files should be labeled 1a, 1b, 2a, 2b, 3a, 3b, 4a, and 4b. This requires at least four disks, plus one disk for archived files.

[Figure 15-1](#) illustrates how redo members should be distributed across disks to minimize contention.

Figure 15–1 Distributing Redo Members Across Disks

In this example, LGWR switches out of log group 1 (member 1a and 1b) and writes to log group 2 (2a and 2b). Concurrently, the archiver process reads from group 1 and writes to its archive destination. Note how the redo log files are isolated from contention.

Note: Mirroring redo log files, or maintaining multiple copies of each redo log file on separate disks, does not slow LGWR considerably. LGWR writes to each disk in parallel and waits until each part of the parallel write is complete. Hence, a parallel write does not take longer than the longest possible single-disk write.

Because redo logs are written serially, drives dedicated to redo log activity generally require limited head movement. This significantly accelerates log writing.

Three Sample Configurations

This section contains three high-level examples of configuring I/O systems. These examples include sample calculations that define the disk topology, stripe depths, and so on.

Stripe Everything Across Every Disk

The simplest approach to I/O configuration is to build one giant volume, striped across all available disks. To account for recoverability, the volume is mirrored

(RAID 1). The striping unit for each disk should be larger than the maximum I/O size for the frequent I/O operations. This provides adequate performance for most cases.

Move Archive Logs to Different Disks

If archive logs are striped on the same set of disks as other files, then any I/O requests on those disks could suffer when redo logs are being archived. Moving archive logs to separate disks provides the following benefits:

- The archive can be performed at very high rate (using sequential I/O).
- Nothing else is affected by the degraded response time on the archive destination disks.

The number of disks for archive logs is determined by the rate of archive log generation and the amount of archive storage required.

Move Redo Logs to Separate Disks

In high-update OLTP systems, the redo logs are write-intensive. Moving the redo log files to disks that are separate from other disks and from archived redo log files has the following benefits:

- Writing redo logs is performed at the highest possible rate. Hence, transaction processing performance is at its best.
- Writing of the redo logs is not impaired with any other I/O.

The number of disks for redo logs is mostly determined by the redo log size, which is generally small compared to current technology disk sizes. Typically, a configuration with two disks (possibly mirrored to four disks for fault tolerance) is adequate. In particular, by having the redo log files alternating on two disks, writing redo log information to one file does not interfere with reading a completed redo log for archiving.

Oracle-Managed Files

For systems where a file system can be used to contain all Oracle data, database administration is simplified by using Oracle-managed files. Oracle internally uses standard file system interfaces to create and delete files as needed for tablespaces, temp files, online logs, and control files. Administrators only specify the file system directory to be used for a particular type of file. You can specify one default location for datafiles and up to five multiplexed locations for the control and online redo log files.

Oracle ensures that a unique file is created and then deleted when it is no longer needed. This reduces corruption caused by administrators specifying the wrong file, reduces wasted disk space consumed by obsolete files, and simplifies creation of test and development databases. It also makes development of portable third-party tools easier, because it eliminates the need to put operating-system specific file names in SQL scripts.

New files can be created as managed files, while old ones are administered in the old way. Thus, a database can have a mixture of Oracle-managed and manually managed files.

Note: Oracle-managed files cannot be used with raw devices.

Tuning Oracle-Managed Files

Several points should be considered when tuning Oracle-managed files.

- Because Oracle-managed files require the use of a file system, DBAs give up control over how the data is laid out. Therefore, it is important to correctly configure the file system.
- The Oracle-managed file system should be built on top of an LVM that supports striping. For load balancing and improved throughput, the disks in the Oracle-managed file system should be striped.
- Oracle-managed files work best if used on an LVM that supports dynamically extensible logical volumes. Otherwise, the logical volumes should be configured as large as possible.
- Oracle-managed files work best if the file system provides large extensible files.

See Also: *Oracle9i Database Administrator's Guide* for detailed information on using Oracle-managed files

Choosing Data Block Size

This section lists considerations when choosing database block size for optimal performance.

Reads Regardless of the size of the data, the goal is to minimize the number of reads required to retrieve the desired data.

- If the rows are small and access is predominantly random, then choose a smaller block size.


- If the rows are small and access is predominantly sequential, then choose a larger block size.
- If the rows are small and access is both random and sequential, then it might be effective to choose a larger block size.
- If the rows are large—for example, if they contain large object (LOB) data—then choose a larger block size.

Writes For high-concurrency OLTP systems, consider appropriate values for `INITRANS`, `MAXTRANS`, and `FREELISTS` when using a larger block size. These parameters affect the degree of update concurrency allowed within a block. However, you do not need to specify the value for `FREELISTS` when using automatic segment-space management.

If you are uncertain about which block size to choose, then try a database block size of 8 KB for most systems that process a large number of transactions. This represents a good compromise and is usually effective. Only systems processing LOB data need more than 8 KB.

[Figure 15-2](#) illustrates the suitability of various block sizes to OLTP or DSS applications.

Figure 15-2 *Block Size (in KB) and Application Type*



See Also: The Oracle documentation specific to your operating system for information on the minimum and maximum block size on your platform

Block Size Advantages and Disadvantages

[Table 15-4](#) lists the advantages and disadvantages of different block sizes.

Table 15–4 *Block Size Advantages and Disadvantages*

Block Size	Advantages	Disadvantages
Smaller	<p>Good for small rows with lots of random access.</p> <p>Reduces block contention.</p>	<p>Has relatively large space overhead due to metadata (that is, block header).</p> <p>Not recommended for large rows. There might only be a few rows stored for each block, or worse, row chaining if a single row does not fit into a block,</p>
Larger	<p>Has lower overhead, so there is more room to store data.</p> <p>Permits reading a number of rows into the buffer cache with a single I/O (depending on row size and block size).</p> <p>Good for sequential access or very large rows (such as LOB data).</p>	<p>Wastes space in the buffer cache, if you are doing random access to small rows and have a large block size. For example, with an 8 KB block size and 50 byte row size, you waste 7,950 bytes in the buffer cache when doing random access.</p> <p>Not good for index blocks used in an OLTP environment, because they increase block contention on the index leaf blocks.</p>

Understanding Operating System Resources

This chapter explains how to tune the operating system for optimal performance of the Oracle database server.

This chapter contains the following sections:

- [Understanding Operating System Performance Issues](#)
- [Solving Operating System Problems](#)
- [Understanding CPU](#)
- [Finding System CPU Utilization](#)

See Also:

- Your Oracle platform-specific documentation and your operating system vendor's documentation
- *Oracle9i Database Performance Planning* for a discussion of the importance of operating system statistics

Understanding Operating System Performance Issues

Operating system performance issues commonly involve process management, memory management, and scheduling. If you have tuned the Oracle instance and you still need better performance, then verify your work or try to reduce system time. Make sure that there is enough I/O bandwidth, CPU power, and swap space. Do not expect, however, that further tuning of the operating system will have a significant effect on application performance. Changes in the Oracle configuration or in the application are likely to make a more significant difference in operating system efficiency than simply tuning the operating system.

For example, if an application experiences excessive buffer busy waits, then the number of system calls increases. If you reduce the buffer busy waits by tuning the application, then the number of system calls decreases.

See Also: Your Oracle platform-specific documentation and your operating system vendor's documentation

Using Operating System Caches

Operating systems and device controllers provide data caches that do not directly conflict with Oracle's own cache management. Nonetheless, these structures can consume resources while offering little or no benefit to performance. This is most noticeable on a UNIX system that has the database files in the UNIX file store; by default all database I/O goes through the file system cache. On some UNIX systems, direct I/O is available to the filestore. This arrangement allows the database files to be accessed within the UNIX file system, bypassing the file system cache. It saves CPU resources and allows the file system cache to be dedicated to non-database activity, such as program texts and spool files.

This problem does not occur on NT. All file requests by the database bypass the caches in the file system.

Although the operating system cache is often redundant because the Oracle buffer cache buffers blocks, there are a number of cases where Oracle does not use the Oracle buffer cache. In these cases, using direct I/O which bypasses the Unix or operating system cache or using raw devices which do not use the operating system cache may yield worse performance than using operating system buffering. Some examples of this include the following:

- Reads or writes to the `TEMPORARY` tablespace
- Data stored in `NOCACHE` LOBs
- Parallel Query slaves reading data

You may want a mix with some files cached at the operating system level and others not.

Hardware Cache

Some underlying I/O subsystems implement hardware-level caching of disk reads and writes to speed up their response times to I/O requests. It is important to ensure that such subsystems are configured to acknowledge write requests only after the written data is guaranteed to be safe.

For example, consider a subsystem that implements a RAM cache that acknowledges writes as soon as the written data is in the RAM cache. The Oracle database server considers this data to be safe. Then a power failure occurs and the data is lost. This can lead to corruption, because the Oracle server has no way of knowing that the write was lost. In order to overcome this problem, most I/O subsystems have a mechanism for ensuring data in the RAM cache is not lost across power failures. If the subsystem cannot guarantee this, then it is generally best to configure the system in such a way that a write is acknowledged only after the data has been written to disk, rather than only to the RAM cache.

Asynchronous I/O

With synchronous I/O, when an I/O request is submitted to the operating system, the writing process blocks until the write is confirmed as complete. It can then continue processing.

Asynchronous I/O allows a process to submit an I/O request but to then continue processing. It may then check on the result of the I/O at a later time. It is also possible to submit several I/O requests and then collect the status of those requests at a later time, thus allowing the operating system to parallelize any of those I/O operations, where possible. Parallel processing can reduce the overall time to complete an operation.

Consider an extreme example: Imagine you want to write out four data blocks to four different files. With synchronous I/O you must submit block 1, wait, submit block 2, wait, submit block 3, wait, submit block 4, and wait. With asynchronous I/O, you can submit blocks 1, 2, 3, and 4 and then wait for all four blocks to complete. Because you gave the operating system all four I/O requests at once, it can act on all the requests in parallel. The total response time is only the duration of the longest I/O of the four, rather than the sum of all four I/O durations.

Some platforms support asynchronous I/O by default, others need special configuration, and some only support asynchronous I/O for certain underlying file system types.

FILESYSTEMIO_OPTIONS Initialization Parameter

With Oracle9i Release 2 (9.2), you can use the `FILESYSTEMIO_OPTIONS` initialization parameter to enable or disable asynchronous I/O or direct I/O on file system files. This parameter is platform-specific and has a default value that is best for a particular platform. It can be dynamically changed to update the default setting.

`FILESYSTEMIO_OPTIONS` can be set to one of the following values:

- `ASYNCH`: enable asynchronous I/O on file system files
- `DIRECTIO`: enable direct I/O on file system files
- `SETALL`: enable both asynchronous and direct I/O on file system files
- `NONE`: disable both asynchronous and direct I/O on file system files

See Also: Your platform-specific documentation for more details

Memory Usage

Memory usage is affected by both buffer cache limits and initialization parameters.

Buffer Cache Limits

The UNIX buffer cache consumes operating system memory resources. Although in some versions of UNIX the UNIX buffer cache may be allocated a set amount of memory, it is common today for more sophisticated memory management mechanisms to be used. Typically these will allow free memory pages to be used to cache I/O. In such systems it is common for operating system reporting tools to show that there is no free memory which is not generally a problem. If processes require more memory, the memory caching I/O data is usually released to allow the process memory to be allocated.

Parameters Affecting Memory Usage

The memory required by any one Oracle session depends on many factors. Typically the major contributing factors are:

- Number of open cursors
- Memory used by PLSQL, such as PLSQL tables
- `SORT_AREA_SIZE` initialization parameter

In Oracle9i, the `PGA_AGGREGATE_TARGET` initialization parameter gives greater control over a session's memory usage.

Using Process Schedulers

Many processes, or threads on NT systems, are involved in the operation of Oracle. They all access the shared memory resources in the SGA.

Be sure that all Oracle processes, both background and user processes, have the same process priority. When you install Oracle, all background processes are given the default priority for the operating system. Do not change the priorities of background processes. Verify that all user processes have the default operating system priority.

Assigning different priorities to Oracle processes might exacerbate the effects of contention. The operating system might not grant processing time to a low-priority process if a high-priority process also requests processing time. If a high-priority process needs access to a memory resource held by a low-priority process, then the high-priority process can wait indefinitely for the low-priority process to obtain the CPU, process the request, and release the resource.

Additionally, do not bind Oracle background processes to CPUs. This can cause the bound processes to be CPU-starved. This is especially the case when binding processes that fork off operating system threads. In this case, the parent process and all its threads bind to the CPU.

Using Operating System Resource Managers

Some platforms provide operating system resource managers. These are designed to reduce the impact of peak load use patterns by prioritizing access to system resources. They usually implement administrative policies that govern which resources users can access and how much of those resources each user is permitted to consume.

Operating system resource managers are different from domains or other similar facilities. Domains provide one or more completely separated environments within one system. Disk, CPU, memory, and all other resources are dedicated to each domain and cannot be accessed from any other domain. Other similar facilities completely separate just a portion of system resources into different areas, usually separate CPU or memory areas. Like domains, the separate resource areas are dedicated only to the processing assigned to that area; processes cannot migrate across boundaries. Unlike domains, all other resources (usually disk) are accessed by all partitions on a system.

Oracle runs within domains, as well as within these other less complete partitioning constructs, as long as the allocation of partitioned memory (RAM) resources is fixed, not dynamic.

Note: Oracle is not supported in any resource partitioned environment in which memory resources are assigned dynamically.

Operating system resource managers prioritize resource allocation within a global pool of resources, usually a domain or an entire system. Processes are assigned to groups, which are in turn assigned resources anywhere within the resource pool.

Note: Oracle is not supported for use with any operating system resource manager's memory management and allocation facility. Oracle Database Resource Manager, which provides resource allocation capabilities within an Oracle instance, cannot be used with any operating system resource manager.

Caution: When running under operating system resource managers, Oracle is supported only when each instance is assigned to a dedicated operating system resource manager group or managed entity. Also, the dedicated entity running all the instance's processes must run at one priority (or resource consumption) level. Management of individual Oracle processes at different priority levels is *not* supported. Severe consequences, including instance crashes, can result.

See Also:

- For a complete list of operating system resource management and resource allocation and deallocation features that work with Oracle and Oracle Database Resource Manager, see your systems vendor and your Oracle representative. Oracle does not certify these system features for compatibility with specific release levels.
- *Oracle9i Database Concepts* and *Oracle9i Database Administrator's Guide* for more information about Oracle Database Resource Manager

Solving Operating System Problems

This section provides hints for tuning various systems by explaining the following topics:

- [Performance Hints on UNIX-Based Systems](#)
- [Performance Hints on NT Systems](#)
- [Performance Hints on Midrange and Mainframe Computers](#)

Familiarize yourself with platform-specific issues so that you know what performance options the operating system provides.

See Also: Your Oracle platform-specific documentation and your operating system vendor's documentation

Performance Hints on UNIX-Based Systems

On UNIX systems, try to establish a good ratio between the amount of time the operating system spends fulfilling system calls and doing process scheduling and the amount of time the application runs. The goal should be to run 60% to 75% of the time in application mode (often called user mode) and 25% to 40% of the time in operating system mode. If you find that the system is spending 50% of its time in each mode, then determine what is wrong.

The ratio of time spent in each mode is only a symptom of the underlying problem, which might involve the following:

- Paging or swapping
- Executing too many operating system calls
- Running too many processes

If such conditions exist, then there is less time available for the application to run. The more time you can release from the operating system side, the more transactions an application can perform.

Performance Hints on NT Systems

On NT systems, as with UNIX-based systems, establish an appropriate ratio between time in application mode and time in system mode. On NT you can easily monitor many factors with Performance Monitor: CPU, network, I/O, and memory are all displayed on the same graph to assist you in avoiding bottlenecks in any of these areas.

Performance Hints on Midrange and Mainframe Computers

Consider the paging parameters on a mainframe, and remember that Oracle can exploit a very large working set.

Free memory in VAX or VMS environments is actually memory that is not mapped to any operating system process. On a busy system, free memory likely contains a page belonging to one or more currently active process. When that access occurs, a soft page fault takes place, and the page is included in the working set for the process. If the process cannot expand its working set, then one of the pages currently mapped by the process must be moved to the free set.

Any number of processes might have pages of shared memory within their working sets. The sum of the sizes of the working sets can thus markedly exceed the available memory. When the Oracle server is running, the SGA, the Oracle kernel code, and the Oracle Forms runtime executable are normally all sharable and account for perhaps 80% or 90% of the pages accessed.

Understanding CPU

To address CPU problems, first establish appropriate expectations for the amount of CPU resources your system should be using. Then, determine whether sufficient CPU resources are available and recognize when your system is consuming too many resources. Begin by determining the amount of CPU resources the Oracle instance utilizes with your system in the following three cases:

- System is idle, when little Oracle and non-Oracle activity exists
- System at average workloads
- System at peak workloads


You can capture various workload snapshots using Statspack or the UTLBSTAT/UTLESTAT utility. Operating system tools, such as `vmstat`, `sar`, and `iostat` on UNIX and Performance Monitor on NT, should be run during the same time interval as UTLBSTAT/UTLESTAT to provide a complimentary view of the overall statistics.

Note: Chapter 21, "Using Statspack" for more information on Statspack and UTLBSTAT/UTLESTAT

Workload is an important factor when evaluating your system's level of CPU utilization. During peak workload hours, 90% CPU utilization with 10% idle and

waiting time can be acceptable. Even 30% utilization at a time of low workload can be understandable. However, if your system shows high utilization at normal workload, then there is no room for a peak workload. For example, [Figure 16-1](#) illustrates workload over time for an application having peak periods at 10:00 AM and 2:00 PM.

Figure 16-1 Average Workload and Peak Workload



This example application has 100 users working 8 hours a day. Each user entering one transaction every 5 minutes translates into 9,600 transactions daily. Over an 8-hour period, the system must support 1,200 transactions an hour, which is an average of 20 transactions a minute. If the demand rate were constant, then you could build a system to meet this average workload.

However, usage patterns are not constant and in this context, 20 transactions a minute can be understood as merely a minimum requirement. If the peak rate you need to achieve is 120 transactions a minute, then you must configure a system that can support this peak workload.

For this example, assume that at peak workload, Oracle uses 90% of the CPU resource. For a period of average workload, then, Oracle uses no more than about 15% of the available CPU resource, as illustrated in the following equation:

$$20 \text{ tpm} / 120 \text{ tpm} * 90\% = 15\% \text{ of available CPU resource}$$

where t_{pm} is transactions a minute.

If the system requires 50% of the CPU resource to achieve 20 tpm, then a problem exists: the system cannot achieve 120 transactions a minute using 90% of the CPU. However, if you tuned this system so that it achieves 20 tpm using only 15% of the CPU, then, assuming linear scalability, the system might achieve 120 transactions a minute using 90% of the CPU resources.

As users are added to an application, the workload can rise to what had previously been peak levels. No further CPU capacity is then available for the new peak rate, which is actually higher than the previous.

CPU capacity issues can be addressed with the following:

- Tuning, or the process of detecting and solving CPU problems from excessive consumption

See Also: ["Finding System CPU Utilization"](#)

- Increasing hardware capacity, including changing the system architecture

See Also: *Oracle9i Database Performance Planning* for information about improving your system architecture

- Reducing the impact of peak load use patterns by prioritizing CPU resource allocation. Oracle's Database Resource Manager does this by allocating and managing CPU resources among database users and applications.

See Also: *Oracle9i Database Concepts* and *Oracle9i Database Administrator's Guide* for more information about Oracle's Database Resource Manager

Context Switching

Oracle has the several features for context switching, described in this section.

Post-wait Driver

An Oracle process needs to be able to post another Oracle process (give it a message) and also needs to be able to wait to be posted.

For example, a foreground process may need to post LGWR to tell it to write out all blocks up to a given point so that it can acknowledge a commit.

Often this post-wait mechanism is implemented through UNIX Semaphores, but these can be resource intensive. Therefore, some platforms supply a post-wait driver, typically a kernel device driver that is a lightweight method of implementing a post-wait interface.

Memory-mapped System Timer

Oracle often needs to query the system time for timing information. This can involve an operating system call that incurs a relatively costly context switch. Some platforms implement a memory-mapped timer that uses an address within the processes virtual address space to contain the current time information. Reading the time from this memory-mapped timer is less expensive than the overhead of a context switch for a system call.

List I/O Interfaces to Submit Multiple Asynchronous I/Os in One Call

List I/O is an application programming interface that allows several asynchronous I/O requests to be submitted in a single system call, rather than submitting several I/O requests through separate system calls. The main benefit of this feature is to reduce the number of context switches required.

Finding System CPU Utilization

Oracle statistics report CPU use by Oracle sessions only, whereas every process running on your system affects the available CPU resources. Therefore, tuning non-Oracle factors can also improve Oracle performance.

Use operating system monitoring tools to determine what processes are running on the system as a whole. If the system is too heavily loaded, check the memory, I/O, and process management areas described later in this section.

Tools such as `sar -u` on many UNIX-based systems let you examine the level of CPU utilization on your entire system. CPU utilization in UNIX is described in statistics that show user time, system time, idle time, and time waiting for I/O. A CPU problem exists if idle time and time waiting for I/O are both close to zero (less than 5%) at a normal or low workload.

On NT, use Performance Monitor to examine CPU utilization. Performance Manager provides statistics on processor time, user time, privileged time, interrupt time, and DPC time. (NT Performance Monitor is not the same as Performance Manager, which is an Oracle Enterprise Manager tool.)

Note: This section describes how to check system CPU utilization on most UNIX-based and NT systems. For other platforms, see your operating system documentation.

Checking Memory Management

Check the following memory management areas:

Paging and Swapping

Use tools such as `sar` or `vmstat` on UNIX or Performance Monitor on NT to investigate the cause of paging and swapping.

Oversize Page Tables

On UNIX, if the processing space becomes too large, then it can result in the page tables becoming too large. This is not an issue on NT.

Checking I/O Management

Thrashing is an I/O management issue. Ensure that your workload fits into memory, so the machine is not **thrashing** (swapping and paging processes in and out of memory). The operating system allocates fixed portions of time during which CPU resources are available to your process. If the process wastes a large portion of each time period checking to be sure that it can run and ensuring that all necessary components are in the machine, then the process might be using only 50% of the time allotted to actually perform work.

See Also: [Chapter 15, "I/O Configuration and Design"](#)

Checking Network Management

Check client/server round trips. There is an overhead in processing messages. When an application generates many messages that need to be sent through the network, the latency of sending a message can result in CPU overload. To alleviate this problem, bundle multiple messages together rather than perform lots of round trips. For example, you can use array inserts, array fetches, and so on.

Checking Process Management

Several process management issues discussed in this section should be checked.

Scheduling and Switching

The operating system can spend excessive time scheduling and switching processes. Examine the way in which you are using the operating system, because you could be using too many processes. On NT systems, do not overload your server with too many non-Oracle processes.

Context Switching

Due to operating system specific characteristics, your system could be spending a lot of time in context switches. Context switching can be expensive, especially with a large SGA. Context switching is not an issue on NT, which has only one process for each instance. All threads share the same page table.

Starting New Operating System Processes

There is a high cost in starting new operating system processes. Programmers often create single-purpose processes, exit the process, and create a new one. Doing this re-creates and destroys the process each time. Such logic uses excessive amounts of CPU, especially with applications that have large SGAs. This is because you need to build the page tables each time. The problem is aggravated when you pin or lock shared memory, because you have to access every page.

For example, if you have a 1 gigabyte SGA, then you might have page table entries for every 4 KB, and a page table entry might be 8 bytes. You could end up with $(1G / 4 \text{ KB}) * 8 \text{ byte}$ entries. This becomes expensive, because you need to continually make sure that the page table is loaded.

Configuring Instance Recovery Performance

This chapter offers guidelines for configuring the time to perform instance recovery.

This chapter contains the following sections:

- [Understanding Instance Recovery](#)
- [Checkpointing and Cache Recovery](#)
- [Reducing Checkpoint Frequency to Optimize Runtime Performance](#)
- [Configuring the Duration of Cache Recovery](#)
- [Monitoring Cache Recovery](#)
- [MTTR Advisory](#)
- [Tuning Transaction Recovery](#)

Understanding Instance Recovery

Instance and crash recovery are the automatic application of redo log records to Oracle data blocks after a crash or system failure. During normal operation, if an instance is shutdown cleanly as when using a `SHUTDOWN IMMEDIATE` statement, rather than terminated abnormally, then the in-memory changes that have not already been written to the datafiles on disk are written to disk as part of the checkpoint performed during shutdown.

However, if a single instance database crashes or if all instances of an Oracle Real Application Cluster configuration crash, then Oracle performs crash recovery at the next startup. If one or more instances of an Oracle Real Application Cluster configuration crash, then a surviving instance performs instance recovery automatically. Instance and crash recovery occur in two steps: cache recovery followed by transaction recovery.

Cache Recovery (Rolling Forward) During the cache recovery step, Oracle applies all committed and uncommitted changes in the redo log files to the affected data blocks. The work required for cache recovery processing is proportional to the rate of change to the database (update transactions each second) and the time between checkpoints.

Transaction Recovery (Rolling Back) To make the database consistent, the changes that were not committed at the time of the crash must be undone (in other words, rolled back). During the transaction recovery step, Oracle applies the rollback segments to undo the uncommitted changes. The work required to do transaction recovery is proportional to the number and size of uncommitted transactions when the system fault occurred.

Checkpointing and Cache Recovery

Periodically, Oracle records a **checkpoint**. A checkpoint is the highest system change number (SCN) such that all data blocks less than or equal to that SCN are known to be written out to the data files. If a failure occurs, then only the redo records containing changes at SCNs higher than the checkpoint need to be applied during recovery. The duration of cache recovery processing is determined by two factors: the number of data blocks that have changes at SCNs higher than the SCN of the checkpoint, and the number of log blocks that need to be read to find those changes.

How Checkpoints Affect Performance

Frequent checkpointing writes dirty buffers to the datafiles more often than otherwise, and so reduces cache recovery time in the event of an instance failure. If checkpointing is frequent, then applying the redo records in the redo log between the current checkpoint position and the end of the log involves processing relatively few data blocks. This means that the cache recovery phase of recovery is fairly short.

However, in a high-update system, frequent checkpointing can reduce runtime performance, because checkpointing causes *DBWn* processes to perform writes.

Fast Instance Recovery Trade-offs

To minimize the duration of instance recovery, you must force Oracle to checkpoint often, thus keeping the number of redo log records to be applied during recovery to a minimum. However, in a high-update system, frequent checkpointing increases the overhead for normal database operations.

If daily operational efficiency is more important than minimizing recovery time, then decrease the frequency of writes to data files due to checkpoints. This should improve operational efficiency, but also increase instance recovery time.

See Also:

- ["Reducing Checkpoint Frequency to Optimize Runtime Performance"](#) for information on how to maximize runtime performance
- ["Configuring the Duration of Cache Recovery"](#) for information on how to minimize instance recovery time

Reducing Checkpoint Frequency to Optimize Runtime Performance

To reduce the checkpoint frequency and optimize runtime performance, you can do the following:

- Size your online redo log files according to the amount of redo your system generates. A rough guide is to switch logs at most once every twenty minutes. Small log files can increase checkpoint activity and reduce performance. You should make all logs the same size.
- Set the value of the `LOG_CHECKPOINT_INTERVAL` initialization parameter (in multiples of physical block size) to zero. This value eliminates interval checkpoints.

- Set the value of the `LOG_CHECKPOINT_TIMEOUT` initialization parameter to zero. This value eliminates time-based checkpoints.
- Set the value of `FAST_START_MTTR_TARGET` (and `FAST_START_IO_TARGET`) to zero to disable fast-start checkpointing.
 - **See Also:** *Oracle9i Database Concepts* for a complete discussion of checkpoints

Configuring the Duration of Cache Recovery

There are several methods for tuning cache recovery to keep the duration of recovery within user-specified bounds. These include the following:

- [Initialization Parameters that Influence Cache Recovery Time](#)
- [Use Fast-Start Checkpointing to Limit Instance Recovery Time](#)
- [Set LOG_CHECKPOINT_TIMEOUT to Influence the Amount of Redo](#)
- [Set LOG_CHECKPOINT_INTERVAL to Influence the Amount of Redo](#)
- [Use Parallel Recovery to Speed up Redo Application](#)

Initialization Parameters that Influence Cache Recovery Time

The initialization parameters in [Example 17-1](#) influence cache recovery time.

Table 17-1 Initialization Parameters Influencing Cache Recovery

Parameter	Purpose
<code>FAST_START_MTTR_TARGET</code>	Lets you specify in seconds the expected mean time to recover (MTTR), which is the expected amount of time Oracle takes to perform recovery and startup the instance.
<code>FAST_START_IO_TARGET</code>	This initialization parameter has been deprecated in favour of <code>FAST_START_MTTR_TARGET</code> . This parameter specifies the upper limit on the number of dirty buffers.
<code>LOG_CHECKPOINT_TIMEOUT</code>	Limits the number of seconds between the most recent redo record and the checkpoint.
<code>LOG_CHECKPOINT_INTERVAL</code>	Limits the number of redo blocks generated between the most recent redo record and the checkpoint.
<code>RECOVERY_PARALLELISM</code>	Specifies the number of concurrent recovery processes to be used in instance or crash recovery.

Table 17–1 (Cont.) Initialization Parameters Influencing Cache Recovery

Parameter	Purpose
LOG_PARALLELISM	Specifies the level of concurrency for redo allocation within Oracle.

Note: Oracle recommends using the `FAST_START_MTTR_TARGET` initialization parameter to control the duration of startup after instance failure. Fast-start checkpointing is only available with Enterprise Edition.

The `FAST_START_IO_TARGET` initialization parameter has been deprecated in favor of the `FAST_START_MTTR_TARGET` parameter.

The initialization parameter `DB_BLOCK_MAX_DIRTY_TARGET` has been removed.

Use Fast-Start Checkpointing to Limit Instance Recovery Time

Oracle Enterprise Edition features include a fast-start fault recovery functionality to control instance recovery. This reduces the time required for cache recovery and makes the recovery bounded and predictable by limiting the number of dirty buffers and the number of redo records generated between the most recent redo record and the last checkpoint.

The foundation of fast-start recovery is the fast-start checkpointing architecture. Instead of the conventional event driven (that is, log switching) checkpointing, which does bulk writes, fast-start checkpointing occurs incrementally. Each DBWn process periodically writes buffers to disk to advance the checkpoint position. The oldest modified blocks are written first to ensure that every write lets the checkpoint advance. Fast-start checkpointing eliminates bulk writes and the resultant I/O spikes that occur with conventional checkpointing.

Administrators specify a target (bounded) time to complete the cache recovery phase of recovery with the `FAST_START_MTTR_TARGET` initialization parameter, and Oracle automatically varies the incremental checkpoint writes to meet that target. The `FAST_START_MTTR_TARGET` initialization parameter lets you specify in seconds the expected mean time to recover (MTTR), which is the expected amount of time Oracle takes to perform crash or instance recovery for a single instance.

FAST_START_MTTR_TARGET

The `FAST_START_MTTR_TARGET` initialization parameter simplifies the configuration of recovery time from instance or system failure. This parameter lets you specify the number of seconds crash or instance recovery is expected to take. The `FAST_START_MTTR_TARGET` is internally converted to a set of parameters that modify the operation of Oracle such that recovery time is as close to this estimate as possible.

Note: You should disable or remove the `FAST_START_IO_TARGET`, `LOG_CHECKPOINT_INTERVAL`, and `LOG_CHECKPOINT_TIMEOUT` initialization parameters when using `FAST_START_MTTR_TARGET`. Setting these parameters to active values interferes with `FAST_START_MTTR_TARGET`, resulting in a different than expected value in the `TARGET_MTTR` column of the `V$INSTANCE_RECOVERY` view.

The maximum value for `FAST_START_MTTR_TARGET` is 3600, or one hour. If you set the value to more than 3600, then Oracle rounds it to 3600. There is no minimum value for `FAST_START_MTTR_TARGET`. However, this does not mean that you can target the recovery time as low as you want. The time to do a crash recovery is limited by the low limit of the target number of dirty buffers, which is 1000, as well as factors such as how long initialization and file open take.

If you set the value of `FAST_START_MTTR_TARGET` too low, then the effective mean time to recover (MTTR) target will be the best MTTR target the system can achieve. If you set the value of `FAST_START_MTTR_TARGET` to such a high value that even in the worst-case recovery would not take that long, then the effective MTTR target will be the estimated MTTR in the worst-case scenario when the whole buffer cache is dirty. Use the `TARGET_MTTR` column in the `V$INSTANCE_RECOVERY` view to see the effective MTTR.

Note: The `TARGET_MTTR` column in `V$INSTANCE_RECOVERY` could be different than `FAST_START_MTTR_TARGET` if the latter is set too low or too high. Periodically check the `MTTR_TARGET` column in the `V$INSTANCE_RECOVERY` view and compare it with the parameter setting. Adjust the parameter setting if it is consistently different from the value in the target.

See Also: ["Monitoring Estimated MTTR: Example Scenario"](#) on page 17-10 for more information on setting `FAST_START_MTTR_TARGET`

Set `LOG_CHECKPOINT_TIMEOUT` to Influence the Amount of Redo

Set the initialization parameter `LOG_CHECKPOINT_TIMEOUT` to an integer value n to require that the latest checkpoint position follow the most recent redo block by no more than n seconds. In other words, at most, n seconds worth of logging activity can occur between the most recent checkpoint position and the last block written to the redo log. This forces the checkpoint position to keep pace with the most recent redo block.

You can also interpret `LOG_CHECKPOINT_TIMEOUT` as specifying an upper bound on the time a buffer can be dirty in the cache before `DBWn` must write it to disk. For example, if you set `LOG_CHECKPOINT_TIMEOUT` to 60, then no buffers remain dirty in the cache for more than 60 seconds. The default value for `LOG_CHECKPOINT_TIMEOUT` is 1800, or 30 minutes.

Set `LOG_CHECKPOINT_INTERVAL` to Influence the Amount of Redo

Set the initialization parameter `LOG_CHECKPOINT_INTERVAL` to a value n (where n is an integer) to require that the checkpoint position never follow the most recent redo block by more than n blocks. In other words, at most n redo blocks can exist between the checkpoint position and the last block written to the redo log. In effect, you are limiting the amount of redo blocks that can exist between the checkpoint and the end of the log.

Oracle limits the maximum value of `LOG_CHECKPOINT_INTERVAL` to 90% of the smallest log to ensure that the checkpoint advances into the current log before that log fills and a log switch is attempted.

`LOG_CHECKPOINT_INTERVAL` is specified in redo blocks. Redo blocks are the same size as operating system blocks. Use the `LOG_FILE_SIZE_REDO_BKLS` column in `V$INSTANCE_RECOVERY` to see the number of redo blocks corresponding to 90% of the size of the smallest log file.

See Also:

- ["Calculating Performance Overhead"](#) on page 17-12
- [Chapter 15, "I/O Configuration and Design"](#) for more information on tuning checkpoints

Use Parallel Recovery to Speed up Redo Application

Use parallel recovery to tune the cache recovery phase of recovery. Parallel recovery uses a division of labor approach to allocate different processes to different data blocks during the cache recovery phase of recovery.

For example, during recovery the redo log is read, and blocks that require redo application are parsed out. These blocks are subsequently distributed evenly to all recovery processes to be read into the buffer cache. Crash, instance, and media recovery of datafiles on different disk drives are good candidates for parallel recovery.

Use the `RECOVERY_PARALLELISM` initialization parameter to specify the number of concurrent recovery processes for instance or crash recovery. To use parallel processing, the value of `RECOVERY_PARALLELISM` must be greater than 1 and cannot exceed the value of the `PARALLEL_MAX_SERVERS` initialization parameter.

The `LOG_PARALLELISM` initialization parameter allows the parallel generation of redo and can increase the throughput of certain update-intensive workloads. If you are using Oracle on high-end servers that have more than 16 processors, and you are experiencing very high contention on the redo allocation latch, then you should consider enabling parallel redo.

Note: The `RECOVERY_PARALLELISM` initialization parameter specifies the number of concurrent recovery processes for instance or crash recovery *only*.

Media recovery is not affected by this parameter. Use the `PARALLEL` clause in the `RECOVER DATABASE` statement for media recovery.

Recovery is usually I/O bound on reads to data blocks. Consequently, parallelism at the block level can only help recovery performance if it speeds up total I/Os. Performance on systems with efficient asynchronous I/O typically does not improve significantly with parallel recovery, unless the recovery is CPU-bound.

See Also: *Oracle9i Database Reference* for more information on initialization parameters

Monitoring Cache Recovery

Use the `V$INSTANCE_RECOVERY` view to see the current recovery parameter settings. You can also use statistics from this view to calculate which parameter has the greatest influence on checkpointing. `V$INSTANCE_RECOVERY` contains the columns shown in [Table 17-2](#).

Note: The last three fields in `V$INSTANCE_RECOVERY` are new with Oracle9i, and they are the most important. With the initialization parameter `FAST_START_MTTR_TARGET`, the other seven fields of `V$INSTANCE_RECOVERY` are less useful.

Table 17-2 `V$INSTANCE_RECOVERY` View

Column	Description
<code>RECOVERY_ESTIMATED_IOS</code>	Contains the number of dirty buffers in the buffer cache. (In Standard Edition, the value of this field is always NULL).
<code>ACTUAL_REDO_BKLS</code>	Current number of redo blocks required to be read for recovery.
<code>TARGET_REDO_BKLS</code>	Goal for the maximum number of redo blocks to be processed during recovery. This value is the minimum of the next three columns (<code>LOG_FILE_SIZE_REDO_BKLS</code> , <code>LOG_CHKPT_TIMEOUT_REDO_BKLS</code> , <code>LOG_CHKPT_INTERVAL_REDO_BKLS</code>).
<code>LOG_FILE_SIZE_REDO_BKLS</code>	Number of redo blocks to be processed during recovery corresponding to 90% of the size of the smallest log file.
<code>LOG_CHKPT_TIMEOUT_REDO_BKLS</code>	Number of redo blocks that must be processed during recovery to satisfy <code>LOG_CHECKPOINT_TIMEOUT</code> .
<code>LOG_CHKPT_INTERVAL_REDO_BKLS</code>	Number of redo blocks that must be processed during recovery to satisfy <code>LOG_CHECKPOINT_INTERVAL</code> .
<code>FAST_START_IO_TARGET_REDO_BKLS</code>	This field is obsolete. It is retained for backward compatibility. The value of this field is always NULL.

Table 17–2 (Cont.) V\$INSTANCE_RECOVERY View

Column	Description
TARGET_MTTR	Effective mean time to recover (MTTR) target in seconds. Usually, it should be equal to the value of the FAST_START_MTTR_TARGET initialization parameter. If FAST_START_MTTR_TARGET is set to such a small value that it is impossible to do a recovery within its time frame, then the TARGET_MTTR field contains the effective MTTR target, which is larger than FAST_START_MTTR_TARGET. If FAST_START_MTTR_TARGET is set to such a high value that even in the worst-case (the whole buffer cache is dirty) recovery would not take that long, then the TARGET_MTTR field contains the estimated MTTR in the worst-case scenario. This field is 0 if FAST_START_MTTR_TARGET is not specified.
ESTIMATED_MTTR	The current estimated mean time to recover (MTTR) in the number of seconds based on the number of dirty buffers and log blocks (gives the current estimated MTTR even if FAST_START_MTTR_TARGET is not specified).
CKPT_BLOCK_WRITES	Number of writes to disk that would be avoided if checkpointing has been disabled.

See Also: *Oracle9i Database Reference* for more information on the V\$INSTANCE_RECOVERY view

Monitoring Estimated MTTR: Example Scenario

The TARGET_MTTR field of V\$INSTANCE_RECOVERY contains the MTTR target in effect. The ESTIMATED_MTTR field of V\$INSTANCE_RECOVERY contains the estimated MTTR should a crash happen right away. Query these two fields to see if the system can keep up with your specified MTTR target.

For example, assume the initialization parameter setting is as follows:

```
FAST_START_MTTR_TARGET = 6          # seconds
```

Execute the following query after database open:

```
SELECT TARGET_MTTR, ESTIMATED_MTTR, CKPT_BLOCK_WRITES
FROM V$INSTANCE_RECOVERY;
```

Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR CKPT_BLOCK_WRITES
18          15             0
```


You see that `TARGET_MTTR` is 18 seconds, which is higher than the value of `FAST_START_MTTR_TARGET` specified (6 seconds). This means that it is impossible to recover the database within 6 seconds. 18 seconds is the minimum MTTR target that the system can achieve.

The 18 second minimum is calculated based on the absolute low limit of 1000 blocks on the target of number of dirty buffers (The deprecated initialization parameter `FAST_START_IO_TARGET` follows this low limit; that is, you cannot set `FAST_START_IO_TARGET` to less than 1000). The `ESTIMATED_MTTR` field contains the estimated mean time to recovery. Because the database has just opened, the system contains few dirty buffers. That is why `ESTIMATED_MTTR` can be lower than the minimum possible `TARGET_MTTR`.

Now assume that you use the following statement to modify `FAST_START_MTTR_TARGET`:

```
ALTER SYSTEM SET FAST_START_MTTR_TARGET = 30;
```

Reissue the query to `V$INSTANCE_RECOVERY`, and Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR CKPT_BLOCK_WRITES
30          15              0
```

The `ESTIMATED_MTTR` field is still 15 seconds, which means that the estimated MTTR at the current time (should a crash happen immediately) is still 15 seconds. This is because no new redo is written, and no data block has become dirty.

Assume that heavy update activity occurs in the database and then you query `V$INSTANCE_RECOVERY` immediately afterward. Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR CKPT_BLOCK_WRITES
30          36             54367
```

You see that the effective MTTR target is 30 seconds. The estimated MTTR at the current time (should a crash happen immediately) is 36 seconds. This is fine. This means that some checkpoints writes might not have finished yet, so the buffer cache contains more dirty buffers than targeted.

Assume that you wait for one minute and reissue the query to `V$INSTANCE_RECOVERY`. Oracle responds with the following:

```
TARGET_MTTR ESTIMATED_MTTR CKPT_BLOCK_WRITES
30          31             55230
```

The estimated MTTR at this time has dropped to 31 seconds. This is because more dirty buffers have been written out during this period. This is shown by the increase of CKPT_BLOCK_WRITES field of V\$INSTANCE_RECOVERY.

Note: The number of physical writes minus the number of physical writes non checkpoint (from V\$SYSSTAT) equals the field CKPT_BLOCK_WRITES in V\$INSTANCE_RECOVERY.

Calculating Performance Overhead

To calculate performance overhead, use the V\$SYSSTAT view. For example, assume that you execute the following query:

```
SELECT NAME, VALUE
FROM V$SYSSTAT
WHERE NAME IN ('physical reads','physical writes',
              'physical writes non checkpoint');
```

Oracle responds with the following:

NAME	VALUE
physical reads	2376
physical writes	14932
physical writes non checkpoint	11165

The first row shows the number of data blocks retrieved from disk. The second row shows the number of data blocks written to disk. The last row shows the number of writes to disk that would occur if you turned off checkpointing.

Use this data to calculate the overhead imposed by setting the FAST_START_MTTR_TARGET initialization parameter. To effectively measure the percentage of extra writes, mark the values for these statistics at different times, t_1 and t_2 .

Calculate the percentage of extra I/Os generated by fast-start checkpointing using the following formula:

$$[(PW_2 - PW_1) - (PWNC_2 - PWNC_1)] / ((PR_2 - PR_1) + (PW_2 - PW_1)) \times 100\% = EIO$$

where the variables are described in [Table 17-3](#).

Table 17-3 Variable Definitions

Variable	Definition
*_1	Value of prefixed variable at time $t_{_1}$, which is any time after the database has been running for a while
*_2	Value of prefixed variable at time $t_{_2}$, which is later than $t_{_1}$ and not immediately after changing any of the checkpoint parameters
PWNC	physical writes non checkpoint
PW	physical writes
PR	physical reads
EIO	Percentage of estimated extra I/Os generated by enabling checkpointing

It can take some time for database statistics to stabilize after instance startup or dynamic initialization parameter modification. After such events, wait until all blocks age out of the buffer cache at least once before taking measurements. If the percentage of extra I/Os is too high, then increase the value of `FAST_START_MTTR_TARGET`.

The number of extra writes caused by setting `FAST_START_MTTR_TARGET` to a nonzero value is application-dependent; it is not dependent on cache size.

Calculating Performance Overhead: Example Scenario

As an example, assume the initialization parameter setting is as follows:

```
FAST_START_MTTR_TARGET = 90 # 90 seconds
```

After the statistics stabilize, you issue this query on `V$SYSSTAT`:

```
SELECT NAME, VALUE
FROM V$SYSSTAT
WHERE NAME IN ('physical reads','physical writes',
              'physical writes non checkpoint');
```

Oracle responds with the following:

```
NAME                                VALUE
physical reads                       2376
physical writes                      14932
physical writes non checkpoint       11165
```

The physical write checkpoint statistics can also be found in the CKPT_BLOCK_WRITES field of the V\$INSTANCE_RECOVERY view. For example:

```
SELECT CKPT_BLOCK_WRITES
FROM V$INSTANCE_RECOVERY;
```

Oracle responds with the following:

```
CKPT_BLOCK_WRITES      3767
```

It is consistent with the result from V\$SYSSTAT: 3767 = 14932 - 11165.

After making updates for a few hours, you reissue the query. Oracle responds with the following:

NAME	VALUE
physical reads	3011
physical writes	17467
physical writes non checkpoint	13231

Substitute the values from the SELECT statements in the formula on page 17-12 to determine how much performance overhead you are incurring:

$$(((17467 - 14932) - (13231 - 11165)) / ((3011 - 2376) + (17467 - 14932))) \times 100\% = 14.8\%$$

As the result indicates, enabling fast-start checkpointing generates about 15% more I/O than required had you not enabled fast-start checkpointing. After calculating the extra I/O, you decide you can afford more system overhead if you decrease recovery time.

To decrease recovery time, reduce the value for the parameter FAST_START_MTTR_TARGET to 60. After items in the buffer cache age out, calculate V\$SYSSTAT statistics across a second interval to determine the new performance overhead.

Query V\$SYSSTAT:

```
SELECT NAME, VALUE FROM V$SYSSTAT
WHERE NAME IN ('physical reads', 'physical writes',
'physical writes non checkpoint');
```

Oracle responds with the following:

NAME	VALUE
physical reads	4652
physical writes	28864
physical writes non checkpoint	21784

After making updates, reissue the query. Oracle responds with the following:

NAME	VALUE
physical reads	6000
physical writes	35394
physical writes non checkpoint	26438

Calculate how much performance overhead you are incurring using the values from the two SELECT statements:

$$[(35394 - 28864) - (26438 - 21784)] / ((6000 - 4652) + (35394 - 28864)) \times 100\% = 23.8\%$$

After changing the parameter, the percentage of I/Os performed by Oracle is now about 24% more than it would be if you disabled fast-start checkpointing.

Calibrating the MTTR

The `FAST_START_MTTR_TARGET` initialization parameter calculates internal system trigger values to limit the length of the redo log and the number of dirty data buffers in the data cache. This calculation uses estimated times to read a redo block and to read and write a data block.

Initially, internal defaults are used. These defaults are replaced by execution time estimates during system operation. However, the best values are obtained from measurements taken from an actual recovery from a failure.

Note: To effectively align `FAST_START_MTTR_TARGET`, make sure that you perform several instance recoveries to ensure that the time to read a redo block and the time to read and write a data block are recorded accurately.

Before doing instance recoveries to calibrate the `FAST_START_MTTR_TARGET`, decide whether `FAST_START_MTTR_TARGET` is being calibrated for a database crash or a hardware crash. This is a consideration if your database files are stored in a file system or if your I/O subsystem has a memory cache, because there is a considerable difference in the read and write time to disk depending on whether or not the files are cached. The workload being run during the instance recovery should be a very good representation of the average workload on the system to ensure that the amount of redo records generated are similar.

MTTR Advisory

Starting with Oracle9i Release 2 (9.2), MTTR advisory is available to help you evaluate the effect of different MTTR settings on system performance in terms of extra physical writes.

How MTTR Advisory Works

When MTTR advisory is enabled, after the system runs a typical workload for a while, you can query `V$MTTR_TARGET_ADVICE`, which tells you the ratio of estimated number of cache writes under other MTTR settings to the number of cache writes under the current MTTR. For instance, a ratio of 1.2 indicates 20% more cache writes.

By looking at the different MTTR settings and their corresponding cache write ratio, you can decide which MTTR value fits your recovery and performance needs. `V$MTTR_TARGET_ADVICE` also gives the ratio on total physical writes (including direct writes), and the ratio on total I/Os (including reads).

Enabling MTTR Advisory

Enabling MTTR Advisory involves setting two initialization parameters:

- `STATISTICS_LEVEL`
- `FAST_START_MTTR_TARGET`

`STATISTICS_LEVEL`

Make sure that `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`.

`FAST_START_MTTR_TARGET`

To enable MTTR advisory, set the initialization parameter `FAST_START_MTTR_TARGET` to a nonzero value. If `FAST_START_MTTR_TARGET` is not specified, then MTTR advisory will be `OFF`.

When MTTR advisory is `ON`, it simulates checkpoint queue behavior under five different MTTR settings:

- Current `FAST_START_MTTR_TARGET` setting
- Current setting times 0.1, 0.5, 1.5, and 2

Note: Whenever you set `FAST_START_MTTR_TARGET` to a nonzero value, and while MTTR advisory is ON, Oracle Corporation recommends that you disable (set to 0) the following parameters:

- `LOG_CHECKPOINT_TIMEOUT`
- `LOG_CHECKPOINT_INTERVAL`
- `FAST_START_IO_TARGET`

Because these initialization parameters either override `FAST_START_MTTR_TARGET` or potentially drive checkpoints more aggressively than `FAST_START_MTTR_TARGET` does, they can interfere with the simulation.

Viewing MTTR Advisory

Oracle9i Release 2 (9.2) provides a dynamic performance view for viewing statistics or advisories collected by MTTR advisory.

V\$MTTR_TARGET_ADVICE

If MTTR advisory has been turned on, `V$MTTR_TARGET_ADVICE` shows the advisory information collected. Usually this view show five rows, corresponding to the current MTTR, 0.1 times the current MTTR, 0.5 times the current MTTR, 1.5 times the current MTTR and 2 times the current MTTR. However, if one or more of the 5 values are less than the smallest MTTR target the system can sustain, their corresponding rows are replaced with a single row corresponding to the smallest MTTR target the system can have. Similarly, if one or more of the 5 values are larger than the worst-case MTTR target the system can have, their corresponding rows are replaced with a single row corresponding to the worst-case MTTR target the system can have.

If MTTR advisory is currently OFF, the view shows information collected the last time MTTR advisory was on.

See Also: "[V\\$MTTR_TARGET_ADVICE](#)" on page 24-21 for column details of these views

Tuning Transaction Recovery

During the second phase of instance recovery, Oracle rolls back uncommitted transactions. Oracle uses two features, fast-start on-demand rollback and fast-start parallel rollback, to increase the efficiency of this recovery phase.

Note: These features are part of fast-start fault recovery and are only available in the Oracle9i Enterprise Edition.

This section contains the following topics:

- [Using Fast-Start On-Demand Rollback](#)
- [Using Fast-Start Parallel Rollback](#)

Using Fast-Start On-Demand Rollback

Using the fast-start on-demand rollback feature, Oracle automatically allows new transactions to begin as soon as the database opens, which is usually a very short time after cache recovery completes. If a user attempts to access a row that is locked by a terminated transaction, Oracle rolls back only those changes necessary to complete the transaction; in other words, it rolls them back *on demand*. Consequently, new transactions do not have to wait until all parts of a long transaction are rolled back.

Note: Oracle does this automatically. You do not need to set any parameters or issue statements to use this feature.

Using Fast-Start Parallel Rollback

In fast-start parallel rollback, the background process SMON acts as a coordinator and rolls back a set of transactions in parallel using multiple server processes. Essentially, fast-start parallel rollback is to transaction recovery what parallel recovery is to cache recovery.

Fast-start parallel rollback is mainly useful when a system has transactions that run a long time before committing, especially parallel `INSERT`, `UPDATE`, and `DELETE` operations. SMON automatically decides when to begin parallel rollback and disperses the work among several parallel processes: process one rolls back one transaction, process two rolls back a second transaction, and so on.

One special form of fast-start parallel rollback is intra-transaction recovery. In intra-transaction recovery, a single transaction is divided among several processes. For example, assume eight transactions require recovery with one parallel process assigned to each transaction. The transactions are all similar in size except for transaction five, which is quite large. This means it takes longer for one process to roll this transaction back than for the other processes to roll back their transactions.

In this situation, Oracle automatically begins intra-transaction recovery by dispersing transaction five among the processes: process one takes one part, process two takes another part, and so on.

You control the number of processes involved in transaction recovery by setting the initialization parameter `FAST_START_PARALLEL_ROLLBACK` to one of three values listed in [Table 17-4](#).

Table 17-4 *FAST_START_PARALLEL_ROLLBACK Parameter Values*

Value	Meaning
FALSE	Turns off fast-start parallel rollback.
LOW	Specifies that the number of recovery servers cannot exceed twice the value of the <code>CPU_COUNT</code> initialization parameter.
HIGH	Specifies that the number of recovery servers cannot exceed four times the value of the <code>CPU_COUNT</code> initialization parameter.

Parallel Rollback in an Oracle Real Application Clusters Configuration

In Oracle Real Application Clusters, you can perform fast-start parallel rollback on each instance. Within each instance, you can perform parallel rollback on transactions that are:

- Online on a given instance
- Offline and not being recovered on instances other than the given instance

After a rollback segment is online for a given instance, only this instance can perform parallel rollback on transactions on that segment.

Monitoring Progress of Fast-Start Parallel Rollback

Monitor the progress of fast-start parallel rollback by examining the `V$FAST_START_SERVERS` and `V$FAST_START_TRANSACTIONS` views. `V$FAST_START_SERVERS` provides information about all recovery processes performing fast-start parallel rollback. `V$FAST_START_TRANSACTIONS` contains data about the progress of the transactions.

See Also:

- *Oracle9i Database Administrator's Guide* for more information on managing undo space
- *Oracle9i Real Application Clusters Deployment and Performance* for more information on fast-start parallel rollback in an Oracle Real Application Clusters environment
- *Oracle9i Database Reference* for more information about initialization parameters
- *Oracle9i Net Services Administrator's Guide* for information on transparent application failure (TAF)

Configuring Undo and Temporary Segments

There are performance considerations when configuring undo and temporary segments.

This chapter contains the following topics:

- [Configuring Undo Segments](#)
- [Configuring Temporary Tablespaces](#)

Configuring Undo Segments

Oracle provides automatic undo management, which completely automates the management of undo data. A database running in automatic undo management mode transparently creates and manages undo segments. Oracle Corporation strongly recommends using automatic undo management, because it significantly simplifies database management and removes the need for any manual tuning of undo (rollback) segments. Manual undo management using rollback segments is supported for backward compatibility reasons.

Configuring Automatic Undo Management

To configure automatic undo, you can simply include the following initialization parameter:

```
UNDO_MANAGEMENT=AUTO
```

You can also create an undo tablespace, and determine the maximum retention time for undo data kept in that tablespace.

See Also: *Oracle9i Database Administrator's Guide* for detailed information on how to configure automatic undo

Configuring Rollback Segments

Automatic undo management is the preferred way of handling rollback space. Automatic undo management lets you allocate undo space in a single undo tablespace, instead of distributing undo space in a set of statically allocated rollback segments. The creation and allocation of space among the undo segments is handled automatically by the Oracle server.

With rollback segments, one or more tablespaces are created; rollback segments are manually created in those tablespaces. The number and size of the rollback segments must be determined by the DBA.

Determining the Number and Size of Rollback Segments

The size of rollback segments can affect performance. Rollback segment size is determined by the rollback segment's storage parameter values. Your rollback segments must be large enough to hold the rollback entries for your transactions. As with other objects, avoid dynamic space management in rollback segments.

Table 18–1 shows some general guidelines for choosing how many rollback segments to allocate based on the number of concurrent transactions on your database. These guidelines are appropriate for most application mixes.

Table 18–1 *Choosing the Number of Rollback Segments*

Number of Current Transactions (<i>n</i>)	Number of Rollback Segments Recommended
$n < 16$	4
$16 \leq n < 32$	8
$32 \leq n$	$n/4$

Use the `SET TRANSACTION` statement to assign transactions to rollback segments of the appropriate size based on the recommendations in the following sections. If you do not explicitly assign a transaction to a rollback segment, then Oracle automatically assigns it to a rollback segment.

For example, the following statement assigns the current transaction to the rollback segment `oltp_13`:

```
SET TRANSACTION USE ROLLBACK SEGMENT oltp_13
```

Note: If you are running multiple concurrent copies of the same application, then be careful not to assign the transactions for all copies to the same rollback segment. This leads to contention for that rollback segment.

Also, monitor the shrinking, or dynamic deallocation, of rollback segments based on the `OPTIMAL` storage parameter.

See Also: *Oracle9i Database Administrator's Guide* for information on choosing values for this parameter, monitoring rollback segment shrinking, and adjusting the `OPTIMAL` parameter

For Long Queries Assign large rollback segments to transactions that modify data that is concurrently selected by long queries. Such queries might require access to rollback segments to reconstruct a read-consistent version of the modified data. The rollback segments must be large enough to hold all the rollback entries for the data while the query is running.

For Long Transactions Assign large rollback segments to transactions that modify large amounts of data. A large rollback segment can improve the performance of such a transaction, because the transaction generates large rollback entries. If a rollback entry does not fit into a rollback segment, then Oracle extends the segment. Dynamic extension reduces performance; avoid it whenever possible.

For OLTP Transactions OLTP applications are characterized by frequent concurrent transactions, each of which modifies a small amount of data. Assign OLTP transactions to small rollback segments, as long as their data is not concurrently queried. Small rollback segments are more likely to remain stored in the buffer cache where they can be accessed quickly. A typical OLTP rollback segment might have two extents, each approximately 10 kilobytes in size. To best avoid contention, create many rollback segments and assign each transaction to its own rollback segment.

Configuring Temporary Tablespaces

Configuring the temporary tablespace helps optimize disk sort performance. This involves choosing good storage clauses and the correct type of tablespace to use for sorting.

Choosing the default storage clause for the sort tablespace includes the following:

- Setting `PCTINCREASE` to zero
- Setting `INITIAL` and `NEXT` to the same size and a factor of `SORT_AREA_SIZE`

Choosing the correct type of tablespace makes disk sorting more efficient. The various tablespaces that could be used for disk sorting include the following:

- [Temporary Tablespaces](#)
- [Tablespaces of Type TEMPORARY](#)
- [Permanent Tablespaces](#)

Temporary Tablespaces These are the most efficient tablespaces for disk sorts. Characteristics of a temporary tablespace include the following:

- Space management (extent allocation and deallocation) is locally managed. Therefore, use of the `ST-enqueue` is avoided.
- The sort segment created for each instance is reused (it is only dropped if the tablespace is dropped). With Oracle Real Application Clusters, one sort segment is required for each instance.

- All processes performing sorts reuse existing sort extents of the sort segment, rather than allocating a segment (and potentially many extents) for each sort. If an insufficient number of extents exist for the number of sorts currently in operation, then the required extents are added once for each instance startup. They are recycled thereafter.
- Temporary tablespaces can be striped, as described in "[Basic I/O Configuration](#)" on page 15-6. Only when the tablespaces is creating an I/O bottleneck should the tablespace be segregated from others.
- To create temporary tablespaces, use the `CREATE TEMPORARY TABLESPACE` statement.

See Also:

- *Oracle9i Database Concepts* for more information on temporary tablespaces
- *Oracle9i SQL Reference* for more information on using the `CREATE TEMPORARY TABLESPACE` statement

Tablespaces of Type TEMPORARY After temporary tablespaces, these are next best tablespaces to use for sort operations. Characteristics of tablespaces of type `TEMPORARY` include the following:

- Space management is dictionary managed so they use the ST-enqueue.
- The sort segment for an instance is dropped on instance startup and re-created when the first sort is performed. Extents are then allocated as needed.
- Although space management is dictionary-managed, the frequency of space allocation and deallocation is reduced. All processes performing sorts reuse the existing sort extents of a single sort segment, rather than allocating and deallocating extents and segments for each sort. With Oracle Real Application Clusters, one sort segment is required for each instance. If insufficient extents exist for the number of sorts, then the required extents are added once for each instance startup, and are recycled thereafter. This in turn reduces the load on the ST enqueue considerably.
- To create tablespaces of type `TEMPORARY`, use the `CREATE TABLESPACE` or `ALTER TABLESPACE` statements with the `TEMPORARY` clause.

See Also: *Oracle9i SQL Reference* for more information on using the `TEMPORARY` clause

Permanent Tablespaces Permanent tablespaces (which are not of type `TEMPORARY`) are least efficient for performance of disk sorts. This is because of the following reasons:

- The ST-enqueue is used for allocation and de-allocation of each extent allocated to a sort segment.
- Sort-segments are not reused. Each process performing a disk sort creates then drops it's own sort segment. In addition, a single sort operation can require the allocation and deallocation of many extents, and each extent allocation requires the ST-enqueue.

For optimal performance, the best choice is to use temporary tablespaces. They bypass the need of using the ST enqueue for space management and have the additional benefit of reuse of sort extents.

Table 18–2 Sort Tablespaces

Tablespace Type	Space Management	Sort Segment Reuse	Sort Extent Reuse	Creation Statement
Temporary (uses tempfiles)	Locally Managed	Yes. Sort segment is not dropped	Always	CREATE TEMPORARY TABLESPACE
Type <code>TEMPORARY</code> (uses datafiles)	Dictionary Managed	Sort segment is re-created on instance startup	Yes. Extents are reused until sort segment is dropped on instance startup	CREATE TABLESPACE ... TYPE TEMPORARY
Permanent	Dictionary Managed	No. Sort segments are not reused	No. Sort segments are not reused	CREATE TABLESPACE

Temporary tablespaces, and tablespaces of type `TEMPORARY`, cannot contain permanent objects, such as tables or rollback segments.

See Also: *Oracle9i SQL Reference* for more information about the syntax of the `CREATE TABLESPACE` and `ALTER TABLESPACE` statements

Configuring Shared Servers

Proper configuration of shared servers can result in significant performance improvement.

This chapter contains the following topic:

- [Introduction to Shared Server Performance](#)
- [Configuring the Number of Shared Servers](#)

Introduction to Shared Server Performance

Using shared servers enables you to reduce the number of processes and the amount of memory consumed on the server machine. Shared servers are beneficial for systems where there are many OLTP users performing intermittent transactions.

Using shared servers rather than dedicated servers is also generally better for systems that have a high connection rate to the database. With shared servers, when a connect request is received, a dispatcher is already available to handle concurrent connection requests. With dedicated servers, on the other hand, a connection-specific dedicated server is sequentially initialized for each connection request.

Performance of certain database features can improve when a shared server architecture is used, and performance of certain database features can degrade slightly when a shared server architecture is used. For example, a session can be prevented from migrating to another shared server while parallel execution is active.

A session can remain nonmigratable even after a request from the client has been processed, because not all the user information has been stored in the UGA. If a server were to process the request from the client, then the part of the user state that was not stored in the UGA would be inaccessible. To avoid this, individual shared servers often need to remain bound to a user session.

Configuring the Number of Shared Servers

When using some features, you may need to configure more shared servers, because some servers might be bound to sessions for an excessive amount of time.

This section discusses how to reduce contention for processes used by Oracle's architecture:

- [Identifying Contention Using the Dispatcher-Specific Views](#)
- [Reducing Contention for Dispatcher Processes](#)
- [Reducing Contention for Shared Servers](#)
- [Determining the Optimal Number of Dispatchers and Shared Servers](#)

Identifying Contention Using the Dispatcher-Specific Views

The following views provide dispatcher performance statistics:

- `V$DISPATCHER` - general information about dispatcher processes
- `V$DISPATCHER_RATE` - dispatcher processing statistics

See Also:

- *Oracle9i Database Reference* for detailed information about these views
- *Oracle Enterprise Manager Concepts Guide* for information about Oracle Tuning Pack applications that monitor statistics

Analyzing `V$DISPATCHER_RATE` Statistics

The `V$DISPATCHER_RATE` view contains current, average, and maximum dispatcher statistics for several categories. Statistics with the prefix `CUR_` are statistics for the current sample. Statistics with the prefix `AVG_` are the average values for the statistics since the collection period began. Statistics with the prefix `MAX_` are the maximum values for these categories since statistics collection began.

To assess dispatcher performance, query the `V$DISPATCHER_RATE` view and compare the current values with the maximums. If your present system throughput provides adequate response time and current values from this view are near the average and less than the maximum, then you likely have an optimally tuned shared server environment.

If the current and average rates are significantly less than the maximums, then consider reducing the number of dispatchers. Conversely, if current and average rates are close to the maximums, then you might need to add more dispatchers. A general rule is to examine `V$DISPATCHER_RATE` statistics during both light and heavy system use periods. After identifying your shared server load patterns, adjust your parameters accordingly.

If needed, you can also mimic processing loads by running system stress tests and periodically polling the `V$DISPATCHER_RATE` statistics. Proper interpretation of these statistics varies from platform to platform. Different types of applications also can cause significant variations on the statistical values recorded in `V$DISPATCHER_RATE`.

Reducing Contention for Dispatcher Processes

This section discusses how to add dispatcher processes and how to enable connection pooling.

Adding Dispatcher Processes

Add dispatcher processes while Oracle is running with the `SET` option of the `ALTER SYSTEM` statement to increase the value for the `DISPATCHERS` initialization parameter.

The total number of dispatcher processes is limited by the value of the initialization parameter `MAX_DISPATCHERS`. You might need to increase this value before adding dispatcher processes. The default value of this parameter is five, and the maximum value varies depending on your operating system.

See Also: *Oracle9i Database Administrator's Guide* and *Oracle9i Net Services Administrator's Guide* for more information on adding dispatcher processes

Enabling Connection Pooling

When system load increases and dispatcher throughput is maximized, it is not necessarily a good idea to immediately add more dispatchers. Instead, consider configuring the dispatcher to support more users with connection pooling.

`DISPATCHERS` lets you enable various attributes for each dispatcher. Oracle supports a name-value syntax to let you specify attributes in a position-independent, case-insensitive manner. For example:

```
DISPATCHERS = " (PROTOCOL=TCP) (POOL=ON) "
```

The optional attribute `POOL` enables the Oracle Net connection pooling feature. `TICK` is the size of a network `TICK` in seconds. The `TICK` default is 1 second.

See Also: *Oracle9i Database Reference* and the *Oracle9i Net Services Administrator's Guide* for more information about the `DISPATCHERS` parameter and its options

Enabling Session Multiplexing

Multiplexing is used by a connection manager process to establish and maintain network sessions from multiple users to individual dispatchers. For example, several user processes can connect to one dispatcher by way of a single connection from a connection manager process.

The connection manager manages communication from users to the dispatcher by way of a shared connection. At any one time, zero, one, or a few users might need the connection, while other user processes linked to the dispatcher by way of the connection manager process are idle. This way, session multiplexing is beneficial because it maximizes use of the dispatcher process connections.

Multiplexing is also useful for multiplexing database link sessions between dispatchers. The limit on the number of sessions for each dispatcher is platform dependent. For example:

```
DISPATCHERS=" (PROTOCOL=TCP) (MULTIPLEX=ON) "
```

Reducing Contention for Shared Servers

This section discusses how to identify contention for shared servers and how to increase the maximum number of shared servers.

Identifying Contention for Shared Servers

Steadily increasing wait times in the requests queue indicate contention for shared servers. To examine wait time data, use the dynamic performance view `V$QUEUE`. This view contains statistics showing request queue activity for shared servers. By default, this view is available only to the user `SYS` and to other users with `SELECT ANY TABLE` system privilege, such as `SYSTEM`. [Table 19–1](#) lists the columns showing the wait times for requests and the number of requests in the queue.

Table 19–1 *Wait Time and Request Columns in V\$QUEUE*

Column	Description
WAIT	Displays the total waiting time, in hundredths of a second, for all requests that have ever been in the queue
TOTALQ	Displays the total number of requests that have ever been in the queue

Monitor these statistics occasionally while your application is running by issuing the following SQL statement:

```
SELECT DECODE(TOTALQ, 0, 'No Requests',
             WAIT/TOTALQ || ' HUNDREDTHS OF SECONDS')
       "AVERAGE WAIT TIME PER REQUESTS"
FROM V$QUEUE
WHERE TYPE = 'COMMON';
```

This query returns the results of a calculation that show the following:

```
AVERAGE WAIT TIME PER REQUEST
-----
.090909 HUNDREDTHS OF SECONDS
```

From the result, you can tell that a request waits an average of 0.09 hundredths of a second in the queue before processing.

You can also determine how many shared servers are currently running by issuing the following query:

```
SELECT COUNT(*) "Shared Server Processes"
FROM V$SHARED_SERVER
WHERE STATUS != 'QUIT';
```

The result of this query could look like the following:

```
Shared Server Processes
-----
10
```

If you detect resource contention with shared servers, then first make sure that this is not a memory contention issue by examining the shared pool and the large pool. If performance remains poor, then you might want to create more resources to reduce shared server process contention. Do this by modifying the optional server process parameters, as explained in the following section.

Setting and Modifying Shared Server Processes

This section explains how to set optional parameters affecting processes for the shared server architecture. This section also explains how and when to modify these parameters to tune performance.

The following static initialization parameters are discussed in this section:

- MAX_DISPATCHERS
- MAX_SHARED_SERVERS

This section also describes the following initialization/session parameters:

- DISPATCHERS
- SHARED_SERVERS

Values for the initialization parameters MAX_DISPATCHERS and MAX_SHARED_SERVERS define upper limits for the number of dispatchers and servers running on

an instance. These parameters are static and cannot be changed after your database is running. You can create as many dispatcher and server processes as you need, but the total number of processes cannot exceed the host operating system's limit for the number of running processes.

Note: Setting `MAX_DISPATCHERS` sets the limit on the number of dispatchers for all `DISPATCHERS`' dispatcher values.

You can also define starting values for the number of dispatchers and servers by setting the `DISPATCHERS` parameter's `DISPATCHER` attribute and the `SHARED_SERVERS` parameter. After system startup, you can dynamically reset values for these parameters to change the number of dispatchers and servers using the `SET` option of the `ALTER SYSTEM` statement. If you enter values for these parameters in excess of limits set by the static parameters, then Oracle uses the static parameter values.

The default value of `MAX_SHARED_SERVERS` is dependent on the value of `SHARED_SERVERS`. If `SHARED_SERVERS` is less than or equal to 10, then `MAX_SHARED_SERVERS` defaults to 20. If `SHARED_SERVERS` is greater than 10, then `MAX_SHARED_SERVERS` defaults to two times the value of `SHARED_SERVERS`.

Self-adjusting Shared Server Architecture Features

When the database starts, `SHARED_SERVERS` is the number of shared servers created. Oracle does not allow the number of shared servers to be less than this minimum. During processing, Oracle automatically adds shared servers up to the limit defined by `MAX_SHARED_SERVERS` if Oracle perceives that the load based on the activity of the requests on the common queue warrant additional shared servers. Therefore, you are unlikely to improve performance by explicitly adding shared servers. However, you might need to adjust your system to accommodate certain resource issues.

If the number of shared server processes has reached the limit set by the initialization parameter `MAX_SHARED_SERVERS` and the average wait time in the request queue is still unacceptable, then you might improve performance by increasing the `MAX_SHARED_SERVERS` value.

If resource demands exceed expectations, then you can either allow Oracle to automatically add shared server processes or you can add shared processes by altering the value for `SHARED_SERVERS`. You can change the value of this parameter in the initialization parameter file, or alter it using the `SHARED_SERVERS`

parameter of the `ALTER SYSTEM` statement. Experiment with this limit and monitor shared servers to determine an ideal setting for this parameter.

Increasing the Maximum Number of Shared Servers

The shared servers are the processes that perform data access and pass back this information to the dispatchers.

The dispatchers then forward the data to the client process. If there are not enough shared servers to handle all the requests, then the queue backs up (`V$QUEUE`), and requests take longer to process. However, before you check the `V$QUEUE` statistics, it is best to first check if you are running out of shared servers.

Find out the amount of free RAM in the system. Examine `ps` or any other operating system utility to find out the amount of memory a shared server uses. Divide the amount of free RAM by the size of a shared server. This gives you the maximum number of shared servers you can add to your system.

The best way to proceed is to increase the `MAX_SHARED_SERVERS` parameter gradually until you begin to swap. If swapping occurs due to the shared server, then reduce the number until swapping stops, or increase the amount of physical RAM. Because each operating system and application is different, the only way to find out the ideal setting for `MAX_SHARED_SERVERS` is through trial and error.

To change the `MAX_SHARED_SERVERS`, first edit the initialization parameter file. Save the file and restart the instance. Remember that setting `SHARED_SERVERS` to `MAX_SHARED_SERVERS` should only be done if you are sure that you want to fix the number of shared server processes. Keep in mind the following rules:

- `SHARED_SERVERS` should be set for slightly greater than the expected number of shared servers that will be needed when the database is at an average load.
- `MAX_SHARED_SERVERS` should be set for slightly greater than the expected number of shared servers that will be needed when the database is at an peak load.

Determining the Optimal Number of Dispatchers and Shared Servers

As mentioned, `SHARED_SERVERS` determines the number of shared servers activated at instance startup. The default setting for `SERVER_SERVERS` is one when `DISPATCHERS` is specified.

To determine the optimal number of dispatchers and shared servers, consider the number of users typically accessing the database and how much processing each requires. Also consider that user and processing loads vary over time. For example,

a customer service system's load might vary drastically from peak OLTP-oriented daytime use to DSS-oriented nighttime use. System use can also predictably change over longer time periods, such as the loads experienced by an accounting system that vary greatly from mid-month to month-end.

If each user makes relatively few requests over a given period of time, then each associated user process is idle for a large percentage of time. In this case, one shared server process can serve 10 to 20 users. If each user requires a significant amount of processing, then establish a higher ratio of servers to user processes.

In the beginning, it is best to allocate fewer shared servers. Additional shared servers start automatically as needed and are deallocated automatically if they remain idle too long. However, the initial servers always remain allocated, even if they are idle.

If you set the initial number of servers too high, then your system might incur unnecessary overhead. Experiment with the number of initial shared servers and monitor shared servers until you achieve ideal system performance for your typical database activity.

Estimating the Maximum Number of Dispatcher Processes

Use values for `MAX_DISPATCHERS` and `DISPATCHERS` that are at least equal to the maximum number of concurrent sessions divided by the number of connections for each dispatcher. For most systems, a value of 1,000 connections for each dispatcher provides good performance.

Disallowing Further Shared Server Use with Concurrent Shared Server Use

You can use the `SET` option of the `ALTER SYSTEM` statement to alter the number of active, shared servers. To prevent additional users from accessing shared servers, set `SHARED_SERVERS` to zero. This temporarily disables additional use of shared servers. Resetting `SHARED_SERVERS` to a positive value enables shared servers for all current users.

To prevent a dispatcher from being used, issue the statement `ALTER SYSTEM SHUTDOWN [IMMEDIATE]`. To bring the dispatcher back online, use the `ALTER SYSTEM SET` command for the `DISPATCHERS` parameter.

See Also:

- *Oracle9i Database Reference* for information about dispatchers, specifically the description of the V\$DISPATCHER and V\$DISPATCHER_RATE views
- *Oracle9i SQL Reference* for more information about the ALTER SYSTEM statement
- *Oracle9i Database Administrator's Guide* for more information on changing the number of shared servers

Part IV

System-Related Performance Tools

Part IV provides information about Oracle's system-related performance tools.

The chapters in this part are:

- [Chapter 20, "Oracle Tools to Gather Database Statistics"](#)
- [Chapter 21, "Using Statspack"](#)

Oracle Tools to Gather Database Statistics

This chapter explains why performance data gathering is important, and it describes how to use available Oracle tools.

This chapter contains the following sections:

- [Overview of Tools](#)
- [Principles of Data Gathering](#)
- [Interpreting Statistics](#)
- [Oracle Enterprise Manager Diagnostics Pack](#)
- [Statspack](#)
- [V\\$ Performance Views](#)

Overview of Tools

Effective data collection and analysis is essential for identifying and correcting system performance problems. Oracle provides a number of tools that allow a performance engineer to gather information regarding instance and database performance.

- Oracle Enterprise Manager Diagnostics Pack, with a graphical user interface, is the most feature-rich performance tool. It provides data analysis and collection of operating system statistics.
- Statspack and BSTAT/ESTAT are command-line interface tools that gather instance related performance data. Statspack is the successor of BSTAT/ESTAT, with significantly increased functionality over the BSTAT/ESTAT tool.
- V\$ views can be queried using SQL. They contain dynamic performance data related to an Oracle instance's performance. The data is lost when the instance is shut down.

Principles of Data Gathering

To effectively diagnose a performance problem, it is vital to have an established performance baseline for later comparison when the system is running poorly. Without a baseline data point, it can be very difficult to identify new problems. For example, perhaps the volume of transactions on the system has increased, or the transaction profile or application has changed, or the number of users has increased.

Although Oracle Enterprise Manager, Statspack and BSTAT/ESTAT, and the V\$ views have different interfaces (GUI, command line, SQL), the majority of data they collect and report on is extracted from the V\$ views. Because the V\$ views are based on memory-resident data, when an instance is shut down, the data related to the instance is lost.

In order to perform analysis of data from one day to the next, data visible through the V\$ views must be saved. On each instance startup, the memory resident V\$ views are reinitialized; hence, to determine what has changed within any particular period, you must calculate the difference in the performance data. This is done by subtracting the statistic values at the beginning of the period from the statistic values at the end of the period. This gives you the activity of the instance during that period, or the *delta*. The delta of each statistic can then be normalized (for example, over seconds or over transactions).

The delta of *all* statistics for a period of time can be considered a baseline, assuming the response time and operations performed on the instance were representative of some typical load on your system (that is, batch, online, or both).

Each of the tools Oracle provides (except for the `v$` views themselves) has a mechanism for saving this data and for determining the delta.

It is also important to gather operating system and network statistics. These can then be correlated with the Oracle performance data. If you are using Statspack, `BSTAT/ESTAT`, or your own tool, you should devise a mechanism for collecting operating system statistics. With Oracle Enterprise Manager, this capability is built-in.

Interpreting Statistics

When initially examining performance data, you can formulate potential theories by examining your statistics. One way to ensure that your interpretation of the statistics is correct is to perform cross-checks with other data. This establishes whether a statistic or event is really of interest.

Some pitfalls are discussed in the following sections:

- Hit ratios

When tuning, it is common to compute a ratio that helps determine whether there is a problem. Such ratios include the buffer cache hit ratio, the soft-parse ratio, and the latch hit ratio. These ratios should not be used as 'hard and fast' identifiers of whether there is or is not a performance bottleneck. Rather, they should be used as indicators. In order to identify whether there is a bottleneck, other related evidence should be examined.

- Wait events with timed statistics

Setting `TIMED_STATISTICS` to true at the instance level directs the Oracle server to gather wait time for events, in addition to wait counts already available. This data is useful for comparing the total wait time for an event to the total elapsed time between the performance data collections. For example, if the wait event accounts for only 30 seconds out of a two hour period, then there is probably little to be gained by investigating this event, even though it may be the highest ranked wait event when ordered by time waited. However, if the event accounts for 30 minutes of a 45 minute period, then the event is worth investigating.

Note: Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

- Comparing Oracle statistics with other factors

When looking at statistics, it is important to consider other factors that influence whether the statistic is of value. Such factors include the user load and the hardware capability. Even an event that had a wait of 30 minutes in a 45 minute snapshot might not be indicative of a problem if you discover that there were 2000 users on the system, and the host hardware was a 64 node machine.

- Wait events without timed statistics

If `TIMED_STATISTICS` is false, then the amount of time waited for an event is not available. Therefore, it is only possible to order wait events by the number of times each event was waited for. Although the events with the largest number of waits might indicate the potential bottleneck, they might not be the main bottleneck. This can happen when an event is waited for a large number of times, but the total time waited for that event is small. The converse is also true: an event with fewer waits might be a problem if the wait time is a significant proportion of the total wait time. Without having the wait times to use for comparison, it is difficult to determine whether a wait event is really of interest.

- Idle wait events

Oracle uses some wait events to indicate if the Oracle server process is idle. Typically, these events are of no value when investigating performance problems, and they should be ignored when examining the wait events.

- Computed statistics

When interpreting computed statistics (such as rates, statistics normalized over transactions, or ratios), it is important to cross-verify the computed statistic with the actual statistic counts. This confirms whether the derived rates are

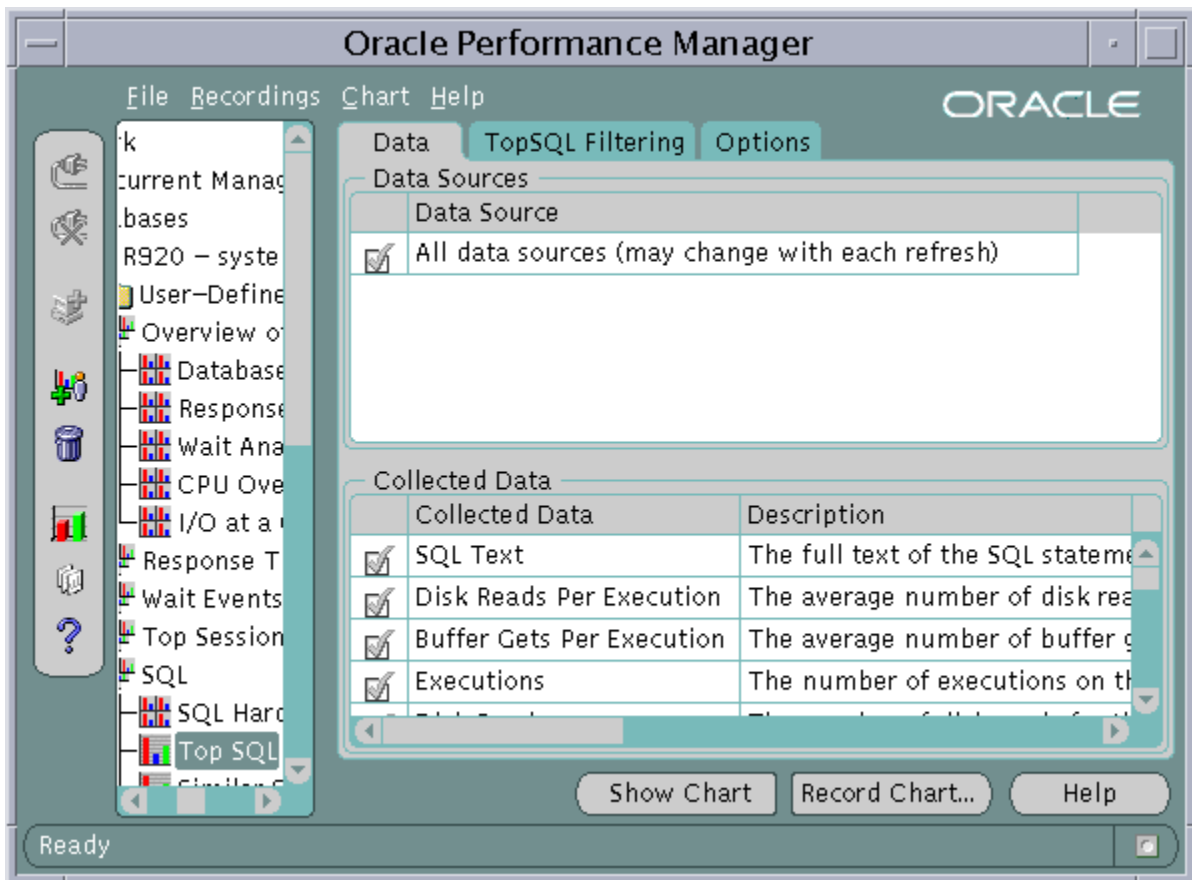
really of interest: small statistic counts usually can discount an unusual ratio. For example, on initial examination, a soft-parse ratio of 50% generally indicates a potential tuning area. If, however, there was only one hard parse and one soft parse during the data collection interval, then the soft-parse ratio would be 50%, even though the statistic counts show this is not an area of concern. In this case, the ratio is not of interest due to the low raw statistic counts.

See Also: ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Oracle Enterprise Manager Diagnostics Pack

Oracle Enterprise Manager (EM) Diagnostics Pack captures related operating system, middle-tier, and application performance data, in addition to instance performance data, allowing end-to-end diagnostics.

The Diagnostics Pack can automatically analyze this performance data, display it in a graphical interface, and use alerts to immediately direct you to any performance problems. You can be alerted automatically through email or page when a problem is detected. Oracle Enterprise Manager also includes an integrated diagnostics methodology that uses guided drilldowns and expert advice to help you quickly resolve performance issues. [Figure 20-1](#) is a screen of the Oracle Performance Manager.

Figure 20–1 Oracle Performance Manager

EM also lets you store the captured data in a separate performance repository database. You can store the performance data for multiple databases in the same repository.

The EM Intelligent Agent data gathering service collects performance data on a scheduled basis. A single agent can manage the data collections for all Oracle databases and the operating system of the target node. The data is automatically stored in an historical data repository for performance reporting. Data stored in the repository can be used to analyze many facets of database performance, such as database load, cache allocations and efficiency, resource contention, and high-impact SQL.

Performance data collections can be initiated directly from the EM Console or through the Capacity Planner application of the Oracle Enterprise Manager Diagnostics Pack. HTML reports of historical performance data can be generated from the EM Console. These reports provide a comprehensive analysis of database system usage and performance, which can be easily accessed and navigated from a browser. EM also provides a graphical real-time Performance Overview for monitoring a subset of these performance metrics using line charts, bar graphs, and so forth.

The Performance Overview charts let you troubleshoot existing performance problems by drilling into performance data to track down the source of a performance bottleneck. For example, a decline in the memory sort percentage can be immediately investigated by drilling down to the sessions and corresponding SQL responsible for high volume sort activity. High-impact SQL statements discovered through this process can be further investigated by launching SQL diagnostic tools in the context of the problem.

See Also:

- *Oracle Enterprise Manager Concepts Guide*
- *Getting Started with the Oracle Diagnostics Pack*

Statspack

Statspack fundamentally differs from the well known UTLBSTAT/UTLESTAT tuning scripts by collecting more information, and by storing the performance statistics data permanently in Oracle tables, which can later be used for reporting and analysis. The data collected can be analyzed using the report provided, which includes an instance health and load summary page, high resource SQL statements, as well as the traditional wait events and initialization parameters.

See Also: [Chapter 21, "Using Statspack"](#)

V\$ Performance Views

The v\$ views are the performance information sources used by all Oracle performance tuning tools. The v\$ views are based on memory structures initialized at instance startup. The memory structures, and the views that represent them, are automatically maintained by Oracle throughout the life of the instance.

See Also: [Chapter 24, "Dynamic Performance Views for Tuning"](#)

If you choose not to use an Oracle tool to gather your performance data, then you need to develop your own. You need to save data from the required performance views to a structure on disk, so that the data can be analyzed and compared with other data collected. Because there are multiple collections, you need a key to identify each collection. This method is very similar to the method used by Statspack.

The following is an example of how to save performance data from a single V\$ view to an Oracle table. To implement your own collection tool, you must perform a similar collection mechanism for all essential V\$ views.

Note: Oracle recommends using Oracle Enterprise Manager Diagnostics Pack or Statspack to gather performance data. These tools have been designed to capture all of the data needed for performance analysis.

Example - Saving File I/O Data

The following example creates a table that stores the collection data. The table has all V\$FILESTAT columns and also includes the collection ID column and the collection date column. A sample SQL statement that reports the delta in this table for the first two collections has also been included.

Create the table, and insert the first data collection:

```
CREATE TABLE coll_filestat AS
  SELECT 1      coll_id      -- collection number
         , sysdate coll_date  -- collection date
         , fs.*
  FROM V$FILESTAT fs;

ALTER TABLE coll_filestat add
  (CONSTRAINT coll_filestat PRIMARY KEY (coll_id, file#));
```

At the end of the interval, insert second collection:

```
INSERT INTO coll_filestat
  SELECT 2      -- collection number
         , sysdate -- collection date
         , fs.*
  FROM V$FILESTAT fs;
```

Note: In order to insert any other collections, you need to keep the collection ID key unique. This could be done using a sequence number. (The sequence number should have value for all v\$ views captured in one collection.)

To query for high I/O tablespaces:

```
SELECT t.tablespace_name
      , SUM(fs2.phyrds-fs1.phyrds)
        / MAX(86400*(fs2.coll_date-fs1.coll_date)) "Rd/sec"
      , SUM(fs2.phyblkrd-fs1.phyblkrd)
        / MAX(86400*(fs2.coll_date-fs1.coll_date)) "Blk/sec"
      , SUM(fs2.phywrt-fs1.phywrt)
        / MAX(86400*(fs2.coll_date-fs1.coll_date)) "Wr/sec"
      , SUM(fs2.phyblkwrt-fs1.phyblkwrt)
        / MAX(86400*(fs2.coll_date-fs1.coll_date)) "Blk/sec"
  FROM coll_filestat fs1, coll_filestat fs2, dba_data_files t
 WHERE fs2.file#   = fs1.file#
       AND fs2.coll_id = fs1.coll_id + 1
       AND t.file_id = fs2.file#
 GROUP BY t.tablespace_name
 ORDER BY sum(fs2.phyblkrd+fs2.phyblkwrt-fs1.phyblkrd-fs1.phyblkwrt) DESC;
```

The following is a sample output from the preceding select statement:

TABLESPACE_N	Rd/sec	Blk/sec	Wr/sec	Blk/sec
AP_T_02	287.1	2245.7	.0	.0
PO_T_01	313.5	650.6	.2	.2
RECEIVABLE_T	401.0	613.8	2.4	2.4
INV_T_01	154.3	155.3	.0	.0
APPLSYS_T_01	63.3	139.6	.4	.4
PA_T_01	102.3	102.3	.0	.0
SO_I_01	63.4	63.4	34.5	34.5
TEMP	2.3	45.0	1.9	47.0
RECEIVABLE_I	73.0	73.0	.1	.1
AP_T_03	69.3	69.3	.0	.0
RECEIVABLE_I	65.1	65.1	1.9	1.9
SO_T_01	54.0	57.8	2.9	2.9
SYSTEM	45.2	59.0	.3	.3
PER_T_01	48.0	58.7	.0	.0
AP_T_01	12.9	51.0	.2	.2
SO_T_03	43.0	43.0	1.2	1.2

V\$ Performance Views

PER_I_01	30.8	30.8	.0	.0
FA_T_01	22.3	22.3	.0	.0
INV_I_01	20.7	20.7	.7	.7
PO_I_01	19.5	19.5	.7	.7
GSR_T_01	19.2	19.2	.4	.4
INV_I_03	18.3	18.3	.0	.0
ROLL_01	1.4	1.4	14.7	14.7
PA_I_01	14.3	14.3	.2	.2

Using Statspack

This chapter explains how to install, configure, and use Statspack.

The chapter contains the following sections:

- [Introduction to Statspack](#)
- [Statspack Compared with BSTAT/ESTAT](#)
- [How Statspack Works](#)
- [Configuring Database Space Requirements for Statspack](#)
- [Installing Statspack](#)
- [Using Statspack](#)
- [Removing Statspack](#)
- [Statspack Supplied Scripts and Documentation](#)

Note: Statspack is not supported with releases earlier than 8.1.6. Storing data from multiple databases in one `PERFSTAT` user account is currently not supported.

See Also: "[Statspack Documentation](#)" on page 21-28 for information about the `SPDOC.TXT` file, installed with the Oracle database

Introduction to Statspack

When tuning a database, it is important to have an established baseline for later comparison when the system is running poorly. A baseline data point helps identify the factors to check when diagnosing new performance problems. Some factors to check are:

- Has the volume of transactions on the system increased?
- Has the transaction profile or application changed?
- Has the number of users increased?

The Statspack package is a set of SQL, PL/SQL, and SQL*Plus scripts that allow the collection, automation, storage, and viewing of performance data. Statspack stores the performance statistics permanently in Oracle tables, which can later be used for reporting and analysis. The data collected can be analyzed using Statspack reports, which includes an instance health and load summary page, high resource SQL statements, and the traditional wait events and initialization parameters.

See Also: Oracle provides a diagnostics pack which contains GUI tools for collecting and analyzing statistics. See "[Oracle Enterprise Manager Diagnostics Pack](#)" on page 20-5.

Statspack Compared with BSTAT/ESTAT

Statspack differs from the existing UTLBSTAT/UTLESTAT performance scripts in the following ways:

- Statspack collects more data, including high-resource SQL.
- Statspack precalculates many ratios useful when performance tuning, such as cache hit ratios, rates, and transaction statistics. Many of these ratios must be calculated manually when using BSTAT/ESTAT.
- Statspack uses permanent tables owned by the PERFSTAT user to store performance statistics. Rather than creating and dropping tables each time, data is inserted into the existing tables, making it easier to compare historical data.
- Statspack separates data collection from report generation. Data is collected when a snapshot is taken. The performance engineer then runs the performance report and views the data collected.

Note: The term **snapshot** as used here denotes a set of performance statistics gathered at a single time, identified by a unique ID that includes the snapshot number (or `SNAP_ID`). This kind of snapshot has nothing to do with Oracle's snapshot replication technology.

- Statspack makes data collection easy to automate using either `DBMS_JOB` or an operating system utility to schedule collection tasks.
- Statspack considers a transaction to finish either with a `COMMIT` or a `ROLLBACK`, and so calculates the number of transactions as `user commits + user rollbacks`. `BSTAT/ESTAT` considers a transaction to complete with a `COMMIT` only, and so assumes that `transactions = user commits`. For this reason, comparing statistics for each transaction between Statspack and `BSTAT/ESTAT` can result in significantly different ratios.

Caution: If you choose to run `BSTAT/ESTAT` in conjunction with Statspack, do not run both as the same user. There is a name conflict with the `STATS$WAITSTAT` table.

How Statspack Works

When you run the Statspack installation script, the `PERFSTAT` user is created automatically. `PERFSTAT` owns all objects needed by the Statspack package and is granted limited query-only privileges on the `V$` views required for performance tuning.

Statspack users become familiar with the concept of a **snapshot**, a single collection of performance data. Each snapshot taken is identified by a snapshot ID, which is a unique number generated at the time the snapshot is taken. Each time a new collection is taken, a new `SNAP_ID` is generated.

The `SNAP_ID`, along with the database identifier (`DBID`) and instance number (`INSTANCE_NUMBER`), comprise the unique key for a snapshot. Use of this unique combination allows storage of multiple instances of an Oracle Real Application Clusters database in the same tables.

After snapshots are taken, you can run the performance report. The report prompts you for start and end snapshot IDs and then calculates activity on the instance between the two snapshots, much like a `BSTAT/ESTAT` report. To compare, the first

`SNAP_ID` supplied can be considered the equivalent of running `BSTAT`; the second `SNAP_ID` specified can be considered the equivalent of `ESTAT`. Unlike `BSTAT/ESTAT`, which can by its nature only compare two static data points, the report can compare any two snapshots specified.

Configuring Database Space Requirements for Statspack

The default initial and next extent sizes are 100K, 1MB, or 5MB for all Statspack tables and indexes. Approximately 64MB is required to install Statspack.

The amount of database space required by the Statspack package depends on the frequency of snapshots, the size of the database and instance, and the amount of data collected, which can be configured. It is therefore difficult to provide general storage clauses and space utilization predictions that are accurate at each site.

- If you install the package in a dictionary-managed tablespace, then you should monitor the space used by the objects created and, if required, adjust the storage clauses of the segments.
- If you install the package in a locally managed tablespace, then storage clauses are not required, because the storage characteristics are automatically managed.

Installing Statspack

There are two ways to install Statspack:

- [Interactive Statspack Installation](#)
- [Batch Mode Statspack Installation](#)

Batch mode is useful when you do not want to be prompted for the `PERFSTAT` user's password, default tablespace, and temporary tablespace.

Interactive Statspack Installation

The first step in the installation is the creation of the `PERFSTAT` user, which owns all PL/SQL code and database objects created, including the Statspack tables, constraints, and the Statspack package. During installation, you are prompted for the `PERFSTAT` user's password, default tablespace, and temporary tablespace. The default tablespace is used to create all Statspack objects, such as tables and indexes. The temporary tablespace is used for sort-type activities.

See Also: *Oracle9i Database Concepts* for more information on temporary tablespaces

Note:

- A password is mandatory and should be kept confidential.
 - Do not specify the `SYSTEM` tablespace for the `PERFSTAT` user's `DEFAULT` or `TEMPORARY` tablespaces. If `SYSTEM` is specified, the installation aborts with an error specifying the problem. Oracle Corporation does not recommend using the `SYSTEM` tablespace to store statistics data or for sorting. Use a `TOOLS` tablespace to store the data, and use your instance's `TEMP` tablespace for sorting. To recover from this error, run the `de-install (SPDROP.SQL)` script, then rerun the installation.
 - During installation, the `DBMS_SHARED_POOL` and `DBMS_JOB` PL/SQL packages are created. `DBMS_SHARED_POOL` pins the Statspack package in the shared pool. `DBMS_JOB` is created on the assumption that you want to schedule periodic snapshots automatically using `DBMS_JOB`.
-
-

When installing the Statspack package, you can either change to the `ORACLE_HOME/rdbms/admin` directory, or fully specify the `ORACLE_HOME/rdbms/admin` directory when calling the installation script, `SPCREATE.SQL`.

To install Statspack, perform the following:

- Start SQL*Plus.
- Connect as a user with `SYSDBA` privilege. For example:

```
SQL> CONNECT / AS SYSDBA
```
- Run the `SPCREATE.SQL` script.
 - On UNIX platforms, enter the following:

```
SQL> @?/rdbms/admin/spcreate
```
 - On Windows platforms, enter the following:

```
SQL> @%ORACLE_HOME%\rdbms\admin\spcreate
```
- Enter appropriate information when prompted for the `PERFSTAT` user's password, default tablespace, and temporary tablespace.

The `SPCREATE.SQL` install script runs the following scripts automatically:

- `SPCUSR.SQL`: Creates the user and grants privileges
- `SPCTAB.SQL`: Creates the tables
- `SPCPKG.SQL`: Creates the package

To ensure that no errors were encountered during the installation, check the `SPCUSR.LIS`, `SPCTAB.LIS`, and `SPCPKG.LIS` output files. For example:

```
ORACLE_HOME/bin/spcusr.lis
ORACLE_HOME/bin/spctab.lis
ORACLE_HOME/bin/spcpkg.lis
```

Batch Mode Statspack Installation

To install Statspack in batch mode, you must assign values to the SQL*Plus variables that specify the default and temporary tablespaces before running `SPCREATE.SQL`. The variables are:

- `DEFAULT_TABLESPACE`: For the default tablespace
- `TEMPORARY_TABLESPACE`: For the temporary tablespace
- `PERFSTAT_PASSWORD`: For the `PERFSTAT` user

For example, on UNIX:

```
SQL> CONNECT / AS SYSDBA
SQL> define default_tablespace='TOOLS'
SQL> define temporary_tablespace='TEMP'
SQL> define perfstat_password='my_perfstat_password'
SQL> @?/rdbms/admin/spcreate
```

When `SPCREATE.SQL` is run, it does not prompt for the information provided by the variables.

Note: After the setup is complete, change the password of the `PERFSTAT` user for security purposes.

Using Statspack

The following topics are discussed in this section:

- [Taking a Statspack Snapshot](#)
- [Automating Statistics Gathering](#)

- [Running a Statspack Performance Report](#)
- [Configuring the Amount of Data Captured in Statspack](#)
- [Time Units Used for Wait Events](#)
- [Event Timings](#)[Managing and Sharing Statspack Performance Data](#)
- [Oracle Real Application Clusters Considerations with Statspack](#)

Taking a Statspack Snapshot

The simplest interactive way to take a snapshot is to login to SQL*Plus as the PERFSTAT user and run the procedure STATSPACK.SNAP. For example:

```
SQL> CONNECT perfstat/my_perfstat_password
SQL> EXECUTE statspack.snap;
```

Note: In an Oracle Real Application Clusters environment, you must connect to the instance for which you want to collect data.

Taking such a snapshot stores the current values for the performance statistics in the Statspack tables. This snapshot can be used as a baseline for comparison with another snapshot taken at a later time.

For better performance analysis, set the initialization parameter TIMED_STATISTICS to TRUE. Statspack will then include important timing information in the data it collects. You can change the TIMED_STATISTICS parameter dynamically by using the ALTER SYSTEM statement. Timing data is important and is usually required by Oracle support to diagnose performance problems.

Note: Timed statistics are automatically collected for the database if the initialization parameter STATISTICS_LEVEL is set to TYPICAL or ALL. If STATISTICS_LEVEL is set to BASIC, then you must set TIMED_STATISTICS to TRUE to enable collection of timed statistics.

If you explicitly set DB_CACHE_ADVICE, TIMED_STATISTICS, or TIMED_OS_STATISTICS, either in the initialization parameter file or by using ALTER_SYSTEM or ALTER_SESSION, the explicitly set value overrides the value derived from STATISTICS_LEVEL.

See Also: ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Typically, to automate the gathering and reporting phases (during a benchmark, for example), you might need to know the `snap_id` of the snapshot just taken. To take a snapshot and display the `snap_id`, call the `STATSPACK.SNAP` function.

Example 21–1 Calling the `snap` Function in SQL*Plus

Using an anonymous PL/SQL block,

```
SQL> variable snap number;
SQL> begin :snap := statspack.snap; end;
2 /
PL/SQL procedure successfully completed.
SQL> print snap
        SNAP
-----
        12
```

Automating Statistics Gathering

To make performance comparisons from one day, week, or year to the next, you need multiple snapshots taken over a period of time. The best method to gather snapshots is to automate the collection at regular intervals. You have the following options:

- Within the database, use the Oracle `DBMS_JOB` procedure to schedule snapshots.
- Use an operating system utility, such as `cron` on UNIX, to schedule snapshots. On Windows platforms, you can use the `at` utility on Windows NT or `System Tools > Scheduled Tasks` on Windows 2000.

Using `DBMS_JOB` to Collect Statistics

The `DBMS_JOB` package is an Oracle-automated method for scheduling and running different tasks, such as collecting statistics. A sample script on how to do this is supplied in `SPAUTO.SQL`, which schedules a snapshot every hour, on the hour.

You might want to schedule snapshots at regular times each day to reflect your system's OLTP or batch peak loads. For example, you could take snapshots at 9 a.m.

10 a.m, 11 a.m, 12 noon, and 6 p.m. for the OLTP load, and then take a snapshot at 12 midnight and 6 a.m for the batch window.

In order to use `DBMS_JOB` to schedule snapshots, you must set the `JOB_QUEUE_PROCESSES` initialization parameter to greater than 0 in the initialization file, so that the job can be run automatically.

Note: If you are using `SPAUTO.SQL` in an Oracle Real Application Clusters environment, you must run the `SPAUTO.SQL` script once on each instance in the cluster. Similarly, you must set the `JOB_QUEUE_PROCESSES` parameter for each instance.

Changing the Interval of Statistics Collection

Use the `DBMS_JOB.INTERVAL` procedure to change the interval of statistics collection. For example:

```
EXECUTE DBMS_JOB.INTERVAL( job_number, 'SYSDATE+(1/48)');
```

Where `'SYSDATE+(1/48)'` results in the statistics being gathered each 1/48 hours, or every half hour, and `job_number` refers to the specific job that you want to run.

To force a job to run immediately:

```
EXECUTE DBMS_JOB.RUN( job_number);
```

To remove a specified job:

```
EXECUTE DBMS_JOB.REMOVE( job_number);
```

See Also: *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on the `DBMS_JOB` package

Running a Statspack Performance Report

After snapshots are taken, you can generate performance reports. The SQL scripts that generate the reports prompts you for a beginning snapshot ID, an ending snapshot ID, and a report name. The Statspack package includes two reports.

- First, run a Statspack report, `SPREPORT.SQL`, which is a general instance health report that covers all aspects of instance performance. This reports calculates and prints ratios and differences for all statistics between the two snapshots, similar to the `BSTAT/ESTAT` report.

- After examining the instance report, run a SQL report, `SPREPSQL.SQL`, on a single SQL statement (identified by its hash value). The SQL report only reports on data relating to the single SQL statement.

Note: It is not correct to specify begin and end snapshots where the begin snapshot and end snapshot were taken from different instance startups. In other words, the instance must not have been shutdown between the times that the begin and end snapshots were taken.

This is necessary because the database's dynamic performance tables, which Statspack queries to gather the data, reside in memory. Hence, shutting down the database resets the values in the performance tables to 0. Because Statspack subtracts the begin-snapshot statistics from the end-snapshot statistics, the resulting output is invalid. If begin and end snapshots taken between shutdowns are specified in the report, then the report shows an appropriate error to indicate this.

Because data gathering is separate from report production, you have flexibility to base a report on any data points you select. For example, as DBA you might want to use the supplied automation script to automate data collection every hour, on the hour. If, at some later point, a performance issue arose that might be better investigated by looking at a three-hour data window, all you have to do is specify the required start point and end point when running the report.

Running the Statspack Report

To examine the change in instance-wide statistics between two time periods, the `SPREPORT.SQL` script is run while connected to the `PERFSTAT` user. The `SPREPORT.SQL` script is located in the `rdbms/admin` directory of the Oracle home.

Note: In an Oracle Real Application Clusters environment, you must connect to the instance on which you want to report.

When the report is run, you are prompted for the following:

- The beginning snapshot ID
- The ending snapshot ID

- The name of the report text file to be created

Note: Both the serial number and the session_id (SID) must be the same for the begin and end snapshots. Blank lines between lines of snapshot IDs means that the instance has been restarted (shutdown/startup) between those times, changing the serial number. The blank lines thus identify begin and end snapshots that cannot be used together when running a Statspack report.

[Example 21-2](#) shows the SQL commands to run the report and an example of the partial report output.

Example 21-2 Creating a Statspack Report with Prompts

```
SQL> connect perfstat/my_perfstat_password
SQL> @?/rdbms/admin/spreport
```

On Windows platforms, the command to run the report is:

```
SQL> @%ORACLE_HOME%\rdbms\admin\spreport
```

Sample output:

```
SQL> connect perfstat/my_perfstat_password
Connected.
SQL> @?/rdbms/admin/spreport
```

```
DB Id          DB Name          Inst Num Instance
-----
2618106428 PRD1                1 prd1
Completed Snapshots

Instance      DB Name          Snap
              DB Name          Id   Snap Started   Snap
-----
prd1          PRD1              1 11 May 2000 12:07   5
              PRD1              2 11 May 2000 12:08   5
```

```
Specify the Begin and End Snapshot Ids
```

```
~~~~~
```

```
Enter value for begin_snap: 1
Begin Snapshot Id specified: 1
```

```
Enter value for end_snap: 2
```

End Snapshot Id specified: 2

Specify the Report Name

~~~~~

The default report file name is sp\_1\_2. To use this name, press <return> to continue, otherwise enter an alternative. Enter value for report\_name: <press return or enter a new name>

Using the report name sp\_1\_2

The report now scrolls past and is also written to the file specified. For example:

ORACLE\_HOME/bin/sp\_1\_2.lis

To run a report without being prompted, assign values to the SQL\*Plus variables that specify the begin snap ID, the end snap ID, and the report name before running SPREPORT.

The variables are:

- BEGIN\_SNAP: Specifies the begin snapshot ID
- END\_SNAP: Specifies the end snapshot ID
- REPORT\_NAME: Specifies the report output name

### **Example 21-3 Creating a Statspack Report Without Prompts (UNIX)**

```
SQL> connect perfstat/my_perfstat_password
SQL> define begin_snap=1
SQL> define end_snap=2
SQL> define report_name=batch_run
SQL> @?/rdbms/admin/spreport
```

When SPREPORT.SQL is run, it does not prompt for the information provided by the variables.

## **Running the SQL Report**

When you examine the instance report, you often find high-load SQL statements that you want to examine more closely. The SQL report, SPREPSQL.SQL, displays statistics, the complete SQL text, and (if a level six snapshot has been taken), information on any SQL plan(s) associated with that statement.

**See Also:** ["Snapshot Levels"](#) on page 21-17 for information about levels

The SQL statement to be reported on is identified by a hash value, which is a numerical representation of the statement's SQL text. The hash value for each statement is displayed for each statement in the SQL sections of the instance report.

You run the `SPREPSQL.SQL` script while connected as the `PERFSTAT` user. The report is located in the `rdbms/admin` directory of the Oracle home.

---



---

**Note:** In an Oracle Real Application Clusters environment, you must connect to the instance on which you want to report.

---



---

The `SPREPSQL.SQL` report prompts you for the following:

- Beginning snapshot ID
- Ending snapshot ID
- Hash value for the SQL statement
- Name of the report text file to be created

**Example 21–4 Sample Output of SPREPORT.SQL**

```
SQL> connect perfstat/my_perfstat_password
Connected.
SQL> @?/rdbms/admin/sprepsql
```

```

      DB Id      DB Name      Inst Num Instance
-----
2618106428 PRD1                1 prd1

```

Completed Snapshots

| Instance | DB Name | Snap Id | Snap Started      | Snap Level | Comment |
|----------|---------|---------|-------------------|------------|---------|
| prd1     | PRD1    | 37      | 02 Mar 2001 11:01 | 6          |         |
|          |         | 38      | 02 Mar 2001 12:01 | 6          |         |
|          |         | 39      | 08 Mar 2001 09:01 | 5          |         |
|          |         | 40      | 08 Mar 2001 10:02 | 5          |         |

Specify the Begin and End Snapshot Ids

```

~~~~~
Enter value for begin_snap: 39
Begin Snapshot Id specified: 39

```

```
Enter value for end_snap: 40
End Snapshot Id specified: 40
```

```
Specify the Hash Value
```

```
~~~~~
Enter value for hash_value: 1988538571
Hash Value specified is: 1988538571
```

```
Specify the Report Name
```

```
~~~~~
The default report file name is sp_39_40_1988538571. To use this name,
press <return> to continue, otherwise enter an alternative.
Enter value for report_name:
```

```
Using the report name sp_39_40_1988538571
```

The report scrolls past and is also written to the file you specified. For example:

```
sp_39_40_1988538571.lis
```

The `SPREPSQL.SQL` script can run the SQL report in batch mode. To run a report without being prompted, assign values to the SQL\*Plus variables that specify the begin snap ID, the end snap ID, the hash value, and the report name before running the `SPREPSQL.SQL` script. The variables are:

- `BEGIN_SNAP`: specifies the begin snapshot ID
- `END_SNAP`: specifies the end snapshot ID
- `HASH_VALUE`: specifies the hash value
- `REPORT_NAME`: specifies the report output name

**Example 21–5 Running SPREPORT.SQL in Batch Mode**

```
SQL> connect perfstat/my_perfstat_password
SQL> define begin_snap=39
SQL> define end_snap=40
SQL> define hash_value=1988538571
SQL> define report_name=batch_sql_run
SQL> @?/rdbms/admin/sprepsql
```

When `SPREPSQL.SQL` is run, it does not prompt for the information provided by the variables.

## Gathering Optimizer Statistics on the PERFSTAT Schema

For best performance when running the performance reports, collect optimizer statistics for tables and indexes owned by PERFSTAT. You should do this whenever there is significant change in data volumes in PERFSTAT's tables.

To collect optimizer statistics on the PERFSTAT schema, use DBMS\_STATS or DBMS\_UTILITY, and specify the PERFSTAT user. For example:

```
EXECUTE DBMS_STATS.GATHER_SCHEMA_STATS(OWNNAME=>'PERFSTAT', CASCADE=>TRUE);
```

or

```
EXECUTE DBMS_UTILITY.ANALYZE_SCHEMA('PERFSTAT', 'COMPUTE');
```

## Configuring the Amount of Data Captured in Statspack

Both the snapshot level and the thresholds specified affect the amount of data Statspack captures. You can change the amount of information gathered by specifying a different snapshot level. The higher the snapshot level, the more data is gathered. The default level set at installation is level 5.

For typical usage, level 5 snapshot is effective on most sites. There are certain situations when using a level 6 snapshot is beneficial. These include the following:

- When taking your first baseline
- When a new application or an application change is installed to take a new baseline
- After gathering optimizer statistics

**See Also:** ["Snapshot Levels"](#) on page 21-17

### Snapshot SQL Thresholds

There are other parameters you can configure, in addition to the snapshot level. These parameters are used as thresholds when collecting data on SQL statements; data is captured on any SQL statements that breach the specified thresholds. Snapshot level and threshold information used by the package is stored in the STATS\$STATSPACK\_PARAMETER table.

### Changing the Default Values for Snapshot Levels and SQL Thresholds

You can change the default parameters used for taking snapshots so that they are tailored to the instance's workload. Simply use the appropriate parameter and the

new value with the Statspack `MODIFY_STATSPACK_PARAMETER` or `SNAP` procedure. For example:

```
SQL> EXECUTE STATSPACK.SNAP(i_ucomment=>'this is a temporary comment');
SQL> EXECUTE STATSPACK.MODIFY_STATSPACK_PARAMETER(i_ucomment=>'this is a
comment that is saved');
```

The parameters that can be passed into the `MODIFY_STATSPACK_PARAMETER` and `SNAP` procedures are listed in [Table 21-1](#) on page 21-20.

**Temporarily Using New Values** To temporarily use a snapshot level or threshold that is different from the instance's default snapshot values, you specify the required threshold or snapshot level when taking the snapshot. This value is used only for the immediate snapshot taken; the new value is not saved as the default.

For example, to take a single level 6 snapshot:

```
SQL> EXECUTE STATSPACK.SNAP(i_snap_level=>6);
```

**Saving New Defaults** You can save the new value as the instance's default in either of two ways:

- Take a snapshot and specify the new defaults to be saved to the database, using `STATSPACK.SNAP` the `I_MODIFY_PARAMETER` input variable.

```
SQL> EXECUTE STATSPACK.SNAP(i_snap_level=>10, i_modify_parameter=>'true');
```

Setting the `I_MODIFY_PARAMETER` value to `TRUE` saves the new thresholds in the `STATS$STATSPACK_PARAMETER` table. These thresholds are used for all subsequent snapshots.

If the `I_MODIFY_PARAMETER` was set to `FALSE` or omitted, then the new parameter values are not saved. Only the snapshot taken at that point uses the specified values. Any subsequent snapshots use the preexisting values in the `STATS$STATSPACK_PARAMETER` table.

- Change the defaults immediately without taking a snapshot, using the `STATSPACK.MODIFY_STATSPACK_PARAMETER` procedure. For example, the following statement changes the snapshot level to 10 and modifies the SQL thresholds for `BUFFER_GETS` and `DISK_READS`:

```
SQL> EXECUTE STATSPACK.MODIFY_STATSPACK_PARAMETER
 (i_snap_level=>10, i_buffer_gets_th=>10000, i_disk_reads_th=>1000);
```

This procedure changes the values permanently, but does not take a snapshot.

## Snapshot Levels

This section discusses the snapshot levels.

**Levels  $\geq$  0 General Performance Statistics** Any level greater than 0 collects general performance statistics, such as wait statistics, system events, system statistics, rollback segment data, row cache, SGA, background events, session events, lock statistics, buffer pool statistics, and parent latch statistics.

**Levels  $\geq$  5 Additional Data: SQL Statements** This level includes all statistics gathered in the lower levels, as well as performance data on SQL statements with high resource usage. In a level 5 snapshot (or higher), the time required for the snapshot to complete depends on the `SHARED_POOL_SIZE` and on the number of SQL statements in the shared pool at the time of the snapshot. The larger the shared pool, the longer it takes to complete the snapshot.

SQL statements gathered by Statspack are those that exceed one of six predefined threshold parameters:

- Number of executions of the SQL statement. The default 100.
- Number of disk reads performed by the SQL statement. The default 1,000.
- Number of parse calls performed by the SQL statement. The default 1,000.
- Number of buffer gets performed by the SQL statement. The default 10,000.
- Size of sharable memory used by the SQL statement. The default 1 Mb.
- Version count for the SQL statement. The default 20.

The values of each of these threshold parameters are used when deciding which SQL statements to collect. If a SQL statement's resource usage exceeds any one of these threshold values, then it is captured during the snapshot.

The SQL threshold levels used are either those stored in the table `STATS$STATSPACK_PARAMETER` or by the thresholds specified when the snapshot is taken.

**Levels  $\geq$  6 Additional Data: SQL Plans and SQL Plan Usage** This level includes all statistics gathered in the lower levels, as well as SQL plans and plan usage data for each of the high-resource SQL statements captured.

A level 6 snapshot gathers valuable information for determining whether the execution plan used for a SQL statement has changed. Therefore, level 6 snapshots should be used whenever a plan might have changed.

To gather the plan for a SQL statement, the statement must be in the shared pool at the time the snapshot is taken, and it must exceed one of the SQL thresholds. To gather plans for all statements in the shared pool, specify the executions threshold to be zero (0) for those snapshots.

**See Also:** ["Changing the Default Values for Snapshot Levels and SQL Thresholds"](#) on page 21-15 for information on how to do this

**Levels  $\geq$  7 Additional data: Segment Level Statistics** This level includes all statistics gathered in the lower levels, and additionally gathers the performance data on highly used segments. RAC specific segment level statistics are also captured with level 7.

A level 7 snapshot gathers information which determines what segments are more heavily accessed and contended. With this information, you can decide to modify the physical layout of some segments or of the tablespaces they reside in. For example, to better spread the segment IO load, you can add files residing on different disks to a tablespace storing a heavily accessed segment or you can partition a segment. This information can also help decide on changing segment attributes values such as `PCTFREE` or `INITTRANS`. On a RAC environment, this information allows us to easily spot the segments responsible for much of the cross-instance traffic.

Level 7 includes the following segment statistics:

- Logical reads
- Db block changes
- Physical reads
- Physical writes
- Physical reads direct
- Physical writes direct
- Global cache consistent read blocks served (RAC specific)
- Global cache current blocks served (RAC specific)
- Buffer busy waits
- ITL waits
- Row lock waits



Although Statspack captures all segment statistics, it reports only the following statistics that exceed one of the predefined threshold parameters:

- Number of logical reads on the segment. The default is 10,000.
- Number of physical reads on the segment. The default is 1,000.
- Number of buffer busy waits on the segment. The default is 100.
- Number of row lock waits on the segment. The default is 100.
- Number of ITL waits on the segment. The default is 100.
- Number of global cache consistent read blocks served (RAC only). The default is 1,000.
- Number of global cache current blocks served (RAC only). The default is 1,000.

The values of the threshold parameters are used when deciding which segment statistics to collect. If a segment's statistic exceeds a threshold value, all statistics regarding this segment are captured during the snapshot. The threshold levels used are either those stored in the table `stats$statspack_parameter`, or by the thresholds specified when the snapshot is taken.

**Levels >= 10 Additional Statistics: Parent and Child Latches** This level includes all statistics gathered in the lower levels, as well as parent and child latch information. Sometimes data gathered at this level can cause the snapshot to take longer to complete. This level can be resource-intensive, and it should only be used when advised by Oracle personnel.

## Specifying a Session ID

If you want to gather session statistics and wait events for a particular session (in addition to the instance statistics and wait events), specify the session ID in the call to Statspack. The statistics gathered for the session include session statistics, session events, and lock activity. The default behavior is to not gather session level statistics.

For example:

```
SQL> EXECUTE STATSPACK.SNAP(i_session_id=>3);
```

## Parameters for SNAP and MODIFY\_STATSPACK\_PARAMETER Procedures

Parameters that can be passed to the `STATSPACK.SNAP` and `STATSPACK.MODIFY_STATSPACK_PARAMETER` procedures are as follows:

**Table 21–1 Parameters for SNAP and MODIFY\_STATSPACK\_PARAMETER Procedures**

| Parameter Name     | Range of Valid Values     | Default Value  | Meaning                                                                                                     |
|--------------------|---------------------------|----------------|-------------------------------------------------------------------------------------------------------------|
| i_snap_level       | 0, 5, 6, 7, 10            | 5              | Snapshot level                                                                                              |
| i_ucomment         | Text                      | Blank          | Comment to be stored with snapshot                                                                          |
| i_executions_th    | Integer >=0               | 100            | SQL threshold: number of times statement was executed                                                       |
| i_disk_reads_th    | Integer >=0               | 1000           | SQL threshold: number of disk reads the statement made                                                      |
| i_parse_calls_th   | Integer >=0               | 1000           | SQL threshold: number of parse calls the statement made                                                     |
| i_buffer_gets_th   | Integer >=0               | 10000          | SQL threshold: number of buffer gets the statement made                                                     |
| i_sharable_mem_th  | Integer >=0               | 1048576        | SQL threshold: amount of sharable memory                                                                    |
| i_version_count_th | Integer >=0               | 20             | SQL threshold: number of versions of a SQL statement                                                        |
| i_seg_phy_reads_th | Integer >=0               | 1000           | Segment statistic threshold: number of physical reads on a segment                                          |
| i_seg_log_reads_th | Integer >=0               | 10000          | Segment statistic threshold: number of logical reads on a segment                                           |
| i_seg_buff_busy_th | Integer >=0               | 100            | Segment statistic threshold: number of buffer busy waits for a segment                                      |
| i_seg_rowlock_w_th | Integer >=0               | 100            | Segment statistic threshold: number of row lock waits for a segment                                         |
| i_seg_itl_waits_th | Integer >=0               | 100            | Segment statistic threshold: number of ITL waits for a segment                                              |
| i_seg_cr_bks_sd_th | Integer >=0               | 1000           | Segment statistic threshold: number of consistent reads blocks served by the instance for the segment (RAC) |
| i_seg_cu_bks_sd_th | Integer >=0               | 1000           | Segment statistic threshold: number of current blocks served by the instance for the segment (RAC)          |
| i_session_id       | Valid SID from V\$SESSION | 0 (no session) | Session ID of the Oracle session for which to capture session granular statistics                           |
| i_modify_parameter | TRUE, FALSE               | FALSE          | Determines whether the parameters specified are used for future snapshots                                   |

## Time Units Used for Wait Events

Oracle supports capturing certain performance data with microsecond granularity. Views that include microsecond timing include the following:

- V\$SYSTEM\_EVENT, V\$SESSION\_EVENT (TIME\_WAITED\_MICRO column)
- V\$SQL (CPU\_TIME, ELAPSED\_TIME columns)
- V\$LATCH (WAIT\_TIME column)
- V\$SQL\_WORKAREA, V\$SQL\_WORKAREA\_ACTIVE (ACTIVE\_TIME column)

---



---

**Note:** Existing columns in other views continue to use centisecond times.

---



---

Because microsecond timing might not be appropriate for rolled-up data, Statspack displays most times in seconds. It displays average times in milliseconds for easier comparison with operating system monitoring utilities, which often report timing in milliseconds.

For clarity, the time units used are specified in the column headings of each timed column in the Statspack report. The following convention are used:

- second (s)
- centisecond (cs) - 100th of a second
- millisecond (ms) - 1,000th of a second
- microsecond (us) - 1,000,000th of a second

## Event Timings

If timings are available, the Statspack report orders wait events by time. These are listed in the Top-5 and background and foreground wait events sections.

If `TIMED_STATISTICS` is `FALSE` for the instance, but a subset of users or programs set `TIMED_STATISTICS` to `TRUE` dynamically, then the Statspack report output can look inconsistent, where some events have timings (those which the individual programs/users waited for) and some do not. The Top-5 section also looks unusual in this situation.

Optimally, `TIMED_STATISTICS` should be set to `true` at the instance level, for ease of diagnosing performance problems.

---

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

---

**See Also:** ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

## Managing and Sharing Statspack Performance Data

This section discusses the following topics:

- [Sharing Data Through Export](#)
- [Removing Unnecessary Data](#)
- [Truncating All Statspack Data](#)

### Sharing Data Through Export

If you want to share data with other sites (for example, if Oracle Support requires the raw statistics), then you can export the `PERFSTAT` user. An export parameter file (`SPUEXP.PAR`) is supplied for this purpose. To use this file, supply the export command with the `userid` parameter, along with the export parameter file name. For example:

```
exp userid=perfstat/my_perfstat_password parfile=spuexp.par
```

This creates a file called `SPUEXP.DMP` and the log file `SPUEXP.LOG`. If you want to load the data into another database, use the import command.

**See Also:** *Oracle9i Database Utilities* for more information on using export and import

## Removing Unnecessary Data

Purge unnecessary data from the `PERFSTAT` schema using the `SPPURGE.SQL` script. This deletes snapshots that fall between the begin and end snapshot IDs you specify.

---



---

**Note:** You should export the schema as a backup before running this script, either using your own export parameters or those provided in `SPUEXP.PAR`.

---



---

Purging can require the use of a large rollback segment, because all data relating to each snapshot ID to be purged is deleted. You can avoid rollback segment extension errors in one of two ways:

- Specify a smaller range of snapshot IDs to purge.
- Explicitly use a large rollback segment, by executing the `SET TRANSACTION USE ROLLBACK SEGMENT` statement before running the `SPPURGE.SQL` script.

**See Also:** *Oracle9i SQL Reference*

When you run `SPPURGE.SQL`, it displays the instance to which you are connected and the available snapshots. It then prompts you for the low snap ID and high snap ID. All snapshots that fall within this range are purged.

### Example 21-6 Sample Run of `SPPURGE.SQL`

```
SQL> CONNECT perfstat/my_perfstat_password
SQL> SET TRANSACTION USE ROLLBACK SEGMENT rbig;
SQL> @?/rdbms/admin/sppurge
```

```
Database Instance currently connected to
```

```
=====
 Instance
 DB Id DB Name Inst Num Name

 720559826 PERF 1 perf
```

```
Snapshots for this database instance
```

```
=====
 Snap
 Snap Id Level Snapshot Started Host Comment

```

```
1 5 30 Feb 2000 10:00:01 perfhost
2 5 30 Feb 2000 12:00:06 perfhost
3 5 01 Mar 2000 02:00:01 perfhost
4 5 01 Mar 2000 06:00:01 perfhost
```

---

---

**Caution:** SPPURGE.SQL deletes all snapshots ranging between the lower and upper bound snapshot IDs specified for the database instance connected to. You might want to export this data before continuing.

---

---

Specify the Low Snap ID and High Snap ID range to purge

```
~~~~~
Enter value for losnapid: 1
Using 1 for lower bound.
```

```
Enter value for hisnapid: 2
Using 2 for upper bound.
```

```
Deleting snapshots 1 - 2
```

Purge of specified snapshot range complete. If you want to rollback the purge, it is still possible to do so. Exiting from SQL\*Plus automatically commits the purge.

```
SQL> -- end of example output
```

To purge in batch mode, you must assign values to the SQL\*Plus variables that specify the low and high snapshot IDs to purge. The variables are:

- LOSNAPID: Begin snapshot ID
- HISNAPID: End snapshot ID

**Example 21-7 Running SPPURGE.SQL in Batch Mode**

```
SQL> CONNECT perfstat/my_perfstat_password
SQL> DEFINE losnapid=1
SQL> DEFINE hisnapid=2
SQL> @?/rdbms/admin/sppurge
```

When SPPURGE . SQL is run, it does not prompt for the information provided by the variables.

## Truncating All Statspack Data

To truncate all performance data indiscriminately, use `SPTRUNC.SQL`. This script truncates all statistics data gathered.

---



---

**Note:** Oracle Corporation recommends that you export the schema as a backup before running this script, either using your own export parameters or those provided in `SPUEXP.PAR`.

---



---

### *Example 21–8 Sample Run of SPTRUNC.SQL*

```
SQL> CONNECT perfstat/my_perfstat_password
SQL> @?/rdbs/admin/sptrunc
~~~~~
```

---



---

**Note:** Running `SPTRUNC.SQL` removes all data from Statspack tables. You might want to export the data before continuing.

---



---

```
If you would like to continue, enter any string, followed by <return>.
Enter value for anystring:
entered - starting truncate operation
Table truncated.
<etc>
Truncate operation complete.
```

## Oracle Real Application Clusters Considerations with Statspack

The unique identifiers for a database instance used by Statspack are the `DBID` and the `INSTANCE_NUMBER`. When you use Oracle Real Application Clusters, the `INSTANCE_NUMBER` could change between startups (either because the `INSTANCE_NUMBER` parameter is set in the initialization file or because the instances are started in a different order).

Statspack uses the `INSTANCE_NUMBER` and the `DBID` to identify the instance's snapshot preferences, so you could have a different set of levels or thresholds being used when taking snapshots of an instance. This happens only under the following set of conditions:

- The instance numbers have switched between startups.
- The DBA has modified the default Statspack parameters used for at least one of the instances.

- The parameters used (for example, thresholds and snapshot level) are not the same on all instances.

The parameters differ only if the DBA explicitly modified them after installation, either by saving the specified values or by using the `MODIFY_STATSPACK_PARAMETER` procedure. To check whether any of the Statspack snapshot parameters are different for the instances, query the `STATS$STATSPACK_PARAMETER` table.

---

---

**Note:** If you have changed the default Statspack parameters, you can avoid encountering this problem by hard-coding the `INSTANCE_NUMBER` in the initialization parameter file for each of the instances in the Oracle Real Application Clusters database.

---

---

**See Also:** *Oracle9i Real Application Clusters Administration* for recommendations and issues with setting the `INSTANCE_NUMBER` initialization parameter

## Removing Statspack

To deinstall Statspack, connect as a user with `SYSDBA` privilege and run the following `SPDROP` script from SQL\*Plus. For example:

```
SQL> CONNECT / AS SYSDBA
SQL> @?/rdbms/admin/spdrop
```

The `SPDROP.SQL` script calls the following scripts:

- `SPDTAB.SQL` - drops tables and public synonyms
- `SPDUSR.SQL` - drops the user

Check each of two output files produced (`SPDTAB.LIS`, `SPDUSR.LIS`) to ensure that the package was completely deinstalled.

## Statspack Supplied Scripts and Documentation

This section discusses the following:

- [Scripts for Statspack Installation and Removal](#)
- [Scripts for Statspack Reporting and Automation](#)
- [Scripts for Upgrading Statspack](#)



- [Scripts for Statspack Performance Data Maintenance](#)
- [Statspack Documentation](#)

## Scripts for Statspack Installation and Removal

The Statspack installation and removal scripts must be run as a user with the SYSDBA privilege.

- `SPCREATE.SQL`: Creates entire Statspack environment, calling the following scripts:
  - `SPCURS.SQL`: Creates the Statspack user (`PERFSTAT`)
  - `SPCTAB.SQL`: Creates Statspack tables
  - `SPCPKG.SQL`: Creates the Statspack package
- `SPDROP.SQL`: Drops entire Statspack environment, calling the following scripts:
  - `SPDTAB.SQL`: Drops Statspack tables
  - `SPDUSR.SQL`: Drops the Statspack user (`PERFSTAT`)

## Scripts for Statspack Reporting and Automation

The Statspack reporting and automation scripts must be run as the `PERFSTAT` user.

- `SPREPORT.SQL`: Generates a Statspack report
- `SPREPSQL.SQL`: Generates a Statspack SQL report for the specific SQL hash value specified
- `SPREPINS.SQL`: Generates a Statspack report for the database and instance specified
- `SPAUTO.SQL`: Automates Statspack statistics collection (using `DBMS_JOB`)

## Scripts for Upgrading Statspack

The Statspack upgrade scripts must be run as a user with the SYSDBA privilege.

- `SPUP90.SQL`: Converts data from the 9.0 schema to the 9.2 schema. Back up the existing schema before running the upgrade.
- `SPUP817.SQL`: If upgrading from Statspack 8.1.7, this script must be run first.
- `SPUP816.SQL`: If upgrading from Statspack 8.1.6, this script must be run first, followed by `SPUP817.SQL`.

## Scripts for Statspack Performance Data Maintenance

The Statspack data maintenance scripts must be run as the `PERFSTAT` user.

- `SPPURGE.SQL`: Purges a limited range of Snapshot IDs for a given database instance.
- `SPTRUNC.SQL`: Truncates all performance data in Statspack tables

---

---

**Caution:** Do not use this script unless you want to remove all data in the schema you are using. You can choose to export the data as a backup before using this script.

---

---

- `SPEXP.PAR`: An export parameter file supplied for exporting the whole `PERFSTAT` user.

## Statspack Documentation

The `SPDOC.TXT` file in the `ORACLE_HOME/rdbms/admin` directory contains instructions and documentation on the Statspack package.

# Part V

---

## Optimizing Instance Performance

Part V describes how to tune various elements of your database system to optimize performance of an Oracle instance.

The chapters in this part are:

- [Chapter 22, "Instance Tuning"](#)
- [Chapter 23, "Tuning Networks"](#)



---

## Instance Tuning

After the initial configuration of a database, tuning an instance is important to eliminate any performance bottlenecks.

This chapter contains the following sections:

- [Performance Tuning Principles](#)
- [Performance Tuning Steps](#)
- [Interpreting Oracle Statistics](#)
- [Wait Events](#)
- [Idle Wait Events](#)

## Performance Tuning Principles

Performance tuning requires a different, although related, method to the initial configuration of a system. Configuring a system involves allocating resources in an ordered manner so that the initial system configuration is functional.

Tuning is driven by identifying the most significant bottleneck and making the appropriate changes to reduce or eliminate the effect of that bottleneck. Usually, tuning is performed reactively, either while the system is preproduction or after it is live.

---

---

**Note:** Before using this performance tuning reference, make sure you have read *Oracle9i Database Performance Planning*. Oracle Corporation has designed a new performance methodology, based on years of Oracle designing and performance experience. This brief book explains clear and simple activities that can dramatically improve system performance. It discusses the following topics:

- Investment Options
  - Scalability
  - System Architecture
  - Application Design Principles
  - Workload Testing, Modeling, and Implementation
  - Deploying New Applications
- 
- 

## Baselines

The most effective way to tune is to have an established performance baseline that can be used for comparison if a performance issue arises. Most DBAs know their system well and can easily identify peak usage periods. For example, the peak periods could be between 10.00am and 12.00pm and also between 1.30pm and 3.00pm. This could include a batch window of 12.00am midnight to 6am.

It is important to identify these high-load times at the site and install a monitoring tool that gathers performance data for those times. Optimally, data gathering should be configured from when the application is in its initial trial phase during the QA cycle. Otherwise, this should be configured when the system is first in production.

---

---

**Note:** Oracle recommends using the Enterprise Manager (EM) Diagnostics Pack for systems monitoring and tuning due to its extended feature list. However, if your site does not have EM, then Statspack can be used to gather Oracle instance statistics.

For illustration purposes, a combination of Statspack report output and direct queries from the v\$ views are used in examples, because they are available on all installations.

---

---

**See Also:** [Chapter 20, "Oracle Tools to Gather Database Statistics"](#) for detailed information on Oracle instance performance tuning tools

Ideally, baseline data gathered should include the following:

- Application statistics (transaction volumes, response time)
- Database statistics
- operating system statistics
- Disk I/O statistics
- Network statistics

## The Symptoms and the Problems

A common pitfall in performance tuning is to mistake the symptoms of a problem for the actual problem itself. It is important to recognize that many performance statistics indicate the symptoms, and that identifying the symptom is not sufficient data to implement a remedy. For example:

- Slow physical I/O  
Generally, this is caused by poorly-configured disks. However, it could also be caused by a significant amount of unnecessary physical I/O on those disks issued by poorly-tuned SQL.
- Latch contention  
Rarely is latch contention tunable by reconfiguring the instance. Rather, latch contention usually is resolved through application changes.
- Excessive CPU usage

Excessive CPU usage usually means that there is little idle CPU on the system. This could be caused by an inadequately-sized system, by untuned SQL statements, or by inefficient application programs.

**See Also:** [Table 22-1, " Wait Events and Potential Causes"](#) on page 22-19

## When to Tune

There are two distinct types of tuning: proactive monitoring and bottleneck elimination.

### Proactive Monitoring

Proactive monitoring usually occurs on a regularly scheduled interval, where a number of performance statistics are examined to identify whether the system behavior and resource usage has changed. Proactive monitoring also can be called proactive tuning.

Usually, monitoring does not result in configuration changes to the system, unless the monitoring exposes a serious problem that is developing. In some situations, experienced performance engineers can identify potential problems through statistics alone, although accompanying performance degradation is usual.

'Tweaking' a system when there is no apparent performance degradation as a proactive action can be a dangerous activity, resulting in unnecessary performance drops. Tweaking a system should be considered reactive tuning, and the steps for reactive tuning should be followed.

Monitoring is usually part of a larger capacity planning exercise, where resource consumption is examined to see the changes in the way the application is being used and the way the application is using the database and host resources.

### Bottleneck Elimination: Tuning

Tuning usually implies fixing a performance problem. However, tuning should be part of the lifecycle of an application, through the analysis, design, coding, production, and maintenance stages. Many times, the tuning phase is left until the system is in production. At this time, tuning becomes a reactive fire-fighting exercise, where the most important bottleneck is identified and fixed.

Usually, the purpose for tuning is to reduce resource consumption or to reduce the elapsed time for an operation to complete. Either way, the goal is to improve the effective use of a particular resource. In general, performance problems are caused



by the over-use of a particular resource. That resource is the *bottleneck* in the system. There are a number of distinct phases in identifying the bottleneck and the potential fixes. These are discussed in the sections that follow.

Remember that the different forms of contention are symptoms that can be fixed by making changes in the following places:

- Changes in the application, or the way the application is used
- Changes in Oracle
- Changes in the host hardware configuration

Often, the most effective way of resolving a bottleneck is to change the application.

## Performance Tuning Steps

These are the main steps in the Oracle Performance Method:

1. Get candid feedback from users about the scope of the performance problem. This step is to [Define the Problem](#).
2. Obtain a full set of operating system, database, and application statistics. Then [Examine the Host System](#) and [Examine the Oracle Statistics](#) for any evidence.
3. Consider the list of common performance errors to see whether the data gathered suggests that they are contributing to the problem.
4. Build a conceptual model of what is happening on the system using the performance data gathered.
5. Propose changes to be made and the expected result of implementing the changes. Then, [Implement and Measure Change](#) in application performance.
6. Determine whether the performance objective defined in step 1 has been met. If not, then repeat steps 5 and 6 until the performance goals are met.

**See Also:** *Oracle9i Database Performance Planning* for a list of common errors and for a theoretical description of this performance method

The remainder of this chapter covers the steps of the Oracle performance method in detail.

## Define the Problem

It is vital to develop a good understanding of the purpose of the tuning exercise and the nature of the problem before attempting to implement a solution. Without this understanding, it is virtually impossible to implement effective changes. The data gathered during this stage helps determine the next step to take and what evidence to examine.

Gather the following data:

1. Identify the performance objective.

What is the measure of acceptable performance? How many transactions an hour, or seconds, response time will meet the required performance level?

2. Identify the scope of the problem.

What is affected by the slowdown? For example, is the whole instance slow? Is it a particular application, program, specific operation, or a single user?

3. Identify the time frame when the problem occurs.

Is the problem only evident during peak hours? Does performance deteriorate over the course of the day? Was the slowdown gradual (over the space of months or weeks) or sudden?

4. Quantify the slowdown.

This helps identify the extent of the problem and also acts as a measure for comparison when deciding whether changes implemented to fix the problem have actually made an improvement. Find a consistently reproducible measure of the response time or job run time. How much worse are the timings than when the program was running well?

5. Identify any changes.

Identify what has changed since performance was acceptable. This may narrow the potential cause quickly. For example, has the operating system software, hardware, application software, or Oracle release been upgraded? Has more data been loaded into the system, or has the data volume or user population grown?

At the end of this phase, you should have a good understanding of the symptoms. If the symptoms can be identified as local to a program or set of programs, then the problem is handled in a different manner than instance-wide performance issues.

**See Also:** [Chapter 6, "Optimizing SQL Statements"](#) for information on solving performance problems specific to an application or user

## Examine the Host System

Look at the load on the database server, as well as the database instance. Consider the operating system, the I/O subsystem, and network statistics, because examining these areas helps determine what might be worth further investigation. In multitier systems, also examine the application server middle-tier hosts.

Examining the host hardware often gives a strong indication of the bottleneck in the system. This determines which Oracle performance data could be useful for cross-reference and further diagnosis.

Data to examine includes the following:

### CPU Usage

If there is a significant amount of idle CPU, then there could be an I/O, application, or database bottleneck. Note that wait I/O should be considered as idle CPU.

If there is high CPU usage, then determine whether the CPU is being used effectively. Is the majority of CPU usage attributable to a small number of high-CPU using programs, or is the CPU consumed by an evenly distributed workload?

If the CPU is used by a small number of high-usage programs, then look at the programs to determine the cause.

**Non-Oracle Processes** If the programs are not Oracle programs, then identify whether they are legitimately requiring that amount of CPU. If so, then can their execution can be delayed to off-peak hours?

**Oracle Processes** If a small number of Oracle processes consumes most of the CPU resources, then use `SQL_TRACE` and `TKPROF` to identify the SQL or PL/SQL statements to see if a particular query or PL/SQL program unit can be tuned. For example, a `SELECT` statement could be CPU-intensive if its execution involves many reads of data in cache (logical reads) that could be avoided with better SQL optimization.

**Oracle CPU Statistics** Oracle CPU statistics are available in three `V$` views:

- `V$SYSSTAT` shows Oracle CPU usage for all sessions. The statistic "CPU used by this session" shows the aggregate CPU used by all sessions.

- `V$SESSTAT` shows Oracle CPU usage for each session. Use this view to determine which particular session is using the most CPU.
- `V$RSRC_CONSUMER_GROUP` shows CPU utilization statistics for each consumer group when the Oracle Database Resource Manager is running.

**Interpreting CPU Statistics** It is important to recognize that CPU time and real time are distinct. With eight CPUs, for any given minute in real time, there are eight minutes of CPU time available. On NT and UNIX, this can be either user time or system time (*privileged* mode on NT). Thus, average CPU time utilized by all processes (threads) on the system could be greater than one minute for every one minute real time interval.

At any given moment, you know how much time Oracle has used on the system. So, if eight minutes are available and Oracle uses four minutes of that time, then you know that 50% of all CPU time is used by Oracle. If your process is not consuming that time, then some other process is. Identify the processes that are using CPU time, figure out why, and then attempt to tune them.

**See Also:** [Chapter 10, "Using SQL Trace and TKPROF"](#)

If the CPU usage is evenly distributed over many Oracle server processes, then examine the Statspack report for other evidence.

### Detecting I/O Problems

An overly active I/O system can be evidenced by disk queue lengths greater than two, or disk service times that are over 20-30ms. If the I/O system is overly active, then check for potential hot spots that could benefit from distributing the I/O across more disks. Also identify whether the load can be reduced by lowering the resource requirements of the programs using those resources.

Use operating system monitoring tools to determine what processes are running on the system as a whole and to monitor disk access to all files. Remember that disks holding datafiles and redo log files can also hold files that are not related to Oracle. Reduce any heavy access to disks that contain database files. Access to non-Oracle files can be monitored only through operating system facilities, rather than through the `V$` views.

Tools, such as `sar -d` (or `iostat`) on many UNIX systems and Performance Monitor on Windows 2000 systems, examine I/O statistics for the entire system.

**See Also:** Your operating system documentation for the tools available on your platform

Check the Oracle wait event data in `V$SYSTEM_EVENT` to see whether the top wait events are I/O related. I/O related events include `db file sequential read`, `db file scattered read`, `db file single write`, and `db file parallel write`. These are all events corresponding to I/Os performed against the data file headers, control files, or data files. If any of these wait events correspond to high average time, then investigate the I/O contention.

Cross reference the host I/O system data with the I/O sections in the Statspack report to identify hot datafiles and tablespaces. Also compare the I/O times reported by the operating system with the times reported by Oracle to see if they are consistent.

Before investigating whether the I/O system should be reconfigured, determine if the load on the I/O system can be reduced. To reduce Oracle I/O load, look at SQL statements that perform many physical reads by querying the `V$SQLAREA` view or by reviewing the 'SQL ordered by physical reads' section of the Statspack report. Examine these statements to see how they can be tuned to reduce the number of I/Os.

If there are Oracle-related I/O problems caused by SQL statements, then tune them. If the Oracle server is not consuming the available I/O resources, then identify the process that is using up the I/O. Determine why the process is using up the I/O, and then tune this process.

**See Also:**

- [Chapter 6, "Optimizing SQL Statements"](#)
- ["V\\$SQLAREA"](#) on page 24-57
- [Chapter 15, "I/O Configuration and Design"](#)
- ["db file scattered read"](#) on page 22-29 and ["db file sequential read"](#) on page 22-31 for the difference between a scattered read and a sequential read, and how this affects I/O

## Network

Using operating system utilities, look at the network round-trip ping time and the number of collisions. If the network is causing large delays in response time, then investigate possible causes.

Network load can be reduced by scheduling large data transfers to off-peak times, or by coding applications to batch requests to remote hosts, rather than accessing remote hosts once (or more) for one request.

**See Also:** *Oracle9i Database Performance Planning* for a description of important operating system statistics

## Examine the Oracle Statistics

Oracle statistics should be examined and cross-referenced with operating system statistics to ensure a consistent diagnosis of the problem. Operating system statistics can indicate a good place to begin tuning. However, if the goal is to tune the Oracle instance, then look at the Oracle statistics to identify the resource bottleneck from Oracle's perspective before implementing corrective action.

**See Also:** ["Interpreting Oracle Statistics"](#) on page 22-15

### Setting the Level of Statistics Collection

Oracle9i Release 2 (9.2) provides the initialization parameter `STATISTICS_LEVEL`, which controls all major statistics collections or advisories in the database. This parameter sets the statistics collection level for the database.

Depending on the setting of `STATISTICS_LEVEL`, certain advisories and statistics are collected, as follows:

**BASIC:** No advisories or statistics are collected.

**TYPICAL:** The following advisories or statistics are collected:

- Buffer cache advisory
- MTRR advisory
- Shared Pool sizing advisory
- Segment level statistics
- PGA target advisory
- Timed statistics

**ALL:** All of the preceding advisories or statistics are collected, plus the following:

- Timed operating system statistics
- Row source execution statistics

The default level is `TYPICAL`. `STATISTICS_LEVEL` is a dynamic parameter and can be altered at the system or the session level.

When modified by `ALTER SYSTEM`, all advisories or statistics in the preceding list are dynamically turned on or off, depending on the new value of `STATISTICS_LEVEL`.

When modified by `ALTER SESSION`, only the following advisories or statistics are turned on or off in the local session only. Their system-wide state is not changed.

- Timed statistics
- Timed operating system statistics
- Row source execution statistics

**V\$STATISTICS\_LEVEL** This view lists the status of the statistics or advisories controlled by `STATISTICS_LEVEL`.

---

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

---

**See Also:** "[V\\$STATISTICS\\_LEVEL](#)" on page 24-61 for column details of this view

The following sections discuss the common Oracle data sources used while tuning. The sources can be divided into two types of statistics: wait events and system statistics.

## Wait Events

Wait events are statistics that are incremented by a server process/thread to indicate that it had to wait for an event to complete before being able to continue processing. Wait event data reveals various symptoms of problems that might be impacting performance, such as latch contention, buffer contention, and I/O contention. Remember that these are only *symptoms* of problems—not the actual causes.

A server process can wait for the following:

- A resource to become available, such as a buffer or a latch
- An action to complete, such as an I/O
- More work to do, such as waiting for the client to provide the next SQL statement to execute. Events that identify that a server process is waiting for more work are known as *idle events*.

Wait event statistics include the number of times an event was waited for and the time waited for the event to complete. The views `V$SESSION_WAIT`, `V$SESSION_EVENT`, and `V$SYSTEM_EVENT` can be queried for wait event statistics. If the configuration parameter `TIMED_STATISTICS` is set to `true`, then you can also see how long each resource was waited for. To minimize user response time, reduce the time spent by server processes waiting for event completion. Not all wait events have the same wait time. Therefore, it is more important to examine events with the most total time waited rather than wait events with a high number of occurrences. Usually, it is best to set the dynamic parameter `TIMED_STATISTICS` to `true` at least while monitoring performance.

---

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER_SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

---

**See Also:** ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Investigate wait events and related timing data when performing reactive performance tuning. The events with the most time listed against them are often strong indications of the performance bottleneck. For example, by looking at `V$SYSTEM_EVENT`, you might notice lots of `buffer busy waits`. It might be that many processes are inserting into the same block and must wait for each other before they can insert. The solution could be to introduce freelists for the object in question.



**See Also:** See [Wait Events](#) on page 22-24 for a description of the differences between the views `V$SESSION_WAIT`, `V$SESSION_EVENT`, and `V$SYSTEM_EVENT`

## System Statistics

System statistics are typically used in conjunction with wait event data to find further evidence of the cause of a performance problem.

For example, if `V$SYSTEM_EVENT` indicates that the largest wait event (in terms of wait time) is the event `buffer busy waits`, then look at the specific buffer wait statistics available in the view `V$WAITSTAT` to see which block type has the highest wait count and the highest wait time. After the block type has been identified, also look at `V$SESSION_WAIT` real-time while the problem is occurring to identify the contended-for object(s) using the file number and block number indicated. The combination of this data indicates the appropriate corrective action.

Statistics are available in many `V$` views. Some common views include the following:

**V\$SYSSTAT** This contains overall statistics for many different parts of Oracle, including rollback, logical and physical I/O, and parse data. Data from `V$SYSSTAT` is used to compute ratios, such as the buffer cache hit ratio.

**V\$FILESTAT** This contains detailed file I/O statistics for each file, including the number of I/Os for each file and the average read time.

**V\$ROLLSTAT** This contains detailed rollback and undo segment statistics for each segment.

**V\$ENQUEUE\_STAT** This contains detailed enqueue statistics for each enqueue, including the number of times an enqueue was requested and the number of times an enqueue was waited for, and the wait time.

**V\$LATCH** This contains detailed latch usage statistics for each latch, including the number of times each latch was requested and the number of times the latch was waited for.

**See Also:** [Chapter 24, "Dynamic Performance Views for Tuning"](#) for detailed descriptions of the `V$` views used in tuning

## Segment-Level Statistics

With Oracle9i Release 2 (9.2) and higher, you can gather segment-level statistics to help you spot performance problems associated with individual segments. Collecting and viewing segment-level statistics is a good way to effectively identify the hot table or index in an instance.

After viewing wait events or system statistics to identify the performance problem, you can use segment-level statistics to find specific tables or indexes that are causing the problem. Consider, for example, that `V$SYSTEM_EVENT` indicates that buffer busy waits cause a fair amount of wait time. You can select from `V$SEGMENT_STATISTICS` the top segments that cause the buffer busy waits. Then you can focus your effort on eliminating the problem in those segments.

You can query segment-level statistics through the following dynamic performance views:

- `V$SEGSTAT_NAME` This view lists the segment statistics being collected, as well as the properties of each statistic (for instance, if it is a sampled statistic).
- `V$SEGSTAT` This is a highly efficient, real-time monitoring view that shows the statistic value, statistic name, and other basic information.
- `V$SEGMENT_STATISTICS` This is a user-friendly view of statistic values. In addition to all the columns of `V$SEGSTAT`, it has information about such things as the segment owner and table space name. It makes the statistics easy to understand, but it is more costly.

**See Also:** [Chapter 24, "Dynamic Performance Views for Tuning"](#) for detailed column information, starting on page 24-31

## Implement and Measure Change

Often at the end of a tuning exercise, it is possible to identify two or three changes that could potentially alleviate the problem. To identify which change provides the most benefit, it is recommended that only one change be implemented at a time. The effect of the change should be measured against the baseline data measurements found in the problem definition phase.

Typically, most sites with dire performance problems implement a number of overlapping changes at once, and thus cannot identify which changes provided any benefit. Although this is not immediately an issue, this becomes a significant hindrance if similar problems subsequently appear, because it is not possible to know which of the changes provided the most benefit and which efforts to prioritize.

If it is not possible to implement changes separately, then try to measure the effects of dissimilar changes. For example, measure the effect of making an initialization change to optimize redo generation separately from the effect of creating a new index to improve the performance of a modified query. It is *impossible* to measure the benefit of performing an operating system upgrade if SQL is tuned, the operating system disk layout is changed, and the initialization parameters are also changed at the same time.

Performance tuning is an iterative process. It is unlikely to find a 'silver bullet' that solves an instance-wide performance problem. In most cases, excellent performance requires iteration through the performance tuning phases, because solving one bottleneck often uncovers another (sometimes worse) problem.

Knowing when to stop tuning is also important. The best measure of performance is user perception, rather than how close the statistic is to an ideal value.

## Interpreting Oracle Statistics

Gather statistics that cover the time when the instance had the performance problem. If you previously captured baseline data for comparison, then you can compare the current data to the data from the baseline that most represents the problem workload.

When comparing two reports, ensure that the two reports are from times where the system was running comparable workloads.

**See Also:** ["Principles of Data Gathering"](#) on page 20-2

## Examine Load

Usually, wait events are the first data examined. However, if you have a baseline report, then check to see if the load has changed. Regardless of whether you have a baseline, it is useful to see whether the resource usage rates are high.

Load-related statistics to examine include redo size, session logical reads, db block changes, physical reads, physical writes, parse count (total), parse count (hard), and user calls. This data is queried from V\$SYSSTAT. It is best to normalize this data over seconds and over transactions.

In the Statspack report, look at the Load Profile section. The data has been normalized over transactions and over seconds.

### Changing Load

The load profile statistics over seconds show the changes in throughput (that is, whether the instance is performing more work each second). The statistics over transactions identify changes in the application characteristics by comparing these to the corresponding statistics from the baseline report.

### High Rates of Activity

Examine the statistics normalized over seconds to identify whether the rates of activity are very high. It is difficult to make blanket recommendations on high values, because the thresholds are different on each site and are contingent on the application characteristics, the number and speed of CPUs, the operating system, the I/O system, and the Oracle release.

The following are some generalized examples (acceptable values vary at each site):

- A hard parse rate of more than 100 a second indicates that there is a very high amount of hard parsing on the system. High hard parse rates cause serious performance issues and must be investigated. Usually, a high hard parse rate is accompanied by latch contention on the shared pool and library cache latches. Check whether waits for 'latch free' appear in the top-5 wait events, and if so, examine the latching sections of the Statspack report.
- A high soft parse rate could be in the rate of 300a second or more. Unnecessary soft parses also limit application scalability. Optimally, a SQL statement should be soft parsed once in each session and executed many times.

## Using Wait Event Statistics to Drill Down to Bottlenecks

Whenever an Oracle process waits for something, it records the wait using one of a set of predefined wait events. (See `V$EVENT_NAME` for a list of all wait events.) Some of these events are termed *idle* events, because the process is idle, waiting for work to perform. Non-idle events indicate nonproductive time spent waiting for a resource or action to complete.

---

---

**Note:** Not all symptoms can be evidenced by wait events. See "[Additional Statistics](#)" on page 22-20 for the statistics that can be checked.

---

---

The most effective way to use wait event data is to order the events by the wait time. This is only possible if `TIMED_STATISTICS` is set to `true`. Otherwise, the

wait events can only be ranked by the number of times waited, which is often not the ordering that best represents the problem.

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

**See Also:** ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

To get an indication of where time is spent, follow these steps:

1. Examine the data collection for `V$SYSTEM_EVENT`. The events of interest should be ranked by wait time.

Identify the wait events that have the most significant percentage of wait time. To determine the percentage of wait time, add the total wait time for all wait events, excluding idle events (such as `Null event`, `SQL*Net message from client`, `SQL*Net message to client`, `SQL*Net more data`). Calculate the relative percentage of the five most prominent events by dividing each event's wait time by the total time waited for all events.

**See Also:** ["Idle Wait Events"](#) for the complete list of idle events on page 22-49

Alternatively, look at the Top 5 Wait Events section on the front page of the Statspack report; this section automatically orders the wait events (omitting idle events), and calculates the relative percentage:

```
Top 5 Wait Events
~~~~~
Event                               Waits      Wait      % Total
                                     Time (cs)  Wt Time
-----
latch free                           217,224      65,056    63.55
```

|                             |        |        |       |
|-----------------------------|--------|--------|-------|
| db file sequential read     | 39,836 | 31,844 | 31.11 |
| db file scattered read      | 3,679  | 2,846  | 2.78  |
| SQL*Net message from dblink | 1,186  | 870    | .85   |
| log file sync               | 830    | 775    | .76   |

In the previous example, the highest ranking wait event is the `latch free` event. In some situations, there might be a few events with similar percentages. This can provide extra evidence if all the events are related to the same type of resource request (for example, all I/O related events).

2. Look at the number of waits for these events, and the average wait time. For example, for I/O related events, the average time might help identify whether the I/O system is slow. The following example of this data is taken from the Wait Event section of the Statspack report:

| Event                      | Waits     | Timeouts  | Total Wait Time (s) | Avg wait (ms) | Waits /txn |
|----------------------------|-----------|-----------|---------------------|---------------|------------|
| latch free                 | 5,560,989 | 2,705,969 | 26,117              | 5             | 827.7      |
| db file sequential read    | 137,027   | 0         | 2,129               | 16            | 20.4       |
| SQL*Net break/reset to cli | 1,566     | 0         | 1,707               | 1091          | 0.2        |

3. The top wait events identify the next places to investigate. A table of common wait events is listed in [Table 22-1](#). For the previous example, the appropriate data to check would be latch-related. (It is usually a good idea to also have quick look at high-load SQL).
4. Examine the related data indicated by the wait events to see what other information this data provides. Determine whether this information is consistent with the wait event data. In most situations, there is enough data to begin developing a theory about the potential causes of the performance bottleneck.
5. To determine whether this theory is valid, cross-check data you have already examined with other statistics available for consistency. (The appropriate statistics vary depending on the problem, but usually include load profile-related data in `V$SYSSTAT`, operating system statistics, and so on). Perform cross-checks with other data to confirm or refute the developing theory.

## Table of Wait Events and Potential Causes

[Table 22-1](#) links wait events to possible causes and gives an overview of the Oracle data that could be most useful to review next.

**Table 22-1** *Wait Events and Potential Causes*

| Wait Event              | General Area              | Possible Causes                                                                                                                      | Look for / Examine                                                                                                                                                                                                           |
|-------------------------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffer busy waits       | Buffer cache, DBWR        | Depends on buffer type. For example, waits for an index block may be caused by a primary key that is based on an ascending sequence. | Examine <code>V\$SESSION_WAIT</code> while the problem is occurring to determine the type of block in contention.                                                                                                            |
| free buffer waits       | Buffer cache, DBWR, I/O   | Slow DBWR (possibly due to I/O?)<br>Cache too small                                                                                  | Examine write time using operating system statistics. Check buffer cache statistics for evidence of too small cache.                                                                                                         |
| db file scattered read  | I/O, SQL statement tuning | Poorly tuned SQL<br>Slow I/O system                                                                                                  | Investigate <code>V\$SQLAREA</code> to see whether there are SQL statements performing many disk reads. Cross-check I/O system and <code>V\$FILESTAT</code> for poor read time.                                              |
| db file sequential read | I/O, SQL statement tuning | Poorly tuned SQL<br>Slow I/O system                                                                                                  | Investigate <code>V\$SQLAREA</code> to see whether there are SQL statements performing many disk reads. Cross-check I/O system and <code>V\$FILESTAT</code> for poor read time.                                              |
| enqueue                 | Locks                     | Depends on type of enqueue                                                                                                           | Look at <code>V\$ENQUEUE_STAT</code> .                                                                                                                                                                                       |
| latch free              | Latch contention          | Depends on latch                                                                                                                     | Check <code>V\$LATCH</code> .                                                                                                                                                                                                |
| log buffer space        | Log buffer, I/O           | Log buffer small<br>Slow I/O system                                                                                                  | Check the statistic redo buffer allocation retries in <code>V\$SYSSTAT</code> . Check configuring log buffer section in configuring memory chapter. Check the disks that house the online redo logs for resource contention. |
| log file sync           | I/O, over-committing      | Slow disks that store the online logs<br>Un-batched commits                                                                          | Check the disks that house the online redo logs for resource contention. Check the number of transactions ( <code>commits + rollbacks</code> ) each second, from <code>V\$SYSSTAT</code> .                                   |

**See Also:**

- ["Wait Events"](#) on page 22-24 for detailed information on each event listed in [Table 22-1](#) and for other information to cross-check
- [Chapter 24, "Dynamic Performance Views for Tuning"](#) for detailed information on querying V\$ views

## Additional Statistics

There are a number of statistics that can indicate performance problems that do not have corresponding wait events.

### Redo Log Space Requests Statistic

The V\$SYSSTAT statistic `redo log space requests` indicates how many times a server process had to wait for space in the online redo log, not for space in the redo log buffer. A significant value for this statistic and the wait events should be used as an indication that checkpoints, DBWR, or archiver activity should be tuned, not LGWR. Increasing the size of log buffer does not help.

### Read Consistency

Your system might spend excessive time rolling back changes to blocks in order to maintain a consistent view. Consider the following scenarios:

- If there are many small transactions and an active long-running query is running in the background on the same table where the changes are happening, then the query might need to roll back those changes often, in order to obtain a read-consistent image of the table. Compare the following V\$SYSSTAT statistics to determine whether this is happening:
  - `consistent changes` statistic indicates the number of times a database block has rollback entries applied to perform a consistent read on the block. Workloads that produce a great deal of `consistent changes` can consume a great deal of resources.
  - `consistent gets` statistic counts the number of logical reads in consistent mode.
- If there are few very, large rollback segments, then your system could be spending a lot of time rolling back the transaction table during delayed block cleanout in order to find out exactly which SCN a transaction was committed. The ratio of the following V\$SYSSTAT statistics should be close to 1:



```
ratio = transaction tables consistent reads undo records applied /
        transaction tables consistent read rollbacks
```

A solution is to create more, smaller rollback segments, or to use automatic undo management.

- If there are insufficient rollback segments, then there is rollback segment (header or block) contention. Evidence of this problem is available by the following:
  - Comparing the number of `WAITS` to the number of `GETS` in `V$ROLLSTAT`; the proportion of `WAITS` to `GETS` should be small.
  - Examining `V$WAITSTAT` to see whether there are many `WAITS` for buffers of `CLASS 'undo header'`.

A solution is to create more rollback segments.

### Table Fetch by Continued Row

You can detect migrated or chained rows by checking the number of `table fetch continued row` statistic in `V$SYSSTAT`. A small number of chained rows (less than 1%) is unlikely to impact system performance. However, a large percentage of chained rows can affect performance.

Chaining on rows larger than the block size is inevitable. You might want to consider using tablespaces with larger block size for such data.

However, for smaller rows, you can avoid chaining by using sensible space parameters and good application design. For example, do *not* insert a row with key values filled in and nulls in most other columns, then update that row with the real data, causing the row to grow in size. Rather, insert rows filled with data from the start.

If an `UPDATE` statement increases the amount of data in a row so that the row no longer fits in its data block, then Oracle tries to find another block with enough free space to hold the entire row. If such a block is available, then Oracle moves the entire row to the new block. This is called migrating a row. If the row is too large to fit into any available block, then Oracle splits the row into multiple pieces and stores each piece in a separate block. This is called chaining a row. Rows can also be chained when they are inserted.

Migration and chaining are especially detrimental to performance with the following:

- `UPDATE` statements that cause migration and chaining to perform poorly

- Queries that select migrated or chained rows because these must perform additional input and output

Identify migrated and chained rows in a table or cluster using the `ANALYZE` statement with the `LIST CHAINED ROWS` clause. This statement collects information about each migrated or chained row and places this information in a specified output table.

---

---

**Note:** Oracle Corporation strongly recommends that you use the `DBMS_STATS` package rather than `ANALYZE` to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by `DBMS_STATS`. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the `ANALYZE` statement rather than `DBMS_STATS` for statistics collection not related to the cost-based optimizer, such as:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
  - To collect information on freelist blocks
- 
- 

The definition of a sample output table named `CHAINED_ROWS` appears in a SQL script available on your distribution medium. The common name of this script is `UTLCHN1.SQL`, although its exact name and location varies depending on your platform. Your output table must have the same column names, datatypes, and sizes as the `CHAINED_ROWS` table.

Increasing `PCTFREE` can help to avoid migrated rows. If you leave more free space available in the block, then the row has room to grow. You can also reorganize or re-create tables and indexes that have high deletion rates. If tables frequently have rows deleted, then data blocks can have partially free space in them. If rows are inserted and later expanded, then the inserted rows might land in blocks with deleted rows but still not have enough room to expand. Reorganizing the table ensures that the main free space is totally empty blocks.

---

---

**Note:** `PCTUSED` is not the opposite of `PCTFREE`.

---

---

**See Also:**

- *Oracle9i Database Concepts* for more information on PCTUSED
- *Oracle9i Database Administrator's Guide* for information on reorganizing tables

**Parse-Related Statistics**

The more your application parses, the more potential for contention exists, and the more time your system spends waiting. If `parse time CPU` represents a large percentage of the CPU time, then time is being spent parsing instead of executing statements. If this is the case, then it is likely that the application is using literal SQL and so SQL cannot be shared, or the shared pool is poorly configured.

**See Also:** [Chapter 14, "Memory Configuration and Use"](#)

There are a number of statistics available to identify the extent of time spent parsing by Oracle. Query the parse related statistics from `V$SYSSTAT`. For example:

```
SELECT NAME, VALUE
       FROM V$SYSSTAT
      WHERE NAME IN ( 'parse time cpu', 'parse time elapsed'
                    , 'parse count (hard)', 'CPU used by this session' );
```

There are various ratios that can be computed to assist in determining whether parsing may be a problem:

- `parse time CPU / parse time elapsed`

This ratio indicates how much of the CPU time spent parsing was due to the parse operation itself, rather than waiting for resources, such as latches. A ratio of one is good, indicating that the elapsed time was not spent waiting for highly contended resources.

- `parse time CPU / CPU used by this session`

This ratio indicates how much of the total CPU used by Oracle server processes was spent on parse-related operations. A ratio closer to zero is good, indicating that the majority of CPU is not spent on parsing.

## Wait Events

The views `V$SESSION_WAIT`, `V$SESSION_EVENT` and `V$SYSTEM_EVENT` provide information on what resources were waited for, and, if the configuration parameter `TIMED_STATISTICS` is set to `true`, how long each resource was waited for.

---

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

---

**See Also:** ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

Investigate wait events and related timing data when performing reactive performance tuning. The events with the most time listed against them are often strong indications of the performance bottleneck.

The three views contain related, but different, views of the same data:

- `V$SESSION_WAIT` is a current state view. It lists either the event currently being waited for or the event last waited for on each session
- `V$SESSION_EVENT` lists the cumulative history of events waited for on each session. After a session exits, the wait event statistics for that session are removed from this view.
- `V$SYSTEM_EVENT` lists the events and times waited for by the whole instance (that is, all session wait events data rolled up) since instance startup.

Because `V$SESSION_WAIT` is a current state view, it also contains a finer-granularity of information than `V$SESSION_EVENT` or `V$SYSTEM_EVENT`. It includes additional identifying data for the current event in three parameter columns: `P1`, `P2`, and `P3`.

For example, `V$SESSION_EVENT` can show that session 124 (`SID=124`) had many waits on the `db file scattered read` event, but it does not show which file and

block number. However, `V$SESSION_WAIT` shows the file number in P1, the block number read in P2, and the number of blocks read in P3 (P1 and P2 let you determine for which segments the wait event is occurring).

This chapter concentrates on examples using `V$SESSION_WAIT`. However, Oracle recommends capturing performance data over an interval and keeping this data for performance and capacity analysis. This form of rollup data is queried from the `V$SYSTEM_EVENT` view by tools such as Enterprise Manager Diagnostics Pack and Statspack.

Most commonly encountered events are described in this chapter, listed in case-sensitive alphabetical order. Other event-related data to examine is also included. The case used for each event name is that which appears in the `V$SYSTEM_EVENT` view.

**See Also:** *Oracle9i Database Reference* for a complete list of wait events

## SQL\*Net

The following events signify that the database process is waiting for acknowledgment from a database link or a client process:

- SQL\*Net break/reset to client
- SQL\*Net break/reset to dblink
- SQL\*Net message from client
- SQL\*Net message from dblink
- SQL\*Net message to client
- SQL\*Net message to dblink
- SQL\*Net more data from client
- SQL\*Net more data from dblink
- SQL\*Net more data to client
- SQL\*Net more data to dblink

If these waits constitute a significant portion of the wait time on the system or for a user experiencing response time issues, then the network or the middle-tier could be a bottleneck.

Events that are client-related should be diagnosed as described for the event `SQL*Net message from client`. Events that are `dblink`-related should be diagnosed as described for the event `SQL*Net message from dblink`.

### **SQL\*Net message from client**

Although this is an idle event, it is important to explain when this event can be used to diagnose what is not the problem. This event indicates that a server process is waiting for work from the client process. However, there are several situations where this event could accrue most of the wait time for a user experiencing poor response time. The cause could be either a network bottleneck or a resource bottleneck on the client process.

**Network Bottleneck** A network bottleneck can occur if the application causes a lot of traffic between server and client and the network latency (time for a round-trip) is high. Symptoms include the following:

- Large number of waits for this event
- Both the database and client process are idle (waiting for network traffic) most of the time

To alleviate network bottlenecks, try the following:

- Tune the application to reduce round trips.
- Explore options to reduce latency (for example, terrestrial lines opposed to VSAT links).
- Change system configuration to move higher traffic components to lower latency links.

**See Also:** *Oracle9i Database Performance Planning*

**Resource Bottleneck on the Client Process** If the client process is using most of the resources, then there is nothing that can be done in the database. Symptoms include the following:

- Number of waits might not be large, but the time waited might be significant
- Client process has a high resource usage

In some cases, you can see the wait time for a waiting user tracking closely with the amount of CPU used by the client process. The term client here refers to any process other than the database process (middle-tier, desktop client) in the n-tier architecture.

## SQL\*Net message from dblink

This event signifies that the session has sent a query to the remote node and is waiting for a response from the database link. This time could go up because of the following:

- Network bottleneck
  - For information, see ["SQL\\*Net message from client"](#) on page 22-26.
- Time taken to run the query on the remote node
  - It is useful to see the query being run on the remote node. Login to the remote database, find the session created by the database link, and examine the SQL statement being run by it.

## buffer busy waits

This wait indicates that there are some buffers in the buffer cache that multiple processes are attempting to access concurrently. Query `V$WAITSTAT` for the wait statistics for each class of buffer. Common buffer classes that have buffer busy waits include data block, segment header, undo header, and undo block.

Check the following `V$SESSION_WAIT` parameter columns:

- P1 - File ID
- P2 - Block ID

## Causes

To determine the possible causes, identify the type of class contended for by querying `V$WAITSTAT`:

```
SELECT class, count
   FROM V$WAITSTAT
  WHERE count > 0
   ORDER BY count DESC;
```

Example output:

| CLASS          | COUNT |
|----------------|-------|
| -----          | ----- |
| data block     | 43383 |
| undo header    | 10680 |
| undo block     | 5237  |
| segment header | 785   |

To identify the segment and segment type contended for, query `DBA_EXTENTS` using the values for File Id and Block Id returned from `V$SESSION_WAIT` (p1 and p2 columns):

```
SELECT segment_owner, segment_name
       FROM DBA_EXTENTS
       WHERE file_id = <&p1>
             AND <&p2> BETWEEN block_id AND block_id + blocks - 1;
```

## Actions

The action required depends on the class of block contended for and the actual segment.

**segment header** If the contention is on the segment header, then this is most likely freelist contention.

Automatic segment-space management in locally managed tablespaces eliminates the need to specify the `PCTUSED`, `FREELISTS`, and `FREELIST GROUPS` parameters. If possible, switch from manual space management to automatic segment-space management.

The following information is relevant if you are *unable* to use automatic segment-space management (for example, because the tablespace uses dictionary space management).

A freelist is a list of free data blocks that usually includes blocks existing in a number of different extents within the segment. Blocks in freelists contain free space greater than `PCTFREE`. This is the percentage of a block to be reserved for updates to existing rows. In general, blocks included in process freelists for a database object must satisfy the `PCTFREE` and `PCTUSED` constraints. Specify the number of process freelists with the `FREELISTS` parameter. The default value of `FREELISTS` is one. The maximum value depends on the data block size.

To find the current setting for freelists for that segment, run the following:

```
SELECT SEGMENT_NAME, FREELISTS
       FROM DBA_SEGMENTS
       WHERE SEGMENT_NAME = segment name
             AND SEGMENT_TYPE = segment type;
```

Set freelists, or increase of number of freelists. If adding more freelists does not alleviate the problem, then use freelist groups (even in single instance this can make a difference). If using Oracle Real Application Clusters, then ensure that each instance has its own freelist group(s).



**See Also:**

- *Oracle9i Database Concepts* for information on automatic segment-space management, freelists, PCTFREE, and PCTUSED
- *Oracle9i Real Application Clusters Setup and Configuration* for information about using freelist groups in an Oracle Real Application Clusters environment

**data block** If the contention is on tables or indexes (not the segment header):

- Check for SQL statements using unselective indexes.
- Check for 'right-hand-indexes' (that is, indexes that are inserted into at the same point by many processes; for example, those which use sequence number generators for the key values).
- Consider using automatic segment-space management, or increasing freelists to avoid multiple processes attempting to insert into the same block.

**undo header** For contention on rollback segment header:

- If you are not using automatic undo management, then add more rollback segments.

**undo block** For contention on rollback segment block:

- If you are not using automatic undo management, then consider making rollback segment sizes larger.

## db file scattered read

This event signifies that the user process is reading buffers into the SGA buffer cache and is waiting for a physical I/O call to return. A `db file scattered read` issues a scatter-read to read the data into multiple discontinuous memory locations. A scattered read is usually a multiblock read. It can occur for a fast full scan (of an index) in addition to a full table scan.

The `db file scattered read` wait event identifies that a full table scan is occurring. When performing a full table scan into the buffer cache, the blocks read are read into memory locations that are not physically adjacent to each other. Such reads are called scattered read calls, because the blocks are scattered throughout memory. This is why the corresponding wait event is called 'db file scattered read'. Multiblock (up to `DB_FILE_MULTIBLOCK_READ_COUNT` blocks) reads due to full table scans into the buffer cache show up as waits for 'db file scattered read'.

Check the following `V$SESSION_WAIT` parameter columns:

- P1 - The absolute file number
- P2 - The block being read
- P3 - The number of blocks (should be greater than 1)

### Actions

On a healthy system, physical read waits should be the biggest waits after the idle waits. However, also consider whether there are direct read waits (signifying full table scans with parallel query) or `db file scattered read` waits on an operational (OLTP) system that should be doing small indexed accesses.

Other things that could indicate excessive I/O load on the system include the following:

- Poor buffer cache hit ratio
- These wait events accruing most of the wait time for a user experiencing poor response time

### Managing Excessive I/O

There are several ways to handle excessive I/O waits. In the order of effectiveness, these are as follows:

1. Reduce the I/O activity by SQL tuning.
2. Reduce the need to do I/O by managing the workload.
3. Add more disks to reduce the number of I/Os for each disk.
4. Alleviate I/O hot spots by redistributing I/O across existing disks.

**See Also:** [Chapter 15, "I/O Configuration and Design"](#)

The first course of action should be to find opportunities to reduce I/O. Examine the SQL statements being run by sessions waiting for these events, as well as statements causing high physical I/Os from `V$SQLAREA`. Factors that can adversely affect the execution plans causing excessive I/O include the following:

- Improperly optimized SQL
- Missing indexes
- High degree of parallelism for the table (skewing the optimizer toward scans)

- Lack of accurate statistics for the optimizer

### Inadequate I/O Distribution

Besides reducing I/O, also examine the I/O distribution of files across the disks. Is I/O distributed uniformly across the disks, or are there hot spots on some disks? Are the number of disks sufficient to meet the I/O needs of the database?

See the total I/O operations (reads and writes) by the database, and compare those with the number of disks used. Remember to include the I/O activity of LGWR and ARCH processes.

### Finding the SQL Statement executed by Sessions Waiting for I/O

Use the following query to see, at a point in time, which sessions are waiting for I/O:

```
SELECT s.sql_address, s.sql_hash_value
       FROM V$SESSION s, V$SESSION_WAIT w
       WHERE w.event LIKE 'db file%read'
              AND w.sid = s.sid ;
```

### Finding the Object Requiring I/O

Use the following query to find the object being accessed:

```
SELECT segment_owner, segment_name
       FROM DBA_EXTENTS
       WHERE file_id = &p1
              AND &p2 between block_id AND block_id + blocks - 1 ;
```

## db file sequential read

This event signifies that the user process is reading buffers into the SGA buffer cache and is waiting for a physical I/O call to return. This call differs from a scattered read, because a sequential read is reading data into contiguous memory space. A sequential read is usually a single-block read.

Single block I/Os are usually the result of using indexes. Rarely, full table scan calls could get truncated to a single block call due to extent boundaries, or buffers already present in the buffer cache. These waits would also show up as 'db file sequential read'.

Check the following V\$SESSION\_WAIT parameter columns:

- P1 - The absolute file number

- P2 - The block being read
- P3 - The number of blocks (should be 1)

**See Also:** "[db file scattered read](#)" on page 22-29 for information on managing excessive I/O, inadequate I/O distribution, and finding the SQL causing the I/O and the segment the I/O is performed on

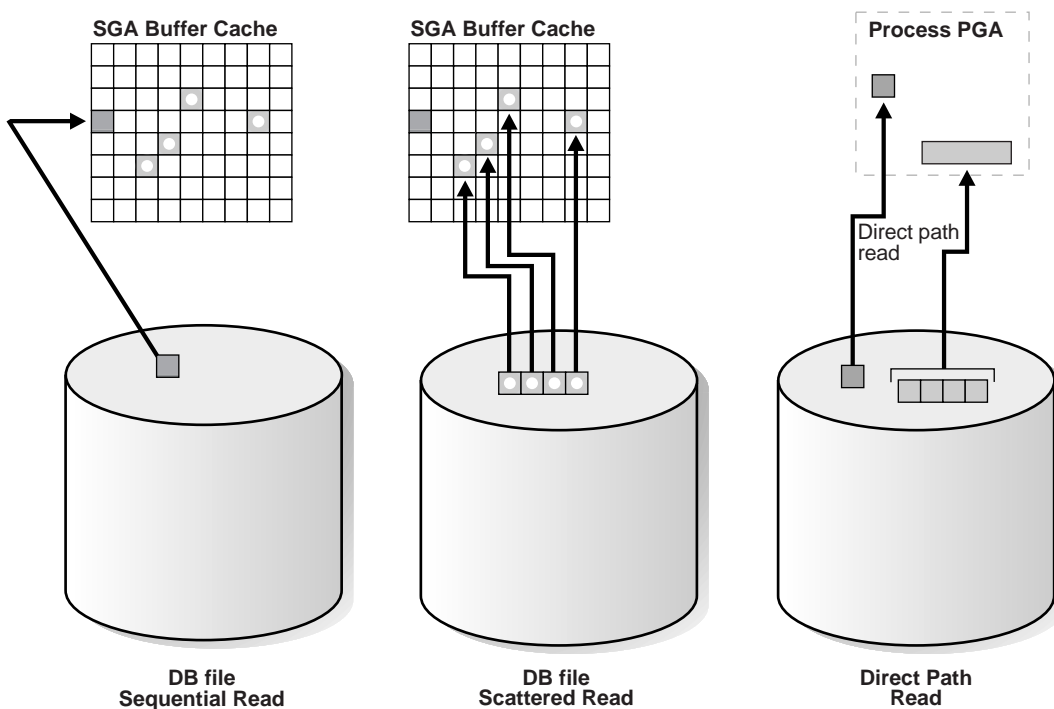
### Actions

On a healthy system, physical read waits should be the biggest waits after the idle waits. However, also consider whether there are `db file sequential read` on a large data warehouse that should be seeing mostly full table scans with parallel query.

[Figure 22-1](#) depicts the differences between the following wait events:

- `db file sequential read` (single block read into one SGA buffer)
- `db file scattered read` (multiblock read into many discontinuous SGA buffers)
- `direct read` (single or multiblock read into the PGA, bypassing the SGA)

Figure 22–1 Scattered Read, Sequential Read, and Direct Path Read



## direct path read and direct path read (lob)

When a session is reading buffers from disk directly into the PGA (opposed to the buffer cache in SGA), it waits on this event. If the I/O subsystem does not support asynchronous I/Os, then each wait corresponds to a physical read request.

If the I/O subsystem supports asynchronous I/O, then the process is able to overlap issuing read requests with processing the blocks already existing in the PGA. When the process attempts to access a block in the PGA that has not yet been read from disk, it then issues a wait call and updates the statistics for this event. Hence, the number of waits is not necessarily the same as the number of read requests (unlike 'db file scattered read' and 'db files sequential read').

Check the following `V$SESSION_WAIT` parameter columns:

- P1 - File\_id for the read call

- P2 - Start block\_id for the read call
- P3 - Number of blocks in the read call

### Causes

This happens in the following situations:

- The sorts are too large to fit in memory and go to disk. If a sort does not fit into memory, then some of the sort data is written out directly to disk. This data is later read back in, using direct reads.
- Parallel slaves are used for scanning data.
- The server process is processing buffers faster than the I/O system can return the buffers. This can indicate an overloaded I/O system.

### Actions

The `file_id` shows if the reads are for an object in `TEMP` tablespace (sorts to disk) or full table scans by parallel slaves. This is the biggest wait for large data warehouse sites. However, if the workload is not a DSS workload, then examine why this is happening.

**Sorts to Disk** Examine the SQL statement currently being run by the session experiencing waits to see what is causing the sorts. Query `V$TEMPSEG_USAGE` to find the SQL statement that is generating the sort. Also query the statistics from `V$SESSTAT` for the session to determine the size of the sort. See if it is possible to reduce the sorting by tuning the SQL statement. If `WORKAREA_SIZE_POLICY` is `MANUAL`, then consider increasing the `SORT_AREA_SIZE` for the system (if the sorts are not too big) or for individual processes. If `WORKAREA_SIZE_POLICY` is `AUTO`, then investigate whether to increase `PGA_AGGREGATE_TARGET`.

**See Also:** ["Configuring the PGA Working Memory"](#) on page 14-48

**Full Table Scans** If tables are defined with a high degree of parallelism, then this could skew the optimizer to use full table scans with parallel slaves. Check the object being read into using the direct path reads, as well as the SQL statement being run by the query-coordinator. If the full table scans are a valid part of the workload, then ensure that the I/O subsystem is sized adequately for the degree of parallelism.

**Hash Area Size** For query plans that call for a hash join, excessive I/O could result from having `HASH_AREA_SIZE` too small. If `WORKAREA_SIZE_POLICY` is `MANUAL`,

then consider increasing the `HASH_AREA_SIZE` for the system or for individual processes. If `WORKAREA_SIZE_POLICY` is `AUTO`, then investigate whether to increase `PGA_AGGREGATE_TARGET`.

**See Also:**

- ["Managing Excessive I/O"](#) on page 22-30
- ["Configuring the PGA Working Memory"](#) on page 14-48

## direct path write

When a process is writing buffers directly from PGA (as opposed to the DBWR writing them from the buffer cache), the process waits on this event for the write call to complete. Operations that could perform direct path writes include when a sort goes to disk, during parallel DML operations, direct-path `INSERTS`, parallel create table as select, and some LOB operations.

Like direct path reads, the number of waits is not the same as number of write calls issued if the I/O subsystem supports asynchronous writes. The session waits if it has processed all buffers in the PGA and is unable to continue work until an I/O request completes.

Check the following `V$SESSION_WAIT` parameter columns:

- `P1` - File\_id for the write call
- `P2` - Start block\_id for the write call
- `P3` - Number of blocks in the write call

### Causes

This happen in the following situations:

- Sorts are too large to fit in memory and are going to disk
- Parallel DML are issued to create/populate objects

### Actions

For large sorts see ["Sorts to Disk"](#) on page 22-34.

For parallel DML, check the I/O distribution across disks and make sure that the I/O subsystem is adequately sized for the degree of parallelism.

## enqueue

Enqueues are locks that serialize access to database resources. This event indicates that the session is waiting for a lock that is held by another session.

Check the following `V$SESSION_WAIT` parameter columns:

- P1 - Lock TYPE (or name) and MODE
- P2 - Resource identifier ID1 for the lock
- P3 - Resource identifier ID2 for the lock

Check the comparison with `V$LOCK` columns:

- `V$LOCK.ID1 = P2`
- `V$LOCK.ID2 = P3`

Performing the following SQL transformation of the P1 column results in the same value displayed in `V$LOCK.TYPE`:

```
V$LOCK.TYPE = chr(bitand(P1,-16777216)/16777215)||  
chr(bitand(P1,16711680)/65535)
```

To obtain the mode in which the enqueue is being requested, issue the following statement:

```
request = mod(P1, 65536);
```

### Finding Locks and Lock Holders

Query `V$LOCK` to find the sessions holding the lock. For every session waiting for the event `enqueue`, there is a row in `V$LOCK` with `REQUEST <> 0`. Therefore, use either of the following two queries to find the sessions holding the locks and waiting for the locks.

If there are enqueue waits, you can see these using the following statement:

```
SELECT * FROM V$LOCK WHERE request > 0;
```

To show only holders and waiters for locks being waited on, use the following:

```
SELECT DECODE(request,0,'Holder: ','Waiter: ')|| sid sess, id1, id2, lmode,  
request, type  
FROM V$LOCK  
WHERE (id1, id2, type) IN (SELECT id1, id2, type FROM V$LOCK WHERE request>0)  
ORDER BY id1, request;
```



**See Also:**

- [Chapter 24, "Dynamic Performance Views for Tuning"](#) for more information on using V\$LOCK
- *Oracle9i Database Reference* for more information on enqueues

**Actions**

The appropriate action depends on the type of enqueue.

**ST enqueue** If the contended-for enqueue is the ST enqueue, then the problem is most likely to be dynamic space allocation. Oracle dynamically allocates an extent to a segment when there is no more free space available in the segment. This enqueue is only used for dictionary managed tablespaces.

To solve contention on this resource:

- Check to see whether the temporary (that is, sort) tablespace uses TEMPFILES. If not, then switch to using TEMPFILES.
- Switch to using locally managed tablespaces if the tablespace that contains segments that are growing dynamically.

**See Also:** *Oracle9i Database Concepts* for detailed information on TEMPFILES and locally managed tablespaces

- If it is not possible to switch to locally managed tablespaces, then ST enqueue resource usage can be decreased by changing the next extent sizes of the growing objects to be large enough to avoid constant space allocation. To determine which segments are growing constantly, monitor the EXTENTS column of the DBA\_SEGMENTS view for all SEGMENT\_NAMES over time to identify which segments are growing and how quickly.
- Preallocate space in the segment (for example, by allocating extents using the ALTER TABLE ALLOCATE EXTENT SQL statement).

**HW enqueue** The HW enqueue is used to serialize the allocation of space beyond the high-water mark of a segment.

- V\$SESSION\_WAIT.P2 / V\$LOCK.ID1 is the tablespace number.
- V\$SESSION\_WAIT.P2 / V\$LOCK.ID2 is the relative dba of segment header of the object for which space is being allocated.

If this is a point of contention for an object, then manual allocation of extents solves the problem.

**TM enqueue** The most common reason for waits on TM locks tend to involve foreign key constraints where the constrained columns are not indexed. Index the foreign key columns to avoid this problem.

**TX enqueue** These are acquired exclusive when a transaction initiates its first change and held until the transaction does a `COMMIT` or `ROLLBACK`.

- Waits for TX in mode 6: occurs when a session is waiting for a row level lock that is already held by another session. This occurs when one user is updating or deleting a row, which another session wishes to update or delete.

The solution is to have the first session already holding the lock perform a `COMMIT` or `ROLLBACK`.

- Waits for TX in mode 4 can occur if the session is waiting for an ITL (interested transaction list) slot in a block. This happens when the session wants to lock a row in the block but one or more other sessions have rows locked in the same block, and there is no free ITL slot in the block. Usually, Oracle dynamically adds another ITL slot. This may not be possible if there is insufficient free space in the block to add an ITL. If so, the session waits for a slot with a TX enqueue in mode 4.

The solution is to increase the number of ITLs available, either by changing the `INITRANS` or `MAXTRANS` for the table (either by using an `ALTER` statement, or by re-creating the table with the higher values).

- Waits for TX in mode 4 can also occur if a session is waiting due to potential duplicates in `UNIQUE` index. If two sessions try to insert the same key value the second session has to wait to see if an `ORA-0001` should be raised or not.

The solution is to have the first session already holding the lock perform a `COMMIT` or `ROLLBACK`.

- Waits for TX in mode 4 is also possible if the session is waiting due to shared bitmap index fragment. Bitmap indexes index key values and a range of `ROWIDs`. Each 'entry' in a bitmap index can cover many rows in the actual table. If two sessions want to update rows covered by the same bitmap index fragment, then the second session waits for the first transaction to either `COMMIT` or `ROLLBACK` by waiting for the TX lock in mode 4.
- Waits for TX in Mode 4 can also occur waiting for a `PREPARED` transaction.

---

**See Also:** *Oracle9i Application Developer's Guide - Fundamentals* for more information about referential integrity and locking data explicitly

## free buffer waits

This wait event indicates that a server process was unable to find a free buffer and has posted the database writer to make free buffers by writing out dirty buffers. A dirty buffer is a buffer whose contents have been modified. Dirty buffers are freed for reuse when DBWR has written the blocks to disk.

### Causes

DBWR may not be keeping up with writing dirty buffers in the following situations:

- The I/O system is slow.
- There are resources it is waiting for, such as latches.
- The buffer cache is so small that DBWR spends most of its time cleaning out buffers for server processes.
- The buffer cache is so big that one DBWR process is not enough to free enough buffers in the cache to satisfy requests.

### Actions

If this event occurs frequently, then examine the session waits for DBWR to see whether there is anything delaying DBWR.

**Writes** If it is waiting for writes, then determine what is delaying the writes and fix it. Check the following:

- Examine `V$FILESTAT` to see where most of the writes are happening.
- Examine the host operating system statistics for the I/O system. Are the write times acceptable?

If I/O is slow:

- Consider using faster I/O alternatives to speed up write times.
- Spread the I/O activity across large number of spindles (disks) and controllers.

**See Also:** [Chapter 15, "I/O Configuration and Design"](#) for information on balancing I/O

**Cache is Too Small** It is possible DBWR is very active because of the cache is too small. Investigate whether this is a probable cause by looking to see if the buffer cache hit ratio is low. Also use the `V$DB_CACHE_ADVICE` view to determine whether a larger cache size would be advantageous.

**See Also:** ["Sizing the Buffer Cache"](#) on page 14-6

**Cache Is Too Big for One DBWR** If the cache size is adequate and the I/O is already evenly spread, then you can potentially modify the behavior of DBWR by using asynchronous I/O or by using multiple database writers.

### **Consider Multiple Database Writer (DBWR) Processes or I/O Slaves**

Configuring multiple database writer processes, or using I/O slaves, is useful when the transaction rates are high or when the buffer cache size is so large that a single DBWR process cannot keep up with the load.

**DB\_WRITER\_PROCESSES** The `DB_WRITER_PROCESSES` initialization parameter lets you configure multiple database writer processes (from DBW0 to DBW9 and from DBWa to DBWj). Configuring multiple DBWR processes distributes the work required to identify buffers to be written, and it also distributes the I/O load over these processes. Multiple db writer processes are highly recommended for systems with multiple CPUs (at least one db writer for every 8 CPUs) or multiple processor groups (at least as many db writers as processor groups).

Based upon the number of CPUs and the number of processor groups, Oracle either selects an appropriate default setting for `DB_WRITER_PROCESSES` or adjusts a user-specified setting.

**DBWR\_IO\_SLAVES** If it is not practical to use multiple DBWR processes, then Oracle provides a facility whereby the I/O load can be distributed over multiple slave processes. The DBWR process is the only process that scans the buffer cache LRU list for blocks to be written out. However, the I/O for those blocks is performed by the I/O slaves. The number of I/O slaves is determined by the parameter `DBWR_IO_SLAVES`.

`DBWR_IO_SLAVES` is intended for scenarios where you cannot use multiple `DB_WRITER_PROCESSES` (for example, where you have a single CPU). I/O slaves are also useful when asynchronous I/O is not available, because the multiple I/O slaves simulate nonblocking, asynchronous requests by freeing DBWR to continue identifying blocks in the cache to be written. Asynchronous I/O at the operating system level, if you have it, is generally preferred.

---

DBWR I/O slaves are allocated immediately following database open when the first I/O request is made. The DBWR continues to perform all of the DBWR-related work, apart from performing I/O. I/O slaves simply perform the I/O on behalf of DBWR. The writing of the batch is parallelized between the I/O slaves.

---

**Note:** Implementing `DBWR_IO_SLAVES` requires that extra shared memory be allocated for I/O buffers and request queues. Multiple DBWR processes cannot be used with I/O slaves. Configuring I/O slaves forces only one DBWR process to start.

---

**Choosing Between Multiple DBWR Processes and I/O Slaves** Configuring multiple DBWR processes benefits performance when a single DBWR process is unable to keep up with the required workload. However, before configuring multiple DBWR processes, check whether asynchronous I/O is available and configured on the system. If the system supports asynchronous I/O but it is not currently used, then enable asynchronous I/O to see if this alleviates the problem. If the system does not support asynchronous I/O, or if asynchronous I/O is already configured and there is still a DBWR bottleneck, then configure multiple DBWR processes.

---

**Note:** If asynchronous I/O is not available on your platform, then asynchronous I/O can be disabled by setting the `DISK_ASYNC_IO` initialization parameter to `FALSE`.

---

Using multiple DBWRs parallelizes the gathering and writing of buffers. Therefore, multiple DBWR processes should deliver more throughput than one DBWR process with the same number of I/O slaves. For this reason, the use of I/O slaves has been deprecated in favor of multiple DBWR processes. I/O slaves should only be used if multiple DBWR processes cannot be configured.

**See Also:** [Chapter 17, "Configuring Instance Recovery Performance"](#) for details on tuning checkpoints

## latch free

A latch is a low-level internal lock used by Oracle to protect memory structures. The latch free event is updated when a server process attempts to get a latch, and the latch is unavailable on the first attempt.

**See Also:** *Oracle9i Database Concepts* for more information on latches and internal locks

### **Actions**

This event should only be a concern if latch waits are a significant portion of the wait time on the system as a whole, or for individual users experiencing problems.

- To help determine the cause of this wait event, identify the latch(es) contended for. There are many types of latches used for different purposes. For example, the shared pool latch protects certain actions in the shared pool, and the cache buffers LRU chain protects certain actions in the buffer cache.
- Examine the resource usage for related resources. For example, if the library cache latch is heavily contended for, then examine the hard and soft parse rates.
- Examine the SQL statements for the sessions experiencing latch contention to see if there is any commonality.

Check the following V\$SESSION\_WAIT parameter columns:

- P1 - Address of the latch
- P2 - Latch number
- P3 - Number of times process has already slept, waiting for the latch

### **Example: Find Latches Currently Waiting For**

```
SELECT n.name, SUM(w.p3) Sleeps
   FROM V$SESSION_WAIT w, V$LATCHNAME n
  WHERE w.event = 'latch free'
        AND w.p2 = n.latch#
  GROUP BY n.name;
```

**Table 22–2 Latch Free Wait Event**

| Latch                      | SGA Area               | Possible Causes                                                                                                                                                                                                                                                                                                                             | Look For:                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Shared pool, library cache | Shared pool            | <p>Lack of statement reuse</p> <p>Statements not using bind variables</p> <p>Insufficient size of application cursor cache</p> <p>Cursors closed explicitly after each execution</p> <p>Frequent logon/logoffs</p> <p>Underlying object structure being modified (for example truncate)</p> <p>Shared pool too small</p>                    | <p>Sessions (in V\$SESSTAT) with high:</p> <p>parse time CPU</p> <p>parse time elapsed</p> <p>Ratio of parse count (hard) / execute count</p> <p>Ratio of parse count (total) / execute count</p> <p>Cursors (in V\$SQLAREA/V\$SQL) with:</p> <p>High ratio of PARSE_CALLS / EXECUTIONS</p> <p>EXECUTIONS = 1 differing only in literals in the WHERE clause (that is, no bind variables used)</p> <p>High RELOADS</p> <p>High INVALIDATIONS</p> <p>Large (&gt; 1mb) SHARABLE_MEM</p> |
| cache buffers lru chain    | Buffer cache LRU lists | <p>Excessive buffer cache throughput. For example, inefficient SQL that accesses incorrect indexes iteratively (large index range scans) or many full table scans</p> <p>DBWR not keeping up with the dirty workload; hence, foreground process spends longer holding the latch looking for a free buffer</p> <p>Cache may be too small</p> | <p>Statements with very high logical I/O or physical I/O, using unselective indexes</p>                                                                                                                                                                                                                                                                                                                                                                                               |
| cache buffers chains       | Buffer cache buffers   | <p>Repeated access to a block (or small number of blocks), known as a <b>hot block</b></p>                                                                                                                                                                                                                                                  | <p>Sequence number generation code that updates a row in a table to generate the number, rather than using a sequence number generator</p> <p>Index leaf chasing from very many processes scanning the same unselective index with very similar predicate</p> <p>Identify the segment the hot block belongs to</p>                                                                                                                                                                    |

### Shared Pool and Library Cache Latch Contention

A main cause of shared pool or library cache latch contention is parsing. There are a number of techniques that can be used to identify unnecessary parsing and a number of types of unnecessary parsing:

**Unshared SQL** This method identifies similar SQL statements that could be shared if literals were replaced with bind variables. The idea is to either:

- Manually inspect SQL statements that have only one execution to see whether they are similar:

```
SELECT sql_text
       FROM V$SQLAREA
      WHERE executions < 4
      ORDER BY sql_text;
```

- Or, automate this process by grouping together what may be similar statements. Do this by estimating the number of bytes of a SQL statement which will likely be the same, and group the SQL statements by that many bytes. For example, the following example groups together statements that differ only after the first 60 bytes.

```
SELECT SUBSTR(sql_text,1, 60), COUNT(*)
       FROM V$SQLAREA
      WHERE executions < 4
      GROUP BY SUBSTR(sql_text, 1, 60)
      HAVING COUNT(*) > 1;
```

**Reparsed Sharable SQL** check the V\$SQLAREA view. Enter the following query:

```
SELECT SQL_TEXT, PARSE_CALLS, EXECUTIONS
       FROM V$SQLAREA
      ORDER BY PARSE_CALLS;
```

When the PARSE\_CALLS value is close to the EXECUTIONS value for a given statement, you might be continually reparsing that statement. Tune the statements with the higher numbers of parse calls.

**By Session** Identify unnecessary parse calls by identifying the session in which they occur. It might be that particular batch programs or certain types of applications do most of the reparsing. To do this, run the following query:

```
SELECT pa.sid, pa.value "Hard Parses", ex.value "Execute Count"
       FROM v$sesstat pa, v$sesstat ex
      WHERE pa.sid=ex.sid
      AND pa.statistic#=(select statistic#
                        FROM v$statname where name='parse count (hard)')
      AND ex.statistic#=(select statistic#
                        FROM v$statname where name='execute count')
      AND pa.value>0;
```



The result is a list of all sessions and the amount of reparsing they do. For each system identifier (SID), go to `V$SESSION` to find the name of the program that causes the reparsing.

---

**Note:** Because this query counts all parse calls since instance startup, it is best to look for sessions with high *rates* of parse. For example, a connection which has been up for 50 days might show a high parse figure, but a second connection might have been up for 10 minutes and be parsing at a much faster rate.

---

The output is similar to the following:

| SID | Hard Parses | Execute Count |
|-----|-------------|---------------|
| 7   | 1           | 20            |
| 8   | 3           | 12690         |
| 6   | 26          | 325           |
| 11  | 84          | 1619          |

**cache buffer lru chain** The `cache buffer lru chain` latches protect the lists of buffers in the cache. When adding, moving, or removing a buffer from a list, a latch must be obtained.

For symmetric multiprocessor (SMP) systems, Oracle automatically sets the number of LRU latches to a value equal to one half the number of CPUs on the system. For non-SMP systems, one LRU latch is sufficient.

Contention for the LRU latch can impede performance on SMP machines with a large number of CPUs. LRU latch contention is detected by querying `V$LATCH`, `V$SESSION_EVENT`, and `V$SYSTEM_EVENT`. To avoid contention, consider tuning the application, bypassing the buffer cache for DSS jobs, or redesigning the application.

**cache buffers chains** The `cache buffers chains` latches are used to protect a buffer list in the buffer cache. These latches are used when searching for, adding, or removing a buffer from the buffer cache. Contention on this latch usually means that there is a block that is greatly contended for (known as a hot block).

To identify the heavily accessed buffer chain, and hence the contended for block, look at latch statistics for the `cache buffers chains` latches using the view `V$LATCH_CHILDREN`. If there is a specific `cache buffers chains` child latch that

has many more GETS, MISSES, and SLEEPS when compared with the other child latches, then this is the contended for child latch.

This latch has a memory address, identified by the ADDR column. Use the value in the ADDR column joined with the V\$BH view to identify the blocks protected by this latch. For example, given the address (V\$LATCH\_CHILDREN.ADDR) of a heavily contended latch, this queries the file and block numbers:

```
SELECT file#, dbablk, class, state, TCH
FROM X$BH
WHERE HLADDR='address of latch';
```

X\$BH.TCH is a touch count for the buffer. A high value for X\$BH.TCH indicates a hot block.

Many blocks are protected by each latch. One of these buffers will probably be the hot block. Any block with a high TCH value is a potential hot block. Perform this query a number of times, and identify the block that consistently appears in the output. After you have identified the hot block, query DBA\_EXTENTS using the file number and block number, to identify the segment.

**See Also:** ["Finding the Object Requiring I/O"](#) on page 22-31 for instructions on how to do this

## log buffer space

This event occurs when server processes are waiting for free space in the log buffer, because you are writing redo to the log buffer faster than LGWR can write it out.

### Actions

Modify the redo log buffer size. If the size of the log buffer is already reasonable, then ensure that the disks on which the online redo logs reside do not suffer from I/O contention. The log buffer space wait event could be indicative of either disk I/O contention on the disks where the redo logs reside, or of a too-small log buffer. Check the I/O profile of the disks containing the redo logs to investigate whether the I/O system is the bottleneck. If the I/O system is not a problem, then the redo log buffer could be too small. Increase the size of the redo log buffer until this event is no longer significant.

## log file switch

There are two wait events commonly encountered:

- log file switch (archiving needed)
- log file switch (checkpoint incomplete)

In both of the events, the LGWR is unable to switch into the next online redo log, and all the commit requests wait for this event.

### Actions

For the `log file switch (archiving needed)` event, examine why the archiver is unable to archive the logs in a timely fashion. It could be due to the following:

- Archive destination is running out of free space.
- Archiver is not able to read redo logs fast enough (contention with the LGWR).
- Archiver is not able to write fast enough (contention on the archive destination, or not enough ARCH processes). If you have ruled out other possibilities (such as slow disks or a full archive destination) consider increasing the number of ARCH processes. The default is 2.
- If you have mandatory remote shipped archive logs, check whether this process is slowing down because of network delays or the write is not completing because of errors.

Depending on the nature of bottleneck, you might need to redistribute I/O or add more space to the archive destination to alleviate the problem. For the `log file switch (checkpoint incomplete)` event:

- Check if DBWR is slow, possibly due to an overloaded or slow I/O system. Check the DBWR write times, check the I/O system, and distribute I/O if necessary.

**See Also:** [Chapter 15, "I/O Configuration and Design"](#)

- Check if there are too few, or too small redo logs. If you have a few redo logs or small redo logs (for example two x 100k logs), and your system produces enough redo to cycle through all of the logs before DBWR has been able to complete the checkpoint, then increase the size or number of redo logs.

**See Also:** ["Sizing Redo Log Files"](#) on page 13-5

## log file sync

When a user session commits (or rolls back), the session's redo information must be flushed to the redo logfile by LGWR. The server process performing the `COMMIT` or `ROLLBACK` waits under this event for the write to the redo log to complete.

### Actions

If this event's waits constitute a significant wait on the system or a significant amount of time waited by a user experiencing response time issues or on a system, then examine the average time waited.

If the average time waited is low, but the number of waits are high, then the application might be committing after every `INSERT`, rather than batching `COMMITs`. Applications can reduce the wait by committing after 50 rows, rather than every row.

If the average time waited is high, then examine the session waits for the log writer and see what it is spending most of its time doing and waiting for. If the waits are because of slow I/O, then try the following:

- Reduce other I/O activity on the disks containing the redo logs, or use dedicated disks.
- Alternate redo logs on different disks to minimize the effect of the archiver on the log writer.
- Move the redo logs to faster disks or a faster I/O subsystem (for example, switch from RAID 5 to RAID 1).
- Consider using raw devices (or simulated raw devices provided by disk vendors) to speed up the writes.
- Depending on the type of application, it might be possible to batch `COMMITs` by committing every *N* rows, rather than every row, so that fewer log file syncs are needed.

## rdbms ipc reply

This event is used to wait for a reply from one of the background processes.

## Idle Wait Events

These events indicate that the server process is waiting because it has no work. This usually implies that if there is a bottleneck, then the bottleneck is not for database resources.

The majority of the idle events should be ignored when tuning, because they do not indicate the nature of the performance bottleneck. Some idle events can be useful in indicating what the bottleneck is not. An example of this type of event is the most commonly encountered idle wait-event `SQL Net message from client`. This and other idle events (and their categories) are listed in [Table 22-3](#).

**Table 22-3** Idle Wait Events

| Wait Name                            | Background Process Idle Event | User Process Idle Event | Parallel Query Idle Event | Shared Server Idle Event | Oracle Real Application Clusters Idle Event |
|--------------------------------------|-------------------------------|-------------------------|---------------------------|--------------------------|---------------------------------------------|
| dispatcher timer                     | .                             | .                       | .                         | X                        | .                                           |
| lock manager wait for remote message | .                             | .                       | .                         | .                        | X                                           |
| pipe get                             | .                             | X                       | .                         | .                        | .                                           |
| pmon timer                           | X                             | .                       | .                         | .                        | .                                           |
| PX Idle Wait                         | .                             | .                       | X                         | .                        | .                                           |
| PX Deq Credit: need buffer           | .                             | .                       | X                         | .                        | .                                           |
| PX Deq Credit: send blkd             | .                             | .                       | X                         | .                        | .                                           |
| rdbms ipc message                    | X                             | .                       | .                         | .                        | .                                           |
| smon timer                           | X                             | .                       | .                         | .                        | .                                           |
| SQL*Net message from client          | .                             | X                       | .                         | .                        | .                                           |
| virtual circuit status               | .                             | .                       | .                         | X                        | .                                           |

---

**Note:** If Statspack is installed, then it is also possible to query the `STATS$IDLE_EVENT` table, which contains a list of idle events.

---

**See Also:** *Oracle9i Database Reference* for explanations of each idle wait event



---

## Tuning Networks

This chapter describes different connection models and introduces networking issues that affect tuning.

This chapter contains the following sections:

- [Understanding Connection Models](#)
- [Detecting Network Problems](#)
- [Solving Network Problems](#)

## Understanding Connection Models

The techniques used to determine the source of problems vary depending on the configuration. You can have a shared server configuration or a dedicated server configuration.

- If you have a shared server configuration, then LSNRCTL services lists dispatchers.
- If you have a dedicated server configuration, then LSNRCTL services lists dedicated servers.

It is possible to connect to dedicated server with a database configured for shared servers by placing the parameter (`SERVER = DEDICATED`) in the connect descriptor.

### Shared Server Configuration

This section discusses the setups for the shared server configuration.

**Registering the Dispatchers** The LSNRCTL control utility's `services` statement lists every dispatcher registered with it. This list includes the dispatchers process ID. You can check the alert log to confirm that the dispatchers have been started successfully.

---

---

**Note:** Remember that PMON can take a minute to register the dispatcher with the listener.

---

---

```
LSNRCTL> services
Connecting to
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=helios)(PORT=1521)))
Services Summary...
Service "sales.us.acme.com" has 1 instance(s).
  Instance "sales", status READY, has 3 handler(s) for this service...
  Handler(s):
    "DEDICATED" established:0 refused:0 state:ready
      LOCAL SERVER
    "D000" established:0 refused:0 current:0 max:10000 state:ready
      DISPATCHER <machine: helios, pid: 1689>
      (ADDRESS=(PROTOCOL=tcp)(HOST=helios)(PORT=52414))
    "D001" established:0 refused:0 current:0 max:10000 state:ready
      DISPATCHER <machine: helios, pid: 1691>
      (ADDRESS=(PROTOCOL=tcp)(HOST=helios)(PORT=52415))
The command completed successfully.
```



**See Also:** *Oracle9i Net Services Administrator's Guide* for information on setting the output mode

### Configuring the Initialization Parameter File

- Make sure that the DISPATCHERS line is correctly set. For example:

```
DISPATCHERS = "(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
                (HOST=hostname)(PORT=1492)(queue_size=32)))
                (DISPATCHERS = 1)
                (LISTENER = alias)
                (SERVICE = servicename)
                (SESSIONS = 1000)
                (CONNECTIONS = 1000)
                (MULTIPLEX = ON)
                (POOL = ON)
                (TICK = 5)"
```

One, and only one, of the following attributes is required:

- PROTOCOL
- ADDRESS
- DESCRIPTION

ADDRESS and DESCRIPTION provide support for the specification of additional network attributes beyond PROTOCOL. In the previous example, the entire DISPATCHERS line can be (PROTOCOL=TCP). The attributes DISPATCHERS, LISTENER, SERVICE, SESSIONS, CONNECTIONS, MULTIPLEX, POOL, and TICKS are all optional.

- Make sure that the optional MAX\_DISPATCHERS line is correctly set. For example:

```
MAX_DISPATCHERS = 4
```

This line should reflect the total number of dispatchers you want to start.

- Make sure that the optional MAX\_SHARED\_SERVERS line is correctly set. For example:

```
MAX_SHARED_SERVERS = 5
```

This line sets the upper bound on the total number of shared servers PMON can create, based on the peak load of the system. This should be set high enough so that all requests can be serviced, but not so high that the system swaps if they

are reached. The purpose of this parameter is to prevent the server from swapping. Run the following script to see what the highwater mark is for the number of servers running, and then set `MAX_SHARED_SERVERS` to more than this.

```
SELECT maximum_connections "MAX CONN", servers_started "STARTED", servers_
terminated "TERMINATED", servers_highwater "HIGHWATER" FROM V$SHARED_SERVER_
MONITOR;
```

- **Make sure that the optional `SHARED_SERVERS` line is correctly set. For example:**

```
SHARED_SERVERS = 5
```

This is the total number of shared servers started when the database is started. It also represents the total number of shared servers PMON tries to keep. It should be the total number of servers expected to be used when the database is active. `MAX_SHARED_SERVERS` is intended to handle peak load.

**Checking the Connections** Use the `LSNRCTL` control utility's `services` command to see if there are excessive connection refusals. Check the listener's log file to see if this is a connection problem. For example:

```
LSNRCTL> services
Connecting to
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=helios)(PORT=1521)))
Services Summary...
Service "sales.us.acme.com" has 1 instance(s).
  Instance "sales", status READY, has 2 handler(s) for this service...
  Handler(s):
    "DEDICATED" established:11 refused:0 state:ready
      LOCAL SERVER
    "D000" established:565 refused:4 current:155 max:10000 state:ready
      DISPATCHER <machine: helios, pid: 5673>
        (ADDRESS=(PROTOCOL=tcp)(HOST=helios)(PORT=38411))
The command completed successfully.
```

Under normal conditions, the number refused should be zero. Shut down the listener and restart it to erase these statistics. If the refused count is increasing after the listener restarts, then the connections are being refused. If the refused count stays at zero, and if the problem you are troubleshooting is occurring, then your problem is not with the connections being refused.

**Checking the Connect/Second Rate** Connection refusals can occur for many reasons. Examine the listener log to see what the connect rate is. Run the listener log analyzer script to check.

The listener is a queue-based process. It receives connect requests from the lower level protocol stack. It has a limited queue stack which is configurable to the operating system maximum. It can only process one connection at a time, and there is a limit to the number of connections a second the process can handle.

If the rate at which the connect requests arrive exceeds that limit, then the requests are queued. The queue stack is also limited, but you can configure it. If there are more listener processes, then the requests made against each individual process are fewer and are handled more quickly.

Increasing the listener queue is done in the `listener.ora` file. The `listener.ora` file can contain many listeners, each by a different name. It is assumed that only one of those listed is having a problem. If not, then apply this method to all applicable listeners. To increase the listener queue, add `(queuesize = number)` to the `listener.ora` file. For example:

```
listener =
  (address =
    (protocol = tcp)
    (host = sales-pc)
    (port = 1521)
    (queuesize = 20)
  )
```

**See Also:** *Oracle9i Net Services Administrator's Guide*

Stop and restart the listener to initialize this new parameter. If you are not currently running a shared server configuration, then consider doing so. It is faster for the listener to handle a client request in a shared server configuration than it is in a dedicated server configuration.

---

---

**Note:** Shared server dispatchers also receive connect requests and can also benefit from tuning the queue size.

The maximum queue size is subject to the maximum size possible for a particular operating system.

---

---

## Detecting Network Problems

This section encompasses local area network (LAN) and wide area network (WAN) troubleshooting methods.

### Using Dynamic Performance Views for Network Performance

Networks entail overhead that adds a certain amount of delay to processing. To optimize performance, you must ensure that your network throughput is fast, and you should try to reduce the number of messages that must be sent over the network. It can be difficult to measure the delay the network adds.

Three dynamic performance views are useful for measuring the network delay:

- V\$SESSION\_EVENT
- V\$SESSION\_WAIT
- V\$SESSTAT

In V\$SESSION\_EVENT, the AVERAGE\_WAIT column indicates the amount of time that Oracle waits between messages. You can use this statistic as a yardstick to evaluate the effectiveness of the network.

In V\$SESSION\_WAIT, the EVENT column lists the events for which active sessions are waiting. The "sqlnet message from client" wait event indicates that the shared or foreground process is waiting for a message from a client. If this wait event has occurred, then you can check to see whether the message has been sent by the user or received by Oracle.

You can investigate hang-ups by looking at V\$SESSION\_WAIT to see what the sessions are waiting for. If a client has sent a message, then you can determine whether Oracle is responding to it or is still waiting for it.

In V\$SESSTAT you can see the number of bytes that have been received from the client, the number of bytes sent to the client, and the number of calls the client has made.

### Understanding Latency and Bandwidth

The most critical aspects of a network that contribute to performance are latency and bandwidth.

- Latency refers to a time delay; for example, the gap between the time a device requests access to a network and the time it receives permission to transmit.

- Bandwidth is the throughput capacity of a network medium or protocol. Variations in the network signals can cause degradation on the network. Sources of degradation can be cables that are too long or wrong cable type. External noise sources, such as elevators, air handlers, or florescent lights, can also cause problems.

## Common Network Topologies

### Local Area Network Topologies:

- Ethernet
- Fast Ethernet
- 1 Gigabit Ethernet
- Token Ring
- FDDI
- ATM

### Wide Area Network Topologies:

- DSL
- ISDN
- Frame Relay
- T-1, T-3, E-1, E-3
- ATM
- SONAT

Table 23–1 lists the most common ratings for various topologies.

**Table 23–1 Bandwidth Ratings**

| Topology or Carrier | Bandwidth           |
|---------------------|---------------------|
| Ethernet            | 10 Megabits/second  |
| Fast Ethernet       | 100 Megabits/second |
| 1 Gigabit Ethernet  | 1 Gigabits/second   |
| Token Ring          | 16 Megabits/second  |
| FDDI                | 100 Megabits/second |

**Table 23–1 (Cont.) Bandwidth Ratings**

| <b>Topology or Carrier</b> | <b>Bandwidth</b>                                                                            |
|----------------------------|---------------------------------------------------------------------------------------------|
| ATM                        | 155 Megabits/second (OC3), 622 Megabits/second (OC12)                                       |
| T-1 (US only)              | 1.544 Megabits/second                                                                       |
| T-3 (US only)              | 44.736 Megabits/second                                                                      |
| E-1 (non-US)               | 2.048 Megabits/second                                                                       |
| E-3 (non-US)               | 34.368 Megabits/second                                                                      |
| Frame Relay                | Committed Information Rate, which can be up to the carrier speed, but usually is not.       |
| DSL                        | This can be up to the carrier speed.                                                        |
| ISDN                       | This can be up to the carrier speed. Usually, it is used with slower modems.                |
| Dial Up Modems             | 56 Kilobits/second. Usually, it is accompanied with data compression for faster throughput. |

## Solving Network Problems

This section describes several techniques for enhancing performance and solving network problems.

- [Finding Network Bottlenecks](#)
- [Dissecting Network Bottlenecks](#)
- [Using Array Interfaces](#)
- [Adjusting Session Data Unit Buffer Size](#)
- [Using TCP.NODELAY](#)
- [Using Connection Manager](#)

**See Also:** *Oracle9i Net Services Administrator's Guide*

### Finding Network Bottlenecks

The first step in solving network problem is to understand the overall topology. Gather as much information about the network that you can. This kind of information usually manifests itself as a network diagram. Your diagram should contain the types of network technology used in the Local Area Network and the

Wide Area Network. It should also contain addresses of the various network segments.

Examine this information. Obvious network bottlenecks include the following:

- Using a dial-up modem (normal modem or ISDN) to access time critical data.
- A frame relay link is running on a T-1, but has a 9.6 Kilobits CIR so that it only reliably transmits up to 9.6 Kilobit's a second and if the rest of the bandwidth is used, then there is a possibly that the data will be lost.
- Data from high speed networks channels through low speed networks.
- There are too many network hops. A router constitutes one hop.
- A 10 Megabit network for a Web site.

There are many problems that can cause a performance breakdown. Follow this checklist:

- Get a network sniffer trace.
- Check the following:
  - Is the bandwidth being exceeded on the network, the client, or the server?
  - Ethernet collisions.
  - Token ring or FDDI ring beacons.
  - Are there many runt frames?
  - The stability of the WAN links.
- Get a bandwidth utilization chart for frame relay, and see if CIR is being exceeded.
- Is any quality of service or packet prioritizing going on?
- Is a firewall in the way somewhere?

If nothing is revealed, then find the network route from the client to the data server. Understanding the travel times on a network gives you an idea as to the time a transaction will take. Client-server communication requires many small packets. High latency on a network slows the transaction down due to the time interval between sending a request and getting the response.

Use trace route (`tracroute` or equivalent) from the client to the server to get address information for each device in the path. For example:

```
tracert usmail05
```

```
Tracing route to usmail05.us.oracle.com [144.25.88.200] over a maximum of 30 hops:
```

```
  1  <10 ms  <10 ms  10 ms  whq1davis-rtr-749-f1-0-a.us.oracle.com
[144.25.216.1]
  2  <10 ms  <10 ms  <10 ms  whq4op3-rtr-723-f0-0.us.oracle.com
[144.25.252.23]
  3  220 ms  210 ms  231 ms  usmail05.us.oracle.com [144.25.88.200]
```

Trace complete.

**Ping each device in turn to get the timings. Use large packets to get the slowest times. Make sure you set the "don't fragment bit" so that routers do not spend time disassembling and reassembling the packet. Also note that the packet size is 1472. This is for Ethernet. Ethernet packets are 1536 octets (actual 8 bit bytes) in size. ICPM packets (this is what ping is designed to use) have 64 octets of header. Evaluate the area where the slowness seems to occur. For example:**

```
ping -l 1472 -n 1 -f 144.25.216.1
Pinging 144.25.216.1 with 1472 bytes of data:
Reply from 144.25.216.1: bytes=1472 time<10ms TTL=255
```

```
ping -l 1472 -n 1 -f 144.25.252.23
Pinging 144.25.252.23 with 1472 bytes of data:
Reply from 144.25.252.23: bytes=1472 time=10ms TTL=254
```

```
ping -l 1472 -n 1 -f 144.25.88.200
Pinging 144.25.88.200 with 1472 bytes of data:
Reply from 144.25.88.200: bytes=1472 time=271ms TTL=253
```

The previous example validates trace route. Ideally, you ping from the workstation to 144.25.216.1, from 144.25.216.1 to 144.25.252.23, then from 144.25.252.23 to 144.25.88.200. This would show the exact latency on each segment traveled.

## Dissecting Network Bottlenecks

This section helps you determine the problem with your network bottleneck.

### Determining if the Problem is with Oracle Net or the Network

Oracle Net tracing reveals whether an error is Oracle-specific or due to conditions that the operating system is passing to the Transparent Network Substrate (Oracle TNS layer).

Enable Oracle Net tracing at the Oracle server, the listener, and at a client suspected of having the problem you are trying to resolve.



To enable tracing at the server, find the `sqlnet.ora` file for the server and create the following lines in it:

```
TRACE_TIMESTAMP_SERVER = ON
TRACE_LEVEL_SERVER = 16
TRACE_UNIQUE_SERVER = ON
```

To enable tracing at the client, find the `sqlnet.ora` file for the client and create the following lines in it:

```
TRACE_TIMESTAMP_CLIENT = ON
TRACE_LEVEL_CLIENT = 16
TRACE_UNIQUE_CLIENT = ON
```

To enable tracing at the listener, find the `listener.ora` file and create the following line in it:

```
TRACE_TIMESTAMP_listener_name = ON
TRACE_LEVEL_listener_name = 16
```

---

---

**Note:** The `TRACE_TIMESTAMP_x` parameters are optional, but they should be included for better debugging

---

---

Reproduce the problem, so that you generate traces on the client and server. Now analyze the traces generated.

**See Also:**

- *Oracle9i Net Services Administrator's Guide* for detailed directions on enabling Oracle Net tracing
- *Oracle9i Database Error Messages* for definitions to Oracle Net errors noted in the trace file

If the problem is with the network and not Oracle Net, then you must determine the following:

- Does the problem only occur in one location on the local network?
- Does the problem only occur in one area on the WAN?

For example, perhaps the system is fine in the building where the Data Center is located, but it is slow in other buildings that are several miles away.

Not all Oracle error codes represent pure Oracle troubles. ORA-3113 is the most common error that points to an underlying network problem.

---

---

**Note:** Enabling tracing on the server can generate a large amount of trace files. To prevent this, set up a separate environment that traces itself. This configuration works for dedicated connections. First, log in to the server's operating system as the Oracle software owner. Create a temporary directory to keep configuration files and trace files that will be created. Copy the `sqlnet.ora`, `listener.ora`, and `tnsnames.ora` to that directory. Edit the `sqlnet.ora` file to enable tracing. Add to the `sqlnet.ora` file the following line:

```
TRACE_DIRECTORY_SERVER = temporary_directory_just_created
```

Now, modify the `listener.ora` file and change the listening port (for TCP, other protocols, use a similar technique) to an unused port. You need to make a similar modification to the client's `tnsnames.ora` file for the connect string you will be using for this test.

Set the `TNS_ADMIN` environment to point to the temporary directory. Start the listener. Now all new connections to the new listener send Server traces to this directory. Reproduce the problem.

---

---

If you are getting an Oracle error message, then look into the trace file to find the error. For troubleshooting bugs, Oracle Net trace analysis takes some time to fully find the problem. However, high-level simple trace analysis is rather simple.

### **Determining if the Problem is on the Client or the Server (on Oracle Net)**

If the problem is with Oracle Net, then use Oracle Net tracing to show you where the problem lies. If there are errors in the trace files, then do they appear in only the client traces, only in the server traces, or in both?

#### **Errors Only in the Client Trace**

The problem is on the client. However, if you are getting ORA-3113 or ORA-3114 errors, then the problem is on the server.

#### **Errors Only in the Server Trace or Listener Trace**

The problem is on the server. However, if you are getting ORA-3113 or ORA-3114 errors, then the problem is on the client.

### Errors in All: Client, Server, and Listener Trace

If you are getting ORA-3113 or ORA-3114 errors, then the problem is on the Network. Troubleshoot the server first. If it is fine, then the client is at fault.

### Checking if the Server is Configured for Shared Servers

The shared server architecture can be more complex to troubleshoot. Check the initialization parameter file for any shared server parameters. Look at the operating system to see if any of the shared server processes are present.

Check for dispatchers by looking for names such as ora\_d000, ora\_d001, and so on. For example:

```
ps -ef | grep ora_d
```

Check for shared servers by looking for names such as ora\_s000, ora\_s001, and so on. For example:

```
ps -ef | grep ora_s
```

#### See Also:

- ["Shared Server Configuration"](#) on page 23-2 for more information on tuning the shared server
- *Oracle9i Database Concepts* and *Oracle9i Net Services Administrator's Guide* for more information on shared server concepts and parameters

## Using Array Interfaces

Reduce network calls by using array interfaces. Instead of fetching one row at a time, it is more efficient to fetch 10 rows with a single network round trip.

**See Also:** *Oracle Call Interface Programmer's Guide* for more information on array interfaces

## Adjusting Session Data Unit Buffer Size

Before sending data across the network, Oracle Net buffers data into the Session Data Unit (SDU). It sends the data stored in this buffer when the buffer is full or when an application tries to read the data. When large amounts of data are being

retrieved and when packet size is consistently the same, it might speed retrieval to adjust the default SDU size.

Optimal SDU size depends on the normal transport size. Use a sniffer to find out the frame size, or set tracing on to its highest level to check the number of packets sent and received and to determine whether they are fragmented. Tune your system to limit the amount of fragmentation.

Use Oracle Net Configuration Assistant to configure a change to the default SDU size on both the client and the server; SDU size is generally the same on both.

**See Also:** *Oracle9i Net Services Administrator's Guide*

## Using TCP.NODELAY

When a session is established, Oracle Net packages and sends data between server and client using packets. The `TCP.NODELAY` parameter, which causes packets to be flushed on to the network more frequently, is enabled by default. Although Oracle Net supports many networking protocols, TCP tends to have the best scalability.

**See Also:** Your platform-specific Oracle documentation for more information on `TCP.NODELAY`

## Using Connection Manager

In Oracle Net, you can use the Connection Manager to conserve system resources by multiplexing. *Multiplexing* means funneling many client sessions through a single transport connection to a server destination. This way, you can increase the number of sessions that a process can handle. This applies only to shared server configurations. Alternately, you can use Connection Manager to control client access to dedicated servers. Connection Manager provides multiple protocol support allowing a client and server with different networking protocols to communicate.

**See Also:** *Oracle9i Net Services Administrator's Guide* for more information on Connection Manager

# Part VI

---

## Performance-Related Reference Information

Part VI provides reference information regarding dynamic performance views and wait events.

The chapters in this part are:

- [Chapter 24, "Dynamic Performance Views for Tuning"](#)



---

## Dynamic Performance Views for Tuning

This chapter provides detailed information on some of the dynamic views that can help you tune your system and investigate performance problems.

The topics discussed in this chapter are:

- [Dynamic Performance Tables](#)
- [Description of Dynamic Performance Views](#)

**See Also:** *Oracle9i Database Reference* for a complete list of the dynamic performance views and their columns

## Dynamic Performance Tables

Throughout its operation, Oracle maintains a set of virtual tables that record current database activity. These tables are created by Oracle and are called dynamic performance tables.

Database administrators can query and create views on the tables and grant access to those views to other users. These views are called fixed views because they cannot be altered or removed by the database administrator.

`SYS` owns the dynamic performance tables. By default, they are available only to the user `SYS` and to users granted `SELECT ANY TABLE` system privilege, such as `SYSTEM`. Their names all begin with `V_`. Views are created on these tables, and then public synonyms are created for the views. The synonym names begin with `V`.

Each view belongs to one of the following categories:

- [Current State Views](#)
- [Counter/Accumulator Views](#)
- [Information Views](#)

### Current State Views

The views listed in [Table 24–1](#) give a picture of what is currently happening on the system.

**Table 24–1** *Current State Views*

| Fixed View                      | Description                                            |
|---------------------------------|--------------------------------------------------------|
| <a href="#">V\$LOCK</a>         | Locks currently held/requested on the instance         |
| <a href="#">V\$LATCHHOLDER</a>  | Sessions/processes holding a latch                     |
| <a href="#">V\$OPEN_CURSOR</a>  | Cursors opened by sessions on the instance             |
| <a href="#">V\$SESSION</a>      | Sessions currently connected to the instance           |
| <a href="#">V\$SESSION_WAIT</a> | Different resources sessions are currently waiting for |

### Counter/Accumulator Views

These views keep track of how many times some activity has occurred since instance/session startup. Select from the view directly to see activity since startup.



If you are interested in activity happening in a given time interval, then take a snapshot before and after the time interval, and the delta between the two snapshots provides the activity during that time interval. This is similar to how operating system utilities like `sar`, `vmstat`, and `iostat` work. Tools provided by Oracle, like `Statspack` and `BSTAT/ESTAT`, do this delta to provide a report of activity in a given interval.

---



---

**Note:** Snapshots should be taken during steady-state, not immediately after system startup. Extra overhead is incurred during system ramp-up, which may not accurately reflect the performance of the system at steady-state.

---



---

**Table 24–2 Summary Since Session Startup**

| Fixed View                           | Description                                                                          |
|--------------------------------------|--------------------------------------------------------------------------------------|
| <code>V\$DB_OBJECT_CACHE</code>      | Object level statistics in shared pool                                               |
| <code>V\$FILESTAT</code>             | File level summary of the I/O activity                                               |
| <code>V\$LATCH</code>                | Latch activity summary                                                               |
| <code>V\$LATCH_CHILDREN</code>       | Latch activity for child latches                                                     |
| <code>V\$LIBRARYCACHE</code>         | Namespace level summary for shared pool                                              |
| <code>V\$LIBRARY_CACHE_MEMORY</code> | Summary of the current memory use of the library cache, by library cache object type |
| <code>V\$MYSTAT</code>               | Resource usage summary for your own session                                          |
| <code>V\$ROLLSTAT</code>             | Rollback segment activity summary                                                    |
| <code>V\$ROWCACHE</code>             | Data dictionary activity summary                                                     |
| <code>V\$SEGMENT_STATISTICS</code>   | User-friendly DBA view for real-time monitoring of segment-level statistics          |
| <code>V\$SEGSTAT</code>              | High-efficiency view for real-time monitoring of segment-level statistics            |
| <code>V\$SESSION_EVENT</code>        | Session-level summary of all the waits for current sessions                          |
| <code>V\$SESSTAT</code>              | Session-level summary of resource usage since session startup                        |
| <code>V\$LIBRARY_CACHE_MEMORY</code> | Simulation of the shared pool's LRU list mechanism                                   |

**Table 24–2 (Cont.) Summary Since Session Startup**

| Fixed View      | Description                                                        |
|-----------------|--------------------------------------------------------------------|
| V\$SQL          | Child cursor details for V\$SQLAREA                                |
| V\$SQLAREA      | Shared pool details for statements/anonymous blocks                |
| V\$SYSSTAT      | Summary of resource usage                                          |
| V\$SYSTEM_EVENT | Instance wide summary of resources waited for                      |
| V\$UNDOSTAT     | Histogram of undo usage. Each row represents a 10-minute interval. |
| V\$WAITSTAT     | Break down of buffer waits by block class                          |

## Information Views

In information views, the information is not as dynamic as in the current state view. Hence, it does not need to be queried as often as the current state views.

**Table 24–3 Information Views**

| Fixed View                           | Description                                                                                                       |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| V\$MTR_TARGET_ADVICE                 | Advisory information collected by MTR advisory, when FAST_START_MTR_TARGET is set                                 |
| V\$PARAMETER and V\$SYSTEM_PARAMETER | Parameters values for your session<br>Instance wide parameter values                                              |
| V\$PROCESS                           | Server processes (background and foreground)                                                                      |
| V\$SEGSTAT_NAME                      | Statistics property view for segment-level statistics                                                             |
| V\$SQL_PLAN                          | Execution plan for cursors that were recently executed                                                            |
| V\$SQL_PLAN_STATISTICS               | Execution statistics of each operation in the execution plan                                                      |
| V\$SQL_PLAN_STATISTICS_ALL           | Concatenates information in V\$SQL_PLAN with execution statistics from V\$SQL_PLAN_STATISTICS and V\$SQL_WORKAREA |
| V\$SQLTEXT                           | SQL text of statements in the shared pool                                                                         |
| V\$STATISTICS_LEVEL                  | Status of the statistics or advisories controlled by the STATISTICS_LEVEL initialization parameter                |

## Description of Dynamic Performance Views

This section discusses the details of some of the dynamic performance views.

### V\$DB\_OBJECT\_CACHE

This view provides object level statistics for objects in the library cache (shared pool). This view provides more details than V\$LIBRARYCACHE and is useful for finding active objects in the shared pool.

#### Useful Columns for V\$DB\_OBJECT\_CACHE

Most of the columns of this table provide current state information.

- OWNER: Object owner
- NAME: Object name (First 1000 characters of SQL text for anonymous blocks/cursors)
- TYPE: Type of object (for example, sequence, procedure, function, package, package body, trigger)
- KEPT: Tells if the object is pinned in the shared pool (yes, no)
- SHARABLE\_MEM: Amount of sharable memory used
- PINS: Sessions currently executing this object
- LOCKS: Sessions currently locking this object

#### Instantaneous State Columns

The following columns keep statistics on the object since its first load:

- LOADS: Number of times this object had to be loaded
- INVALIDATIONS: Number of times this object was invalidated

#### *Example 24-1 Summary of Shared Pool Executions and Memory Usage*

The following query shows the distribution of shared pool memory across different type of objects. It also shows if any of the objects have been pinned in the shared pool using the procedure DBMS\_SHARED\_POOL.KEEP().

```
SELECT type, kept, COUNT(*), SUM(sharable_mem)
FROM V$DB_OBJECT_CACHE
GROUP BY type, kept;
```

**Example 24–2 Finding Objects with Large Number of Loads**

```
SELECT owner, name sharable_mem, kept, loads
FROM V$DB_OBJECT_CACHE
WHERE loads > 1
OR invalidations > 0
ORDER BY loads DESC;
```

**Example 24–3 Finding Large Unpinned Objects**

The following query finds all objects using large amounts of memory. They can be pinned using `DBMS_SHARED_POOL.KEEP()`.

```
SELECT owner, name, sharable_mem, kept
FROM V$DB_OBJECT_CACHE
WHERE sharable_mem > 102400
AND kept = 'NO'
ORDER BY sharable_mem DESC;
```

## V\$FILESTAT

This view keeps information on physical I/O activity for each file. This is useful in isolating where the I/O activity is happening if the bottleneck is I/O related. `V$FILESTAT` shows the following information for database I/O (but not for log file I/O):

- Number of physical reads and writes
- Number of blocks read and written
- Total I/O time for reads and writes

The numbers reflect activity since the instance startup. If two snapshots are taken, then the differences in the statistics provides the I/O activity for the time interval.

**Useful Columns for V\$FILESTAT**

- `FILE#`: Number of the file
- `PHYRDS`: Number of physical reads done
- `PHYBLKRD`: Number of physical blocks read
- `PHYWRTS`: Number of physical writes done
- `PHYBLKWRT`: Number of physical blocks written

## Notes on V\$FILESTAT

- Physical reads and blocks read can be different because of multiblock read calls.
- Physical writes and blocks written can differ because of direct writes by processes.
- Sum(Physical blocks read) should correlate closely with `physical reads` from `V$SYSSTAT`.
- Sum(Physical blocks written) should correlate closely with `physical writes` from `V$SYSSTAT`.
- Reads (into buffer cache as well as direct reads) are done by server processes. Writes from buffer cache are handled only by the DBWR. The direct writes are handled by the server processes.

## Join Columns for V\$FILESTAT

Table 24–4 lists the join columns for `V$FILESTAT`.

**Table 24–4 Join Columns for V\$FILESTAT**

| Column | View           | Joined Column(s) |
|--------|----------------|------------------|
| FILE#  | DBA_DATA_FILES | FILE_ID          |

### Example 24–4 Checking Oracle Datafile I/O

The following query monitors the values of physical reads and physical writes over some period of time while your application is running:

```
SELECT NAME, PHYRDS, PHYWRTS
FROM V$DATAFILE df, V$FILESTAT fs
WHERE df.FILE# = fs.FILE#;
```

The preceding query also retrieves the name of each datafile from the dynamic performance view `V$DATAFILE`. Sample output might look like the following:

```
NAME                                PHYRDS    PHYWRTS
-----
/oracle/ora81/dbs/ora_system.dbf      7679      2735
/oracle/ora81/dbs/ora_temp.dbf        32        546
```

The `PHYRDS` and `PHYWRTS` columns of `V$FILESTAT` can also be obtained through SNMP.

The total I/O for a single disk is the sum of PHYRDS and PHYWRTS for all the database files managed by the Oracle instance on that disk. Determine this value for each of your disks. Also, determine the rate at which I/O occurs for each disk by dividing the total I/O by the interval of time over which the statistics were collected.

---



---

**Note:** Although Oracle records read and write times accurately, a database that is running on Unix file system (UFS) might not reflect true disk accesses. For example, the read times might not reflect a true disk read, but rather a UFS cache hit. However, read and write times should be accurate for raw devices. Additionally, write times are only recorded for each batch, with all blocks in the same batch given the same time after the completion of the write I/O.

---



---

**Example 24-5 Finding the Files with Large Numbers of Multiblock Reads**

The following example is useful for finding tablespaces that might be getting hit by large number of scans.

```
SELECT t.tablespace_name
      ,SUM(a.phyrds-b.phyrds)
      /MAX(86400*(a.snap_date-b.snap_date)) "Rd/sec"
      ,SUM(a.phyblkrd-b.phyblkrd)
      /greatest(SUM(a.phyrds-b.phyrds),1) "Blk/rd"
      ,SUM(a.phywrts-b.phywrts)
      /MAX(86400*(a.snap_date-b.snap_date)) "Wr/sec"
      ,SUM(a.phyblkwrt-b.phyblkwrt)
      /greatest(SUM(a.phywrts-b.phywrts),1) "Blk/wr"
FROM snap_filestat a, snap_filestat b, dba_data_files t
WHERE a.file# = b.file#
      AND a.snap_id = b.snap_id + 1
      AND t.file_id = a.file#
GROUP BY t.tablespace_name
HAVING sum(a.phyblkrd-b.phyblkrd)
      /greatest(SUM(a.phyrds-b.phyrds),1) > 1.1
      OR SUM(a.phyblkwrt-b.phyblkwrt)
      /greatest(SUM(a.phywrts-b.phywrts),1) > 1.1
ORDER BY 3 DESC, 5 DESC;
```

| TABLESPACE_N | Rd/sec | Blk/rd | Wr/sec | Blk/wr |
|--------------|--------|--------|--------|--------|
| TEMP         | 2.3    | 19.7   | 1.9    | 24.7   |
| AP_T_02      | 287.1  | 7.8    | .0     | 1.0    |

|              |       |     |     |     |
|--------------|-------|-----|-----|-----|
| AP_T_01      | 12.9  | 4.0 | .2  | 1.0 |
| APPLSYS_T_01 | 63.3  | 2.2 | .4  | 1.0 |
| PO_T_01      | 313.5 | 2.1 | .2  | 1.0 |
| RECEIVABLE_T | 401.0 | 1.5 | 2.4 | 1.0 |
| SHARED_T_01  | 9.2   | 1.3 | .4  | 1.0 |
| SYSTEM       | 45.2  | 1.3 | .3  | 1.0 |
| PER_T_01     | 48.0  | 1.2 | .0  | .0  |
| DBA_T_01     | .2    | 1.0 | .4  | 1.4 |

You can see that most of the multiblock reads and writes are going to TEMP tablespace, due to large sorts going to disk. Other tablespaces are getting multiblock reads due to full table scans.

**See Also:** [Chapter 20, "Oracle Tools to Gather Database Statistics"](#) for an example of how to gather file I/O data.

## V\$LATCH

This view keeps a summary of statistics for each type of latch since instance startup. It is useful for identifying the area within SGA experiencing problems when latch contention is observed in V\$SESSION\_WAIT.

### Useful Columns for V\$LATCH

- NAME: Latch name
- IMMEDIATE\_GETS: Requests for the latch in immediate mode
- IMMEDIATE\_MISSES: IMMEDIATE\_GETS that failed
- GETS: Requests for the latch in a willing to wait mode
- MISSES: GETS that did not obtain the latch on first try
- SPIN\_GETS: GETS that got the latch within SPIN\_GET tries and did not have to sleep
- SLEEP1-SLEEP3: GETS that succeeded only after sleeping one to three times
- SLEEP4: GETS that only succeeded after sleeping four or more times
- WAIT\_TIME: Elapsed time spent waiting for this latch
- MISSES\_WAITLIST: Number of latch misses on the associated wait list latch
- SLEEPS\_WAITLIST: Number of sleeps while trying to acquire the associated wait list latch

- `WAIT_TIME_WAITLIST`: Wait time for the associated wait list latch

## Join Columns for V\$LATCH

Table 24–5 lists the join columns for V\$LATCH.

**Table 24–5 Join Columns for V\$LATCH**

| Column | View                                                | Joined Column(s) |
|--------|-----------------------------------------------------|------------------|
| NAME   | V\$LATCH_CHILDREN<br>V\$LATCHHOLDER<br>V\$LATCHNAME | NAME             |
| NAME   | V\$LATCH_MISSES                                     | PARENT_NAME      |
| LATCH# | V\$LATCH_CHILDREN<br>V\$LATCHNAME                   | LATCH#           |

### Example 24–6 Querying V\$LATCH

In the following example, a table is created to hold data queried from V\$LATCH:

```
CREATE TABLE snap_latch as
SELECT 0 snap_id, sysdate snap_date, a.*
  FROM V$LATCH a;
ALTER TABLE snap_latch add
  (constraint snap_filestat primary key (snap_id, name));
```

Initially, the `snap_id` has been set to 0. After some interval of time, the `snap_latch` table is updated with the `snap_id` set to 1:

```
INSERT INTO snap_latch
SELECT 1, sysdate, a.*
  FROM V$LATCH a;
```

Note that you must increment the `snap_id` each time you use the previous SQL statement to insert records.

After you inserted records for consecutive intervals, use the following SELECT statement to displays statistics. Note that zero is substituted when there is an attempt to divide by zero.

```
SELECT SUBSTR(a.name,1,20) NAME, (a.gets-b.gets)/1000 "Gets(K)",
  (a.gets-b.gets)/(86400*(a.snap_date-b.snap_date)) "Get/s",
  DECODE ((a.gets-b.gets), 0, 0, (100*(a.misses-b.misses)/(a.gets-b.gets))) MISS,
  DECODE ((a.misses-b.misses), 0, 0,
```



```

        (100*(a.spin_gets-b.spin_gets)/(a.misses-b.misses))) SPIN,
    (a.immediate_gets-b.immediate_gets)/1000 "Iget(K)",
    (a.immediate_gets-b.immediate_gets)/ (86400*(a.snap_date-b.snap_date)) "IGet/s",
    DECODE ((a.immediate_gets-b.immediate_gets), 0, 0,
    (100*(a.immediate_misses-b.immediate_misses)/ (a.immediate_gets-b.immediate_gets))) IMISS
FROM snap_latch a, snap_latch b
WHERE a.name = b.name
    AND a.snap_id = b.snap_id + 1
    AND ( (a.misses-b.misses) > 0.001*(a.gets-b.gets)
        or (a.immediate_misses-b.immediate_misses) >
        0.001*(a.immediate_gets-b.immediate_gets))
ORDER BY 2 DESC;

```

Before running the previous SQL statement, you may want to specify various display formatting setups, such as:

```

SET LIN 120
SET PAGES 60
SET NUMFORMAT 999999.9

```

#### **Example 24–7 Sample Latch Statistics**

The following example output shows the latch statistics obtained by doing a delta over a period of one hour as was done with the V\$FILESTAT numbers. Those latches that had misses less than 0.1% of the gets have been filtered out.

| NAME               | Gets(K) | Get/s  | MISS | SPIN  | IGets(K) | IGet/s | IMISS |
|--------------------|---------|--------|------|-------|----------|--------|-------|
| cache buffers chai | 255,272 | 69,938 | 0.4  | 99.9  | 3,902    | 1,069  | 0.0   |
| library cache      | 229,405 | 62,851 | 9.1  | 96.9  | 51,653   | 14,151 | 3.7   |
| shared pool        | 24,206  | 6,632  | 14.1 | 72.1  | 0        | 0      | 0.0   |
| latch wait list    | 1,828   | 501    | 0.4  | 99.9  | 1,836    | 503    | 0.5   |
| row cache objects  | 1,703   | 467    | 0.7  | 98.9  | 1,509    | 413    | 0.2   |
| redo allocation    | 984     | 270    | 0.2  | 99.7  | 0        | 0      | 0.0   |
| messages           | 116     | 32     | 0.2  | 100.0 | 0        | 0      | 0.0   |
| cache buffers lru  | 91      | 25     | 0.3  | 99.0  | 7,214    | 1,976  | 0.3   |
| modify parameter v | 2       | 0      | 0.1  | 100.0 | 0        | 0      | 0.0   |
| redo copy          | 0       | 0      | 92.3 | 99.3  | 1,460    | 400    | 0.0   |

When examining latch statistics, look at the following:

- What is the ratio of misses/gets?
- What percentage of misses are obtained by just spinning?
- How many times was the latch requested?

- How many sleeps were on the latch?

There seems to be a lot of contention for the redo copy latch with a 92.3 percent miss rate. But, look carefully. Redo copy latches are obtained mostly in immediate mode. The numbers for immediate gets look fine, and the immediate gets are several orders of magnitude bigger than the willing to wait gets. So, there is no contention for redo copy latches.

However, there does seem to be contention for the shared pool and library cache latches. Consider running a query that checks the sleeps for these latches to see if there is actually a problem, such as the following output:

| NAME                     | Gets(K) | Get/s  | MISS | SPIN  | SL01 | SL02 | SL03 | SL04 |
|--------------------------|---------|--------|------|-------|------|------|------|------|
| cache buffers chain      | 255,272 | 69,938 | 0.4  | 99.9  | 0.1  | 0.0  | 0.0  | 0.0  |
| library cache            | 229,405 | 62,851 | 9.1  | 96.9  | 3.0  | 0.1  | 0.0  | 0.0  |
| shared pool              | 24,206  | 6,632  | 14.1 | 72.1  | 22.4 | 4.8  | 0.8  | 0.0  |
| latch wait list          | 1,828   | 501    | 0.4  | 99.9  | 0.1  | 0.0  | 0.0  | 0.0  |
| row cache objects        | 1,703   | 467    | 0.7  | 98.9  | 0.6  | 0.0  | 0.4  | 0.0  |
| redo allocation messages | 984     | 270    | 0.2  | 99.7  | 0.1  | 0.0  | 0.2  | 0.0  |
| cache buffers lru        | 116     | 32     | 0.2  | 100.0 | 0.0  | 0.0  | 0.0  | 0.0  |
| modify parameter v       | 91      | 25     | 0.3  | 99.0  | 1.0  | 0.0  | 0.0  | 0.0  |
| redo copy                | 2       | 0      | 0.1  | 100.0 | 0.0  | 0.0  | 0.0  | 0.0  |
|                          | 0       | 0      | 92.3 | 99.3  | 0.0  | 0.7  | 0.0  | 0.0  |

You can see that there is a 14% miss rate on the shared pool latches. 72% of the missed latched without relinquishing the CPU (having to sleep even once) by spinning. There are some misses for which you have to sleep multiple times.

Investigate why the shared pool latch is needed so many times. Look at the SQL being run by sessions holding or waiting for the latch, as well as the resource usage characteristics of the system. Compare them with baselines when there was no problem.

### Tuning Latches

Do not tune latches. If you see latch contention, then it is a symptom of a part of SGA experiencing abnormal resource usage. Latches control access with certain assumptions (for example, a cursor is parsed once and executed many times). To fix the problem, examine the resource usage for the parts of SGA experiencing contention. Merely looking at `V$LATCH` does not address the problem.

**See Also:** *Oracle9i Database Concepts* for more information on latches

## V\$LATCH\_CHILDREN

There are multiple latches in the database for some type of latches. V\$LATCH provides aggregate summary for each type of latch. To look at individual latches, query the V\$LATCH\_CHILDREN view.

### *Example 24-8 Finding the Number of Multiple Latches on the System*

```
SELECT name, count(*)
   FROM v$latch_children
  ORDER BY count(*) desc;
```

| NAME                                 | COUNT(*) |
|--------------------------------------|----------|
| global tx hash mapping               | 2888     |
| global transaction                   | 2887     |
| cache buffers chains                 | 2048     |
| latch wait list                      | 32       |
| Token Manager                        | 23       |
| enqueue hash chains                  | 22       |
| session idle bit                     | 22       |
| redo copy                            | 22       |
| process queue reference              | 20       |
| Checkpoint queue latch               | 11       |
| library cache                        | 11       |
| msg queue latch                      | 11       |
| session queue latch                  | 11       |
| process queue                        | 11       |
| cache buffers lru chain              | 11       |
| done queue latch                     | 11       |
| channel operations parent latch      | 4        |
| session switching                    | 4        |
| message pool operations parent latch | 4        |
| ksfv messages                        | 2        |
| parallel query stats                 | 2        |
| channel handle pool latch            | 1        |
| temp table ageout allocation latch   | 1        |

## V\$LATCHHOLDER

This view is useful to see if the session holding the latch is changing. Most of the time, the latch is held for such a small time that it is impossible to join to some other table to see the SQL statement being executed or the events that latch holder is waiting for.

This latch is useful in finding sessions that might be holding latches for a significant amount of time.

### Join Columns for V\$LATCHHOLDER

[Table 24–6](#) lists the join columns for V\$LATCHHOLDER.

**Table 24–6 Join Columns for V\$LATCHHOLDER**

| Column | View                                            | Joined Column(s) |
|--------|-------------------------------------------------|------------------|
| LADDR  | V\$LATCH_CHILDREN                               | ADDR             |
| NAME   | V\$LATCH,<br>V\$LATCHNAME,<br>V\$LATCH_CHILDREN | NAME             |
| PID    | V\$PROCESS                                      | PID              |
| SID    | V\$SESSION                                      | SID              |

### Example 24–9 Finding the SQL Statement Executed by the Latch Holder

```
SELECT s.sql_hash_value, l.name
   FROM V$SESSION s, V$LATCHHOLDER l
  WHERE s.sid = l.sid;
```

```
SQL_HASH_VALUE NAME
-----
299369270 library cache
1052917712 library cache
3198762001 library cache
SQL> /
```

```
SQL_HASH_VALUE NAME
-----
749899113 cache buffers chains
1052917712 library cache
SQL> /
```

```
SQL_HASH_VALUE NAME
-----
1052917712 library cache
SQL> /
```

```
SQL_HASH_VALUE NAME
-----
```

```

749899113 library cache
1052917712 library cache

```

This example indicates that the SQL statement 1052917712 is using a lot of parsing resources. The next step is to find the resources used by the session and examine the statement.

## V\$LIBRARYCACHE

This view has a namespace level summary for the objects in library cache since instance startup. When experiencing performance issues related to the library cache, this view can help identify the following:

- Specific parts (namespace) of the library cache (shared pool)
- Possible causes of problems

Then use V\$DB\_OBJECT\_CACHE, V\$SQLAREA to get more details.

### Useful Columns for V\$LIBRARYCACHE

- **NAMESPACE:** Class of objects (SQL area, trigger, and so on)
- **GETS:** Handle requests for objects of this namespace
- **GETHITS:** Requests that found handle in the cache
- **PINS:** PIN requests for objects of this namespace
- **PINHITS:** Requests able to reuse an existing PIN
- **RELOADS:** Number of times objects stored in the library cache had to be reloaded into memory because part of the object had been flushed from the cache. If there are a significant number of reloads, then reusable information is being flushed from the library cache. This requires a reload/rebuild of the object before it can again be accessed.
- **INVALIDATIONS:** The number of times objects were invalidated. For example, an object is invalidated automatically by Oracle when it is no longer safe to execute. For example, if the optimizer statistics for a table were recomputed, then all SQL statements currently in the library cache at the time the recompute occurred would be invalidated, because their execution plans may no longer be optimal.

GETHITRATIO (GETHITS/GETS) and GETPINRATIO (PINHITS/PINS) can be used if just examining activity since instance startup. If examining activity over a

specified time interval, it is better to compute these from the differences in snapshots before and after the interval.

**Example 24–10 Querying V\$LIBRARYCACHE**

```
SELECT namespace, gets, 100*gethits/gets gethitratio,
       pins, 100* pinhits/pins getpinratio,
       reloads, invalidations
FROM V$LIBRARYCACHE
ORDER BY gets DESC
```

Look for the following when querying this view:

- High RELOADS or INVALIDATIONS
- Low GETHITRATIO or GETPINRATIO

High number of RELOADS could be due to the following:

- Objects being invalidated (large number of INVALIDATIONS)
- Objects getting swapped out of memory

Low GETHITRATIO could indicate that objects are getting swapped out of memory.

Low PINHITRATIO could indicate the following:

- Session not executing the same cursor multiple times (even though it might be shared across different sessions)
- Session not finding the cursor shared

The next step is to query V\$DB\_OBJECT\_CACHE/V\$SQLAREA to see if problems are limited to certain objects or spread across different objects. If invalidations are high, then it might be worth investigating which of the (invalidated object's) underlying objects are being changed.

## V\$LIBRARY\_CACHE\_MEMORY

This fixed view summarizes the current memory use of the library cache, by library cache object type. The view can be queried often, without increasing library cache latch contention. Column descriptions are listed in [Table 24–7](#).

**Table 24–7 V\$LIBRARY\_CACHE\_MEMORY Column Description**

| Column Name          | Datatype | Column Description                    |
|----------------------|----------|---------------------------------------|
| libcache_object_name | char(24) | Name of the library cache object type |

**Table 24–7 V\$LIBRARY\_CACHE\_MEMORY Column Description**

| Column Name              | Datatype | Column Description                                                                 |
|--------------------------|----------|------------------------------------------------------------------------------------|
| libcache_pinned_memory   | number   | Amount of library cache memory pinned for all library cache objects of this type   |
| libcache_pinned_count    | number   | Number of library cache objects of this type currently pinned                      |
| libcache_unpinned_memory | number   | Amount of library cache memory unpinned for all library cache objects of this type |
| libcache_unpinned_count  | number   | Number of library cache objects of this type currently unpinned                    |

## V\$LOCK

This view has a row for every lock held or requested on the system. You should examine this view if you find sessions waiting for the wait event `enqueue`. If you find sessions waiting for a lock, then the sequence of events could be the following:

1. Use `V$LOCK` to find the sessions holding the lock.
2. Use `V$SESSION` to find the SQL statements being executed by the sessions holding the lock and waiting for the lock.
3. Use `V$SESSION_WAIT` to find what the session holding the lock is blocked on.
4. Use `V$SESSION` to get more details about the program and user holding the lock.

### Useful Columns for V\$LOCK

- `SID`: Identifier of the session holding/requesting the lock
- `TYPE`: Type of lock
- `LMODE`: Lock mode in which the session holds the lock
- `REQUEST`: Lock mode in which the session requests the lock
- `ID1, ID2`: Lock resource identifiers

### Common Lock Types

Several common locks are described in this section.

### **TX: Row Transaction Lock**

- This lock is required in exclusive mode (mode 6) to change data.
- One lock is acquired for each active transaction. It is released when the transaction ends due to a commit or rollback.
- If a block containing the row(s) to be changed does not have any ITL (interested transaction list) entries left, then the session requests the lock in shared mode (mode 4). It is released when the session gets an ITL entry for the block.
- If any of the rows to be changed are locked by another session, then locking session's transaction lock is requested in exclusive mode. When the locking transaction ends, this request ends, and the rows are covered under the requesting session's existing TX lock.
- The lock points to the rollback segment and transaction table entries for the transaction.

Do the following to avoid contention on this enqueue:

- To avoid contention on TX-6 enqueues, review the application.
- To avoid contention on TX-4 enqueues, consider increasing `INITTRANS` for the object.

### **TM: DML Lock**

- This lock is required in exclusive mode for executing any DDL statements on a database object; for example, lock table in exclusive mode, alter table, drop table.
- This lock is also acquired in shared mode when executing DML statements such as `INSERT`, `UPDATE`, or `DELETE`. This prevents other sessions from executing a DDL statement on the same object concurrently.
- For every object whose data is being changed, a TM lock is required.
- The lock points to the object.

To avoid contention on TM enqueues, consider disabling the table lock for the object. Disabling the table lock prevents any DDL from executing on the object.

### **ST - Space Transaction Lock**

- There is only one lock for each database (not instance).
- This lock is required in exclusive mode for any space management activity (creation or dropping any extents) except with locally managed tablespaces.



- Object creation, dropping, extension, and truncation all serialize on this lock.
- Most common causes for contention on this lock are sorting to disk (not using true temporary tablespaces) or rollback segment extension and shrinking.

Do the following to avoid contention on this enqueue:

- Use true temporary tablespaces, utilizing tempfiles. Temporary segments are not created and dropped after every sort to disk.
- Use locally managed tablespaces
- Size rollback segments to avoid dynamic extension and shrinking, or use automatic undo management.
- Avoid application practices that create and drop database objects.

#### **UL - User Defined Locks**

Users can define their own locks.

**See Also:** *Oracle9i Database Concepts* for more information on locks

#### **Common Modes for Request/Lmode**

- 0: None
- 2: Row Share: used for shared DML locks
- 4: Share: used for shared TX when waiting for ITL entry
- 6: Exclusive used for row level, DML locks

Any row in V\$LOCK either has LMODE=0 (indicating it is a request) or REQUEST=0 (indicating it is a held lock).

#### **Resource Identifier ID1**

For DML locks, ID1 is the object\_id.

For TX locks, ID1 points to the rollback segment and transaction table entry.

#### **Join Columns for V\$LOCK**

[Table 24-8](#) lists the join columns for V\$LOCK.

**Table 24–8 Join Columns for V\$LOCK**

| Column          | View        | Joined Column(s) |
|-----------------|-------------|------------------|
| SID             | V\$SESSION  | SID              |
| ID1, ID2, TYPE  | V\$LOCK     | ID1, ID2, TYPE   |
| ID1             | DBA_OBJECTS | OBJECT_ID        |
| TRUNCID1/65536) | V\$ROLLNAME | USN              |

1. This is used to find the session holding the lock, if a session is waiting for a lock.
2. This can be used to find the locked object for DML locks (type = 'TM').
3. This can be used to find the rollback segment in use for row transaction locks (TYPE = 'TX'). However, a less cryptic join might be through V\$TRANSACTION.

**Example 24–11 Finding the Sessions Holding the Lock**

Find the (ID1, ID2, type) for sessions waiting for a lock (LMODE=0).

Find the session holding the lock (REQUEST=0) for that ID1, ID2, type.

```
SELECT lpad(' ',DECODE(request,0,0,1))||sid sess, id1, id2, lmode, request, type
FROM V$LOCK
WHERE id1 IN (SELECT id1 FROM V$LOCK WHERE lmode = 0)
ORDER BY id1,request
```

| SID  | ID1    | ID2    | LMODE | REQUEST TY          |
|------|--------|--------|-------|---------------------|
| 1237 | 196705 | 200493 | 6     | 0 TX <- Lock Holder |
| 1256 | 196705 | 200493 | 0     | 6 TX <- Lock Waiter |
| 1176 | 196705 | 200493 | 0     | 6 TX <- Lock Waiter |
| 938  | 589854 | 201352 | 6     | 0 TX <- Lock Holder |
| 1634 | 589854 | 201352 | 0     | 6 TX <- Lock Waiter |

**Example 24–12 Finding the Statements being Executed by These Sessions**

```
SELECT sid, sql_hash_value
FROM V$SESSION
WHERE SID IN (1237,1256,1176,938,1634);
```

```
SID SQL_HASH_VALUE
-----
938 2078523611 <-Holder
```

```

1176      1646972797 <-Waiter
1237      3735785744 <-Holder
1256      1141994875 <-Waiter
1634      2417993520 <-Waiter

```

### Example 24–13 Finding the Text for These SQL Statements

```

HASH_VALUE  SQL_TEXT
-----
1141994875  SELECT TO_CHAR(CURRENT_MAX_UNIQUE_IDENTIFIER + 1 ) FROM PO_UNI
QUE_IDENTIFIER_CONTROL WHERE TABLE_NAME = DECODE(:b1,'RFQ','PO_
HEADERS_RFQ','QUOTATION','PO_HEADERS_QUOTE','PO_HEADERS') FOR UP
DATE OF CURRENT_MAX_UNIQUE_IDENTIFIER
1646972797  SELECT TO_CHAR(CURRENT_MAX_UNIQUE_IDENTIFIER + 1 ) FROM PO_UNI
QUE_IDENTIFIER_CONTROL WHERE TABLE_NAME = 'PO_HEADERS' FOR UPD
ATE OF CURRENT_MAX_UNIQUE_IDENTIFIER
2078523611  select CODE_COMBINATION_ID, enabled_flag, nvl(to_char(start_da
te_active, 'J'), -1), nvl(to_char(end_date_active, 'J'), -1), S
EGMENT2||'.'||SEGMENT1||'.'||SEGMENT6,detail_posting_allowed_f
lag,summary_flag from GL_CODE_COMBINATIONS where CHART_OF_ACCO
UNTS_ID = 101 and SEGMENT2 in ('000','341','367','388','389','4
52','476','593','729','N38','N40','Q21','Q31','U21') order by S
EGMENT2, SEGMENT1, SEGMENT6
2417993520  select 0 into :b0 from pa_projects where project_id=:b1 for upd
ate
3735785744  begin :X0 := FND_ATTACHMENT_UTIL_PKG.GET_ATCHMT_EXISTS(:L_ENTITY
_NAME, :L_PKEY1, :L_PKEY2, :L_PKEY3, :L_PKEY4, :L_PKEY5, :L_FUNC
TION_NAME, :L_FUNCTION_TYPE); end;

```

The locked sessions' statements show that the sessions 1176 and 1256 are waiting for a lock on the PO\_UNIQUE\_IDENTIFIER\_CONTROL held by session 1237, while session 1634 is waiting for a lock on PA\_PROJECTS held by session 938. Query V\$SESSION\_WAIT, V\$SESSION, and V\$SESSION\_EVENT to get more details about the sessions and users. For example:

- Who is holding the lock?
- Is the session holding the lock active or idle
- Is the session executing long running queries while holding the lock?

## V\$MTTR\_TARGET\_ADVICE

V\$MTTR\_TARGET\_ADVICE contains rows that predict the number of physical I/Os for the MTTR corresponding to each row. The rows also compute a physical I/O

factor, which is the ratio of the number of estimated I/Os to the number of I/Os actually performed by the current MTTR setting during the measurement interval. Column descriptions are listed in [Table 24-9](#).

**Table 24-9 V\$MTTR\_TARGET\_ADVICE Column Description**

| Column Name              | Datatype         | Column Description                                                                                                                                                                            |
|--------------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MTTR_TARGET_FOR_ESTIMATE | NUMBER           | MTTR setting being simulated. It should equal to the current MTTR setting if this is the first row of the view.                                                                               |
| ADVICE_STATUS            | VARCHAR2<br>( 5) | Current status of MTTR simulation (ON   READY   OFF).                                                                                                                                         |
| DIRTY_LIMIT              | NUMBER           | The dirty buffer limit derived from the MTTR being simulated.                                                                                                                                 |
| ESTD_CACHE_WRITES        | NUMBER           | The estimated number of cache physical writes under this MTTR.                                                                                                                                |
| ESTD_CACHE_WRITE_FACTOR  | NUMBER           | The estimated cache physical write ratio under this MTTR. It is the ratio of the estimated number of cache writes to the number of cache writes under current MTTR setting.                   |
| ESTD_TOTAL_WRITES        | NUMBER           | The estimated total number of physical write under this MTTR.                                                                                                                                 |
| ESTD_TOTAL_WRITE_FACTOR  | NUMBER           | The estimated total physical write ratio under this MTTR. It is the ratio of the estimated total number of physical writes to the total number of physical writes under current MTTR setting. |
| ESTD_TOTAL_IOS           | NUMBER           | The estimated total number of I/Os under this MTTR.                                                                                                                                           |
| ESTD_TOTAL_IO_FACTOR     | NUMBER           | The estimated total I/O ratio under this MTTR. It is the ratio of the estimated total number of I/Os to the total number of I/Os under current MTTR setting.                                  |

## V\$MYSTAT

This view is a subset of V\$SESSTAT returning current session's statistics. When auditing resource usage for sessions through triggers, use V\$MYSTAT to capture the resource usage, because it is much cheaper than scanning the rows in V\$SESSTAT.

## V\$OPEN\_CURSOR

This view lists all the cursors opened by the sessions. There are several ways it can be used. For example, you can monitor the number of cursors opened by different sessions.

When diagnosing system resource usage, it is useful to query V\$SQLAREA and V\$SQL for expensive SQL (high logical or physical I/O). In such cases, the next step is to find its source. On applications where users log in to the database as the same generic user (and have the same PARSING\_USER\_ID in V\$SQLAREA), this can get difficult. The statistics in V\$SQLAREA are updated after the statement completes execution (and disappears from V\$SESSION.SQL\_HASH\_VALUE). Therefore, unless the statement is being executed again, you cannot find the session directly. However, if the cursor is still open for the session, then use V\$OPEN\_CURSOR to find the session(s) that have executed the statement.

### Join Columns for V\$OPEN\_CURSOR

Table 24–10 lists the join columns for V\$OPEN\_CURSOR.

**Table 24–10 Join Columns for V\$OPEN\_CURSOR**

| Column                 | View                              | Joined Column(s)       |
|------------------------|-----------------------------------|------------------------|
| HASH_VALUE,<br>ADDRESS | V\$SQLAREA, V\$SQL,<br>V\$SQLTEXT | HASH_VALUE,<br>ADDRESS |
| SID                    | V\$SESSION                        | SID                    |

### Example 24–14 Finding the Session(s) that Executed a Statement

```
SELECT hash_value, buffer_gets, disk_reads
FROM V$SQLAREA
WHERE disk_reads > 1000000
ORDER BY buffer_gets DESC;
```

```
HASH_VALUE  BUFFER_GETS  DISK_READS
-----
1514306888   177649108    3897402
 478652562   63168944     2532721
 360282550   14158750     2482065
 226079402   40458060     1592621
2144648214   1493584      1478953
1655760468   1997868      1316010
 160130138   6609577      1212163
3000880481   2122483      1158608
```

8 rows selected.

```
SQL> SELECT sid FROM V$SESSION WHERE sql_hash_value = 1514306888 ;
```

no rows selected

```
SQL> SELECT sid FROM V$OPEN_CURSOR WHERE hash_Value = 1514306888 ;
```

```

SID
-----
1125
 233
  935
1693
  531

```

5 rows selected.

**Example 24–15 Finding Sessions That Have More Than 400 Cursors Open**

```

SELECT sid, count(*)
FROM v$open_cursor
GROUP BY sid
HAVING COUNT(*) > 400
ORDER BY count(*) desc;

```

```

SID    COUNT(*)
-----
2359    456
1796    449
1533    445
1135    442
1215    442
  810    437
1232    429
  27     426
1954    421
2067    421
1037    416
1584    413
  416    407
  398    406
  307    405
1545    403

```

## V\$PARAMETER and V\$SYSTEM\_PARAMETER

These views list each initialization parameter by name and show the value for that parameter. The V\$PARAMETER view shows the current value for the session performing the query. The V\$SYSTEM\_PARAMETER view shows the instance-wide value for the parameter.

For example, executing the following query shows the SORT\_AREA\_SIZE parameter setting for the session executing the query:

```
SELECT value
   FROM V$PARAMETER
  WHERE name = 'sort_area_size';
```

### Useful Columns for V\$PARAMETER

- NAME: Name of the parameter
- VALUE: Current value for this session (if modified within the session); otherwise, the instance-wide value
- ISDEFAULT: Whether the parameter value is the default value
- ISSES\_MODIFIABLE: Whether this parameter can be modified at the session level
- ISSYS\_MODIFIABLE: Whether this parameter can be modified at an instance-wide level dynamically after the instance has started
- ISMODIFIED: Whether this parameter has been modified after instance startup, and if so, whether it was modified at the session level or at the instance (system) level
- ISADJUSTED: Whether Oracle has adjusted a value specified by the user
- DESCRIPTION: Brief description of the parameter
- UPDATE\_COMMENT: Set if a comment has been supplied by the DBA for this parameter

#### See Also:

- *Oracle9i Database Reference* for more information on the range column values
- *Oracle9i Database Administrator's Guide* for information on server parameter files

## Uses for V\$PARAMETER and V\$SYSTEM\_PARAMETER Data

V\$PARAMETER is queried during performance tuning to determine the current settings for a parameter. For example, if the buffer cache hit ratio is low, then the value for DB\_BLOCK\_BUFFERS (or DB\_CACHE\_SIZE) can be queried to determine the current buffer cache size.

The SHOW PARAMETER statement in SQL\*Plus queries data from V\$PARAMETER.

### Example 24–16 Determining the SORT\_AREA\_SIZE From Within SQL\*Plus

```
column name          format a20
column value         format a10
column isdefault     format a5
column isses_modifiable format a5
```

```
SELECT name, value, isdefault, isses_modifiable, issys_modifiable, ismodified
FROM V$PARAMETER
WHERE name = 'sort_area_size';
```

| NAME           | VALUE   | ISDEF | ISSES | ISSYS_MOD | ISMODIFIED |
|----------------|---------|-------|-------|-----------|------------|
| sort_area_size | 1048576 | TRUE  | TRUE  | DEFERRED  | MODIFIED   |

The preceding example shows that the SORT\_AREA\_SIZE initialization parameter was not set as an initialization parameter on instance startup, but was modified at the session level (indicated by the ISMODIFIED column having the value of MODIFIED) for this session.

---



---

**Note:** Use caution when querying from V\$PARAMETER. If you want to see the instance-wide parameters, use V\$SYSTEM\_PARAMETER view instead of V\$PARAMETER.

---



---

## V\$PROCESS

This view contains information about all Oracle processes running on the system. It is used to relate the Oracle or operating system process ID of the server process to the database session. This is needed in several situations:

- If the bottleneck on the database server is related to an operating system resource (for example, CPU, memory), and if the top resource users are localized within a small set of server processes, then perform the following steps:



1. Find the resource intensive processes.
  2. Find their sessions. You must relate the processes to sessions.
  3. Find out why the session is using so many resources.
- The SQL\*Trace file names are based on the operating system process ID of the server process. To locate the trace file for a session, you must relate the session to the server process.
  - Some events, like `rdbms ipc reply`, identify the Oracle process ID of the process a session is waiting on. To find out what those processes are doing, you must find their sessions.
  - The background processes you see on the server (DBWR, LGWR, PMON, and so on) are all server processes. To see what they are doing in the database, you must find their session.

### Useful Columns for V\$PROCESS

- PID: Oracle process ID of the process
- SPID: Operating system process ID of the process

### Join Columns for V\$PROCESS

[Table 24–11](#) lists the join columns for V\$PROCESS.

**Table 24–11 Join Columns for V\$PROCESS**

| Column | View       | Joined Column(s) |
|--------|------------|------------------|
| ADDR   | V\$SESSION | PADDR            |

### Example 24–17 Finding the Session for Server Process 20143

```
SELECT ' sid, serial#, aud sid : '|| s.sid||' , '||s.serial#||' , '||
s.audsid||chr(10)|| '      DB User / OS User : '||s.username||
' / '||s.osuser||chr(10)|| '      Machine - Terminal : '||
s.machine||' - '|| s.terminal||chr(10)||
'      OS Process Ids : '|| s.process||' (Client) '||
p.spid||' - '||p.pid||' (Server)'|| chr(10)||
'      Client Program Name : '||s.program "Session Info"
FROM V$PROCESS P,V$SESSION S
WHERE p.addr = s.paddr
      AND p.spid = '20143';
```

Session Info

```

-----
Sid, Serial#, Aud sid : 2204 , 5552 , 14478782
  DB User / OS User : APPS / sifapmgr
  Machine - Terminal : finprod3 -
  OS Process Ids : 9095 (Client) 20143 - 1404 (Server)
  Client Program Name : RGRARG@finprod3 (TNS V1-V3)

```

### **Example 24–18 Finding the Session for PMON**

```

SELECT ' sid, serial#, aud sid : ' || s.sid || ' , ' || s.serial# || ' , ' ||
  s.audsid || chr(10) || '      DB User / OS User : ' || s.username ||
  ' / ' || s.osuser || chr(10) || '      Machine - Terminal : ' ||
  s.machine || ' - ' || s.terminal || chr(10) ||
  '      OS Process Ids : ' || s.process || ' (Client) ' ||
  p.spid || ' - ' || p.pid || ' (Server)' || chr(10) ||
  '      Client Program Name : ' || s.program "Session Info"
FROM V$PROCESS p, V$SESSION s
WHERE p.addr = s.paddr
      AND s.program LIKE '%PMON%'

```

Session Info

```

-----
Sid, Serial#, Aud sid : 1 , 1 , 0
  DB User / OS User : / oracle
  Machine - Terminal : finprod7 - UNKNOWN
  OS Process Ids : 20178 (Client) 20178 - 2 (Server)
  Client Program Name : oracle@finprod7 (PMON)

```

You can see that the client and server processes are the same for the background process, which is why we could specify the client program name.

## **V\$ROLLSTAT**

This view keeps a summary of statistics for each rollback segment since startup.

### **Useful Columns for V\$ROLLSTAT**

- USN: Rollback segment number
- RSSIZE: Current size of the rollback segment
- XACTS: Number of active transactions

### **Columns Useful for Doing a Delta Over a Period of Time**

- WRITES: Number of bytes written to the rollback segment

- SHRINKS: Number of times the rollback segment grew past OPTIMAL and shrank back
- EXTENDS: Number of times the rollback segment had to extend because there was an active transaction in the next extent
- WRAPS: Number of times the rollback segment wrapped around
- GETS: Number of header gets
- WAITS: Number of header waits

### Join Columns for V\$ROLLSTAT

Table 24–12 lists the join columns for V\$ROLLSTAT.

**Table 24–12 Join Columns for V\$ROLLSTAT**

| Column | View        | Joined Column(s) |
|--------|-------------|------------------|
| USN    | V\$ROLLNAME | USN              |

### Example 24–19 Querying V\$ROLLSTAT

By dividing the elapsed time by wraps, you can determine the average time taken for a rollback segment to wrap. This is useful in sizing rollback segments for long running queries to avoid 'Snapshot Too Old' errors.

Also, monitor the extends and shrinks to see if the optimal size should be increased.

## V\$ROWCACHE

This view displays statistics for the dictionary cache (also known as the rowcache). Each row contains statistics for the various types of dictionary cache data. Note that there is a hierarchy in the dictionary cache, so the same cache name can appear more than once.

### Useful Columns for V\$ROWCACHE

- PARAMETER: Name of the cache
- COUNT: Number of entries allocated to this cache
- USAGE: Current number of used entries
- GETS: Total number of requests
- GETMISSES: Number of requests resulting in dictionary cache miss

- SCANS: Number of scan requests
- SCANMISSES: Number of times a scan failed to find the required data
- MODIFICATIONS: Number of additions, changes or deletions of cache entries
- DLM\_REQUESTS: Number of DLM Real Application Clusters requests
- DLM\_CONFLICTS: Number of DLM Real Application Clusters conflicts
- DLM\_RELEASES: Number of DLM Real Application Clusters releases

### Uses for V\$ROWCACHE Data

- Determine whether the dictionary cache is adequately sized. If the shared pool is too small, then the dictionary cache is not able to grow to a sufficient size to cache the required information.

**See Also:** ["Sizing the Shared Pool"](#) on page 14-27 for details about tuning the shared pool

- Determine whether the application is accessing the cache efficiently. If the application design uses the dictionary cache inefficiently (in this case, a larger dictionary cache will not alleviate the performance problem). For example, if a large number of GETS appear for the DC\_USERS cache within the sample period, then it is likely that there are large number of distinct users created within the database, and that the application is logging the users on and off frequently. To verify this, check the logon rate and also the number of users in the system. The parse rates will also be high. If this is a large OLTP system with a middle tier, then it might be more efficient to manage individual accounts on the middle tier, allowing the middle tier to logon as a single user: the application owner. Reducing logon/logoff rate by keeping connections active also helps.
- Determine whether dynamic space allocation is occurring. A large number of similarly sized modifications for DC\_SEGMENTS, DC\_USED\_EXTENTS, and DC\_FREE\_EXTENTS can indicate much dynamic space allocation. Possible solutions include sizing the next extents appropriately, or using locally managed tablespaces. If the space allocation is occurring on the temp tablespace, then use a true temporary tablespace for the temp.
- Identify large amounts of sequence number generation occurring. Modifications to dc\_sequences indicates this. Check to see whether the number of cache entries for each sequence number are sufficient for then number of changes.

- Gather evidence for hard parsing. Hard parsing can also be evidenced by many GETS to DC\_COLUMNS, DC\_VIEWS and DC\_OBJECTS caches.

#### **Example 24–20 Querying V\$ROWCACHE Data**

A good way to view dictionary cache statistics is to group the data by the cache name.

```
SELECT parameter
       , sum("COUNT")
       , sum(usage)
       , sum(gets)
       , sum(getmisses)
       , sum(scans)
       , sum(scanmisses)
       , sum(modifications)
       , sum(dlm_requests)
       , sum(dlm_conflicts)
       , sum(dlm_releases)
FROM V$ROWCACHE
GROUP BY parameter;
```

## **V\$SEGMENT\_STATISTICS**

This is a user-friendly view, available with Oracle9i Release 2 (9.2) and higher, that allows real-time monitoring of segment-level statistics, enabling a DBA to identify performance problems associated with an individual table or index.

**Table 24–13 V\$SEGMENT\_STATISTICS View**

| <b>Column</b>   | <b>Datatype</b> | <b>Description</b>                                  |
|-----------------|-----------------|-----------------------------------------------------|
| OWNER           | VARCHAR2(30)    | Owner of the object                                 |
| OBJECT_NAME     | VARCHAR2(30)    | Name of the object                                  |
| SUBOBJECT_NAME  | VARCHAR2(30)    | Name of the sub-object                              |
| TABLESPACE_NAME | VARCHAR2(30)    | Name of the table space to which the object belongs |
| TS#             | NUMBER          | Tablespace number                                   |
| OBJ#            | NUMBER          | Dictionary object number                            |
| DATAOBJ#        | NUMBER          | Data object number                                  |
| OBJECT_TYPE     | VARCHAR2(18)    | Type of the object                                  |

**Table 24–13 (Cont.) V\$SEGMENT\_STATISTICS View**

| Column         | Datatype        | Description           |
|----------------|-----------------|-----------------------|
| STATISTIC_NAME | VARCHAR2 ( 64 ) | Name of the statistic |
| STATISTIC#     | NUMBER          | Statistic number      |
| VALUE          | NUMBER          | Statistic value       |

## V\$SEGSTAT

This is a high-efficiency view for real-time monitoring of segment-level statistics, available with Oracle9i Release 2 (9.2) and higher.

**Table 24–14 V\$SEGSTAT View**

| Column         | Datatype        | Description              |
|----------------|-----------------|--------------------------|
| TS#            | NUMBER          | Tablespace number        |
| OBJ#           | NUMBER          | Dictionary object number |
| DATAOBJ#       | NUMBER          | Data object number       |
| STATISTIC_NAME | VARCHAR2 ( 64 ) | Name of the statistic    |
| STATISTIC#     | NUMBER          | Statistic number         |
| VALUE          | NUMBER          | Statistic value          |

## V\$SEGSTAT\_NAME

This is a statistics property view for segment-level statistics, available with Oracle9i Release 2 (9.2) and higher.

**Table 24–15 V\$SEGSTAT\_NAME View**

| Column     | Datatype        | Description                              |
|------------|-----------------|------------------------------------------|
| STATISTIC# | NUMBER          | Statistic number                         |
| NAME       | VARCHAR2 ( 64 ) | Statistic name                           |
| SAMPLED    | VARCHAR2 ( 3 )  | Whether or not it is a sampled statistic |

## V\$SESSION

This view has one row for every session connected to the database instance. The sessions include user sessions, as well as background processes like DBWR, LGWR, archiver.

**See Also:** *Oracle9i Database Concepts*

### Useful Columns for V\$SESSION

V\$SESSION is basically an information view used for finding the SID or SADDR of a user. However, it has some columns that change dynamically and are useful for examining a user. For example:

SQL\_HASH\_VALUE, SQL\_ADDRESS: These identify the SQL statement currently being executed by the session. If NULL or 0, then the session is not executing any SQL statement. PREV\_HASH\_VALUE and PREV\_ADDRESS identify the previous statement being executed by the session.

---



---

**Note:** When selecting from SQL\*Plus, make sure that you have the column defined with adequate width (11 numbers wide) to see the complete number.

---



---

STATUS: This column identifies if the session is:

- Active: executing a SQL statement (waiting for/using a resource)
- Inactive: waiting for more work (that is, SQL statements)
- Killed: marked to be killed

The following columns provide information about the session and can be used to find a session when a combination (one or more) of the following are known:

#### Session Information

- SID: Session identifier, used to join to other columns
- SERIAL#: Counter, which is incremented each time a SID is reused by another session (when a session ends and another session starts and uses the same SID)
- AUDSID: Auditing session ID uniquely identifies a session over the life of a database. It is also useful when finding the parallel query slaves for a query coordinator (during the PQ execution they have the same AUDSID)

- **USERNAME:** The Oracle user name for the connected session

### Client Information

The database session is initiated by a client process that could be running on the database server or connecting to the database across SQL\*Net from a middle tier server or even a desktop. The following columns provide information about this client process:

- **OSUSER:** Operating system user name for the client process
- **MACHINE:** Machine where the client process is executing
- **TERMINAL:** Terminal (if applicable) where the client process is running
- **PROCESS:** Process ID of the client process
- **PROGRAM:** Client program being executed by the client process

To display **TERMINAL**, **OSUSER** for users connecting from PCs, set the keys **TERMINAL**, **USERNAME** in **ORACLE.INI** or the Windows registry on their PCs if they are not showing up by default.

### Application Information

Call the package **DBMS\_APPLICATION\_INFO** to set some information to identify the user. This shows up in the following columns:

- **CLIENT\_INFO:** Set in **DBMS\_APPLICATION\_INFO**
- **ACTION:** Set in **DBMS\_APPLICATION\_INFO**
- **MODULE:** Set in **DBMS\_APPLICATION\_INFO**

The following **V\$SESSION** columns are also useful:

- **ROW\_WAIT\_OBJ#**
- **ROW\_WAIT\_FILE#**
- **ROW\_WAIT\_BLOCK#**
- **ROW\_WAIT\_ROW#**

### Join Columns for V\$SESSION

[Table 24-16](#) is a list of several columns that can be used to join to other fixed views.



**Table 24–16 Join Columns for V\$SESSION**

| Column                                 | View                                                                            | Joined Column(s)         |
|----------------------------------------|---------------------------------------------------------------------------------|--------------------------|
| SID                                    | V\$SESSION_WAIT,<br>V\$SESSTAT, V\$LOCK,<br>V\$SESSION_EVENT,<br>V\$OPEN_CURSOR | SID                      |
| (SQL_HASH_VALUE,<br>SQL_ADDRESS)       | V\$SQLTEXT,<br>V\$SQLAREA, V\$SQL                                               | (HASH_VALUE,<br>ADDRESS) |
| (PREV_HASH_VALUE,<br>PREV_SQL_ADDRESS) | V\$SQLTEXT,<br>V\$SQLAREA, V\$SQL                                               | (HASH_VALUE,<br>ADDRESS) |
| TADDR                                  | V\$TRANSACTION                                                                  | ADDR                     |
| PADDR                                  | V\$PROCESS                                                                      | ADDR                     |

**Example 24–21 Finding Your Session**

```
SELECT SID, OSUSER, USERNAME, MACHINE, PROCESS
   FROM V$SESSION
  WHERE audsid = userenv('SESSIONID');
```

```
  SID OSUSER      USERNAME      MACHINE      PROCESS
-----
  398 amerora     PERFSTAT     rgmdbs1      26582
```

**Example 24–22 Finding a Session When the Machine Is Known**

```
SELECT SID, OSUSER, USERNAME, MACHINE, TERMINAL
   FROM V$SESSION
  WHERE terminal = 'pts/tl'
     AND machine = 'rgmdbs1';
```

```
  SID OSUSER      USERNAME      MACHINE      TERMINAL
-----
  398 amerora     PERFSTAT     rgmdbs1      pts/tl
```

**Example 24–23 Finding the SQL Statement Currently Being Run by a Session**

It is a common requirement to find the SQL statement currently being executed by a given session. If a session is experiencing or responsible for a bottleneck, then the statement explains what the session might be doing.

```
col hash_value form 9999999999
SELECT sql_hash_value hash_value
```

```
FROM V$SESSION WHERE sid = 406;

HASH_VALUE
-----
4249174653
SQL> /

HASH_VALUE
-----
4249174653
SQL> /

HASH_VALUE
-----
4249174653
SQL> /

HASH_VALUE
-----
4249174653
```

This example waited for five seconds, executed the statement again, and repeated the action couple of times. The same hash\_value comes up again and again, indicating that the statement is being executed by the session. As a next step, find the statement text using the view V\$SQLTEXT and statement statistics from V\$SQLAREA.

## V\$SESSION\_EVENT

This view summarizes wait events for every session. While V\$SESSION\_WAIT shows the current waits for a session, V\$SESSION\_EVENT provides summary of all the events the session has waited for since it started.

### Useful Columns for V\$SESSION\_EVENT

- SID: Identifier for the session
- EVENT: Name of the wait event
- TOTAL\_WAITS: Total number of waits for this event by this session
- TIME\_WAITED: Total time waited for this event (in hundredths of a second)
- AVERAGE\_WAIT: Average amount of time waited for this event by this session (in hundredths of a second)

- `TOTAL_TIMEOUTS`: Number of times the wait timed out

### Join Columns for `V$SESSION_EVENT`

Table 24–17 is a list of join columns for `V$SESSION_EVENT`.

**Table 24–17 Join Columns for `V$SESSION_EVENT`**

| Column | View                    | Joined Column(s) |
|--------|-------------------------|------------------|
| SID    | <code>V\$SESSION</code> | SID              |

### Example 24–24 Finding the Waits for the Database Writer

```
SELECT s.sid, bgp.name
       FROM v$session s, v$bgprocess bgp
       WHERE bgp.name LIKE '%DBW%'
         AND bgp.paddr = s.paddr;
```

```
SELECT event, total_waits waits, total_timeouts timeouts,
       time_waited total_time, average_wait avg
       FROM v$session_event
       WHERE sid = 3
       ORDER BY time_waited DESC;
```

| EVENT                        | WAITS   | TIMEOUTS | TOTAL_TIME | AVG    |
|------------------------------|---------|----------|------------|--------|
| -----                        | -----   | -----    | -----      | -----  |
| rdbms ipc message            | 1684385 | 921495   | 284706709  | 169.03 |
| db file parallel write       | 727326  | 0        | 3012982    | 4.14   |
| latch free                   | 157     | 157      | 281        | 1.78   |
| control file sequential read | 123     | 0        | 61         | 0.49   |
| file identify                | 45      | 0        | 29         | 0.64   |
| direct path read             | 41      | 0        | 5          | 0.12   |
| file open                    | 49      | 0        | 2          | 0.04   |
| db file sequential read      | 2       | 0        | 2          | 1.00   |

## `V$SESSION_WAIT`

This is a key view for finding bottlenecks. It tells what every session in the database is currently waiting for (or the last event waited for by the session if it is not waiting for anything). This view can be used as a starting point to find which direction to proceed in when a system is experiencing performance problems.

`V$SESSION_WAIT` has a row for every session connected to the instance. It indicates if the session is:

- Using a resource
- Waiting for a resource
- Idle (waiting on one of the idle events)

### Useful Columns for V\$SESSION\_WAIT

- SID: Session identifier for the session
- EVENT: Event the session is currently waiting for, or the last event the session had to wait for
- WAIT\_TIME: Time (in hundredths of a second) that the session waited for the event; if the WAIT\_TIME is 0, then the session is currently waiting for the event
- SEQ#: Gets incremented with every wait of the session
- P1, P2, P3: Wait event specific details for the wait
- P1TEXT, P2TEXT, P3TEXT: Description of P1,P2,P3 for the given event

**See Also:** *Oracle9i Database Reference* and ["Wait Events"](#) on page 22-24

**Table 24–18** Wait Time Description

| WAIT_TIME | Meaning                                             | Waiting |
|-----------|-----------------------------------------------------|---------|
| >0        | Time waited in the last wait (in 10 ms clock ticks) | No      |
| 0         | Session is currently waiting for this event         | Yes     |
| -1        | Time waited in the last wait was less than 10 ms    | No      |
| -2        | Timing is not enabled                               | No      |

**Table 24–19** shows an example of how the EVENT, SEQ#, and WAIT\_TIME might change over a period of time:

**Table 24–19** Events Changing Over Time

| Time | Seq # | Event                   | Wait Time | P1       | P2   | P3 | Action              | Waiting |
|------|-------|-------------------------|-----------|----------|------|----|---------------------|---------|
| 0    | 43    | latch free              | 0         | 800043F8 | 31   | 1  | Get LRU latch       | Yes     |
| 10   | 43    | latch free              | 10        | 800043F8 | 31   | 1  | Get free buffer     | No      |
| 20   | 44    | db file sequential read | 0         | 5        | 1345 | 1  | Issue the read call | Yes     |

**Table 24–19 (Cont.) Events Changing Over Time**

| Time | Seq # | Event                   | Wait Time | P1         | P2     | P3   | Action             | Waiting |
|------|-------|-------------------------|-----------|------------|--------|------|--------------------|---------|
| 30   | 44    | db file sequential read | 10        | 5          | 1345   | 1    | Process the buffer | No      |
| 35   | 45    | enqueue                 | 0         | 1415053318 | 196631 | 6355 | Lock the buffer    | Yes     |
| 1040 | 45    | enqueue                 | 1000      | 1415053318 | 196631 | 6355 | Modify the buffer  | No      |

In this example, the session waited for a *latch* from 0-10, waited for *db file sequential read* from 20-30, waited for a *lock* from 35-1040. The times in between have been exaggerated for illustration purposes. Event and Seq# do not change until the session has to wait again. The Wait Time indicates if the session is actually waiting or using a resource.

### Join Columns for V\$SESSION\_WAIT

Table 24–20 is a list of join columns for V\$SESSION\_WAIT.

**Table 24–20 Join Columns for V\$SESSION\_WAIT**

| Column | View       | Joined Column(s) |
|--------|------------|------------------|
| SID    | V\$SESSION | SID              |

### Example 24–25 Finding Current Waits on the System

```
SELECT event,
       sum(decode(wait_time,0,1,0)) "Curr",
       sum(decode(wait_time,0,0,1)) "Prev",
       count(*)"Total"
  FROM v$session_wait
 GROUP BY event
 ORDER BY count(*);
```

| EVENT                         | Prev | Curr | Tot |
|-------------------------------|------|------|-----|
| PL/SQL lock timer             | 0    | 1    | 1   |
| SQL*Net more data from client | 0    | 1    | 1   |
| smon timer                    | 0    | 1    | 1   |
| pmon timer                    | 0    | 1    | 1   |
| SQL*Net message to client     | 2    | 0    | 2   |
| db file scattered read        | 2    | 0    | 2   |
| rdbms ipc message             | 0    | 7    | 7   |
| enqueue                       | 0    | 12   | 12  |

|                             |     |      |      |
|-----------------------------|-----|------|------|
| pipe get                    | 0   | 12   | 12   |
| db file sequential read     | 3   | 10   | 13   |
| latch free                  | 9   | 6    | 15   |
| SQL*Net message from client | 835 | 1380 | 2215 |

This query, which groups the data by event and by wait\_time (0=waiting, nonzero=not waiting), shows the following:

- Most of the sessions are waiting for idle events like SQL\*Net message from client, pipe get, PMON timer, and so on.
- The number of sessions using the CPU can be approximated by the number of sessions not waiting (prev), except for one problem: there seem to be a lot of sessions that are not waiting for anything (hence actively using resources) and whose last wait was SQL\*Net message from client.

The next step should be to check V\$SESSION to see if the session is active or not. Only count the session as actively waiting or using a resource if it is active. Use the following statement to accomplish this. The total column counts the total of all the sessions, however the currently waiting and previously waited (using resource) columns only count active sessions.

```
SELECT event,
       sum(decode(wait_Time,0,0,DECODE(s.status,'ACTIVE',1,0))) "Prev",
       sum(decode(wait_Time,0,1,DECODE(s.status,'ACTIVE',1,0))) "Curr",
       count(*) "Tot"
FROM v$session s, v$session_wait w
WHERE s.sid = w.sid
GROUP BY event
ORDER BY count(*);
```

| EVENT                       | Prev | Curr | Tot                |
|-----------------------------|------|------|--------------------|
| SQL*Net message to client   | 1    | 1    | 1 <- idle event    |
| buffer busy waits           | 1    | 1    | 1                  |
| file open                   | 1    | 1    | 1                  |
| pmon timer                  | 0    | 1    | 1 <- idle event    |
| smon timer                  | 0    | 1    | 1 <- idle event    |
| log file sync               | 0    | 1    | 1                  |
| db file scattered read      | 0    | 2    | 2                  |
| rdbms ipc message           | 0    | 7    | 7 <- idle event    |
| pipe get                    | 0    | 12   | 12 <- idle event   |
| enqueue                     | 0    | 14   | 14                 |
| latch free                  | 10   | 17   | 20                 |
| db file sequential read     | 7    | 22   | 23                 |
| SQL*Net message from client | 0    | 1383 | 2240 <- idle event |

Now sessions are counted as actively waiting or using a resource only if they are active. This highlights the following:

- There are a total of 2324 sessions.
- 20 sessions are actively using resources (active sessions without an active wait).
- 1463 sessions are waiting.
- 58 of these are waiting for non-idle events. The idle events here being `SQL*Net message from client`, `pipe get`, `rdbms ipc message`, `PMON timer`, `SMON timer`, and `SQL*Net message to client`.

**See Also:** ["Wait Events"](#) on page 22-24

- 14 sessions are locked out (and may be experiencing poor performance).
- PMON and SMON are sleeping on their timers.
- 24 sessions are waiting for I/O calls to return (`db file%read`).

## V\$SESSTAT

V\$SESSTAT stores session-specific resource usage statistics, beginning at login and ending at logout.

Similar to V\$SYSSTAT, this view stores the following types of statistics:

- A count of the number of times an action occurred, such as `user commits`
- A running total of volumes of data generated, accessed or manipulated, such as `redo size`
- The cumulative time spent performing some actions, such as CPU used by this session, if `TIMED_STATISTICS` is set to `TRUE`

---

---

**Note:** Timed statistics are automatically collected for the database if the initialization parameter `STATISTICS_LEVEL` is set to `TYPICAL` or `ALL`. If `STATISTICS_LEVEL` is set to `BASIC`, then you must set `TIMED_STATISTICS` to `TRUE` to enable collection of timed statistics.

If you explicitly set `DB_CACHE_ADVICE`, `TIMED_STATISTICS`, or `TIMED_OS_STATISTICS`, either in the initialization parameter file or by using `ALTER_SYSTEM` or `ALTER SESSION`, the explicitly set value overrides the value derived from `STATISTICS_LEVEL`.

---

---

**See Also:** ["Setting the Level of Statistics Collection"](#) on page 22-10 for information about `STATISTICS_LEVEL` settings

The differences between `V$SYSSTAT` and `V$SESSTAT` are the following:

- `V$SESSTAT` only stores data for each session, whereas `V$SYSSTAT` stores the accumulated values for all sessions.
- `V$SESSTAT` is transitory, and is lost after a session logs out. `V$SYSSTAT` is cumulative, and is only lost when the instance is shutdown.
- `V$SESSTAT` does not include the name of the statistic. In order to find the statistic name, this view must be joined to either `V$SYSSTAT` or `V$STATNAME`.

`V$SESSTAT` can be used to find sessions with the following:

- The highest resource usage
- The highest average resource usage rate (ratio of resource usage to logon time)
- The current resource usage rate (delta between two snapshots)

### Useful Statistics in `V$SESSTAT`

The most referenced statistics in `V$SESSTAT` are a subset of those described for `V$SYSSTAT` and include `session logical reads`, `CPU used by this session`, `db block changes`, `redo size`, `physical writes`, `parse count (hard)`, `parse count (total)`, `sorts (memory)`, and `sorts (disk)`.

### Useful Columns for `V$SESSTAT`

- `SID`: Session identifier



- STATISTIC#: Resource identifier
- VALUE: Resource usage

### Join Columns for V\$SESSTAT

Table 24–21 lists the join columns for V\$SESSTAT.

**Table 24–21 Join Columns for V\$SESSTAT**

| Column     | View        | Joined Column(s) |
|------------|-------------|------------------|
| STATISTIC# | V\$STATNAME | STATISTIC#       |
| SID        | V\$SESSION  | SID              |

### **Example 24–26 Finding the Top Sessions with Highest Logical and Physical I/O Rates Currently Connected to the Database**

The following SQL statement shows the logical and physical read rates (each second) for all active sessions connected to the database. Rates for logical and physical I/O are calculated using the elapsed time since logon (from V\$SESSION.LOGON\_TIME). This might not be particularly accurate for sessions connected to the database for long periods, but it is sufficient for this example.

To determine the STATISTIC#'s for the session logical reads and physical reads statistics:

```
SELECT name, statistic#
       FROM V$STATNAME
       WHERE name IN ('session logical reads','physical reads') ;
```

```
NAME                                STATISTIC#
-----                                -
session logical reads                9
physical reads                        40
```

Use these values in the following query, which orders the sessions by resource usage:

```
SELECT ses.sid
       , DECODE(ses.action,NULL,'online','batch')      "User"
       , MAX(DECODE(sta.statistic#,9,sta.value,0))
         /greatest(3600*24*(sysdate-ses.logon_time),1)  "Log IO/s"
       , MAX(DECODE(sta.statistic#,40,sta.value,0))
         /greatest(3600*24*(sysdate-ses.logon_time),1)  "Phy IO/s"
       , 60*24*(sysdate-ses.logon_time)                "Minutes"
```

```

FROM V$SESSION ses
     , V$SESSTAT sta
WHERE ses.status      = 'ACTIVE'
      AND sta.sid      = ses.sid
      AND sta.statistic# IN (9,40)
GROUP BY ses.sid, ses.action, ses.logon_time
ORDER BY
      SUM( DECODE(sta.statistic#,40,100*sta.value,sta.value) )
      / greatest(3600*24*(sysdate-ses.logon_time),1) DESC;

```

| SID  | User   | Log IO/s | Phy IO/s | Minutes |
|------|--------|----------|----------|---------|
| 1951 | batch  | 291      | 257.3    | 1       |
| 470  | online | 6,161    | 62.9     | 0       |
| 730  | batch  | 7,568    | 43.2     | 197     |
| 2153 | online | 1,482    | 98.9     | 10      |
| 2386 | batch  | 7,620    | 35.6     | 35      |
| 1815 | batch  | 7,503    | 35.5     | 26      |
| 1965 | online | 4,879    | 42.9     | 19      |
| 1668 | online | 4,318    | 44.5     | 1       |
| 1142 | online | 955      | 69.2     | 35      |
| 1855 | batch  | 573      | 70.5     | 8       |
| 1971 | online | 1,138    | 56.6     | 1       |
| 1323 | online | 3,263    | 32.4     | 5       |
| 1479 | batch  | 2,857    | 35.1     | 3       |
| 421  | online | 1,322    | 46.8     | 15      |
| 2405 | online | 258      | 50.4     | 8       |

To better show the impact of each individual session on the system, the results were ordered by the total resource usage each second. The resource usage was calculated by adding session logical reads and (a weighted) physical reads.

Physical reads was weighted by multiplying the raw value by a factor of 100, to indicate that a physical I/O is significantly more expensive than reading a buffer already in the cache.

To calculate the physical I/O weighting factor, the following assumptions were made:

- Average wait for a physical I/O (PIO) was 10 ms (queried from V\$SYSTEM\_EVENT.AVERAGE\_WAIT for the events db file sequential read and db file scattered read).
- Average logical I/O rate (LIO) was 13000/second/CPU (queried from V\$SYSSTAT for the statistic name session logical reads. This statistic was divided by the elapsed time in seconds and the number of CPUs on the system).

- This provides a ratio of 130 logical reads for each 10 ms, and 1 physical read for each 10 ms for this configuration. This ratio was rounded to the ballpark number of 100.

## V\$SHARED\_POOL\_ADVICE

V\$SHARED\_POOL\_ADVICE displays information about estimated parse time savings in the shared pool for different sizes. The sizes range from 50% to 200% of the current shared pool size, in equal intervals. The value of the interval depends on the current size of the shared pool.

**Table 24–22 V\$SHARED\_POOL\_ADVICE View**

| Column                        | Datatype | Description                                                                                                                                |
|-------------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------|
| SHARED_POOL_SIZE_FOR_ESTIMATE | NUMBER   | Shared pool size for the estimate (in megabytes)                                                                                           |
| SHARED_POOL_SIZE_FACTOR       | NUMBER   | Size factor with respect to the current shared pool size                                                                                   |
| ESTD_LC_SIZE                  | NUMBER   | Estimated memory in use by the library cache (in megabytes)                                                                                |
| ESTD_LC_MEMORY_OBJECTS        | NUMBER   | Estimated number of library cache memory objects in the shared pool of the specified size                                                  |
| ESTD_LC_TIME_SAVED            | NUMBER   | Estimated elapsed parse time saved (in seconds), owing to library cache memory objects being found in a shared pool of the specified size. |
| ESTD_LC_TIME_SAVED_FACTOR     | NUMBER   | Estimated parse time saved factor with respect to the current shared pool size                                                             |
| ESTD_LC_MEMORY_OBJECT_HITS    | NUMBER   | Estimated number of times a library cache memory object was found in a shared pool of the specified size                                   |

## V\$SQL

A SQL statement can map to multiple cursors, because the objects referred to in the cursor can differ from user to user. If there are multiple cursors (child cursors) present, then V\$SQLAREA provides aggregated information for all the cursors.

For looking at individual cursors, V\$SQL can be used. This view contains cursor level details for the SQL. It can be used when trying to locate the session or person responsible for parsing the cursor.

The PLAN\_HASH\_VALUE column contains the numerical representation of the SQL plan for the cursor and can be used for comparing plans. PLAN\_HASH\_VALUE allows you to easily identify whether or not two plans are the same without comparing the two plans line by line.

## V\$SQL\_PLAN

This view provides a way of examining the execution plan for cursors that were executed and are still cached.

Normally, the information in this view is very similar to the output of an `EXPLAIN PLAN` statement. However, `EXPLAIN PLAN` shows a theoretical plan that might be used if this statement were to be executed, whereas `V$SQL_PLAN` contains the actual plan used. The execution plan obtained by the `EXPLAIN PLAN` statement can be different from the execution plan used to execute the cursor, because the cursor might have been compiled with different values of session parameters (for example, `HASH_AREA_SIZE`).

### Uses for V\$SQL\_PLAN Data

- Determining the current execution plan
- Identifying the effect of creating an index on a table
- Finding cursors containing a certain access path (for example, full table scan or index range scan)
- Identifying indexes that are, or are not, selected by the optimizer
- Determining whether the optimizer selects the particular execution plan (for example, nested loops join) expected by the developer

This view can also be used as a key mechanism in plan comparison. Plan comparison can be useful when the following types of changes occur:

- Dropping or creating indexes
- Running the `ANALYZE` statement on the database objects
- Modifying initialization parameter values
- Switching from the rule-based optimizer to the cost-based optimizer
- After upgrading the application or the database to a new release

If previous plans are kept (for example, selected from `V$SQL_PLAN` and stored in permanent Oracle tables for reference), then it is then possible to identify how changes in the performance of a SQL statement can be correlated with changes in the execution plan for that statement.

---



---

**Note:** Oracle Corporation strongly recommends that you use the `DBMS_STATS` package rather than `ANALYZE` to collect optimizer statistics. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. Further, the cost-based optimizer will eventually use only statistics that have been collected by `DBMS_STATS`. See *Oracle9i Supplied PL/SQL Packages and Types Reference* for more information on this package.

However, you must use the `ANALYZE` statement rather than `DBMS_STATS` for statistics collection not related to the cost-based optimizer, such as:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
  - To collect information on freelist blocks
- 
- 

### Useful Columns for `V$SQL_PLAN`

The view contains almost all `PLAN_TABLE` columns, in addition to new columns. The columns that are also present in the `PLAN_TABLE` have the same values:

- `ADDRESS`: Address of the handle to the parent for this cursor
- `HASH_VALUE`: Hash value of the parent statement in the library cache

The two columns `ADDRESS` and `HASH_VALUE` can be used to join with `V$SQLAREA` to add the cursor-specific information.

- `CHILD_NUMBER`: Child cursor number using this execution plan

The columns `ADDRESS`, `HASH_VALUE` and `CHILD_NUMBER` can be used to join with `V$SQL` to add the child cursor specific information.

- `OPERATION`: Name of the internal operation performed in this step; for example, `TABLE ACCESS`
- `OPTIONS`: A variation on the operation described in the `OPERATION` column; for example, `FULL`
- `OBJECT_NODE`: Name of the database link used to reference the object (a table name or view name); for local queries using parallel execution, this column describes the order in which output from operations is consumed
- `OBJECT#`: Object number of the table or the index

- **OBJECT\_OWNER**: Name of the user who owns the schema containing the table or index
- **OBJECT\_NAME**: Name of the table or index
- **OPTIMIZER**: Current mode of the optimizer for the first row in the plan (statement line); for example, `CHOOSE`. In case the operation is a database access (e.g. `TABLE ACCESS`), it tells whether the object is analyzed or not
- **ID**: A number assigned to each step in the execution plan
- **PARENT\_ID**: ID of the next execution step that operates on the output of the current step
- **DEPTH**: The depth (or level) of the operation in the tree; that is, it is not necessary to do a `CONNECT BY` to get the level information generally used to indent the rows from the `PLAN_TABLE` - the root operation (statement) has level 0.
- **POSITION**: Order of processing for operations that all have the same `PARENT_ID`
- **COST**: Cost of the operation as estimated by the optimizer's cost-based approach; for statements that use the rule-based approach, this column is null
- **CARDINALITY**: The estimate, by the cost-based optimizer, of the number of rows produced by the operation
- **BYTES**: The estimate, by the cost-based optimizer, of the number of bytes produced by the operation
- **OTHER\_TAG**: Describes the contents of the `OTHER` column (see [Chapter 9, "Using EXPLAIN PLAN"](#) for values)
- **PARTITION\_START**: The start partition of a range of accessed partition
- **PARTITION\_STOP**: The stop partition of a range of accessed partitions
- **PARTITION\_ID**: The step that has computed the pair of values of the `PARTITION_START` and `PARTITION_STOP` columns
- **OTHER**: Other information that is specific to the execution step that a user may find useful (see [Chapter 9, "Using EXPLAIN PLAN"](#) for values)
- **DISTRIBUTION**: For parallel query, stores the method used to distribute rows from producer query servers to consumer query servers.

- **CPU\_COST**: The CPU cost of the operation as estimated by the optimizer's cost-based approach; for statements that use the rule-based approach, this column is null
- **IO\_COST**: The I/O cost of the operation as estimated by the optimizer's cost-based approach; for statements that use the rule-based approach, this column is null
- **TEMP\_SPACE**: Temporary space usage of the operation (sort or hash-join) as estimated by the optimizer's cost-based approach; for statements that use the rule-based approach, this column is null
- **ACCESS\_PREDICATES**: Predicates used to locate rows in an access structure; for example, start or stop predicates for an index range scan
- **FILTER\_PREDICATES**: Predicates used to filter rows before producing them

The **DEPTH** column replaces the **LEVEL** pseudo-column produced by the **CONNECT BY** operator, which sometimes is used in SQL scripts to help indent the **PLAN\_TABLE** data.

### Join Columns for V\$SQL\_PLAN

The columns **ADDRESS**, **HASH\_VALUE** and **CHILD\_NUMBER** are used to join with **V\$SQL** or **V\$SQLAREA** to fetch the cursor-specific information; for example, **BUFFER\_GETS**, or with **V\$SQLTEXT** to return the full text of the SQL statement.

[Table 24–23](#) lists the join columns for **V\$SQL\_PLAN**.

**Table 24–23** Join Columns for V\$SQL\_PLAN

| Column                            | View       | Joined Column(s)                  |
|-----------------------------------|------------|-----------------------------------|
| ADDRESS, HASH_VALUE               | V\$SQLAREA | ADDRESS, HASH_VALUE               |
| ADDRESS, HASH_VALUE, CHILD_NUMBER | V\$SQL     | ADDRESS, HASH_VALUE, CHILD_NUMBER |
| ADDRESS, HASH_VALUE               | V\$SQLTEXT | ADDRESS, HASH_VALUE               |

### Determining the Optimizer Plan for a SQL Statement

The following statement shows the **EXPLAIN PLAN** for a specified SQL statement. Looking at the plan for a SQL statement is one of the first steps in tuning a SQL

statement. The SQL statement for which to return the plan is identified by the statement's HASH\_VALUE and address.

**Example 24–27** shows a query and sample output from V\$SQL\_PLAN (assumes only one child cursor):

**Example 24–27 Querying V\$SQL\_PLAN**

```
SELECT /* TAG */ count(*)
      FROM employees e, departments d
      WHERE e.department_id = d.department_id;
```

```
      COUNT(*)
-----
          14
```

```
column operation      format a20
column options        format a20
column object_name    format a20
column cost            format a20
column cost            format 9999
SELECT sql_text, address, hash_value
      FROM v$sql
      WHERE sql_text like '%TAG%';
```

```
SQL_TEXT  ADDRESS  HASH_VALUE
-----  -
          82117BEC  171077025
```

```
SELECT sql_text, address, hash_value
      FROM v$sql
      WHERE sql_text LIKE '%TAG%'
```

```
SELECT /* TAG */ count(*)
      FROM employees e, departments d
      WHERE e.department_id = d.department_id
```

```
SQL_TEXT  ADDRESS  HASH_VALUE
-----  -
          82157784  1224822469
```

```
SELECT operation, options, object_name, cost
      FROM v$sql_plan
```



```
WHERE address = '\ 82157784'
AND hash_value = 1224822469;
```

| OPERATION        | OPTIONS | OBJECT_NAME | COST |
|------------------|---------|-------------|------|
| SELECT STATEMENT |         |             | 5    |
| SORT             |         |             |      |
| AGGREGATE        |         |             |      |
| HASH JOIN        |         |             | 5    |
| TABLE ACCESS     | FULL    | DEPARTMENTS | 2    |
| TABLE ACCESS     | FULL    | EMPLOYEES   | 2    |

6 rows selected.

`V$SQL_PLAN` shows the plan for a cursor, not for a SQL statement. The difference is that a SQL statement can have more than one cursor associated with it, with each cursor further identified by a `CHILD_NUMBER`. The following are a few examples of how a SQL statement can result in more than one cursor:

- When the same table name resolves to two separate tables:

```
User1: SELECT * FROM EMPLOYEES;
```

```
User2: SELECT * FROM EMPLOYEES;
```

Where user2 has his own employee table, and user1 uses the table referenced by a public synonym.

- When the environment for user1 differs from user2. For example, if user2 specified the first rows (`ALTER SESSION SET OPTIMIZER_GOAL = FIRST_ROWS`) in their login script, and user1 did not.

If the results of querying `V$SQL_PLAN` for a `HASH_VALUE` and `ADDRESS` result in more than one plan appearing, it is because this SQL statement has more than one child cursor. In this case, for each child cursor (identified by `CHILD_NUMBER`), look at the plan to identify whether they differ significantly.

**See Also:** ["Identifying and Gathering Data on Resource-Intensive SQL"](#) on page 6-3 for information on how to identify SQL statements to tune

## V\$SQL\_PLAN\_STATISTICS

This view provides, for each cached cursor, the execution statistics of each operation in the execution plan.

To view row source statistics in this view, the DBA must set the parameter `STATISTICS_LEVEL` to `ALL`.

**Table 24–24 V\$SQL\_PLAN\_STATISTICS**

| Column                           | Datatype            | Description                                                                                                                                                                                                                       |
|----------------------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>address</code>             | <code>raw(4)</code> | Address of the handle to the parent for this cursor                                                                                                                                                                               |
| <code>hash_value</code>          | <code>number</code> | Hash value of the parent statement in the library cache. The two columns ( <code>address</code> and <code>hash_value</code> ) can be used to join with <code>v\$sqlarea</code> to locate the parent cursor.                       |
| <code>child_number</code>        | <code>number</code> | Child cursor number using this work area. The columns ( <code>address</code> , <code>hash_value</code> , and <code>child_number</code> ) can be used to join with <code>v\$sql</code> to locate the child cursor using this area. |
| <code>operation_id</code>        | <code>number</code> | Number assigned to each step in the execution plan.                                                                                                                                                                               |
| <code>executions</code>          | <code>number</code> | Number of times this cursor has been executed.                                                                                                                                                                                    |
| <code>last_starts</code>         | <code>number</code> | Number of times this operation has been started, <sup>1</sup> during the last execution.                                                                                                                                          |
| <code>starts</code>              | <code>number</code> | Number of times this operation has been started, accumulated over the past executions.                                                                                                                                            |
| <code>last_output_rows</code>    | <code>number</code> | Number of rows produced by the row source, during the last execution.                                                                                                                                                             |
| <code>output_rows</code>         | <code>number</code> | Number of rows produced by the row source, accumulated over the past executions.                                                                                                                                                  |
| <code>last_cr_buffer_gets</code> | <code>number</code> | Number of buffers retrieved in consistent mode, during the last execution. Buffers are usually retrieved in consistent mode for queries.                                                                                          |
| <code>cr_buffer_gets</code>      | <code>number</code> | Number of buffers retrieved in consistent mode, accumulated over the past executions. Buffers are usually retrieved in consistent mode for queries.                                                                               |
| <code>last_cu_buffer_gets</code> | <code>number</code> | Number of buffers retrieved in current mode, during the last execution. Buffers are retrieved in current mode for statements such as <code>INSERT</code> , <code>UPDATE</code> , and <code>DELETE</code> .                        |
| <code>cu_buffer_gets</code>      | <code>number</code> | Number of buffers retrieved in current mode, accumulated over the past executions. Buffers are retrieved in current mode for statements such as <code>INSERT</code> , <code>UPDATE</code> , and <code>DELETE</code> .             |
| <code>last_disk_reads</code>     | <code>number</code> | Number of physical disk reads performed by the operation, during the last execution.                                                                                                                                              |
| <code>disk_reads</code>          | <code>number</code> | Number of physical disk reads performed by the operation, accumulated over the past executions.                                                                                                                                   |

**Table 24–24 (Cont.) V\$SQL\_PLAN\_STATISTICS**

| Column            | Datatype | Description                                                                                           |
|-------------------|----------|-------------------------------------------------------------------------------------------------------|
| last_disk_writes  | number   | Number of physical disk writes performed by the operation, during the last execution.                 |
| disk_writes       | number   | Number of physical disk writes performed by the operation, accumulated over the past executions.      |
| last_elapsed_time | number   | Elapsed time (in microseconds) corresponding to this operation, during the last execution.            |
| elapsed_time      | number   | Elapsed time (in microseconds) corresponding to this operation, accumulated over the past executions. |

<sup>1</sup> An operation can be started several times if it is after an iterator (PARTITION or INLIST), i.e, the number of starts is in fact the number of iterations (number of partitions or number of elements in the in-list).

## V\$SQL\_PLAN\_STATISTICS\_ALL

This table concatenates information from V\$SQL\_PLAN with execution statistics from V\$SQL\_PLAN\_STATISTICS and V\$SQL\_WORKAREA. V\$SQL\_WORKAREA contains memory usage statistics for row sources that use SQL memory (for example, hash-join and sort).

**Table 24–25 V\$SQL\_PLAN\_STATISTICS\_ALL**

| Column       | Datatype     | Description                                                                                                                                                                         |
|--------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| address      | raw(4)       | Address of the handle to the parent for this cursor                                                                                                                                 |
| hash_value   | number       | Hash value of the parent statement in the library cache. The two columns (address and hash_value) can be used to join with v\$sqlarea to add information specific to the cursor.    |
| child_number | number       | Child cursor number using this execution plan. The columns (address, hash_value, and child_number) can be used to join with v\$sql to add information specific to the child cursor. |
| operation    | varchar2(30) | Name of the internal operation performed in this step—for example, TABLE ACCESS.                                                                                                    |
| options      | varchar2(30) | A variation on the operation described in the OPERATION column—for example, FULL.                                                                                                   |

**Table 24–25 (Cont.) V\$SQL\_PLAN\_STATISTICS\_ALL**

| Column          | Datatype     | Description                                                                                                                                                                                                               |
|-----------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| object_node     | varchar2(10) | Name of the database link used to reference the object (a table name or view name). For local queries using parallel execution, this column describes the order in which output from operations is consumed.              |
| object#         | number       | Object number of the table or the index.                                                                                                                                                                                  |
| object_owner    | varchar2(30) | Name of the user who owns the schema containing the table or index.                                                                                                                                                       |
| object_name     | varchar2(30) | Name of the table or index.                                                                                                                                                                                               |
| optimizer       | varchar2(20) | Current mode of the optimizer for the first row in the plan (statement line)—for example, CHOOSE. In case the operation is a database access (for example, TABLE ACCESS), it tells whether the object is analyzed or not. |
| id              | number       | Number assigned to each step in the execution plan.                                                                                                                                                                       |
| parent_id       | number       | ID of the next execution step that operates on the output of the current step.                                                                                                                                            |
| depth           | number       | Depth (or level) of the operation in the tree; it is not necessary to do a CONNECT BY to get the level information, generally used to indent the rows from the plan_table. The root operation (statement) has level 0.    |
| position        | number       | Order of processing for operations that all have the same PARENT_ID.                                                                                                                                                      |
| cost            | number       | Cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null.                                                                          |
| cardinality     | number       | Estimate, by the cost-based optimizer, of the number of rows produced by the operation.                                                                                                                                   |
| bytes           | number       | Estimate, by the cost-based optimizer, of the number of bytes produced by the operation.                                                                                                                                  |
| other_tag       | varchar2(35) | Describes the contents of the OTHER column; see <a href="#">Table 9–2</a> for values.                                                                                                                                     |
| partition_start | varchar2(5)  | Start partition of a range of accessed partitions.                                                                                                                                                                        |
| partition_stop  | varchar2(5)  | Stop partition of a range of accessed partitions.                                                                                                                                                                         |
| partition_id    | number       | Step that has computed the pair of values of the PARTITION_START and PARTITION_STOP columns.                                                                                                                              |

**Table 24–25 (Cont.) V\$SQL\_PLAN\_STATISTICS\_ALL**

| Column              | Datatype       | Description                                                                                                                                                                           |
|---------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| other               | varchar2(4000) | Other information that is specific to the execution step that a user may find useful; see <a href="#">Table 9–2</a> for values.                                                       |
| distribution        | varchar2(20)   | Method used to distribute rows from producer query servers to consumer query servers; see <a href="#">Table 9–3</a> for values.                                                       |
| cpu_cost            | number         | CPU cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null.                                  |
| io_cost             | number         | I/O cost of the operation as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null.                                  |
| temp_space          | number         | Temporary space usage of the operation (sort or hash-join) as estimated by the optimizer's cost-based approach. For statements that use the rule-based approach, this column is null. |
| access_predicates   | varchar2(4000) | Predicates used to locate rows in an access structure; for example, start or stop predicates for an index range scan.                                                                 |
| filter_predicates   | varchar2(4000) | Predicates used to filter rows before producing them.                                                                                                                                 |
| executions          | number         | Number of times this cursor has been executed.                                                                                                                                        |
| last_starts         | number         | Number of times this operation has been started, during the last execution.                                                                                                           |
| starts              | number         | Number of times this operation has been started, accumulated over the past executions.                                                                                                |
| last_output_rows    | number         | Number of rows produced by the row source, during the last execution.                                                                                                                 |
| output_rows         | number         | Number of rows produced by the row source, accumulated over the past executions.                                                                                                      |
| last_cr_buffer_gets | number         | Number of buffers retrieved in consistent mode, during the last execution. Buffers are usually retrieved in consistent mode for queries.                                              |
| cr_buffer_gets      | number         | Number of buffers retrieved in consistent mode, accumulated over the past executions. Buffers are usually retrieved in consistent mode for queries.                                   |
| last_cu_buffer_gets | number         | Number of buffers retrieved in current mode, during the last execution. Buffers are retrieved in current mode for statements such as INSERT, UPDATE, and DELETE.                      |

**Table 24–25 (Cont.) V\$SQL\_PLAN\_STATISTICS\_ALL**

| Column                 | Datatype     | Description                                                                                                                                                                                        |
|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cu_buffer_gets         | number       | Number of buffers retrieved in current mode, accumulated over the past executions. Buffers are retrieved in current mode for statements such as INSERT, UPDATE, and DELETE.                        |
| last_disk_reads        | number       | Number of physical disk reads performed by the operation, during the last execution.                                                                                                               |
| disk_reads             | number       | Number of physical disk reads performed by the operation, accumulated over the past executions.                                                                                                    |
| last_disk_writes       | number       | Number of physical disk writes performed by the operation, during the last execution.                                                                                                              |
| disk_writes            | number       | Number of physical disk writes performed by the operation, accumulated over the past executions.                                                                                                   |
| last_elapsed_time      | number       | Elapsed time (in microseconds) corresponding to this operation, during the last execution.                                                                                                         |
| elapsed_time           | number       | Elapsed time (in microseconds) corresponding to this operation, accumulated over the past executions.                                                                                              |
| policy                 | varchar2(10) | Sizing policy for this work area. Values are either MANUAL or AUTO.                                                                                                                                |
| estimated_optimal_size | number       | Estimated size (in KB) required by this work area to execute the operation completely in memory (optimal execution). This is either derived from optimizer statistics or from previous executions. |
| estimated_onepass_size | number       | Estimated size (in KB) required by this work area to execute the operation in a single pass. This is either derived from optimizer statistics or from previous executions.                         |
| last_memory_used       | number       | Memory size, in KB, used by this work area during the last execution of the cursor.                                                                                                                |
| last_execution         | varchar(10)  | Whether this work area ran using OPTIMAL, ONE PASS, or under ONE PASS memory requirement (that is, MULTI-PASS), during the last execution of the cursor.                                           |
| last_degree            | number       | Degree of parallelism used, during the last execution of the cursor.                                                                                                                               |
| total_executions       | number       | Number of times this work area was active.                                                                                                                                                         |
| optimal_executions     | number       | Number of times this work area ran in optimal mode.                                                                                                                                                |
| onepass_executions     | number       | Number of times this work area ran in one pass mode.                                                                                                                                               |

**Table 24–25 (Cont.) V\$SQL\_PLAN\_STATISTICS\_ALL**

| Column                | Datatype | Description                                                               |
|-----------------------|----------|---------------------------------------------------------------------------|
| multiplies_executions | number   | Number of times this work area ran under the one-pass memory requirement. |
| active_time           | number   | Time the work area has been active (in microseconds).                     |
| max_tempseg_size      | number   | Maximum temporary space used.                                             |
| tempseg_size          | number   | Temporary space used.                                                     |

## V\$SQLAREA

This view keeps track of all the shared cursors present in the shared pool. It has one row for every SQL statement present in the shared pool. It is an invaluable view for finding the resource usage of a SQL statement.

### Information columns in V\$SQLAREA

- **HASH\_VALUE:** Hash value of the SQL statement
- **ADDRESS:** SGA address for the SQL statement

These two columns are used to identify the SQL statement. Sometimes, two different statements could hash to the same value. In such cases, it is necessary to use the address along with the hash\_value.

- **PARSING\_USER\_ID:** User who parsed the first cursor for the statement
- **VERSION\_COUNT:** Number of cursors for the statement
- **KEPT\_VERSIONS:** Cursors of the statement pinned using `DBMS_SHARED_POOL.KEEP()`
- **SHARABLE\_MEMORY:** Total shared memory used by the cursor
- **PERSISTENT\_MEMORY:** Total persistent memory used by the cursor
- **RUNTIME\_MEMORY:** Total runtime memory used by the cursor
- **SQL\_TEXT:** Up to first 1000 characters of SQL statement
- **MODULE, ACTION:** Information about the session parsing the first cursor if set using `DBMS_APPLICATION_INFO`

### Other Useful Columns in V\$SQLAREA

These columns get incremented with each execution of the statement.

- BUFFER\_GETS: Number of logical reads for this statement
- DISK\_READS: Number of physical reads for this statement
- SORTS: Number of sorts for this statement
- CPU\_TIME: CPU time used for parsing and executing this statement
- ELAPSED\_TIME: Elapsed time for parsing and executing this statement
- PARSE\_CALLS: Number of parse calls (hard and soft) for this statement
- EXECUTIONS: Number of times this statement was executed
- INVALIDATIONS: Number of times the cursors for this statement have been invalidated
- LOADS: Number of loads (and reloads) for this statement
- ROWS\_PROCESSED: Total number of rows this statement returns

### Join Columns in V\$SQLAREA

[Table 24–26](#) lists the join columns for V\$SQLAREA.

**Table 24–26** Join Columns for V\$SQLAREA

| Column                 | View                                  | Joined Column(s)               |
|------------------------|---------------------------------------|--------------------------------|
| HASH_VALUE,<br>ADDRESS | V\$SESSION                            | SQL_HASH_VALUE,<br>SQL_ADDRESS |
| HASH_VALUE,<br>ADDRESS | V\$SQLTEXT, V\$SQL,<br>V\$OPEN_CURSOR | HASH_VALUE,<br>ADDRESS         |
| SQL_TEXT               | V\$DB_OBJECT_CACHE                    | NAME                           |

### Example 24–28 Finding Resource-intensive SQL

There are several costs you can use:

- Total logical I/O (LIO), LIO for each execution
- Total physical I/O (PIO), PIO for each execution
- PIO/LIO (poor cache hit ratio)
- parse\_calls, parse\_calls for each executions

```
SELECT hash_value, executions, buffer_gets, disk_reads, parse_calls
FROM V$SQLAREA
WHERE buffer_gets > 10000000
```



```

OR disk_reads > 1000000
ORDER BY buffer_gets + 100*disk_reads DESC;

```

| HASH_VALUE | EXECUTIONS | BUFFER_GETS | DISK_READS | PARSE_CALLS |
|------------|------------|-------------|------------|-------------|
| 2676594883 | 126        | 7583140     | 6199113    | 126         |
| 4074144966 | 126        | 7264362     | 6195433    | 49          |
| 228801498  | 136        | 236116544   | 2371187    | 136         |
| 360282550  | 5467       | 21102603    | 4476317    | 2355        |
| 1559420740 | 201        | 8197831     | 4537591    | 39          |
| 3213702248 | 28039654   | 364516977   | 44         | 131         |
| 1547710012 | 865        | 7579025     | 3337735    | 865         |
| 3000880481 | 4481       | 3676546     | 2212658    | 2885        |
| 1398193708 | 4946       | 73018658    | 1515257    | 1418        |
| 1052917712 | 8342025    | 201246652   | 38240      | 327462      |
| 371697988  | 7          | 74380777    | 862611     | 7           |
| 1514306888 | 3922461    | 29073852    | 1223482    | 268         |
| 1848522009 | 1          | 1492281     | 1483635    | 1           |
| 1478599096 | 28042103   | 140210513   | 594        | 164         |
| 226079402  | 21473      | 22121577    | 1034787    | 4484        |
| 478652562  | 4468       | 21669366    | 1020370    | 4438        |
| 2054874295 | 73520      | 118272694   | 29987      | 73520       |

---

**Note:** If a statement is executing for the first time on the system and responsible for large fraction of the current resource usage, then this statement does not find that statement, because the BUFFER\_GETS and DISK\_READS statistics do not get updated until the statement finishes execution.

---

#### **Example 24–29 Finding Resources Used by a SQL Statement**

```

SELECT hash_value, buffer_gets, disk_reads, executions, parse_calls
FROM V$SQLAREA
WHERE hash_value = 228801498
AND address = hextoraw('CBD8E4B0');

```

| HASH_VALUE | BUFFER_GETS | DISK_READS | EXECUTIONS | PARSE_CALLS |
|------------|-------------|------------|------------|-------------|
| 228801498  | 236116544   | 2371187    | 136        | 136         |

## V\$SQLTEXT

This view contains the complete SQL text for the SQL statements in the shared pool.

---



---

**Note:** V\$SQLAREA only contains only the first 1000 characters.

---



---

### Useful Columns for V\$SQLTEXT

- HASH\_VALUE: Hash value for the SQL statement
- ADDRESS: Address of the SQL statement cursor in SGA
- SQL\_TEXT: Statement text in 64 character chunks
- PIECE: Ordering information for the SQL statement pieces

### Join Columns for V\$SQLTEXT

Table 24-27 lists the join columns for V\$SQLTEXT.

**Table 24-27** Join Columns for V\$SQLTEXT

| Column                 | View               | Joined Column(s)               |
|------------------------|--------------------|--------------------------------|
| HASH_VALUE,<br>ADDRESS | V\$SQL, V\$SESSION | HASH_VALUE,<br>ADDRESS         |
| HASH_VALUE.<br>ADDRESS | V\$SESSION         | SQL_HASH_VALUE.<br>SQL_ADDRESS |

### Example 24-30 Finding the SQL Statement for a Hash Value

```
SELECT sql_text
   FROM V$SQLTEXT
  WHERE hash_value = 228801498
 ORDER BY piece;
```

SQL\_TEXT

```
-----
select dbsu.primary_flag, i.site_use_code, i.rowid
from ra_customers dbc, ra_addresses dbad, ra_site_uses dbsu, ra_customers_
interface i
where ((((((i.orig_system_customer_ref=dbc.orig_system_reference and
dbad.address_id=dbsu.address_id) and i.site_use_code=dbsu.site_use_code) and
dbsu.status='A') and dbad.customer_id=dbc.customer_id) and i.request_id=:b0) and
nvl(i.validated_flag,'N')<>'Y') and ((i.primary_site_use_flag='Y' and
dbsu.primary_flag='Y') or dbsu.site_use_code in ('STMTS','DUN','LEGAL'))
group by dbsu.primary_flag,i.orig_system_customer_ref,i.site_use_code,i.insert_
update_flag,i.rowid
```

## V\$STATISTICS\_LEVEL

V\$STATISTICS\_LEVEL lists the status of the statistics or advisories controlled by the STATISTICS\_LEVEL initialization parameter. Each row of V\$STATISTICS\_LEVEL represents one of these statistics or advisories.

**Table 24–28 V\$STATISTICS\_LEVEL View**

| Column               | Datatype       | Description                                                                                                                                                                                                                            |
|----------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATISTICS_NAME      | VARCHAR2(64)   | Name of the statistics/advisory.                                                                                                                                                                                                       |
| DESCRIPTION          | VARCHAR2(4000) | Description of what the statistics/advisory does and what it can be used for.                                                                                                                                                          |
| SESSION_STATUS       | VARCHAR2(8)    | ENABLED DISABLED. Status of the statistics/advisory for this session.                                                                                                                                                                  |
| SYSTEM_STATUS        | VARCHAR2(8)    | ENABLED DISABLED. Status of the statistics/advisory system-wide.                                                                                                                                                                       |
| ACTIVATION_LEVEL     | VARCHAR2(7)    | BASIC TYPICAL ALL. What level of STATISTICS_LEVEL enables this statistics/advisory?                                                                                                                                                    |
| STATISTICS_VIEW_NAME | VARCHAR2(64)   | If there is a single view externalizing this statistics/advisory, the name of that view. If there is no such a view, this column is empty. If there are multiple views involved, the DESCRIPTION column should mention the view names. |
| SESSION_SETTABLE     | VARCHAR2(3)    | YES NO. Whether this statistics/advisory can be set at the session level.                                                                                                                                                              |

## V\$SYSSTAT

V\$SYSSTAT stores instance-wide statistics on resource usage, cumulative since the instance was started.

Similar to V\$SESSTAT, this view stores the following types of statistics:

- A count of the number of times an action occurred (user commits)
- A running total of volumes of data generated, accessed, or manipulated (redo size)
- If TIMED\_STATISTICS is true, then the cumulative time spent performing some actions (CPU used by this session)

### Useful Columns in V\$SYSSTAT

- STATISTIC#: Identifier for the statistic
- NAME: Statistic name
- VALUE: Resource usage

The value for each statistic stores the resource usage for that statistic since instance startup. The following are sample column values for the statistic `execute count`.

**Table 24–29 Useful Columns in V\$SYSSTAT**

| Statistic# | Name                       | Value      |
|------------|----------------------------|------------|
| 215        | <code>execute count</code> | 19,003,070 |

---

**Note:** The `STATISTIC#` for a statistic can change between releases. Do not rely on `STATISTIC#` to remain constant. Instead, use the statistic `NAME` column to query the `VALUE`.

---

### Uses for V\$SYSSTAT Data

The data in this view is used for monitoring system performance. Derived statistics, such as the buffer cache hit ratio and soft parse ratio, are computed from `V$SYSSTAT` data.

Data in this view is also used for monitoring system resource usage and how the system's resource usage changes over time. As with most performance data, examine the system's resource usage over an interval. To do this, take a snapshot of the data within the view at the beginning of the interval and another at the end. The difference in the values (end value - begin value) for each statistic is the resource used during the interval. This is the methodology used by Oracle tools such as Statspack and `BSTAT/ESTAT`.

In order to compare one interval's data with another, the data can be normalized (for each transaction, for each execution, for each second, or for each logon). Normalizing the data on both workloads makes identifying the variances between the two workloads easier. This type of comparison is especially useful after patches have been applied, applications have been upgraded, or simply over time to see how increases in user population or data growth affects the resource usage.

You can also use `V$SYSSTAT` data to examine the resource consumption of contended-for resources that were identified by querying the `V$SYSTEM_EVENT` view.

## Useful Statistics for V\$SYSSTAT

This section describes some of the V\$SYSSTAT statistics that are most useful during tuning, along with an explanation of the statistic. This list is in alphabetical order.

**See Also:** *Oracle9i Database Reference* for a complete list of statistics and their description

### Key Database Usage Indicators

- `CPU used by this session`: The total amount of CPU used by all sessions, excluding background processes. This unit for this statistic is hundredths of a second. Calls that complete in less than 10ms are rounded up to this unit.
- `db block changes`: The number of changes made to database blocks in the SGA that were part of an insert, update, or delete operation. This statistic is a rough indication of total database work. On a for each transaction level, this statistic indicates the rate at which buffers are being dirtied.
- `execute count`: The total number of SQL statement executions (including recursive SQL).
- `logons current`: Sessions currently connected to the instance. When using two snapshots across an interval, an average value (rather than the difference) should be used.
- `logons cumulative`: The total number of logons since the instance started. To determine the number of logons in a particular period, subtract the end value from the begin value. A useful derived statistic is to divide the number of connections between a begin and end time, and divide this by the number of seconds the interval covered. This gives the logon rate. Optimally, there should be no more than two logons each second. To contrast, a logon rate of 50 a second is considered very high. Applications that continually connect and disconnect from the database (for example, once for each transaction) do not scale well.
- `parse count (hard)`: The number of parse calls that resulted in a miss in the shared pool. A hard parse occurs when a SQL statement is executed and the SQL statement is either not in the shared pool, or it is in the shared pool but it cannot be shared because part of the metadata for the two SQL statements is different. This can happen if a SQL statement is textually identical to a preexisting SQL statement, but the tables referred to in the two statements resolve to physically different tables. A hard parse is a very expensive operation in terms of CPU and resource use (for example, latches), because it requires

Oracle to allocate memory within the shared pool, then determine the execution plan before the statement can be executed.

- `parse count (total)`: The total number of parse calls, both hard and soft. A soft parse occurs when a session executes a SQL statement, and the statement is already in the shared pool and can be used. For a statement to be used (that is, shared) all data pertaining to the existing SQL statement (including data such as the optimizer execution plan) must be equally applicable to the current statement being issued. These two statistics are used to calculate the soft-parse ratio.
- `parse time cpu`: Total CPU time spent parsing in hundredths of a second. This includes both hard and soft parses.
- `parse time elapsed`: The total elapsed time for the parse call to complete.
- `physical reads`: The number of blocks read from the operating system. It includes physical reads into the SGA buffer cache (a buffer cache miss) and direct physical reads into the PGA (for example, during direct sort operations). This statistic is *not* the number of I/O requests.
- `physical writes`: The number of database blocks written from the SGA buffer cache to disk by DBWR and from the PGA by processes performing direct writes.
- `redo log space requests`: The number of times a server process waited for space in the redo logs, typically because a log switch is needed.
- `redo size`: The total amount of redo generated (and hence written to the log buffer), in bytes. This statistic (normalized over seconds or over transactions) is a good indicator of update activity.
- `session logical reads`: The number of logical read requests that can be satisfied in the buffer cache or by a physical read.
- `sorts (memory)` and `sorts (disk)`: `sorts (memory)` is the number of sort operations that fit inside the `SORT_AREA_SIZE` (and hence did not require an on disk sort). `sorts (disk)` is the number of sort operations that were larger than `SORT_AREA_SIZE` and had to use space on disk to complete the sort. These two statistics are used to compute the in-memory sort ratio.
- `sorts (rows)`: The total number of rows sorted. This statistic can be divided by the 'sorts (total)' statistic to determine rows for each sort. It is an indicator of data volumes and application characteristics.

- `table fetch by rowid`: The number of rows returned using `ROWID` (due to index access or because a SQL statement of the form "where rowid = &rowid" was issued).
- `table scans (rows gotten)`: The total number of rows processed during full table scans.
- `table scans (blocks gotten)`: The number of blocks scanned during full table scans, excluding those for split rows.
- `user commits + user rollbacks`: This provides the total number of transactions on the system. This number is used as the divisor when calculating the ratios for each transaction for other statistics. For example, to calculate the number of logical reads for each transaction, use the following formula:  

$$\text{session logical reads} / (\text{user commits} + \text{user rollbacks}).$$

### Notes on Physical I/O

A physical read as reported by Oracle might not result in an actual physical disk I/O operation. This is possible because most operating systems have an operating system files system cache where the block might be present. Alternatively, the block might also be present in disk or controller level cache, again avoiding an actual I/O. A physical read as reported by Oracle merely indicates that the required block was not in the buffer cache (or in the case of a direct read operation, was required to be read into private memory).

### Instance Efficiency Ratios From V\$SYSSTAT Statistics

The following are typical instance efficiency ratios calculated from `V$SYSSTAT` data. Each ratio's computed value should all be as close as possible to 1:

**Buffer cache hit ratio:** This is a good indicator of whether the buffer cache is too small.

$$1 - ((\text{physical reads} - \text{physical reads direct} - \text{physical reads direct (lob)}) / \text{session logical reads})$$

**Soft parse ratio:** This shows whether there are many hard parses on the system. The ratio should be compared to the raw statistics to ensure accuracy. For example, a soft parse ratio of 0.2 typically indicates a high hard parse rate. However, if the total number of parses is low, then the ratio should be disregarded.

$$1 - (\text{parse count (hard)} / \text{parse count (total)})$$

**In-memory sort ratio:** This shows the proportion of sorts that are performed in memory. Optimally, in an operational (OLTP) system, most sorts are small and can be performed solely as in-memory sorts.

```
sorts (memory) / ( sorts (memory) + sorts (disk) )
```

**Parse to execute ratio:** In an operational environment, optimally a SQL statement should be parsed once and executed many times.

```
1 - (parse count/execute count)
```

**Parse CPU to total CPU ratio:** This shows how much of the total CPU time used was spent on activities other than parsing. When this ratio is low, the system is performing too many parses.

```
1 - (parse time cpu / CPU used by this session)
```

**Parse time CPU to parse time elapsed:** Often, this can indicate latch contention. The ratio calculates whether the time spent parsing is allocated to CPU cycles (that is, productive work) or whether the time spent parsing was not spent on CPU cycles. Time spent parsing not on CPU cycles usually indicates that the time was spent sleeping due to latch contention.

```
parse time cpu / parse time elapsed
```

### **Load Profile Data from V\$SYSSTAT Statistics**

To determine the load profile of the system, normalize the following statistics over seconds and over transactions: `logons cumulative`, `parse count (total)`, `parse count (hard)`, `executes`, `physical reads`, `physical writes`, `block changes`, and `redo size`.

The normalized data can be examined to see if the 'rates' are high, or it can be compared to another baseline data set to identify how the system profile is changing over time. For example, block changes for each transaction is calculated by the following:

```
db block changes / ( user commits + user rollbacks )
```

Additional computed statistics that measure load include the following:

- **Blocks changed for each read:**

This shows the proportion of block changes to block reads. It is an indication of whether the system is predominantly read only or whether the system performs many data changes (inserts/updates/deletes).



db block changes / session logical reads

- Rows for each sort:

sorts (rows) / ( sorts (memory) + sorts (disk) )

### Join Columns for V\$SYSSTAT

Table 24–30 lists the join columns for V\$SYSSTAT.

**Table 24–30 Join Columns for V\$SYSSTAT**

| Column     | View        | Joined Column(s) |
|------------|-------------|------------------|
| STATISTIC# | V\$STATNAME | STATISTIC#       |

## V\$SYSTEM\_EVENT

This view is a summary of waits for an event by an instance. While V\$SESSION\_WAIT shows the current waits on the system, V\$SYSTEM\_EVENT provides a summary of all the event waits on the instance since it started. It is useful to get a historical picture of waits on the system. By taking two snapshots and doing the delta on the waits, you can determine the waits on the system in a given time interval.

### Useful Columns for V\$SYSTEM\_EVENT

- EVENT: Name of the wait event
- TOTAL\_WAITS: Total number of waits for this event
- TIME\_WAITED: Total time waited for this event (in hundredths of a second)
- AVERAGE\_WAIT: Average amount of time waited for this event by this session (in hundredths of a second)
- TOTAL\_TIMEOUTS: Number of times the wait timed out

#### Example 24–31 Finding the Total Waits on the System

```
SELECT event, total_waits waits, total_timeouts timeouts,
       time_waited total_time, average_wait avg
FROM V$SYSTEM_EVENT
ORDER BY 4 DESC;
```

```
EVENT                               WAITS  TIMEOUTS  TOTAL_TIME  AVG
-----
```

## Description of Dynamic Performance Views

---

|                                  |           |           |            |          |
|----------------------------------|-----------|-----------|------------|----------|
| SQL*Net message from client      | 112079628 | 0         | 8622695365 | 76.93    |
| virtual circuit status           | 83559794  | 1168000   | 4275791401 | 51.17    |
| rdbms ipc message                | 131463191 | 115900505 | 2865926648 | 21.80    |
| dispatcher timer                 | 311975975 | 168152330 | 2296760866 | 7.36     |
| PX Idle Wait                     | 7198490   | 7198559   | 1439690729 | 199.99   |
| pmon timer                       | 939711    | 939639    | 287866277  | 306.33   |
| smon timer                       | 9892      | 9114      | 287627013  | 29076.73 |
| lock manager wait for remote mes | 72001548  | 71967858  | 287526387  | 3.99     |
| db file sequential read          | 29419894  | 0         | 32395392   | 1.10     |
| PL/SQL lock timer                | 19725     | 19688     | 29702609   | 1505.83  |
| log file sync                    | 7055611   | 86        | 9550819    | 1.35     |
| log file parallel write          | 7184801   | 4         | 8123534    | 1.13     |
| SQL*Net more data from client    | 991402    | 0         | 3543149    | 3.57     |
| db file parallel write           | 727317    | 0         | 3012928    | 4.14     |
| control file parallel write      | 950531    | 0         | 1975646    | 2.07     |
| log file sequential read         | 1162465   | 0         | 813715     | 0.69     |
| enqueue                          | 9975      | 7692      | 423191     | 42.42    |
| direct path read                 | 453873    | 0         | 298944     | 0.65     |
| db file scattered read           | 347172    | 0         | 292875     | 0.84     |
| row cache lock                   | 472207    | 25        | 169365     | 0.35     |
| direct path write                | 124323    | 0         | 132075     | 1.06     |
| buffer busy due to global cache  | 148122    | 0         | 122381     | 0.82     |
| SQL*Net more data to client      | 17171954  | 52        | 101762     | 0.00     |
| db file parallel read            | 68849     | 0         | 100842     | 1.46     |
| DFS lock handle                  | 18615     | 1080      | 97651      | 5.24     |
| SQL*Net message to client        | 112079756 | 0         | 77604      | 0.00     |
| control file sequential read     | 65793     | 0         | 62560      | 0.95     |
| buffer busy waits                | 132402    | 97        | 60351      | 0.45     |
| latch free                       | 67675     | 57975     | 58365      | 0.86     |
| log file switch completion       | 1449      | 24        | 34244      | 23.63    |
| db file single write             | 10868     | 0         | 25518      | 2.34     |
| SQL*Net break/reset to client    | 19130     | 0         | 9387       | 0.49     |
| LGWR wait for redo copy          | 120199    | 356       | 8613       | 0.07     |
| global cache lock busy           | 4447      | 0         | 7574       | 1.70     |
| undo segment extension           | 5363841   | 5363828   | 6375       | 0.00     |
| log file single write            | 2143      | 0         | 6267       | 2.92     |
| refresh controlfile command      | 2644      | 0         | 4837       | 1.82     |
| library cache load lock          | 49        | 10        | 3859       | 78.75    |
| file open                        | 178566    | 0         | 2930       | 0.01     |
| switch logfile command           | 100       | 0         | 2468       | 24.68    |
| library cache pin                | 9261      | 1         | 1716       | 0.18     |
| pipe get                         | 9         | 3         | 1460       | 162.22   |
| rdbms ipc reply                  | 10296     | 0         | 846        | 0.08     |
| wait for gms registration        | 32        | 32        | 672        | 21.00    |
| process startup                  | 43        | 2         | 662        | 15.39    |

|                                    |      |     |     |        |
|------------------------------------|------|-----|-----|--------|
| file identify                      | 5438 | 0   | 584 | 0.10   |
| control file single write          | 332  | 0   | 475 | 1.43   |
| Null event                         | 17   | 17  | 409 | 24.05  |
| log buffer space                   | 18   | 0   | 209 | 11.61  |
| wait for lock db to unfreeze       | 1    | 1   | 199 | 199.00 |
| local write wait                   | 11   | 0   | 44  | 4.00   |
| LMON wait for LMD to inherit commu | 1    | 1   | 10  | 10.00  |
| wait for lock db to become frozen  | 2    | 2   | 3   | 1.50   |
| instance state change              | 2    | 0   | 0   | 0.00   |
| global cache bg acks               | 2    | 0   | 0   | 0.00   |
| buffer deadlock                    | 141  | 141 | 0   | 0.00   |

**To find the bottlenecks:**

- Statspack lists idle events at the end.
- Examine the time spent waiting for different events.
- Examine the average time for each wait also, because some waits (like log file switch completion) might happen only periodically, but cause a big performance hit when they happen.

**V\$UNDOSTAT**

This view monitors how undo space and transactions are executed in the current instance. Statistics for undo space consumption, transaction concurrency, and length of queries in the instance are available.

**Useful Columns for V\$UNDOSTAT**

- Endtime: End time for each ten minute interval
- UndoBlocksUsed: Total number of undo blocks consumed
- TxnConcurrency: Maximum number of transactions executed concurrently
- TxnTotal: Total number of transactions executed within the interval
- QueryLength: Maximum length of queries, in seconds executed in the instance
- ExtentsStolen: Number of times an undo extent must be transferred from one undo segment to another within the interval
- SSTooOldError: Number of 'Snapshot Too Old' errors that occurred within the interval
- UNDOTSN: undo tablespaces in service during each time period

The first row of the view shows statistics for the current time interval. Each subsequent row represents a ten minute interval. There is a total of 144 rows, spanning a 24 hour cycle.

### **Example 24–32 Querying V\$UNDOSTAT**

This example shows how undo space is consumed in the system for the previous 24 hours from the time 16:07.

```
SELECT * FROM V$UNDOSTAT;
```

| End-Time | UndoBlocks | TxnConcrcy | TxnTotal | QueryLen | ExtentsStolen | SSTooOldError |
|----------|------------|------------|----------|----------|---------------|---------------|
| 16:07    | 252        | 15         | 1511     | 25       | 2             | 0             |
| 16:00    | 752        | 16         | 1467     | 150      | 0             | 0             |
| 15:50    | 873        | 21         | 1954     | 45       | 4             | 0             |
| 15:40    | 1187       | 45         | 3210     | 633      | 20            | 1             |
| 15:30    | 1120       | 28         | 2498     | 1202     | 5             | 0             |
| 15:20    | 882        | 22         | 2002     | 55       | 0             | 0             |

Among the statistics collected, you see that the peak undo consumption happened at the interval of (15:30, 15:40). 1187 undo blocks were consumed in 10 minutes (or about two blocks a second). Also, the highest transaction concurrency occurred during that same period with 45 transactions executing at the same time. The longest query (1202 seconds) was executed (and ended) in the period (15:20, 15:30). Note that the query actually was started in the interval (15:00, 15:10) and continued until around 15:20.

## **V\$WAITSTAT**

This view keeps a summary all buffer waits since instance startup. It is useful for breaking down the waits by class if you see a large number of buffer busy waits on the system.

### **Useful Columns for V\$WAITSTAT**

- `class`: Class of block (data segment header, undo segment header, data block)
- `waits`: Number of waits for this class of blocks
- `time`: Total time waited for this class of block

### **Reasons for Waits**

The following are possible reasons for waits:

- Undo segment header: not enough rollback segments
- Data segment header/freelist: freelist contention
- Data block
- Large number of CR clones for the buffer
- Range scans on indexes with large number of deletions
- Full table scans on tables with large number of deleted rows
- Blocks with high concurrency

**See Also:** [Chapter 22, "Instance Tuning"](#) for more information on wait events



---

---

# Schemas Used in Performance Examples

The tables described in this appendix are used in various examples in this book. The statistics are from representative systems.

This appendix discusses the following:

- [PER\\_ALL\\_PEOPLE\\_F Table](#)
- [RA\\_CUSTOMERS Table](#)
- [SO\\_HEADERS\\_ALL and SO\\_HEADERS Tables](#)
- [MTL\\_SYSTEM\\_ITEMS Table](#)
- [SO\\_LINES\\_ALL and SO\\_LINES Tables](#)

---

---

**Note:** These schemas are used in examples in several chapters, including [Chapter 9, "Using EXPLAIN PLAN"](#).

---

---

## PER\_ALL\_PEOPLE\_F Table

This table stores data for employees on the system. For large corporations, it is common to have 10,000 to 30,000 rows in this table. The unique key is a concatenated index, but `person_id` by itself is quite selective also. Other selective columns are `employee_number` and `full_name`.

The following are indexes on the table:

| Unique | Index Name       | Column Name                                             |
|--------|------------------|---------------------------------------------------------|
| NO     | PER_PEOPLE_F_FK1 | BUSINESS_GROUP_ID                                       |
| NO     | PER_PEOPLE_F_FK2 | PERSON_TYPE_ID                                          |
| NO     | PER_PEOPLE_F_N50 | LAST_NAME                                               |
| NO     | PER_PEOPLE_F_N51 | EMPLOYEE_NUMBER                                         |
| NO     | PER_PEOPLE_F_N52 | APPLICANT_NUMBER                                        |
| NO     | PER_PEOPLE_F_N53 | NATIONAL_IDENTIFIER                                     |
| NO     | PER_PEOPLE_F_N54 | FULL_NAME                                               |
| YES    | PER_PEOPLE_F_PK  | PERSON_ID<br>EFFECTIVE_START_DATE<br>EFFECTIVE_END_DATE |

## RA\_CUSTOMERS Table

This table has a row for every customer in the system. For large companies, this table has several hundred thousand rows. The primary key is `customer_id`. Other selective columns are the following:

- `Customer_number`
- `Customer_name`
- `Orig_system_reference` - tracks the customer identifier for customers imported from another system

The following are indexes on the table:

| Unique | Index Name      | Column Name           |
|--------|-----------------|-----------------------|
| NO     | RA_CUSTOMERS_N1 | CUSTOMER_NAME         |
| NO     | RA_CUSTOMERS_N2 | CREATION_DATE         |
| NO     | RA_CUSTOMERS_N3 | CUSTOMER_KEY          |
| NO     | RA_CUSTOMERS_N4 | JGZZ_FISCAL_CODE      |
| YES    | RA_CUSTOMERS_U1 | CUSTOMER_ID           |
| YES    | RA_CUSTOMERS_U2 | ORIG_SYSTEM_REFERENCE |
| YES    | RA_CUSTOMERS_U3 | CUSTOMER_NUMBER       |



## SO\_HEADERS\_ALL and SO\_HEADERS Tables

This table has a row for every order on the system. For large companies, it is common to have several million rows in this table. The primary key is `header_id`, and there is another unique key on (`order_number`, `order_type_id`). Other selective columns are the following:

- `customer_id` (the customer placing the order)
- `purchase_order_num` (the purchase order for billing)
- `original_system_reference` (tracks the order identifier for orders imported from other systems)

The following are indexes on the table:

| Unique | Index Name     | Column Name                                              |
|--------|----------------|----------------------------------------------------------|
| NO     | SO_HEADERS_N1  | CUSTOMER_ID                                              |
| NO     | SO_HEADERS_N10 | WH_UPDATE_DATE                                           |
| NO     | SO_HEADERS_N2  | OPEN_FLAG                                                |
| NO     | SO_HEADERS_N3  | PURCHASE_ORDER_NUM                                       |
| NO     | SO_HEADERS_N4  | INVOICE_TO_SITE_USE_ID                                   |
| NO     | SO_HEADERS_N5  | ORIGINAL_SYSTEM_REFERENCE                                |
| NO     | SO_HEADERS_N6  | S1                                                       |
| NO     | SO_HEADERS_N7  | S4                                                       |
| NO     | SO_HEADERS_N8  | S6                                                       |
| NO     | SO_HEADERS_N9  | ORIGINAL_SYSTEM_REFERENCE<br>ORIGINAL_SYSTEM_SOURCE_CODE |
| YES    | SO_HEADERS_U1  | HEADER_ID                                                |
| YES    | SO_HEADERS_U2  | ORDER_NUMBER<br>ORDER_TYPE_ID                            |

## MTL\_SYSTEM\_ITEMS Table

This table is the parts master for `so_lines_all`. It has a row for every part in every organization. The primary key is `inventory_item_id`, `organization_id`.

The following are indexes on the table:

| Unique | Index Name          | Column Name                 |
|--------|---------------------|-----------------------------|
| NO     | MTL_SYSTEM_ITEMS_N1 | ORGANIZATION_ID<br>SEGMENT1 |
| NO     | MTL_SYSTEM_ITEMS_N2 | ORGANIZATION_ID             |

|     |                     | DESCRIPTION                                  |
|-----|---------------------|----------------------------------------------|
| NO  | MTL_SYSTEM_ITEMS_N3 | INVENTORY_ITEM_STATUS_CODE                   |
| NO  | MTL_SYSTEM_ITEMS_N4 | ORGANIZATION_ID<br>AUTO_CREATED_CONFIG_FLAG  |
| NO  | MTL_SYSTEM_ITEMS_N5 | WH_UPDATE_DATE                               |
| NO  | MTL_SYSTEM_ITEMS_N6 | ITEM_CATALOG_GROUP_ID<br>CATALOG_STATUS_FLAG |
| NO  | MTL_SYSTEM_ITEMS_N7 | PRODUCT_FAMILY_ITEM_ID<br>ORGANIZATION_ID    |
| NO  | MTL_SYSTEM_ITEMS_N8 | SEGMENT1<br>SEGMENT2<br>SEGMENT3             |
| YES | MTL_SYSTEM_ITEMS_U1 | INVENTORY_ITEM_ID<br>ORGANIZATION_ID         |

## SO\_LINES\_ALL and SO\_LINES Tables

This table has a row for every order line on the system. It joins to the `so_headers_all` table using `header_id`. Because an order has 10 to 12 lines, this table is 10 to 12 times the rows in `so_headers_all`. The primary key is `line_id`. Other selective columns are the following:

- `header_id`
- `parent_line_id`
- `service_parent_line_id`
- `original_system_reference`

The following are indexes on the table:

| Unique | Index Name   | Column Name               |
|--------|--------------|---------------------------|
| -----  |              |                           |
| NO     | SO_LINES_N1  | HEADER_ID                 |
| NO     | SO_LINES_N10 | S5                        |
| NO     | SO_LINES_N11 | S6                        |
| NO     | SO_LINES_N12 | S8                        |
| NO     | SO_LINES_N13 | S9                        |
| NO     | SO_LINES_N14 | S28                       |
| NO     | SO_LINES_N15 | S29                       |
| NO     | SO_LINES_N16 | S30                       |
| NO     | SO_LINES_N17 | PARENT_LINE_ID            |
| NO     | SO_LINES_N18 | SHIPMENT_SCHEDULE_LINE_ID |
| NO     | SO_LINES_N19 | ATO_LINE_ID               |
| NO     | SO_LINES_N2  | LINK_TO_LINE_ID           |

---

|     |              |                                |
|-----|--------------|--------------------------------|
| NO  | SO_LINES_N20 | SERVICE_PARENT_LINE_ID         |
| NO  | SO_LINES_N21 | SHIP_TO_SITE_USE_ID            |
| NO  | SO_LINES_N22 | SOURCE_LINE_ID                 |
| NO  | SO_LINES_N23 | ORIGINAL_SYSTEM_LINE_REFERENCE |
| NO  | SO_LINES_N24 | RETURN_REFERENCE_ID            |
| NO  | SO_LINES_N25 | S27                            |
| NO  | SO_LINES_N26 | CREDIT_INVOICE_LINE_ID         |
| NO  | SO_LINES_N27 | S25                            |
| NO  | SO_LINES_N28 | WH_UPDATE_DATE                 |
| NO  | SO_LINES_N29 | DEMAND_STREAM_ID               |
| NO  | SO_LINES_N3  | OPEN_FLAG                      |
| NO  | SO_LINES_N4  | COMMITMENT_ID                  |
| NO  | SO_LINES_N5  | INVENTORY_ITEM_ID              |
| NO  | SO_LINES_N6  | REQUEST_ID                     |
| NO  | SO_LINES_N7  | S2                             |
| NO  | SO_LINES_N8  | S3                             |
| NO  | SO_LINES_N9  | S4                             |
| YES | SO_LINES_U1  | LINE_ID                        |



---

---

# Glossary

## **asynchronous I/O**

Independent I/O, in which there is no timing requirement for transmission, and other processes can be started before the transmission has finished.

## **Autotrace**

Generates a report on the execution path used by the SQL optimizer and the statement execution statistics. The report is useful to monitor and tune the performance of DML statements.

## **bind variable**

A variable in a SQL statement that must be replaced with a valid value, or the address of a value, in order for the statement to successfully execute.

## **block**

A unit of data transfer between main memory and disk. Many blocks from one section of memory address space form a segment.

## **bottleneck**

The delay in transmission of data, typically when a system's bandwidth cannot support the amount of information being relayed at the speed it is being processed. There are, however, many factors that can create a bottleneck in a system.

## **buffer**

A main memory address in which the buffer manager caches currently and recently used data read from disk. Over time, a buffer can hold different blocks. When a new block is needed, the buffer manager can discard an old block and replace it with a new one.

**buffer pool**

A collection of buffers.

**cache**

Also known as buffer cache. All buffers and buffer pools.

**cache recovery**

The part of instance recovery where Oracle applies all committed and uncommitted changes in the redo log files to the affected data blocks. Also known as the *rolling forward* phase of instance recovery.

**Cartesian product**

A join with no join condition results in a Cartesian product, or a cross product. A Cartesian product is the set of all possible combinations of rows drawn one from each table. In other words, for a join of two tables, each row in one table is matched in turn with every row in the other. A Cartesian product for more than two tables is the result of pairing each row of one table with every row of the Cartesian product of the remaining tables. All other kinds of joins are subsets of Cartesian products effectively created by deriving the Cartesian product and then excluding rows that fail the join condition.

**CBO**

Cost-based optimizer. Generates a set of potential execution plans for SQL statements, estimates the cost of each plan, calls the plan generator to generate the plan, compares the costs, and chooses the plan with the lowest cost. This approach is used when the data dictionary has statistics for at least one of the tables accessed by the SQL statements. The CBO is made up of the query transformer, the estimator, and the plan generator.

**compound query**

A query that uses set operators (UNION, UNION ALL, INTERSECT, or MINUS) to combine two or more simple or complex statements. Each simple or complex statement in a compound query is called a *component query*.

**contention**

When some process has to wait for a resource that is being used by another process.

**dictionary cache**

A collection of database tables and views containing reference information about the database, its structures, and its users. Oracle accesses the data dictionary

frequently during the parsing of SQL statements. Two special locations in memory are designated to hold dictionary data. One area is called the data dictionary cache, also known as the row cache because it holds data as rows instead of buffers (which hold entire blocks of data). The other area is the library cache. All Oracle user processes share these two caches for access to data dictionary information.

#### **distributed statement**

A statement that accesses data on two or more distinct nodes/instances of a distributed database. A *remote statement* accesses data on one remote node of a distributed database.

#### **dynamic performance views**

The views database administrators create on dynamic performance tables (virtual tables that record current database activity). Dynamic performance views are called fixed views because they cannot be altered or removed by the database administrator.

#### **enqueue**

This is another term for a lock.

#### **equijoin**

A join condition containing an equality operator.

#### **estimator**

Uses statistics to estimate the selectivity, cardinality, and cost of execution plans. The main goal of the estimator is to estimate the overall cost of an execution plan.

#### **EXPLAIN PLAN**

A SQL statement that enables examination of the execution plan chosen by the optimizer for DML statements. `EXPLAIN PLAN` causes the optimizer to choose an execution plan and then to put data describing the plan into a database table.

#### **instance recovery**

The automatic application of redo log records to Oracle uncommitted data blocks after a crash or system failure.

#### **join**

A query that selects data from more than one table. A join is characterized by multiple tables in the `FROM` clause. Oracle pairs the rows from these tables using the condition specified in the `WHERE` clause and returns the resulting rows. This

condition is called the join condition and usually compares columns of all the joined tables.

**latch**

A simple, low-level serialization mechanism to protect shared data structures in the System Global Area.

**library cache**

A memory structure containing shared SQL and PL/SQL areas. The library cache is one of three parts of the shared pool.

**LIO**

Logical I/O. A block read which may or may not be satisfied from the buffer cache.

**literal**

A constant value, written at compile-time and read-only at run-time. Literals can be accessed quickly, and are used when modification is not necessary.

**MTBF**

Mean time between failures. A common database statistic important to tuning I/O.

**mirroring**

Maintaining identical copies of data on one or more disks. Typically, mirroring is performed on duplicate hard disks at the operating system level, so that if one of the disks becomes unavailable, the other disk can continue to service requests without interruptions.

**nonequijoin**

A join condition containing something other than an equality operator.

**optimizer**

Determines the most efficient way to execute SQL statements by evaluating expressions and translating them into equivalent, quicker expressions. The optimizer formulates a set of execution plans and picks the best one for a SQL statement. See [CBO](#).

**Oracle Trace**

Used by the Oracle Server to collect performance and resource utilization data, such as SQL parse, execute, fetch statistics, and wait statistics. Oracle Trace provides several SQL scripts that can be used to access server event tables, collect server



event data and stores it in memory, and allows data to be formatted while a collection is occurring.

### **outer join**

A join condition using the outer join operator (+) with one or more columns of one of the tables. Oracle returns all rows that meet the join condition. Oracle also returns all rows from the table without the outer join operator for which there are no matching rows in the table with the outer join operator.

### **paging**

A technique for increasing the memory space available by moving infrequently-used parts of a program's working memory from main memory to a secondary storage medium, usually a disk. The unit of transfer is called a page.

### **parse**

A *hard parse* occurs when a SQL statement is executed, and the SQL statement is either not in the shared pool, or it is in the shared pool but it cannot be shared. A SQL statement is not shared if the metadata for the two SQL statements is different. This can happen if a SQL statement is textually identical as a preexisting SQL statement, but the tables referred to in the two statements resolve to physically different tables, or if the optimizer environment is different.

A *soft parse* occurs when a session attempts to execute a SQL statement, and the statement is already in the shared pool, and it can be used (that is, shared). For a statement to be shared, all data, (including metadata, such as the optimizer execution plan) pertaining to the existing SQL statement must be equally applicable to the current statement being issued.

### **parse call**

A call to Oracle to prepare a SQL statement for execution. This includes syntactically checking the SQL statement, optimizing it, and building (or locating) an executable form of that statement.

### **parser**

Performs syntax analysis and semantic analysis of SQL statements, and expands views (referenced in a query) into separate query blocks.

### **PGA**

Program Global Area. A nonshared memory region that contains data and control information for a server process, created when the server process is started.

**PIO**

Physical I/O. A block read which could not be satisfied from the buffer cache, either because the block was not present or because the I/O is a direct I/O (and bypasses the buffer cache).

**plan generator**

Tries out different possible plans for a given query so that the CBO can choose the plan with the lowest cost. It explores different plans for a query block by trying out different access paths, join methods, and join orders.

**predicate**

A `WHERE` condition in SQL.

**query transformer**

Decides whether to rewrite a user query to generate a better query plan, merges views, and performs subquery unnesting.

**RAID**

Redundant arrays of inexpensive disks. RAID configurations provide improved data reliability with the option of striping (manually distributing data). Different RAID configurations (levels) are chosen based on performance and cost, and are suited to different types of applications, depending on their I/O characteristics.

**RBO**

Rule-based optimizer. Chooses an execution plan for SQL statements based on the access paths available and the ranks of these access paths (if there is more than one way, then the RBO uses the operation with the lowest rank). The RBO is used if no statistics are available, otherwise the CBO is used.

**row source generator**

Receives the optimal plan from the optimizer and outputs the execution plan for the SQL statement. A row source is an iterative control structure that processes a set of rows in an iterated manner and produces a row set.

**segment**

A set of extents allocated for a specific type of database object such as a table, index, or cluster.

**simple statement**

An INSERT, UPDATE, DELETE, or SELECT statement that involves only a single table.

**simple query**

A SELECT statement that references only one table and does not make reference to GROUP BY functions.

**SGA**

System Global Area. A memory region within main memory used to store data for fast access. Oracle uses the shared pool to allocate SGA memory for shared SQL and PL/SQL procedures.

**SQL Compiler**

Compiles SQL statements into a shared cursor. The SQL Compiler is made up of the parser, the optimizer, and the row source generator.

**SQL statements (identical)**

Textually identical SQL statements do not differ in any way.

**SQL statements (similar)**

Similar SQL statements differ only due to changing literal values. If the literal values were replaced with bind variables, then the SQL statements would be textually identical.

**SQL Trace**

A basic performance diagnostic tool to help monitor and tune applications running against the Oracle server. SQL Trace lets you assess the efficiency of the SQL statements an application runs and generates statistics for each statement. The trace files produced by this tool are used as input for TKPROF.

**SQL\*Loader**

Reads and interprets input files. It is the most efficient way to load large amounts of data.

**Statspack**

A set of SQL, PL/SQL, and SQL\*Plus scripts that allow the collection, automation, storage, and viewing of performance data. Statspack supersedes the traditional UTLBSTAT/UTLESTAT tuning scripts.

**striping**

The interleaving of a related block of data across disks. Proper striping reduces I/O and improves performance.

- Stripe depth is the size of the stripe, sometimes called stripe unit.
- Stripe width is the product of the stripe depth and the number of drives in the striped set.

**TKPROF**

A diagnostic tool to help monitor and tune applications running against the Oracle Server. `TKPROF` primarily processes SQL trace output files and translates them into readable output files, providing a summary of user-level statements and recursive SQL calls for the trace files. It can also assess the efficiency of SQL statements, generate execution plans, and create SQL scripts to store statistics in the database.

**transaction recovery**

The part of instance recovery where Oracle applies the rollback segments to undo the uncommitted changes. Also known as the *rolling back* phase of instance recovery.

**UGA**

User Global Area. A memory region in the large pool used for user sessions.

**wait events**

Statistics that are incremented by a server process/thread to indicate that it had to wait for an event to complete before being able to continue processing. Wait events are one of the first places for investigation when performing reactive performance tuning.

**wait events (idle)**

These events indicate that the server process is waiting because it has no work. These events should be ignored when tuning, because they do not indicate the nature of the performance bottleneck.

**work area**

A private allocation of memory used for sorts, hash joins, and other operations that are memory-intensive. A sort operator uses a work area (the sort area) to perform the in-memory sort of a set of rows. Similarly, a hash-join operator uses a work area (the hash area) to build a hash table from its left input.

---

---

# Index

## A

---

access paths

- cluster join, 8-6
- cluster scans, 1-35
- composite index, 8-8
- defined, 1-22
- execution plans, 1-18
- hash cluster key, 8-7
- hash scans, 1-36
- index scans, 1-28
- indexed cluster key, 8-7
- single row by cluster join, 8-4
- single row by hash cluster key (with unique key), 8-5
- single row by rowid, 8-4
- single row by unique or primary key, 8-6

ALL operator, 2-23

ALL\_ROWS

- optimizer mode parameter, 1-7

ALL\_ROWS hint, 1-8, 5-7

allocation

- of memory, 14-2

ALTER INDEX statement, 4-7

ALTER SESSION statement

- examples, 10-5
- SET SESSION\_CACHED\_CURSORS clause, 14-40

ALTER SYSTEM statement

- DISPATCHERS initialization parameter, 19-4

ANALYZE statement, 1-8, 22-22

- creating histograms, 3-22

AND\_EQUAL hint, 4-6, 5-17

anti-joins, 1-43

- transformations not allowed, 1-44

ANY operator, 2-22

APPEND hint, 5-34

APPINFO

- tuning, 11-9

applications

- data warehousing and star queries, 1-44

ApplReg event, 12-15

array interface, 23-13

ARRAYSIZE

- tuning, 11-10

automatic segment-space management, 15-22, 22-28

automatic undo management, 18-2

AUTOTRACE

- settings, 11-2
- system variable, 11-2

autotrace

- SQL\*Plus, 11-1

## B

---

BEGIN\_SNAP variable, 21-12

BETWEEN comparison operator, 2-24

binary files

- formatting using Oracle Trace, 12-3

bind variables, 14-22

- optimization, 1-38

BITMAP CONVERSION row source, 4-18

bitmap indexes, 4-12, 4-17

- compared with B-tree indexes, 4-13
- inlist iterator, 9-19
- maintenance, 4-14
- on index-organized tables, 4-16

- on joins, 4-19
- when to use, 4-12
- BITMAP\_MERGE\_AREA\_SIZE initialization
  - parameter, 4-14, 4-17
- bitmaps
  - mapping to rowids, 4-16
- block sampling, 3-4
- bottlenecks
  - disk I/O, 15-3
  - memory, 14-2
  - resource, 22-26
- broadcast
  - distribution value, 9-26
- B-tree indexes, 4-15, 4-18
- buffer busy wait events, 22-27
  - actions, 22-28
- buffer caches
  - reducing buffers, 14-12, 14-35
- buffer pools
  - default cache, 14-14
  - KEEP cache, 14-14
  - multiple, 14-13
  - RECYCLE cache, 14-14
- BYTES column
  - PLAN\_TABLE table, 9-24

## C

---

- CACHE hint, 5-35
- caching tables
  - automatic caching of small tables, 5-35
- CARDINALITY column
  - PLAN\_TABLE table, 9-24
- cartesian joins, 1-50
- CATALOG.SQL script, 13-5
- CATPROC.SQL script, 13-5
- chained rows, 22-21
- CHAR datatype, 13-3
- character sets
  - database options, 13-3
- checkpoints
  - choosing checkpoint frequency, 17-3
- CHOOSE
  - optimizer mode parameter, 1-7
- CHOOSE hint, 1-8, 5-9

- CLEAR TIMING command
  - SQL\*Plus, 11-7
- client/server applications, 16-12
- CLUSTER hint, 5-11
- clusters, 4-20
  - hash and scans of, 1-36
  - joins and, 8-4, 8-6
  - scans of, 1-35, 8-4
  - scans of hash, 8-5, 8-7
  - scans of index, 8-7
  - scans of joins, 8-6
- collections, 12-8
- columns
  - pseudocolumn ROWNUM, 2-36, 2-45
  - ROWNUM pseudocolumn, 8-15
  - selectivity, 3-2
  - selectivity estimates and histograms, 3-20
  - to index, 4-3
- command files
  - registering, 11-9
- complex view merging, 2-37
- composite indexes, 4-4
- composite partitioning
  - examples of, 9-14
- CONNECT BY clause
  - optimizing view queries, 2-36
- Connection event, 12-15
- connection manager, 23-14
- connection pooling, 19-4
- consistency
  - read, 22-20
- consistent gets statistic, 14-9, 18-3
- consistent mode
  - TKPROF, 10-13
- constants
  - comparisons and, 2-19
  - evaluation of expressions, 2-19
  - when computed, 2-19
- constraints, 4-8
- contention
  - disk, 15-3
  - memory, 14-2, 22-1
  - tuning, 22-1
  - wait events, 22-41
- context switches, 16-13

- CONTROL\_FILES initialization parameter, 13-13
- cost
  - optimizer calculation, 1-10
- COST column
  - PLAN\_TABLE table, 9-24
- cost-based optimizations, 1-10
  - extensible optimization, 1-61
  - histograms, 3-20
  - procedures for plan stability, 7-12
  - selectivity of predicates, 3-2
  - selectivity of predicates and histograms, 3-20
  - selectivity of predicates for user-defined, 1-62
  - star queries, 1-44
  - statistics, 3-2
  - statistics and user-defined, 1-62
  - upgrading to, 7-14
  - user-defined costs, 1-63
- counter/accumulator views, 24-2
- CPU\_COUNT initialization parameter, 17-19
- CPUs
  - utilization, 16-11
- CREATE DATABASE statement, 13-3
- CREATE INDEX statement
  - example, 14-71
  - NOSORT clause, 14-71
  - PARALLEL clause, 13-11
- CREATE OUTLINE statement, 7-5
- CREATE\_BITMAP\_AREA\_SIZE initialization parameter, 4-14, 4-17
- CREATE\_STORED\_OUTLINES parameter, 7-4
- creating databases, 13-2
  - manually, 13-2
  - parameters, 13-2
  - with Installer, 13-2
- cross-facility 3 event, 12-18
- cross-product items
  - See also cross-facility 3 event
- current mode
  - TKPROF, 10-13
- current state views, 24-2
- CURSOR\_NUM column
  - TKPROF\_TABLE table, 10-19
- CURSOR\_SHARING initialization parameter, 1-58, 14-24, 14-44
- CURSOR\_SHARING\_EXACT hint, 5-39

- CURSOR\_SPACE\_FOR\_TIME initialization parameter
  - setting, 14-39

## D

---

- data cache, 16-2
- data dictionary, 14-34
  - CATALOG.SQL scripts, 13-5
  - CATPROC.SQL scripts, 13-5
  - scripts, 13-5
  - statistics in, 3-15
  - views used in optimization, 3-15
- data indexing, 13-10
- data loading, 13-10
- Data Viewer
  - collecting data for specific wait events, 12-35
  - tips on using, 12-35
- data warehousing
  - dimensions, 1-44
  - star queries, 1-44
- Database Connection event, 12-2
- database options, 13-3
- Database Resource Manager, 16-6, 16-10, 22-8
- databases
  - buffers, 14-12, 14-34
  - character set options, 13-3
  - creating, 13-2
  - creating manually, 13-2
  - creation parameters, 13-2
  - creation with Installer, 13-2
  - distributed statement optimization on, 2-13
  - identifier (DBID), 21-3
  - location of initial datafile, 13-4
  - national character set options, 13-4
  - optimization on distributed statement, 2-13
  - SQL.BSQ file options, 13-4
- datatypes
  - CHAR, 13-3
  - NCHAR, 13-4
  - NVARCHAR, 13-4
  - NVARCHAR2, 13-4
  - user-defined and statistics, 1-62
  - VARCHAR, 13-3
  - VARCHAR2, 13-3

- DATE\_OF\_INSERT column
  - TKPROF\_TABLE table, 10-19
- db block gets statistic, 14-9, 18-3
- DB file scattered read wait events, 22-29
  - actions, 22-30
- DB file sequential read wait events
  - actions, 22-32
- DB file sequential/scattered read wait
  - events, 22-29, 22-31
- DB\_BLOCK\_BUFFERS initialization
  - parameter, 14-12, 14-35
- DB\_BLOCK\_SIZE initialization parameter, 13-2, 13-14, 15-14
- DB\_CACHE\_ADVICE parameter, 14-12
- DB\_CACHE\_SIZE initialization parameter, 13-14, 14-13
- DB\_DOMAIN initialization parameter, 13-13
- DB\_FILE\_MULTIBLOCK\_READ\_COUNT
  - initialization parameter, 1-24
- DB\_FILE\_MULTIBLOCK\_READ\_COUNT
  - initialization parameter, 1-58, 15-13, 15-14, 22-29
  - cost-based optimization, 1-43
- DB\_KEEP\_CACHE\_SIZE
  - initialization parameter, 14-17
- DB\_NAME initialization parameter, 13-13
- DB\_nK\_CACHE\_SIZE initialization
  - parameter, 14-12
- DB\_RECYCLE\_CACHE\_SIZE
  - initialization parameter, 14-18
- DB\_WRITER\_PROCESSES initialization
  - parameter, 22-40
- DBA\_OBJECTS view, 14-16
- DBID
  - database identifier, 21-3
  - Statspack, 21-25
- DBMS\_APPLICATION\_INFO package, 11-9
- DBMS\_JOB procedure, 21-8
- DBMS\_JOB.INTERVAL procedure, 21-9
- DBMS\_OUTLN package, 7-4
- DBMS\_OUTLN\_EDIT package, 7-4
- DBMS\_SHARED\_POOL package, 14-42, 14-43
- DBMS\_STATS package, 1-8, 3-5, 3-6
  - creating histograms, 3-22
- default cache, 14-14
- DEFAULT\_TABLESPACE variable, 21-6
- DEFINE OFF
  - tuning, 11-10
- deleting
  - data, 21-23
  - snapshots, 21-23
- DEPTH column
  - TKPROF\_TABLE table, 10-19
- deterministic functions
  - PL/SQL, 2-28
- dictionary managed tablespaces, 21-4
- dimensions
  - star joins, 1-44
  - star queries, 1-44
- direct path read events, 22-33
  - actions, 22-34
  - causes, 22-34
- direct path wait events, 22-35
- direct path write events
  - actions, 22-35
  - causes, 22-35
- direct-path INSERT, 5-34
- disabled constraints, 4-8
- Disconnect event, 12-15
- disk reads and buffer gets
  - monitoring, 11-9
- disks
  - contention, 15-3
  - monitoring operating system file activity, 22-8
- dispatcher processes, 19-4
- DISPATCHERS initialization parameter, 19-4, 23-3
- DISTINCT operator
  - optimizing views, 2-37
- distributed databases
  - statement optimization on, 2-13
- distributed transactions
  - optimizing, 2-13
  - sample table scan not supported, 1-36
- distribution
  - hints for, 5-31
- DISTRIBUTION column
  - PLAN\_TABLE table, 9-25
- DML locks, 24-18
- domain indexes
  - and EXPLAIN PLAN, 9-20



- extensible optimization, 1-61
- user-defined statistics, 1-62
  - using, 4-19
- DRIVING\_SITE hint, 5-27
- duration events
  - in Oracle Trace, 12-2, 12-15
- dynamic performance views, 24-2
- DYNAMIC\_SAMPLING hint, xxxii, 5-39

## E

---

- enabled constraints, 4-8
- END\_SNAP variable, 21-12
- enforced constraints, 4-8
- enqueue wait events
  - actions, 22-37
- EPC\_ERROR.LOG file, 12-37
- equijoins, 6-10
- ErrorStack event, 12-15
- event timings, 21-21
- examples
  - ALTER SESSION statement, 10-5
  - concurrently creating tablespaces, 13-7
  - CREATE DATABASE script, 13-4
  - CREATE INDEX statement, 14-71
  - creating indexes efficiently, 13-12
  - executing required data dictionary scripts, 13-5
  - execution plan, 8-18
  - EXPLAIN PLAN output, 8-18, 10-16
  - full table scan, 8-19
  - indexed query, 8-19
  - minimal initialization file, 13-14
  - NOSORT clause, 14-71
  - SET TRANSACTION statement, 18-3
  - SQL trace facility output, 10-16
  - V\$DB\_OBJECT\_CACHE view, 24-5
  - V\$FILESTAT view, 24-8
  - V\$LATCH view, 24-11
  - V\$LATCH\_CHILDREN view, 24-13
  - V\$LATCHHOLDER view, 24-14
  - V\$LIBRARYCACHE view, 24-16
  - V\$LOCK view, 24-20
  - V\$OPEN\_CURSOR view, 24-23, 24-24
  - V\$PROCESS view, 24-27
  - V\$ROLLSTAT view, 24-29
  - V\$SESSION view, 24-35
  - V\$SESSION\_EVENT view, 24-37
  - V\$SESSION\_WAIT view, 24-39
  - V\$SQLAREA view, 24-58, 24-59
  - V\$SQLTEXT view, 24-60
- Execute event, 12-15
- execution plans
  - accessing views, 2-39, 2-42, 2-43
  - comparing with PLAN\_HASH\_VALUE, 24-45
  - complex statements, 2-34
  - compound queries, 2-48, 2-49, 2-50
  - examples, 2-34, 8-18, 10-7
  - execution sequence of, 1-23
  - joining views, 2-46
  - joins, 1-40
  - optimizer path, 11-3
  - OR operators, 2-31, 8-18
  - overview of, 1-18
  - plan stability, 7-2
  - preserving with plan stability, 7-2
  - table output, 11-3
  - TKPROF, 10-7, 10-11
  - viewing with the utlxpls.sql script, 1-18
- EXPLAIN PLAN statement
  - access paths, 1-36, 8-4, 8-5, 8-6, 8-7, 8-8, 8-9, 8-10, 8-11, 8-12, 8-13, 8-14, 8-15
  - and domain indexes, 9-20
  - and full partition-wise joins, 9-17
  - and partial partition-wise joins, 9-16
  - and partitioned objects, 9-12
  - basic steps, 1-19
  - examples of output, 8-18, 10-16
  - execution order of steps in output, 1-19
  - invoking with the TKPROF program, 10-11
  - PLAN\_TABLE table, 9-4
  - restrictions, 9-22
  - scripts for viewing output, 1-19
  - viewing the output, 1-18
- Export utility
  - copying statistics, 3-2
- exporting data, 21-22
- extensible optimization, 1-61
  - user-defined costs, 1-63
  - user-defined selectivity, 1-62
  - user-defined statistics, 1-62

## F

---

- FACT hint, 5-22
- fact tables
  - star joins, 1-44
  - star queries, 1-44
- fast full index scans, 1-34
- FAST\_START\_IO\_TARGET initialization
  - parameter, 17-4, 17-5
- FAST\_START\_MTTT\_TARGET initialization
  - parameter, 17-4, 17-5, 17-9, 17-12
- FAST\_START\_PARALLEL\_ROLLBACK
  - initialization parameter, 17-19
- FastCGI
  - iSQL\*Plus, 11-14
- fast-start checkpoints
  - FAST\_START\_MTTT\_TARGET initialization
    - parameter, 17-6
  - LOG\_CHECKPOINT\_INTERVAL initialization
    - parameter, 17-7
  - LOG\_CHECKPOINT\_TIMEOUT initialization
    - parameter, 17-7
- fast-start on-demand rollback, 17-18
- fast-start parallel rollback, 17-18
- features, new, xxxi
- Fetch event, 12-15
- FIRST\_ROWS
  - optimizer mode parameter, 1-7
- FIRST\_ROWS hint, 1-8
- FIRST\_ROWS(n) hint, 1-8, 5-7
- FIRST\_ROWS\_n
  - optimizer mode parameter, 1-7
- FLUSH OFF
  - tuning, 11-10
- FORCE\_UNION\_REWRITE hint, xxxii, 5-19
- FORMAT statement
  - in Oracle Trace, 12-3
- formatter tables
  - in Oracle Trace, 12-3
- free buffer wait events, 22-39
- FULL hint, 4-6, 5-10
- full outer joins, 1-54
- full partition-wise joins, 9-17
- full table scans, 8-14, 8-19, 22-34
  - rule-based optimizer, 8-14

- function-based indexes, 4-10
- functions
  - PL/SQL deterministic, 2-28
  - SQL and optimizing view queries, 2-43
  - user-defined and extensible optimization, 1-61

## G

---

- GATHER\_INDEX\_STATS procedure
  - in DBMS\_STATS package, 3-6
- GATHER\_DATABASE\_STATS procedure
  - in DBMS\_STATS package, 3-6
- GATHER\_SCHEMA\_STATS procedure
  - in DBMS\_STATS package, 3-6
- GATHER\_TABLE\_STATS procedure
  - in DBMS\_STATS package, 3-6
- GETMISSES column
  - in VSROWCACHE table, 14-34
- GETS column
  - in VSROWCACHE view, 14-34
- global hints, 5-44
- GLOGIN.SQL
  - site profile, 11-3
- GROUP BY clause
  - NOSORT clause, 14-72
  - optimizing views, 2-37

## H

---

- hash
  - distribution value, 9-26
- hash clusters
  - scans of, 1-36, 8-5, 8-7
- HASH hint, 5-12
- hash joins, 1-47
  - index join, 1-35
- hash partitions, 9-12
  - examples of, 9-12
- HASH\_AJ hint, 1-43, 5-28
- HASH\_AREA\_SIZE initialization parameter, 1-58
- HASH\_JOIN\_ENABLED initialization
  - parameter, 1-59
- HASH\_SJ hint, 1-44, 5-28
- hashing, 4-21
- HIGH\_VALUE statistics, 1-38

## hints, 5-2

- access paths, 5-9, 5-17
- ALL\_ROWS hint, 5-7
- AND\_EQUAL hint, 4-6, 5-17
- as used in outlines, 7-3
- CACHE hint, 5-35
- cannot override sample access path, 1-37
- CHOOSE hint, 5-9
- CLUSTER hint, 5-11
- CURSOR\_SHARING\_EXACT hint, 5-39
- degree of parallelism, 5-29
- EXPAND\_GSET\_TO\_UNION hint, 5-19
- expandable optimization, 1-62
- FACT hint, 5-22
- FIRST\_ROWS hint, 5-7
- FIRST\_ROWS(n) hint, 5-7
- FORCE\_UNION\_REWRITE hint, 5-20
- FULL hint, 4-6, 5-10
  - global, 5-44
- HASH hint, 5-12
- HASH\_AJ hint, 5-28
- HASH\_SJ hint, 5-28
- how to use, 5-2
- INDEX hint, 4-6, 5-12, 5-23
- INDEX\_ASC hint, 5-14
- INDEX\_DESC hint, 5-14, 5-15
- INDEX\_FFS, 1-34
- INDEX\_JOIN, 1-35
- join operations, 5-24
- LEADING hint, 5-27
- MERGE hint, 5-20
- MERGE\_AJ and HASH\_AJ, 1-43
- MERGE\_AJ hint, 5-28
- MERGE\_SJ and HASH\_SJ, 1-44
- MERGE\_SJ hint, 5-28
- NL\_AJ hint, 5-28
- NL\_SJ hint, 5-28
- NO\_EXPAND hint, 5-18
- NO\_FACT hint, 5-22
- NO\_INDEX, 4-6
- NO\_INDEX hint, 5-16
- NO\_MERGE hint, 5-21
- NO\_PUSH\_PRED hint, 5-37
- NO\_PUSH\_SUBQ, 5-38
- NO\_PUSH\_SUBQ hint, 5-38

- NO\_UNNEST hint, 5-37
- NOCACHE hint, 5-35
- NOPARALLEL hint, 5-30
- NOREWRITE hint, 5-20
- optimization approach and goal, 5-6
- ORDERED hint, 1-43, 5-23
- overriding optimizer choice, 1-37
- overriding OPTIMIZER\_MODE, 1-8
- PARALLEL hint, 5-29
- parallel query option, 5-29
- PQ\_DISTRIBUTE hint, 5-31
- PUSH\_PRED hint, 5-37
- PUSH\_SUBQ hint, 5-37
- REWRITE hint, 5-19
- ROWID hint, 5-11
- STAR hint, 5-23
  - syntax, 5-3
- UNNEST hint, 5-36
- USE\_CONCAT hint, 5-18
- USE\_MERGE hint, 5-25
- USE\_NL hint, 5-24
- histograms, 3-20
  - number of buckets, 3-22
- HOLD\_CURSOR clause, 14-27

## I

---

- ID column
  - PLAN\_TABLE table, 9-24
- idle timeout
  - tuning, 11-14
- idle wait events, 22-49
  - SQL\*Net message from client, 22-26
- Import utility
  - copying statistics, 3-2
- IN operator, 2-22
  - merging views, 2-38
- IN subquery, 2-37
- INDEX hint, 4-6, 4-15, 5-12
- index joins, 1-35
- INDEX\_ASC hint, 5-14
- INDEX\_COMBINE hint, 4-6, 4-15
- INDEX\_DESC hint, 5-14, 5-15
- INDEX\_FFS hint, 1-34
- INDEX\_JOIN hint, 1-35

## indexes

- avoiding the use of, 4-6
- bitmap, 4-12, 4-17
- choosing columns for, 4-3
- composite, 4-4, 8-8
- creating, 13-11
- domain, 4-19
- domain indexes and extensible
  - optimization, 1-61
- domain indexes and user-defined statistics, 1-62
- dropping, 4-2
- enforcing uniqueness, 4-8
- ensuring the use of, 4-6
- example, 8-19
- function-based, 4-10
- improving selectivity, 4-4
- index joins, 1-35
- low selectivity, 4-6
- modifying values of, 4-4
- non-unique, 4-8
- optimization and, 2-30, 8-17
- placement on disk, 15-16
- rebuilding, 4-7
- re-creating, 4-7
- restrictions on scans of, 8-14
- scans of, 1-28
- scans of bounded range, 8-10
- scans of cluster key, 8-7
- scans of composite, 8-8
- scans of MAX or MIN, 8-13
- scans of ORDER BY, 8-13
- scans of single-column, 8-8
- scans of unbounded range, 8-11
- selectivity of, 4-3
- statement conversion and, 2-30, 8-17
- statistics gathering, 3-9

indexing data, 13-10

information views, 24-4

initial database creation, 13-2

initialization files, 13-2, 13-13

initialization parameters

- CONTROL\_FILES, 13-13
- CPU\_COUNT, 17-19
- DB\_BLOCK\_SIZE, 13-2, 13-14
- DB\_CACHE\_SIZE, 13-14

- DB\_DOMAIN, 13-13
- DB\_FILE\_MULTIBLOCK\_READ\_COUNT, 1-43
- DB\_NAME, 13-2, 13-13
- FAST\_START\_PARALLEL\_ROLLBACK, 17-19
- in Oracle Trace, 12-7
- INITRANS, 13-8
- JAVA\_POOL\_SIZE, 13-14
- JOB\_QUEUE\_PROCESSES, 21-9
- LOG\_ARCHIVE\_XXX, 13-14
- LOG\_CHECKPOINT\_INTERVAL, 17-7
- LOG\_CHECKPOINT\_TIMEOUT, 17-7
- LOG\_PARALLELISM, 17-8
- OPEN\_CURSORS, 13-13
- OPTIMIZER\_FEATURES\_ENABLE, 1-34, 1-35, 2-37
- OPTIMIZER\_MODE, 1-6, 5-7, 8-2
- PARALLEL\_MAX\_SERVERS, 17-8
- PGA\_AGGREGATE\_TARGET, 13-11
- PROCESSES, 13-14
- RECOVERY\_PARALLELISM, 17-8
- SESSION\_CACHED\_CURSORS, 14-40
- SESSIONS, 13-14
- SHARED\_POOL\_SIZE, 13-14
- SORT\_AREA\_SIZE, 1-42, 13-12
- SQL\_TRACE, 10-6
- TIMED\_STATISTICS, 21-7
- USER\_DUMP\_DEST, 10-4

INIT.ORA file

- ORACLE\_TRACE\_ENABLE parameter, 12-35

INITRANS initialization parameter, 13-8

IN-lists, 5-14, 5-18

input parameters

- SNAP and MODIFY\_STATSPACK\_PARAMETERS, 21-20

INPUT\_IO item, 12-16

INSERT statement

- append, 5-34

instance configuration, 13-13

instance numbers, 21-3

INSTANCE\_NUMBER

- Statpack, 21-25

instrumentation

- of Oracle Server, 12-15

INTERSECT operator

- example, 2-50

- optimizing view queries, 2-36
- intratransaction recovery, 17-19
- I/O
  - and SQL statements, 22-31
  - balancing, 15-4
  - excessive I/O waits, 22-30
  - objects causing I/O waits, 22-31
  - reducing, 4-4
- iSQL\*Plus
  - FastCGI, 11-14
  - idle timeout, 11-14
  - interpreting statistics, 11-13
  - iSQLPlusHashTableSize, 11-13
  - iSQLPlusNumberOfThreads, 11-13
  - iSQLPlusTimeOutInterval, 11-13
  - parameters for tuning, 11-13
  - server statistics report, 11-11
  - statistics report, 11-11
  - tuning statistics, 11-13
- isqlplus.conf file, 11-13
- iSQLPlusHashTableSize
  - tuning, 11-13
- iSQLPlusNumberOfThreads
  - impact on iSQLPlusHashTableSize, 11-13
  - impact on request load, 11-13
  - tuning, 11-13
- iSQLPlusTimeOutInterval
  - tuning, 11-13
- items
  - cross-product, 12-17
  - standard resource utilization, 12-16
  - types of, 12-16

## J

---

- JAVA\_POOL\_SIZE initialization parameter, 13-14
- JOB\_QUEUE\_PROCESSES initialization parameter, 21-9
- joins
  - anti-joins, 1-43
  - cartesian, 1-50
  - cluster, 8-4
  - convert to subqueries, 2-33
  - execution plans and, 1-40
  - full outer, 1-54

- hash, 1-47
- index joins, 1-35
- join order and execution plans, 1-18
- join order and selectivity of predicates, 1-62, 3-2, 3-20
- nested loop, 1-45
- nested loops and cost-based optimization, 1-42
- optimization of, 8-16
- outer, 1-51
- outer and non-null values for nulls, 2-45
- parallel, and PQ\_DISTRIBUTE hint, 5-31
- partition-wise
  - examples of full, 9-17
  - examples of partial, 9-16
  - full, 9-17
- sample table scan not supported, 1-36
- searches on clusters, 8-6
- select-project-join views, 2-35
- semi-joins, 1-43
- sort merge, 1-49
- sort-merge and cost-based optimization, 1-42
- sort-merge example, 8-12
- star joins, 1-44
- star queries, 1-44

## K

---

- KEEP cache, 14-14
- keys
  - searches, 8-5

## L

---

- LARGE\_POOL\_SIZE initialization parameter, 14-36
- latch free wait events
  - actions, 22-42
- latches
  - tuning, 24-12
- LEADING hint, 5-27
- level 7 snapshot
  - Statspack, 21-18
- library cache
  - memory allocation, 14-34
- LIKE operator, 2-22

- Lmode modes, 24-19
- load balancing, 15-4
- loading data, 13-10
- locally managed tablespaces, 21-4
- location of initial datafile
  - database options, 13-4
- lock types
  - common, 24-17
  - ST (space transaction) locks, 24-18
  - TM (DML) locks, 24-18
  - TX (row transaction) locks, 24-18
  - UL (user defined) locks, 24-19
- locking rows, 13-8
- locks and lock holders
  - finding, 22-36
- log buffer tuning, 14-47
- log file switch wait events, 22-46
- log writer processes
  - tuning, 15-18
- LOG\_ARCHIVE\_XXX initialization
  - parameter, 13-14
- LOG\_BUFFER initialization parameter, 14-47
  - setting, 14-48
- LOG\_CHECKPOINT\_INTERVAL initialization
  - parameter, 17-3
  - recovery time, 17-7
- LOG\_CHECKPOINT\_TIMEOUT initialization
  - parameter, 17-4
  - recovery time, 17-7
- LOG\_PARALLELISM initialization
  - parameter, 17-8
- LogicalTX event, 12-15
- lookup tables
  - star queries, 1-44
- LOW\_VALUE statistics, 1-38
- LRU
  - aging policy, 14-13
  - latch contention, 22-45

## M

---

- manual database creation, 13-2
- max session memory statistic, 14-37
- MAX\_DISPATCHERS initialization
  - parameter, 19-4

- MAX\_DUMP\_FILE\_SIZE initialization parameter
  - SQL Trace, 10-4
- MAX\_SHARED\_SERVERS initialization
  - parameter, 19-7
- MAXOPENCURSORS clause, 14-27
- MAXRS\_SIZE item, 12-16
- mean time to recover, 17-4
  - advisory, xxxv
  - See also MTTR
- memory allocation
  - importance, 14-2
  - library cache, 14-34
  - shared SQL areas, 14-34
  - sort areas, 14-69
  - tuning, 14-5
- MERGE hint, 5-20
- MERGE\_AJ hint, 1-43, 5-28
- MERGE\_SJ hint, 1-44, 5-28
- merging complex views, 2-37
- merging views into statements, 2-35
- migrated rows, 22-21
- Migration event, 12-15
- MINUS operator
  - optimizing view queries, 2-36
- mirroring
  - redo logs, 15-19
- modes
  - Lmode, 24-19
  - request, 24-19
- monitoring
  - disk reads and buffer gets, 11-9
- MTBF (mean time between failures)
  - tuning I/O, 15-2
- MTTR
  - initialization parameter, 17-6
  - mean time to recover advisory, xxxv
  - See also mean time to recover
- multiple buffer pools, 14-13

## N

---

- NAMESPACE column
  - VSLIBRARYCACHE view, 14-29
- national character set database option, 13-4
- NCHAR datatype, 13-4

- nested loop joins, 1-45
  - cost-based optimization, 1-42
- network
  - array interface, 23-13
  - detecting performance problems, 23-6
  - problem solving, 23-8
  - Session Data Unit, 23-13
  - tuning, 23-1
- network communication wait events, 22-25
  - DB file sequential/scattered read wait events, 22-29, 22-31
  - SQL\*Net message from Dblink, 22-27
- new features, xxxi
- NL\_AJ hint, 5-28
- NL\_SJ hint, 5-28
- NLS\_SORT initialization parameter
  - ORDER BY access path, 8-13
- NO\_EXPAND hint, 5-18
- NO\_FACT hint, 5-22
- NO\_INDEX hint, 4-6, 5-16
- NO\_MERGE hint, 5-21
- NO\_PUSH\_PRED hint, 5-37
- NO\_UNNEST hint, 5-37
- NOAPPEND hint, 5-34
- NOCACHE hint, 5-35
- NOPARALLEL hint, 5-30
- NOPARALLEL\_INDEX hint, 5-33
- NOREWRITE hint, 5-20
- NOSORT clause, 14-71, 14-72
- NOT IN subquery, 1-43
- NOT operator, 2-24
- NT performance, 16-7
- nulls
  - non-null values for, 2-45
- NUM\_DISTINCT column
  - USER\_TAB\_COLUMNS view, 1-38
- NUM\_ROWS column
  - USER\_TABLES view, 1-38
- NVARCHAR datatype, 13-4
- NVARCHAR2 datatype, 13-4

## O

---

- OBJECT\_INSTANCE column
  - PLAN\_TABLE table, 9-23
- OBJECT\_NAME column
  - PLAN\_TABLE table, 9-23
- OBJECT\_NODE column
  - PLAN\_TABLE table, 9-23
- OBJECT\_OWNER column
  - PLAN\_TABLE table, 9-23
- OBJECT\_TYPE column
  - PLAN\_TABLE table, 9-23
- OPEN\_CURSORS initialization parameter, 13-13
  - increasing cursors for each session, 14-34
- operating system
  - data cache, 16-2
  - monitoring disk I/O, 22-8
- OPERATION column
  - PLAN\_TABLE table, 9-23, 9-27
- OPTIMAL parameter, 18-3
- optimization
  - choosing the approach, 1-6
  - conversion of expressions and predicates, 2-2
  - cost calculation, 1-10
  - cost-based, 1-10
  - cost-based and choosing an access path, 1-37
  - cost-based and histograms, 3-20
  - cost-based and star queries, 1-44
  - cost-based and user-defined costs, 1-63
  - cost-based examples, 1-37
  - cost-based on remote databases, 2-14
  - described, 1-3
  - DISTINCT, 2-37
  - distributed SQL statements, 2-13
  - extensible optimizer, 1-61
  - fast-response method, 1-9
  - GROUP BY views, 2-37
  - hints, 1-8, 1-34, 1-35
  - manual, 1-8
  - merging complex views, 2-37
  - merging views into statements, 2-35
  - non-null values for nulls, 2-45
  - operations performed, 1-5
  - rule-based, 8-2, 8-3, 8-16
  - selectivity of predicates, 3-2
  - selectivity of predicates and histograms, 3-20
  - selectivity of predicates for user-defined, 1-62
  - select-project-join views, 2-35
  - semi-joins, 1-43

- statistics, 3-2
  - statistics for user-defined, 1-62
  - transitivity and, 2-25
  - without merging, 2-45
- optimizer
  - cost calculation, 1-10
  - execution path, 11-3
  - goals, 1-5
  - introduction, 1-3
  - operations, 1-5
  - plan stability, 7-2
  - response time, 1-5
  - throughput, 1-5
- OPTIMIZER column
  - PLAN\_TABLE, 9-24
- optimizer mode parameters
  - ALL\_ROWS, 1-7
  - CHOOSE, 1-7
  - FIRST\_ROWS, 1-7
  - FIRST\_ROWS\_n, 1-7
  - RULE, 1-7
- OPTIMIZER\_FEATURES\_ENABLE initialization
  - parameter, 1-34, 1-35, 1-56, 2-37
- OPTIMIZER\_INDEX\_CACHING initialization
  - parameter, 1-59
- OPTIMIZER\_INDEX\_COST\_ADJ initialization
  - parameter, 1-59
- OPTIMIZER\_MAX\_PERMUTATIONS initialization
  - parameter, 1-59
- OPTIMIZER\_MODE initialization parameter, 1-6, 1-7, 1-60, 5-7, 8-2
  - hints affecting, 1-8
- OPTIONS column
  - PLAN\_TABLE table, 9-23
- Oracle Forms, 10-6
  - control of parsing and private SQL areas, 14-27
- Oracle Net Configuration Assistant, 23-14
- Oracle Performance Manager
  - illustration, 20-5
- Oracle Real Application Clusters
  - and Statspack, 21-25
- Oracle SQL Analyze
  - illustration, 6-3
- Oracle Trace, 12-1
  - accessing collected data, 12-3
  - binary files, 12-3
  - collection results, 12-12
  - collections, 12-8
  - command-line interface, 12-3
  - deleting files, 12-7
  - deprecated, xxxvii
  - deprecated in future release, xxxvii
  - duration events, 12-2
  - events, 12-2
  - FORMAT statement, 12-3
  - formatter tables, 12-3
  - parameters, 12-7
  - point events, 12-2
  - reporting utility, 12-14
  - START statement, 12-3, 12-4
  - STOP statement, 12-3, 12-6
- ORACLE\_TRACE\_COLLECTION\_NAME
  - initialization parameter, 12-7
- ORACLE\_TRACE\_COLLECTION\_PATH
  - initialization parameter, 12-8
- ORACLE\_TRACE\_COLLECTION\_SIZE
  - initialization parameter, 12-8
- ORACLE\_TRACE\_ENABLE initialization
  - parameter, 12-8, 12-35
- ORACLE\_TRACE\_FACILITY\_NAME initialization
  - parameter, 12-8, 12-9
- ORACLE\_TRACE\_FACILITY\_PATH initialization
  - parameter, 12-8
- Oracle-managed files, 15-20
  - tuning, 15-20
- ORDERED hint, 1-43, 5-23
- ORDERED\_PREDICATES hint, 5-38
- OTHER column
  - PLAN\_TABLE table, 9-25
- OTHER\_TAG column
  - PLAN\_TABLE table, 9-24
- outer joins, 1-51
  - non-null values for nulls, 2-45
- Outline Editor
  - illustration, 7-7
- outlines
  - CREATE OUTLINE statement, 7-5
  - creating and using, 7-4
  - execution plans and plan stability, 7-2
  - hints, 7-3



- moving tables, 7-11
- storage requirements, 7-4
- using, 7-6
- using to move to the cost-based optimizer, 7-13
- viewing data for, 7-10

OUTPUT\_IO item, 12-16

overloaded disks, 15-10

## P

---

page table, 16-12

PAGEFAULT\_IO item, 12-16

PAGEFAULTS item, 12-16

paging, 16-12

- reducing, 14-4

parallel broadcast, 1-57

PARALLEL clause

- CREATE INDEX statement, 13-11
- RECOVER statement, 17-8

parallel execution

- hints, 5-29

PARALLEL hint, 5-29

parallel joins

- and PQ\_DISTRIBUTE hint, 5-31

parallel recovery, 17-8

PARALLEL\_MAX\_SERVERS initialization

- parameter, 17-8

parameter files, 13-2

parameters

- iSQL\*Plus tuning, 11-13
- SNAP and MODIFY\_STATSPACK\_PARAMETERS, 21-20

PARENT\_ID column

- PLAN\_TABLE table, 9-24

Parse event, 12-15

parsing

- Oracle Forms, 14-27
- Oracle precompilers, 14-27
- reducing unnecessary calls, 14-26

PARTITION\_ID column

- PLAN\_TABLE table, 9-25

PARTITION\_START column

- PLAN\_TABLE table, 9-24

PARTITION\_STOP column

- PLAN\_TABLE table, 9-25

PARTITION\_VIEW\_ENABLED initialization

- parameter, 1-60

partitioned objects

- and EXPLAIN PLAN statement, 9-12

partitioning

- distribution value, 9-26
- examples of, 9-12
- examples of composite, 9-14
- hash, 9-12
- range, 9-12
- start and stop columns, 9-13

partitions

- statistics, 3-4

partition-wise joins

- full, 9-17
- full, and EXPLAIN PLAN output, 9-17
- partial, and EXPLAIN PLAN output, 9-16

PCTFREE parameter, 13-8, 22-22

PCTINCREASE parameter, 18-4

PCTUSED parameter, 22-22

performance

- generating reports, 21-9
- mainframe, 16-8
- NT, 16-7
- of SQL statements, 11-2
- running reports, 21-3, 21-9
- UNIX-based systems, 16-7
- viewing execution plans, 1-18

Performance Monitor

- NT, 16-12

PERFSTAT user, 21-3, 21-4, 21-15

PGA\_AGGREGATE\_TARGET initialization

- parameter, 13-11, 14-50

physical reads statistic, 14-10

PhysicalTX event, 12-15

plan stability, 7-2

- limitations of, 7-2
- preserving execution plans, 7-2
- procedures for the cost-based optimizer, 7-12
- use of hints, 7-2

PLAN\_HASH\_VALUE

- V\$SQL view column, 24-45

PLAN\_TABLE table

- BYTES column, 9-24
- CARDINALITY column, 9-24

- COST column, 9-24
- creating, 9-4, 11-2
- DISTRIBUTION column, 9-25
- ID column, 9-24
- OBJECT\_INSTANCE column, 9-23
- OBJECT\_NAME column, 9-23
- OBJECT\_NODE column, 9-23
- OBJECT\_OWNER column, 9-23
- OBJECT\_TYPE column, 9-23
- OPERATION column, 9-23
- OPTIMIZER column, 9-24
- OPTIONS column, 9-23
- OTHER column, 9-25
- OTHER\_TAG column, 9-24
- PARENT\_ID column, 9-24
- PARTITION\_ID column, 9-25
- PARTITION\_START column, 9-24
- PARTITION\_STOP column, 9-25
- POSITION column, 9-24
- REMARKS column, 9-23
- SEARCH\_COLUMNS column, 9-24
- STATEMENT\_ID column, 9-23
- TIMESTAMP column, 9-23
- PL/SQL
  - deterministic functions, 2-28
- PLUSTRACE
  - creating role, 11-2
  - granting role, 11-3
- PLUSTRACE role, 11-2
- point events
  - in Oracle Trace, 12-2, 12-15
- POOL attribute, 19-4
- POSITION column
  - PLAN\_TABLE table, 9-24
- PQ\_DISTRIBUTE hint, 5-31
- precompilers
  - control of parsing and private SQL areas, 14-27
- predicates
  - pushing into a view, 2-38, 2-43
  - pushing into a view examples, 2-39, 2-41
  - selectivity, 3-2
  - selectivity estimates and histograms, 3-20
  - selectivity for user-defined, 1-62
- PRIMARY KEY constraint, 4-8
- primary keys

- optimization, 2-34
  - searches, 8-6
- PRIVATE\_SGA variable, 14-38
- procedures
  - DBMS\_JOB, 21-8
  - DBMS\_JOB.INTERVAL, 21-9
  - deterministic functions, 2-28
  - STATSPACK.MODIFY\_STATSPACK\_PARAMETER, 21-16, 21-19
  - STATSPACK.SNAP, 21-7, 21-8, 21-19
- processes
  - dispatcher process configuration, 19-4
  - priority, 16-5
  - scheduler, 16-5
  - scheduling, 16-13
- PROCESSES initialization parameter, 13-14
- program global area (PGA)
  - direct path read, 22-33
  - direct path write, 22-35
  - shared servers, 14-36
- pseudocolumns
  - ROWNUM and optimizing view queries, 2-36, 2-45
  - ROWNUM cannot use indexes, 8-15
- PUSH\_PRED hint, 5-37

## Q

---

- queries
  - avoiding the use of indexes, 4-6
  - compound and optimization of, 2-48
  - compound converted to ORs, 2-30
  - compound with ORs converted to, 8-17
  - ensuring the use of indexes, 4-6
  - optimizing IN subquery, 2-37
  - SAMPLE clause and cost-based optimization, 1-4
  - star queries, 1-44
  - tracing, 11-7
- QUERY\_REWRITE\_ENABLED initialization parameter, 1-60

## R

---

- range
  - distribution value, 9-26
- range partitions, 9-12
  - examples of, 9-12
- read consistency, 22-20
- read events
  - direct path, 22-33
- read wait events
  - scattered, 22-29
- REBUILD clause, 4-7
- RECOVER statement
  - PARALLEL clause, 17-8
- recovery
  - parallel intratransaction recovery, 17-19
  - parallel processes for, 17-8
  - PARALLEL\_MAX\_SERVERS initialization parameter, 17-8
  - setting number of processes to use, 17-8
- RECOVERY\_PARALLELISM initialization parameter, 17-8
- recursive calls, 10-14
- RECYCLE cache, 14-14
- REDO BUFFER ALLOCATION RETRIES
  - statistic, 14-47
- redo logs, 13-5
  - mirroring, 15-19
  - placement on disk, 15-17
  - sizing, 13-5
- reducing
  - contention with dispatchers, 19-4
  - contention with operating system processes, 16-5
  - contention with shared servers, 19-5
  - data dictionary cache misses, 14-34
  - paging and swapping, 14-4
  - rollback segment contention, 18-3
  - unnecessary parse calls, 14-26
- RELEASE\_CURSOR clause, 14-27
- REMARKS column
  - PLAN\_TABLE table, 9-23
- removing
  - data, 21-23
  - snapshots, 21-23
- REPORT\_NAME variable, 21-12
- reports
  - performance, 21-3, 21-9
  - SPREPORT.SQL, 21-9
  - SPREPSQL.SQL, 21-10
  - Statspack, 21-10
- request modes, 24-19
- resource bottlenecks, 22-26
- resource wait events, 22-31
- response time
  - cost-based approach, 1-7
  - optimizer goal, 1-5
  - optimizing, 1-5, 5-7
- REWRITE hint, 5-18
- rollback segments, 22-21
  - assigning to transactions, 18-3
  - choosing how many, 18-3
  - creating, 18-3
- rollback tablespaces, 13-6
- rollbacks
  - fast-start on-demand, 17-18
  - fast-start parallel, 17-18
- round-robin
  - distribution value, 9-26
- row locking, 13-8
- row sampling, 3-4
- row sources, 1-22
- row transaction locks, 24-18
- ROWID hint, 5-11
- rowids
  - mapping to bitmaps, 4-16
  - table access by, 1-27
- ROWNUM pseudocolumn
  - cannot use indexes, 8-15
  - optimizing view queries, 2-36, 2-45
- rows
  - row sources, 1-22
  - rowids used to locate, 1-27, 8-4
  - setting number retrieved at one time, 11-10
- RowSource event, 12-2, 12-15
- RULE
  - optimizer mode parameter, 1-7
- RULE hint
  - OPTIMIZER\_MODE and, 1-8
- rule-based optimization, 8-2, 8-3

## S

---

- SAMPLE BLOCK clause, 1-36
  - access path, 1-36
  - access path and hints cannot override, 1-37
- SAMPLE clause, 1-36
  - access path, 1-36
  - access path and hints cannot override, 1-37
  - cost-based optimization, 1-4
- sample table scans, 1-36
  - hints cannot override, 1-37
- sar UNIX command, 16-12
- scans
  - bounded range, 8-10
  - cluster, 8-4, 8-5, 8-6, 8-7
  - cluster key, 8-7
  - composite index, 8-8
  - full table, 8-14
  - full table using rule-based optimizer, 8-14
  - hash cluster, 8-5, 8-7
  - index, 1-28
  - index bounded range, 8-10
  - index cluster key, 8-7
  - index joins, 1-35
  - index of type bitmap, 1-35
  - index restrictions, 8-14
  - index with ORDER BY, 8-13
  - MAX or MIN of index, 8-13
  - range, 8-8
  - range of MAX or MIN, 8-13
  - range with ORDER BY, 8-13
  - sample table, 1-36
  - sample table and hints cannot override, 1-37
  - single-column index, 8-8
  - unbounded range, 8-11
  - unbounded range index, 8-11
  - unique, 8-6, 8-7
- scattered read wait events, 22-29
  - actions, 22-30
- schemas
  - star schemas, 1-44
- SCPU item, 12-16
- scripts
  - registering automatically, 11-9
  - SPAUTO.SQL, 21-9
  - SPCPKG.SQL, 21-6
  - SPCREATE.SQL, 21-5
  - SPCTAB.SQL, 21-6
  - SPCUSR.SQL, 21-6
  - SPPURGE.SQL, 21-23
  - SPTRUNC.SQL, 21-25
  - Statspack documentation scripts, 21-28
  - Statspack installation scripts, 21-27
  - Statspack performance data maintenance scripts, 21-28
  - Statspack reporting and automation scripts, 21-27
  - Statspack supplied scripts, 21-26
  - upgrading Statspack scripts, 21-27
- SEARCH\_COLUMNS column
  - PLAN\_TABLE table, 9-24
- segment-level statistics, 22-14
- SELECT statement
  - SAMPLE clause, 1-36
  - SAMPLE clause and access path, 1-36, 1-37
  - SAMPLE clause and cost-based optimization, 1-4
- selectivity, 3-2
  - histograms, 3-20
  - improving for an index, 4-4
  - indexes, 4-3, 4-6
  - SQL statement predicate, 3-2
  - user-defined, 1-62
- select-project-join views, 2-35
- semi-joins, 1-43
  - transformations not allowed, 1-44
- sequential read wait events
  - actions, 22-32
- SERVEROUTPUT
  - tuning, 11-10
- Session Data Unit (SDU), 23-13
- session id, 21-19
- session memory statistic, 14-37
- SESSION\_CACHED\_CURSORS initialization parameter, 14-40
- SESSIONS initialization parameter, 13-14
- SET AUTOTRACE, 11-2
- SET command
  - APPINFO variable, 11-9
  - ARRAYSIZE variable, 11-10

- SET TRANSACTION statement, 18-3
- setting
  - system variables for SQL\*PLUS
    - performance, 11-9
- SGA size, 14-47
- shared server
  - performance issues, 19-2
  - reducing contention, 19-2
  - tuning, 19-2
  - tuning memory, 14-35
- shared SQL areas
  - memory allocation, 14-34
- SHARED\_POOL\_RESERVED\_SIZE initialization parameter, 14-42
- SHARED\_POOL\_SIZE initialization parameter, 13-14, 14-34, 14-42
  - allocating library cache, 14-34
  - tuning the shared pool, 14-38
- sharing data, 21-22
- SHOW SGA statement, 14-5
- sizing redo logs, 13-5
- snapshot levels, 21-15, 21-17
- snapshot thresholds, 21-15, 21-17
- snapshots
  - begin and end, 21-10
  - databases identifier (DBID), 21-3
  - deleting, 21-23
  - instance numbers, 21-3
  - levels, 21-15, 21-17
  - removing, 21-23
  - SNAP\_ID, 21-3
  - Statspack, 21-3
  - taken by Statspack, 21-3
  - taking snapshots, 21-7
  - thresholds, 21-15, 21-17
- SOME operator, 2-22
- sort areas
  - memory allocation, 14-69
- sort merge joins, 1-49
  - access path, 8-12
  - cost-based optimization, 1-42
  - example, 8-12
- SORT\_AREA\_SIZE initialization parameter, 1-60, 4-14, 13-12
  - configuring, 14-66
  - cost-based optimization and, 1-42
  - See also PGA\_AGGREGATE\_TARGET
    - initialization parameter
    - tuning sorts, 14-70
  - use PGA\_AGGREGATE\_TARGET, 1-42, 1-60, 13-12
- sorts
  - (disk) statistic, 14-68
  - (memory) statistic, 14-68
  - avoiding on index creation, 14-71
- space transaction locks, 24-18
- SPAUTO.SQL script, 21-8, 21-9, 21-27
- SPCPKG.LIS output file, 21-6
- SPCPKG.SQL script, 21-6, 21-27
- SPCREATE.SQL script, 21-5, 21-27
  - running, 21-6
- SPCTAB.LIS output file, 21-6
- SPCTAB.SQL script, 21-6, 21-27
- SPCUSR.LIS output file, 21-6
- SPCUSR.SQL script, 21-6, 21-27
- SPDOC.TXT
  - Statspack documentation, 21-28
- SPDROP.SQL script, 21-26, 21-27
- SPDTAB.LIS output file, 21-26
- SPDTAB.SQL script, 21-26, 21-27
- SPDUSR.LIS output file, 21-26
- SPDUSR.SQL script, 21-26, 21-27
- SPPURGE.SQL script, 21-23, 21-28
- SPREPINS.SQL script, 21-27
- SPREPORT.SQL script, 21-27
  - performance report, 21-9
  - running the script, 21-10
- SPREPSQL.SQL script, 21-27
  - performance report, 21-10
- SPTRUNC.SQL script, 21-25, 21-28
- SPUEXP.PAR parameter file, 21-28
- SPUP816.SQL script, 21-27
- SPUP817.SQL script, 21-27
- SPUP90.SQL script, 21-27
- SQL functions
  - optimizing view queries, 2-43
- SQL Parse event, 12-2
- SQL statements
  - avoiding the use of indexes, 4-6
  - complex, 2-33

- complex optimizing, 2-33
- converting examples, 2-30, 8-17
- distributed optimization, 2-13
- ensuring the use of indexes, 4-6
- execution plans of, 1-18
- modifying indexed data, 4-4
- optimization of complex statements, 2-33
- optimization of distributed, 2-13
- optimizing complex, 2-33
- processing overview, 1-2
- thresholds, 21-15, 21-17
- waiting for I/O, 22-31
- SQL trace facility, 10-2, 10-6
  - example of output, 10-16
  - output, 10-13
  - statement truncation, 10-15
  - steps to follow, 10-3
  - trace files, 10-5
- SQL\*Loader, 13-11
- SQL\*Net message from client idle events, 22-26
- SQL\*Net message from dblink wait events, 22-27
- SQL\*Plus
  - autotrace, 11-1
  - BEGIN\_SNAP variable, 21-12
  - CLEAR TIMING command, 11-7
  - DEFAULT\_TABLESPACE variable, 21-6
  - END\_SNAP variable, 21-12
  - REPORT\_NAME variable, 21-12
  - statistics, 11-4
  - system variables influencing performance, 11-9
  - TEMPORARY\_TABLESPACE variable, 21-6
  - TIMING command, 11-7
- SQL\_STATEMENT column
- TKPROF\_TABLE, 10-18
- SQL\_TRACE
  - initialization parameter, 10-6
- SQL.BSQ file, 13-4
- SQLSegment event, 12-15
- ST locks, 24-18
- standard resource utilization items, 12-16
- STAR hint, 5-23
- star joins, 1-44
- star query, 1-44
- star transformation, 5-21
- STAR\_TRANSFORMATION hint, 5-21
- STAR\_TRANSFORMATION\_ENABLED
  - initialization parameter, 1-61, 5-22
- start columns
  - in partitioning and EXPLAIN PLAN statement, 9-13
- START statement in Oracle Trace, 12-3, 12-4
- STATEMENT\_ID column
  - PLAN\_TABLE table, 9-23
- statistics
  - automated collecting, 21-8
  - automated gathering, 21-8
  - collecting, 11-7, 21-8
  - collection interval, 21-9
  - consistent gets, 14-9, 18-3
  - database server, 11-4
  - db block gets, 14-9, 18-3
  - estimates and block sampling, 3-4
  - estimates and row sampling, 3-4
  - exporting and importing, 3-2
  - extensible optimization, 1-61
  - from B-tree or bitmap index, 3-9
  - gathering with DBMS\_STATS package, 3-6
  - generating, 3-3
  - generating and managing with DBMS\_STATS, 3-5
  - generating for cost-based optimization, 3-3
  - HIGH\_VALUE and LOW\_VALUE, 1-38
  - max session memory, 14-37
  - optimizer mode, 1-7
  - optimizer use of, 1-10, 3-2
  - partitions and subpartitions, 3-4
  - physical reads, 14-10
  - segment-level, 22-14
  - selectivity of predicates, 3-2
  - selectivity of predicates and histograms, 3-20
  - selectivity of predicates for user-defined, 1-62
  - session memory, 14-37
  - shared server processes, 19-5
  - sorts (disk), 14-68
  - sorts (memory), 14-68
  - SQL\*Plus, 11-4
  - user-defined statistics, 1-62
- Statspack
  - and Oracle Real Application Clusters, 21-25
  - automatic statistics gathering, 21-8

- compared with BSTAT/ESTAT, 20-7, 21-3
- DBID, 21-25
- documentation, 21-28
- exporting data, 21-22
- installation scripts, 21-27
- installing in batch mode, 21-6
- installing interactively, 21-4
- INSTANCE\_NUMBER, 21-25
- level 7 snapshot, 21-18
- performance data maintenance scripts, 21-28
- removing, 21-26
- reporting and automation scripts, 21-27
- running reports, 21-10
- scripts, 21-26
- sharing data, 21-22
- SNAP\_ID, 21-3
- snapshots, 21-3
- space requirements, 21-4
- SPCREATE.SQL, 21-5
- SPDOC.TXT, 21-28
- uninstalling, 21-26
- upgrading scripts, 21-27
- using DBMS\_JOB to gather statistics, 21-8
- STATSPACK.MODIFY\_STATSPACK\_
  - PARAMETER procedure, 21-16, 21-19
- STATSPACK.SNAP procedure, 21-7, 21-8, 21-19
- stop columns
  - in partitioning and EXPLAIN PLAN
    - statement, 9-13
- STOP statement in Oracle Trace, 12-3, 12-6
- STORAGE clause
  - OPTIMAL parameter, 18-3
- stored outlines
  - creating and using, 7-4
  - execution plans and plan stability, 7-2
  - hints, 7-3
  - moving tables, 7-11
  - storage requirements, 7-4
  - using, 7-6
  - viewing data for, 7-10
- striping, 15-4
  - manual, 15-16
- subpartitions
  - statistics, 3-4
- subqueries

- converting to joins, 2-33
  - NOT IN, 1-43
  - optimizing IN subquery, 2-37
- subquery unnesting, 6-21
- substitution variables
  - parsing, 11-10
- swapping, 16-12
  - reducing, 14-4
- switching processes, 16-13
- System Global Area tuning, 14-4
- system statistics
  - gathering, 3-6
- system variables
  - influencing SQL\*Plus performance, 11-9

## T

---

- tables
  - creating, 13-7
  - dimensions and star queries, 1-44
  - fact tables and star queries, 1-44
  - formatter in Oracle Trace, 12-3
  - full scans, 22-34
  - lookup tables, 1-44
  - placement on disk, 15-16
  - setting storage options, 13-7
- tablespaces, 13-6
  - creating, 13-6
  - dictionary managed, 21-4
  - locally managed, 21-4
  - rollback, 13-6
  - temporary, 13-6
- TCP.NODELAY parameter, 23-14
- temporary tablespaces, 13-6
- TEMPORARY\_TABLESPACE variable, 21-6
- thrashing, 16-12
- thread, 16-5
- thresholds
  - SQL statement, 21-17
  - SQL statements, 21-15
- throughput
  - cost-based approach, 1-7
  - optimizer goal, 1-5
  - optimizing, 1-5, 5-7
- TIMED\_STATISTICS initialization parameter, 21-7

- SQL Trace, 10-4
- TIMESTAMP column
  - PLAN\_TABLE table, 9-23
- TIMING command
  - SQL\*Plus, 11-7
- TKPROF program, 10-3, 10-6
  - editing the output SQL script, 10-17
  - example of output, 10-16
  - generating the output SQL script, 10-17
  - syntax, 10-8
  - using the EXPLAIN PLAN statement, 10-11
- TKPROF\_TABLE, 10-18
  - querying, 10-18
- TM locks, 24-18
- Trace, Oracle, 12-1
- tracing
  - queries, 11-7
- tracing statements
  - for performance statistics, 11-5
  - for query execution path, 11-5
  - using a database link, 11-6
  - with parallel query option, 11-7
- transactions
  - assigning rollback segments, 18-3
- TRIMOUT
  - tuning, 11-10
- TRIMSPool
  - tuning, 11-11
- truncating data, 21-25
- tuning
  - DEFINE OFF, 11-10
  - FLUSH OFF, 11-10
  - iSQL\*Plus parameters, 11-13
  - latches, 24-12
  - logical structure, 4-2
  - memory allocation, 14-5
  - resource contention, 22-1
  - SERVEROUTPUT, 11-10
  - SET ARRAYSIZE, 11-10
  - shared server, 19-2
  - SQL statements, 11-2
  - System Global Area (SGA), 14-4
  - TRIMOUT, 11-10
  - TRIMSPool, 11-11
- TX locks, 24-18

## U

---

- UCPU item, 12-16
- UL locks, 24-19
- UNION ALL operator
  - examples, 2-31, 2-33, 2-48
  - optimizing view queries, 2-36
  - transforming OR into, 2-30, 8-17
- UNION operator
  - examples, 2-39, 2-49
  - optimizing view queries, 2-36
- UNIQUE constraint, 4-8
- unique keys
  - optimization, 2-34
  - searches, 8-6
- uniqueness, 4-8
- UNIX system performance, 16-7
- UNNEST hint, 5-36
- upgrade
  - to the cost-based optimizer, 7-14
- USE\_CONCAT hint, 5-18
- USE\_MERGE hint, 5-25
- USE\_NL hint, 5-24
- USE\_STORED\_OUTLINES parameter, 7-6
- user defined locks, 24-19
- user global area (UGA)
  - shared servers, 14-35, 19-2
  - VSESSTAT, 14-37
- USER\_DUMP\_DEST initialization parameter, 10-4
  - SQL Trace, 10-4
- USER\_ID column
  - TKPROF\_TABLE, 10-19
- USER\_OUTLINE\_HINTS view
  - stored outline hints, 7-10
- USER\_OUTLINES view
  - stored outlines, 7-10
- USER\_TAB\_COL\_STATISTICS view, 1-38
- USER\_TAB\_COLUMNS view, 1-38
- USER\_TABLES view, 1-38
- user-defined costs, 1-63
- UTLCHN1.SQL script, 22-22
- utlxplp.sql
  - SQL script for viewing EXPLAIN PLANS, 1-19
- utlxpls.sql
  - SQL script for viewing EXPLAIN PLANS, 1-19



## V

---

VSBH view, 14-15  
VSBUFFER\_POOL\_STATISTICS view, 14-15  
VSDATAFILE view, 24-7  
VSDB\_CACHE\_ADVICE view, 14-6, 14-9, 14-11, 14-12, 14-15  
VSDB\_OBJECT\_CACHE view, 24-5  
VSFAST\_START\_SERVERS view, 17-19  
VSFAST\_START\_TRANSACTIONS view, 17-19  
VSFILESTAT view, 24-6  
VSINSTANCE\_RECOVERY view, 17-9  
VSLATCH view, 24-9  
VSLATCH\_CHILDREN view, 24-13  
VSLATCHHOLDER view, 24-13  
VSLIBRARY\_CACHE\_MEMORY view, 24-16  
VSLIBRARYCACHE view, 24-15  
    NAMESPACE column, 14-29  
VSLOCK view, 24-17  
VSMTTR\_TARGET\_ADVICE view, 24-21  
VSMYSTAT view, 24-22  
VSOPEN\_CURSOR view, 24-23  
VSPARAMETER view, 24-25  
VSPROCESS view, 24-26  
VSQUEUE view, 19-5  
VSROLLSTAT view, 24-28  
V\$ROWCACHE view, 24-29  
    GETMISSES column, 14-34  
    GETS column, 14-34  
    performance statistics, 14-32  
VSRSRC\_CONSUMER\_GROUP view, 22-8  
VSSEGMENT\_STATISTICS view, 24-31  
VSSEGSTAT view, 24-32  
VSSEGSTAT\_NAME view, 24-32  
VSSESSION view, 24-33  
VSSESSION\_EVENT view, 24-36  
    network information, 23-6  
VSSESSION\_WAIT view, 22-13, 24-37  
    network information, 23-6  
VSSESSTAT view, 22-8, 24-41  
    network information, 23-6  
    statistics, 24-42  
    using, 14-37  
VSSHARED\_POOL\_ADVICE view, 24-45  
VSSHARED\_POOL\_RESERVED view, 14-42

VSSQL view, 24-45  
    PLAN\_HASH\_VALUE column, 24-45  
VSSQL\_PLAN view, 24-46  
VSSQL\_PLAN\_STATISTICS view, 24-51  
VSSQL\_PLAN\_STATISTICS\_ALL view, 24-53  
VSSQLAREA view, 24-57  
VSSQLTEXT view, 24-59  
VSSTATISTICS\_LEVEL view, 24-61  
VSSYSSTAT view, 24-61  
    redo buffer allocation, 14-47  
    statistics, 24-63  
    tuning sorts, 14-68  
    using, 14-9  
VSSYSTEM\_EVENT view, 24-67  
VSSYSTEM\_PARAMETER view, 24-25  
VSUNDOSTAT view, 13-15, 24-69  
V\$WAITSTAT view, 22-13, 24-70  
VARCHAR datatype, 13-3  
VARCHAR2 datatype, 13-3  
variables  
    bind variables and optimization, 1-38  
views  
    complex view merging, 2-37  
    counter/accumulator, 24-2  
    current state views, 24-2  
    dynamic performance, 24-2  
    histograms, 3-25  
    information views, 24-4  
    non-null values for nulls, 2-45  
    select-project-join views, 2-35  
    statistics, 3-15  
vmstat UNIX command, 16-12

## W

---

### wait events

- buffer busy waits, 22-27
- contention wait events, 22-41
- direct path, 22-35
- event timings, 21-21
- free buffer waits, 22-39
- idle wait events, 22-49
- log file switch, 22-46
- network communication wait events, 22-25
- reasons for, 24-70
- resource wait events, 22-31
- time units, 21-20