

Oracle® Database

Advanced Security Guide

12c Release 2 (12.2)

E49707-21

April 2017

Oracle Database Advanced Security Guide, 12c Release 2 (12.2)

E49707-21

Copyright © 1996, 2017, Oracle and/or its affiliates. All rights reserved.

Primary Author: Patricia Huey

Contributors: Rahil Mir, Paul Youn, Adam Lee, Preetam Ramakrishna, Gopal Mulagund, Rajbir Chahal, Min-Hank Ho, Michael Hwa, Sudha Iyer, Adam Lindsey, Supriya Kalyanasundaram, Lakshmi Kethana, Andrew Koyfman, Vikram Pesati, Andy Philips, Dinesh Rajasekharan, Saikat Saha, Philip Thornton, Lixia Yuan, Peter Wahl

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	xi
Audience	xi
Documentation Accessibility	xi
Related Documents.....	xi
Conventions.....	xii
Changes in This Release for Oracle Database Advanced Security Guide	xiii
Changes in Oracle Database Advanced Security 12c Release 2 (12.2.0.1)	xiii
New Features	xiii
1 Introduction to Oracle Advanced Security	
1.1 Transparent Data Encryption.....	1-1
1.2 Oracle Data Redaction.....	1-1
Part I Using Transparent Data Encryption	
2 Introduction to Transparent Data Encryption	
2.1 What Is Transparent Data Encryption?	2-1
2.2 Benefits of Using Transparent Data Encryption	2-1
2.3 Who Can Configure Transparent Data Encryption?	2-2
2.4 Types and Components of Transparent Data Encryption.....	2-2
2.4.1 About Transparent Data Encryption Types and Components.....	2-3
2.4.2 How Transparent Data Encryption Column Encryption Works	2-3
2.4.3 How Transparent Data Encryption Tablespace Encryption Works	2-4
2.4.4 How the Keystore for the Storage of TDE Master Encryption Keys Works.....	2-6
2.4.5 Supported Encryption and Integrity Algorithms.....	2-7
3 Configuring Transparent Data Encryption	
3.1 Configuring a Software Keystore.....	3-1
3.1.1 About Configuring a Software Keystore	3-2
3.1.2 Step 1: Set the Keystore Location in the sqlnet.ora File	3-2

3.1.3	Step 2: Create the Software Keystore.....	3-6
3.1.4	Step 3: Open the Software Keystore	3-9
3.1.5	Step 4: Set the Software TDE Master Encryption Key.....	3-11
3.1.6	Step 5: Encrypt Your Data.....	3-13
3.2	Configuring a Hardware Keystore.....	3-13
3.2.1	About Configuring a Hardware (External) Keystore.....	3-14
3.2.2	Step 1: Set the Hardware Keystore Type in the sqlnet.ora File	3-14
3.2.3	Step 2: Configure the Hardware Security Module	3-14
3.2.4	Step 3: Open the Hardware Keystore	3-15
3.2.5	Step 4: Set the Hardware Keystore TDE Master Encryption Key	3-17
3.2.6	Step 5: Encrypt Your Data.....	3-19
3.3	Encrypting Columns in Tables	3-19
3.3.1	About Encrypting Columns in Tables.....	3-20
3.3.2	Data Types That Can Be Encrypted with TDE Column Encryption	3-20
3.3.3	Restrictions on Using Transparent Data Encryption Column Encryption	3-21
3.3.4	Creating Tables with Encrypted Columns	3-22
3.3.5	Encrypting Columns in Existing Tables.....	3-26
3.3.6	Creating an Index on an Encrypted Column	3-27
3.3.7	Adding Salt to an Encrypted Column.....	3-28
3.3.8	Removing Salt from an Encrypted Column	3-28
3.3.9	Changing the Encryption Key or Algorithm for Tables with Encrypted Columns...	3-28
3.4	Encryption Conversions for Tablespaces and Databases	3-29
3.4.1	About Encryption Conversions for Tablespaces and Databases.....	3-29
3.4.2	Restrictions on Using Transparent Data Encryption Tablespace Encryption	3-31
3.4.3	Creating an Encrypted New Tablespace.....	3-32
3.4.4	Encrypting Future Tablespaces	3-36
3.4.5	Encryption Conversions for Existing Offline Tablespaces.....	3-37
3.4.6	Encryption Conversions for Existing Online Tablespaces	3-41
3.4.7	Encryption Conversions for Existing Databases.....	3-48
3.5	Transparent Data Encryption Data Dynamic and Data Dictionary Views	3-51

4 Managing the Keystore and the TDE Master Encryption Key

4.1	Managing the Keystore.....	4-1
4.1.1	Performing Operations That Require a Keystore Password.....	4-2
4.1.2	Changing the Password of a Software Keystore.....	4-3
4.1.3	Changing the Password of a Hardware Keystore	4-4
4.1.4	Backing Up Password-Based Software Keystores.....	4-5
4.1.5	Backups of the Hardware Keystore	4-7
4.1.6	Merging Software Keystores.....	4-8
4.1.7	Moving a Software Keystore to a New Location.....	4-11
4.1.8	Moving a Software Keystore Out of Automatic Storage Management.....	4-12
4.1.9	Migrating Between a Software Password Keystore and a Hardware Keystore	4-13
4.1.10	Migration of Keystores to and from Oracle Key Vault.....	4-20

4.1.11	Closing a Keystore.....	4-20
4.1.12	Using a Software Keystore That Resides on ASM Volumes	4-23
4.1.13	Backup and Recovery of Encrypted Data	4-23
4.1.14	Deletion of Keystores.....	4-24
4.2	Managing the TDE Master Encryption Key.....	4-25
4.2.1	Creating TDE Master Encryption Keys for Later Use.....	4-25
4.2.2	Activation of TDE Master Encryption Keys	4-28
4.2.3	TDE Master Encryption Key Attribute Management	4-31
4.2.4	Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes .	4-32
4.2.5	Setting or Rotating the TDE Master Encryption Key in the Keystore	4-34
4.2.6	Exporting and Importing the TDE Master Encryption Key	4-39
4.2.7	Management of TDE Master Encryption Keys Using Oracle Key Vault	4-45
4.3	Storing Secrets Used by Oracle Database.....	4-45
4.3.1	About Storing Oracle Database Secrets in a Keystore	4-46
4.3.2	Storage of Oracle Database Secrets in a Software Keystore	4-47
4.3.3	Example: Adding an HSM Password to a Software Keystore.....	4-48
4.3.4	Example: Changing an HSM Password Stored as a Secret in a Software Keystore ..	4-49
4.3.5	Example: Deleting an HSM Password Stored as a Secret in a Software Keystore.....	4-49
4.3.6	Storage of Oracle Database Secrets in a Hardware Keystore.....	4-49
4.3.7	Example: Adding an Oracle Database Secret to a Hardware Keystore.....	4-51
4.3.8	Example: Changing an Oracle Database Secret in a Hardware Keystore.....	4-51
4.3.9	Example: Deleting an Oracle Database Secret in a Hardware Keystore	4-51
4.3.10	Configuring Auto-Login Hardware Security Modules	4-52
4.4	Storing Oracle GoldenGate Secrets in a Keystore	4-53
4.4.1	About Storing Oracle GoldenGate Secrets in Keystores.....	4-53
4.4.2	Oracle GoldenGate Extract Classic Capture Mode TDE Requirements.....	4-54
4.4.3	Configuring TDE Keystore Support for Oracle GoldenGate	4-54

5 General Considerations of Using Transparent Data Encryption

5.1	Compression and Data Deduplication of Encrypted Data.....	5-1
5.2	Security Considerations for Transparent Data Encryption	5-2
5.2.1	Transparent Data Encryption General Security Advice	5-2
5.2.2	Transparent Data Encryption Column Encryption-Specific Advice	5-3
5.2.3	Managing Security for Plaintext Fragments.....	5-3
5.3	Performance and Storage Overhead of Transparent Data Encryption.....	5-3
5.3.1	Performance Overhead of Transparent Data Encryption.....	5-4
5.3.2	Storage Overhead of Transparent Data Encryption.....	5-4
5.4	Modifying Your Applications for Use with Transparent Data Encryption	5-5
5.5	How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT	5-6
5.6	Using Transparent Data Encryption with PKI Encryption.....	5-9
5.6.1	Software Master Encryption Key Use with PKI Key Pairs.....	5-10
5.6.2	TDE Tablespace and Hardware Keystores with PKI Encryption	5-10
5.6.3	Backup and Recovery of a PKI Key Pair	5-10

5.7	Data Loads from External Files to Tables with Encrypted Columns	5-11
5.8	Transparent Data Encryption and Database Close Operations	5-11

6 Using Transparent Data Encryption with Other Oracle Features

6.1	How Transparent Data Encryption Works with Export and Import Operations	6-1
6.1.1	About Exporting and Importing Encrypted Data	6-2
6.1.2	Exporting and Importing Tables with Encrypted Columns	6-2
6.1.3	Using Oracle Data Pump to Encrypt Entire Dump Sets	6-3
6.2	How Transparent Data Encryption Works with Oracle Data Guard	6-4
6.3	How Transparent Data Encryption Works with Oracle Real Application Clusters	6-5
6.3.1	About Using Transparent Data Encryption with Oracle Real Application Clusters ..	6-5
6.3.2	Using a Non-Shared File System to Store a Software Keystore in Oracle RAC	6-5
6.4	How Transparent Data Encryption Works with SecureFiles	6-7
6.4.1	About Transparent Data Encryption and SecureFiles	6-7
6.4.2	Example: Creating a SecureFiles LOB with a Specific Encryption Algorithm	6-7
6.4.3	Example: Creating a SecureFiles LOB with a Column Password Specified	6-7
6.5	How Transparent Data Encryption Works in a Multitenant Environment	6-8
6.5.1	About Using Transparent Data Encryption in a Multitenant Environment	6-8
6.5.2	Operations That Must Be Performed in Root	6-9
6.5.3	Operations That Can Be Performed in Root or in a PDB	6-11
6.5.4	Moving PDBs from One CDB to Another	6-11
6.5.5	Exporting and Importing TDE Master Encryption Keys for a PDB	6-12
6.5.6	Unplugging and Plugging a PDB with Encrypted Data in a CDB	6-14
6.5.7	Managing Cloned PDBs with Encrypted Data	6-17
6.5.8	How Keystore Open and Close Operations Work in a Multitenant Environment ...	6-18
6.5.9	Finding the Keystore Status for All of the PDBs in a Multitenant Environment	6-19
6.6	How Transparent Data Encryption Works with Oracle Call Interface	6-20
6.7	How Transparent Data Encryption Works with Editions	6-20
6.8	Configuring Transparent Data Encryption to Work in a Multidatabase Environment	6-21

7 Frequently Asked Questions About Transparent Data Encryption

7.1	Transparency Questions About Transparent Data Encryption	7-1
7.2	Performance Questions About Transparent Data Encryption	7-4

Part II Using Oracle Data Redaction

8 Introduction to Oracle Data Redaction

8.1	What Is Oracle Data Redaction?	8-1
8.2	When to Use Oracle Data Redaction	8-2
8.3	Benefits of Using Oracle Data Redaction	8-2
8.4	Target Use Cases for Oracle Data Redaction	8-2
8.4.1	Oracle Data Redaction Use with Database Applications	8-3
8.4.2	Oracle Data Redaction with Ad Hoc Database Queries Considerations	8-3

9 Oracle Data Redaction Features and Capabilities

9.1 Full Data Redaction to Redact All Data.....	9-1
9.2 Partial Data Redaction to Redact Sections of Data.....	9-2
9.3 Regular Expressions to Redact Patterns of Data	9-3
9.4 Redaction Using Null Values.....	9-4
9.5 Random Data Redaction to Generate Random Values	9-4
9.6 Comparison of Full, Partial, and Random Redaction Based on Data Types	9-5
9.6.1 Oracle Built-in Data Types Redaction Capabilities	9-6
9.6.2 ANSI Data Types Redaction Capabilities	9-6
9.6.3 Built-in and ANSI Data Types Full Redaction Capabilities	9-7
9.6.4 User-Defined Data Types or Oracle Supplied Types Redaction Capabilities	9-9
9.7 No Redaction for Testing Purposes.....	9-9
9.8 Central Management of Named Data Redaction Policy Expressions.....	9-9

10 Configuring Oracle Data Redaction Policies

10.1 About Oracle Data Redaction Policies.....	10-2
10.2 Who Can Create Oracle Data Redaction Policies?	10-3
10.3 Planning an Oracle Data Redaction Policy	10-4
10.4 General Syntax of the DBMS_REDACT.ADD_POLICY Procedure.....	10-4
10.5 Using Expressions to Define Conditions for Data Redaction Policies	10-6
10.5.1 About Using Expressions in Data Redaction Policies.....	10-7
10.5.2 Supported Functions for Data Redaction Expressions	10-7
10.5.3 Applying the Redaction Policy Based on User Environment.....	10-11
10.5.4 Applying the Redaction Policy Based on Database Roles.....	10-11
10.5.5 Applying the Redaction Policy Based on Oracle Label Security Label Dominance	10-12
10.5.6 Applying the Redaction Policy Based on Application Express Session States	10-12
10.5.7 Applying the Redaction Policy to All Users.....	10-13
10.6 Creating and Managing Multiple Named Policy Expressions	10-13
10.6.1 About Data Redaction Policy Expressions to Define Conditions.....	10-13
10.6.2 Creating and Applying a Named Data Redaction Policy Expression	10-15
10.6.3 Updating a Named Data Redaction Policy Expression	10-16
10.6.4 Dropping a Named Data Redaction Expression Policy.....	10-16
10.6.5 Tutorial: Creating and Sharing a Named Data Redaction Policy Expression	10-17
10.7 Creating a Full Redaction Policy and Altering the Full Redaction Value	10-24
10.7.1 Creating a Full Redaction Policy.....	10-24
10.7.2 Altering the Default Full Data Redaction Value.....	10-26
10.8 Creating a DBMS_REDACT.NULLIFY Redaction Policy.....	10-28
10.8.1 About Creating a Policy That Returns Null Values	10-29
10.8.2 Syntax for Creating a Policy That Returns Null Values	10-29
10.8.3 Example: Redaction Policy That Returns Null Values.....	10-29
10.9 Creating a Partial Redaction Policy.....	10-30
10.9.1 About Creating Partial Redaction Policies	10-30

10.9.2	Syntax for Creating a Partial Redaction Policy	10-31
10.9.3	Creating Partial Redaction Policies Using Fixed Character Formats	10-31
10.9.4	Creating Partial Redaction Policies Using Character Data Types.....	10-35
10.9.5	Creating Partial Redaction Policies Using Number Data Types	10-36
10.9.6	Creating Partial Redaction Policies Using Date-Time Data Types	10-38
10.10	Creating a Regular Expression-Based Redaction Policy	10-39
10.10.1	About Creating Regular Expression-Based Redaction Policies.....	10-39
10.10.2	Syntax for Creating a Regular Expression-Based Redaction Policy	10-40
10.10.3	Regular Expression-Based Redaction Policies Using Formats	10-42
10.10.4	Custom Regular Expression Redaction Policies	10-46
10.11	Creating a Random Redaction Policy	10-47
10.11.1	Syntax for Creating a Random Redaction Policy.....	10-48
10.11.2	Example: Random Redaction Policy.....	10-48
10.12	Creating a Policy That Uses No Redaction	10-49
10.12.1	Syntax for Creating a Policy with No Redaction	10-49
10.12.2	Example: Performing No Redaction.....	10-49
10.13	Exemption of Users from Oracle Data Redaction Policies.....	10-50
10.14	Altering an Oracle Data Redaction Policy	10-51
10.14.1	About Altering Oracle Data Redaction Policies.....	10-51
10.14.2	Syntax for the DBMS_REDACT.ALTER_POLICY Procedure	10-51
10.14.3	Parameters Required for DBMS_REDACT.ALTER_POLICY Actions.....	10-52
10.14.4	Tutorial: Altering an Oracle Data Redaction Policy	10-53
10.15	Redacting Multiple Columns	10-56
10.15.1	Adding Columns to a Data Redaction Policy for a Single Table or View.....	10-57
10.15.2	Example: Redacting Multiple Columns	10-57
10.16	Disabling and Enabling an Oracle Data Redaction Policy.....	10-57
10.16.1	Disabling an Oracle Data Redaction Policy	10-57
10.16.2	Enabling an Oracle Data Redaction Policy	10-58
10.17	Dropping an Oracle Data Redaction Policy	10-59
10.18	Tutorial: SQL Expressions to Build Reports with Redacted Values.....	10-60
10.19	Oracle Data Redaction Policy Data Dictionary Views	10-61

11 Using Oracle Data Redaction in Oracle Enterprise Manager

11.1	About Using Oracle Data Redaction in Oracle Enterprise Manager.....	11-1
11.2	Oracle Data Redaction Workflow.....	11-2
11.3	Management of Sensitive Column Types in Enterprise Manager.....	11-2
11.4	Managing Oracle Data Redaction Formats Using Enterprise Manager	11-4
11.4.1	About Managing Oracle Data Redaction Formats Using Enterprise Manager.....	11-5
11.4.2	Creating a Custom Oracle Data Redaction Format Using Enterprise Manager	11-5
11.4.3	Editing a Custom Oracle Data Redaction Format Using Enterprise Manager.....	11-7
11.4.4	Viewing Oracle Data Redaction Formats Using Enterprise Manager	11-8
11.4.5	Deleting a Custom Oracle Data Redaction Format Using Enterprise Manager.....	11-9
11.5	Managing Oracle Data Redaction Policies Using Enterprise Manager	11-9

11.5.1	About Managing Oracle Data Redaction Policies Using Enterprise Manager.....	11-10
11.5.2	Creating an Oracle Data Redaction Policy Using Enterprise Manager.....	11-11
11.5.3	Editing an Oracle Data Redaction Policy Using Enterprise Manager	11-14
11.5.4	Viewing Oracle Data Redaction Policy Details Using Enterprise Manager	11-16
11.5.5	Enabling or Disabling an Oracle Data Redaction Policy in Enterprise Manager...	11-16
11.5.6	Deleting an Oracle Data Redaction Policy Using Enterprise Manager	11-17
11.6	Managing Named Data Redaction Policy Expressions Using Enterprise Manager	11-17
11.6.1	About Named Data Redaction Policy Expressions in Enterprise Manager.....	11-18
11.6.2	Creating a Named Data Redaction Policy Expression in Enterprise Manager	11-18
11.6.3	Editing a Named Data Redaction Policy Expression in Enterprise Manager	11-19
11.6.4	Viewing Named Data Redaction Policy Expressions in Enterprise Manager.....	11-20
11.6.5	Deleting a Named Data Redaction Policy Expression in Enterprise Manager	11-21

12 Oracle Data Redaction Use with Oracle Database Features

12.1	Oracle Data Redaction and DML and DDL Operations	12-2
12.2	Oracle Data Redaction and Nested Functions, Inline Views, and the WHERE Clause	12-2
12.3	Oracle Data Redaction and Database Links.....	12-3
12.4	Oracle Data Redaction and Aggregate Functions.....	12-3
12.5	Oracle Data Redaction and Object Types.....	12-3
12.6	Oracle Data Redaction and XML Generation	12-3
12.7	Oracle Data Redaction and Editions	12-3
12.8	Oracle Data Redaction in a Multitenant Environment.....	12-3
12.9	Oracle Data Redaction and Oracle Virtual Private Database.....	12-4
12.10	Oracle Data Redaction and Oracle Database Real Application Security	12-4
12.11	Oracle Data Redaction and Oracle Database Vault	12-4
12.12	Oracle Data Redaction and Oracle Data Pump	12-5
12.12.1	Oracle Data Pump Security Model for Oracle Data Redaction	12-5
12.12.2	Export of Objects That Have Oracle Data Redaction Policies Defined	12-5
12.12.3	Export of Data Using the EXPDP Utility access_method Parameter.....	12-7
12.12.4	Import of Data into Objects Protected by Oracle Data Redaction.....	12-7
12.13	Oracle Data Redaction and Data Masking and Subsetting Pack	12-8
12.14	Oracle Data Redaction and JSON.....	12-8

13 Security Considerations for Oracle Data Redaction

13.1	Oracle Data Redaction General Usage Guidelines	13-1
13.2	Restriction of Administrative Access to Oracle Data Redaction Policies	13-2
13.3	How Oracle Data Redaction Affects the SYS, SYSTEM, and Default Schemas	13-2
13.4	Policy Expressions That Use SYS_CONTEXT Attributes	13-3
13.5	Oracle Data Redaction Policies on Materialized Views	13-3
13.6	Dropped Oracle Data Redaction Policies When the Recycle Bin Is Enabled	13-3

Glossary

Index

Preface

Welcome to *Oracle Database Advanced Security Guide* for the 12c Release 2 (12.2) of Oracle Advanced Security. This guide describes how to implement, configure, and administer Oracle Advanced Security.

[Audience](#) (page xi)

[Documentation Accessibility](#) (page xi)

[Related Documents](#) (page xi)

[Conventions](#) (page xii)

Audience

Oracle Database Advanced Security Guide is intended for users and systems professionals involved with the implementation, configuration, and administration of Oracle Advanced Security including:

- Implementation consultants
- System administrators
- Security administrators
- Database administrators (DBAs)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

Before you configure Oracle Advanced Security features, you should be familiar with the following guides:

- *Oracle Database Administrator's Guide*
- *Oracle Database Security Guide*

Many books in the documentation set use the sample schemas of the default database. Refer to *Oracle Database Sample Schemas* for information about how these schemas were created and how you can use them.

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

<http://www.oracle.com/technetwork/index.html>

If you already have a user name and password for OTN, then you can go directly to the documentation section of the OTN website at

<http://www.oracle.com/technetwork/documentation/index.html>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Changes in This Release for Oracle Database Advanced Security Guide

This preface contains:

[Changes in Oracle Database Advanced Security 12c Release 2 \(12.2.0.1\)](#)
(page xiii)

Changes in Oracle Database Advanced Security 12c Release 2 (12.2.0.1)

The following are changes in *Oracle Database Advanced Security Guide* for Oracle Database 12c release 2 (12.2.0.1).

[New Features](#) (page xiii)

New Features

The following features are new to this release:

[Ability to Encrypt Existing Tablespaces and Fully Encrypt Databases](#) (page xiv)
Starting with this release, you can encrypt existing tablespaces and fully encrypt databases.

[Additional Supported Encryption Algorithms](#) (page xiv)
You now can use the ARIA, GOST, and SEED encryption algorithms for column and tablespace encryption, and the AES and DES encryption standards.

[Ability to Force Software Keystore Operations](#) (page xv)
You now can force a keystore operation that is prevented because of an in-use auto-login keystore or a closed software or hardware keystore.

[Ability to Use an External Store for Software Keystore Passwords](#) (page xv)
You now can configure a software or a hardware keystore to use an external store for its password.

[New Way to Specify Oracle Key Vault as a Keystore](#) (page xv)
As an alternative to third-party hardware security modules, you now can specify Oracle Key Vault as a keystore.

[Ability to Redact Data Based on Different Runtime Conditions](#) (page xv)
You now can define and associate different Data Redaction policy expressions with different columns within the same table or view.

[Ability to Centrally Manage Data Redaction Policy Expressions within a Database](#) (page xvi)

This new feature applies to named Oracle Data Redaction policy expressions.

[Ability to Use NULL as the Redacted Value](#) (page xvi)

Starting with this release, the redacted value can be NULL.

[Enhanced Support for Redacting Unstructured Data](#) (page xvi)

You now can define regular expression-based redaction (DBMS_REDACT.REGEXP) policies on columns of the CLOB and NCLOB data types.

Ability to Encrypt Existing Tablespaces and Fully Encrypt Databases

Starting with this release, you can encrypt existing tablespaces and fully encrypt databases.

In previous releases, you could only encrypt new tablespaces. However, this new feature enables you to encrypt both offline and online tablespaces. To encrypt a database, you encrypt the Oracle-supplied tablespaces, such as SYSTEM and SYSAUX. Offline tablespace conversion can be used for tablespaces in Oracle Database 11g release 1 (11.1) and Oracle Database 12c release 1 (12.1). You can perform encryption conversion operations in parallel and perform the encryption in an Oracle Data Guard environment. You can configure all future tablespaces to be automatically encrypted, which is beneficial for an Oracle Cloud environment.

Related Topics:

[Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

You can perform encryption operations on both offline and online tablespaces and databases.

Additional Supported Encryption Algorithms

You now can use the ARIA, GOST, and SEED encryption algorithms for column and tablespace encryption, and the AES and DES encryption standards.

The main benefit of these new encryption standards is that they meet the national standards for their respective countries.

- ARIA uses the same block sizes as AES. It is designed for lightweight environments and the implementation of hardware. ARIA meets the standards used in Korea.
- GOST is very similar to DES except that it has a large number of rounds and secret S-boxes. GOST meets the standards used in Russia.
- SEED is used by several standard protocols: S/MIME, TLS/SSL, IPSec, and ISO/IEC. SEED meets the standards used in Korea.

Related Topics:

[Supported Encryption and Integrity Algorithms](#) (page 2-7)

By default, Transparent Data Encryption (TDE) Column encryption uses the Advanced Encryption Standard (AES).

Ability to Force Software Keystore Operations

You now can force a keystore operation that is prevented because of an in-use auto-login keystore or a closed software or hardware keystore.

In previous releases, for many keystore operations, you had to manually open the software keystore before performing the operation. In this release, you can perform these two actions in one `ADMINISTER KEY MANAGEMENT` statement execution by including the `FORCE KEYSTORE` clause.

The operations that you can use the `FORCE KEYSTORE` clause on are as follows: rotating a keystore password; creating, using, rekeying, tagging, importing, exporting, migrating, or reverse migrating encryption keys; opening or backing up keystores; adding, updating, or deleting secret keystores.

Related Topics:

[Performing Operations That Require a Keystore Password](#) (page 4-2)

Many `ADMINISTER KEY MANAGEMENT` operations require access to a keystore password, for both software and hardware keystores.

Ability to Use an External Store for Software Keystore Passwords

You now can configure a software or a hardware keystore to use an external store for its password.

This feature enables you to store the password in a separate location where it can be centrally managed and accessed. To use this functionality, you must first set the `EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION` initialization parameter to a location where the external keystore credential will be stored. Afterward, you can include the `EXTERNAL_STORE` setting in the `IDENTIFIED BY` clause when you run the `ADMINISTER KEY MANAGEMENT` statement for the following operations: opening, closing, backing up the keystore; adding, updating, or deleting a secret keystore; creating, using, rekeying, tagging, importing, exporting encryption keys.

Related Topics:

[Configuring an External Store for a Keystore Password](#) (page 3-4)

An external store for a keystore password stores the software keystore password in a centrally accessed and managed location.

New Way to Specify Oracle Key Vault as a Keystore

As an alternative to third-party hardware security modules, you now can specify Oracle Key Vault as a keystore.

To configure Oracle Key Vault as a keystore, you can edit the `sqlnet.ora` file `METHOD` setting in the `WALLET_LOCATION` parameter to point to OKV.

Related Topics:

[Configuring a Hardware Keystore](#) (page 3-13)

A hardware keystore resides in a hardware security module (HSM), which is designed to store encryption keys.

Ability to Redact Data Based on Different Runtime Conditions

You now can define and associate different Data Redaction policy expressions with different columns within the same table or view.

This feature provides greater flexibility for anyone who creates Data Redaction policies.

For example, this feature enables you to share a single Data Redaction policy expression with multiple Data Redaction policies.

When you create the policy expression, you can apply it to any table or view column that is included in an existing Data Redaction policy. If you change the policy expression, the change is reflected in all Data Redaction policies that redact the associated table or view columns.

Related Topics:

[Creating and Managing Multiple Named Policy Expressions](#) (page 10-13)

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

Ability to Centrally Manage Data Redaction Policy Expressions within a Database

This new feature applies to named Oracle Data Redaction policy expressions.

This feature facilitates the maintenance and administration of policy expressions. When you modify the named policy expression, the changes are automatically applied to all tables and views in the database that use the expression.

Related Topics:

[Creating and Managing Multiple Named Policy Expressions](#) (page 10-13)

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

Ability to Use NULL as the Redacted Value

Starting with this release, the redacted value can be NULL.

For example, you can use this feature to hide data.

When you define an Oracle Data Redaction policy, you can set the `function_type` parameter to `DBMS_REDACT.NULLIFY` to ensure that the redacted value to always be NULL.

Related Topics:

[Creating a DBMS_REDACT.NULLIFY Redaction Policy](#) (page 10-28)

You can create Oracle Data Redaction policies that return null values for the displayed value of the table or view column.

Enhanced Support for Redacting Unstructured Data

You now can define regular expression-based redaction (`DBMS_REDACT.REGEXP`) policies on columns of the CLOB and NCLOB data types.

Related Topics:

[Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40)

The `regexp_*` parameters of the `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression-based redaction policy.

Introduction to Oracle Advanced Security

Two features comprise Oracle Advanced Security: Transparent Data Encryption and Oracle Data Redaction.

[Transparent Data Encryption](#) (page 1-1)

Transparent Data Encryption (TDE) enables you to encrypt data so that only an authorized recipient can read it.

[Oracle Data Redaction](#) (page 1-1)

Oracle Data Redaction enables you to redact (mask) column data using several redaction types.

1.1 Transparent Data Encryption

Transparent Data Encryption (TDE) enables you to encrypt data so that only an authorized recipient can read it.

Use encryption to protect sensitive data in a potentially unprotected environment, such as data you placed on backup media that is sent to an off-site storage location. You can encrypt individual columns in a database table, or you can encrypt an entire tablespace.

To use Transparent Data Encryption, you do not need to modify your applications. TDE enables your applications to continue working seamlessly as before. It automatically encrypts data when it is written to disk, and then automatically decrypts the data when your applications access it. Key management is built-in, eliminating the complex task of managing and securing encryption keys.

1.2 Oracle Data Redaction

Oracle Data Redaction enables you to redact (mask) column data using several redaction types.

The types of redaction that you can perform are as follows:

- **Full redaction.** You redact all of the contents of the column data. The redacted value that is returned to the querying user depends on the data type of the column. For example, columns of the `NUMBER` data type are redacted with a zero (0) and character data types are redacted with a blank space.
- **Partial redaction.** You redact a portion of the column data. For example, you can redact most of a Social Security number with asterisks (*), except for the last 4 digits.
- **Regular expressions.** You can use regular expressions in both full and partial redaction. This enables you to redact data based on a search pattern for the data. For example, you can use regular expressions to redact specific phone numbers or email addresses in your data.

- **Random redaction.** The redacted data presented to the querying user appears as randomly generated values each time it is displayed, depending on the data type of the column.
- **No redaction.** This option enables you to test the internal operation of your redaction policies, with no effect on the results of queries against tables with policies defined on them. You can use this option to test the redaction policy definitions before applying them to a production environment.

Data Redaction performs the redaction at runtime, that is, the moment that the user tries to view the data. This functionality is ideally suited for dynamic production systems in which data constantly changes. While the data is being redacted, Oracle Database is able to process all of the data normally and to preserve the back-end referential integrity constraints. Data redaction can help you to comply with industry regulations such as Payment Card Industry Data Security Standard (PCI DSS) and the Sarbanes-Oxley Act.

Part I

Using Transparent Data Encryption

Part I describes how to use Transparent Data Encryption.

[Introduction to Transparent Data Encryption](#) (page 2-1)

Transparent Data Encryption enables you to encrypt sensitive data, such as credit card numbers or Social Security numbers.

[Configuring Transparent Data Encryption](#) (page 3-1)

You can configure software or hardware keystores, for use on both individual table columns or entire tablespaces.

[Managing the Keystore and the TDE Master Encryption Key](#) (page 4-1)

You can modify settings for the keystore and TDE master encryption key, and store Oracle Database and store Oracle GoldenGate secrets in a keystore.

[General Considerations of Using Transparent Data Encryption](#) (page 5-1)

When you use Transparent Data Encryption, you should consider factors such as security, performance, and storage overheads.

[Using Transparent Data Encryption with Other Oracle Features](#) (page 6-1)

You can use Oracle Data Encryption with other Oracle features, such as Oracle Data Guard or Oracle Real Application Clusters.

[Frequently Asked Questions About Transparent Data Encryption](#) (page 7-1)

Users frequently have questions about transparency and performance issues with Transparent Data Encryption.

Introduction to Transparent Data Encryption

Transparent Data Encryption enables you to encrypt sensitive data, such as credit card numbers or Social Security numbers.

[What Is Transparent Data Encryption?](#) (page 2-1)

Transparent Data Encryption (TDE) enables you to encrypt sensitive data that you store in tables and tablespaces.

[Benefits of Using Transparent Data Encryption](#) (page 2-1)

Transparent Data Encryption (TDE) ensures that sensitive data is encrypted, meets compliance, and provides functionality that streamlines encryption operations.

[Who Can Configure Transparent Data Encryption?](#) (page 2-2)

You must be granted the `ADMINISTER KEY MANAGEMENT` system privilege to configure Transparent Data Encryption (TDE).

[Types and Components of Transparent Data Encryption](#) (page 2-2)

Transparent Data Encryption can be applied to individual columns or entire tablespaces.

2.1 What Is Transparent Data Encryption?

Transparent Data Encryption (TDE) enables you to encrypt sensitive data that you store in tables and tablespaces.

After the data is encrypted, this data is transparently decrypted for authorized users or applications when they access this data. TDE helps protect data stored on media (also called data at rest) in the event that the storage media or data file is stolen.

Oracle Database uses authentication, authorization, and auditing mechanisms to secure data in the database, but not in the operating system data files where data is stored. To protect these data files, Oracle Database provides Transparent Data Encryption (TDE). TDE encrypts sensitive data stored in data files. To prevent unauthorized decryption, TDE stores the encryption keys in a security module external to the database, called a keystore.

You can configure Oracle Key Vault as part of the TDE implementation. This enables you to centrally manage TDE keystores (called TDE wallets in Oracle Key Vault) in your enterprise. For example, you can upload a software keystore to Oracle Key Vault and then make the contents of this keystore available to other TDE-enabled databases. See *Oracle Key Vault Administrator's Guide* for more information.

2.2 Benefits of Using Transparent Data Encryption

Transparent Data Encryption (TDE) ensures that sensitive data is encrypted, meets compliance, and provides functionality that streamlines encryption operations.

Benefits are as follows:

- As a security administrator, you can be sure that sensitive data is encrypted and therefore safe in the event that the storage media or data file is stolen.
- Using TDE helps you address security-related regulatory compliance issues.
- You do not need to create auxiliary tables, triggers, or views to decrypt data for the authorized user or application. Data from tables is transparently decrypted for the database user and application. An application that processes sensitive data can use TDE to provide strong data encryption with little or no change to the application.
- Data is transparently decrypted for database users and applications that access this data. Database users and applications do not need to be aware that the data they are accessing is stored in encrypted form.
- You can encrypt data with zero downtime on production systems by using online table redefinition or you can encrypt it offline during maintenance periods. (See *Oracle Database Administrator's Guide* for more information about online table redefinition.)
- You do not need to modify your applications to handle the encrypted data. The database manages the data encryption and decryption.
- Oracle Database automates TDE master encryption key and keystore management operations. The user or application does not need to manage TDE master encryption keys.

2.3 Who Can Configure Transparent Data Encryption?

You must be granted the `ADMINISTER KEY MANAGEMENT` system privilege to configure Transparent Data Encryption (TDE).

If you must open the keystore at the mount stage, then you must be granted the `SYSKM` administrative privilege, which includes the `ADMINISTER KEY MANAGEMENT` system privilege and other necessary privileges.

When you grant the `SYSKM` administrative privilege to a user, ensure that you create a password file for it so that the user can connect to the database as `SYSKM` using a password. This enables the user to perform actions such as querying the `V$DATABASE` view.

To configure TDE column or tablespace encryption, you do not need the `SYSKM` or `ADMINISTER KEY MANAGEMENT` privileges. You must have the following additional privileges to create TDE policies on tables and tablespaces:

- `CREATE TABLE`
- `ALTER TABLE`
- `CREATE TABLESPACE`

2.4 Types and Components of Transparent Data Encryption

Transparent Data Encryption can be applied to individual columns or entire tablespaces.

About Transparent Data Encryption Types and Components (page 2-3)

You can encrypt sensitive data at the column level or the tablespace level.

How Transparent Data Encryption Column Encryption Works (page 2-3)

Transparent Data Encryption (TDE) column encryption protects confidential data, such as credit card and Social Security numbers, that is stored in table columns.

How Transparent Data Encryption Tablespace Encryption Works (page 2-4)

Transparent Data Encryption (TDE) tablespace encryption enables you to encrypt an entire tablespace.

How the Keystore for the Storage of TDE Master Encryption Keys Works (page 2-6)

To control the encryption, you use a keystore and TDE master encryption key.

Supported Encryption and Integrity Algorithms (page 2-7)

By default, Transparent Data Encryption (TDE) Column encryption uses the Advanced Encryption Standard (AES).

2.4.1 About Transparent Data Encryption Types and Components

You can encrypt sensitive data at the column level or the tablespace level.

At the column level, you can encrypt data using selected table columns. TDE tablespace encryption enables you to encrypt all of the data that is stored in a tablespace.

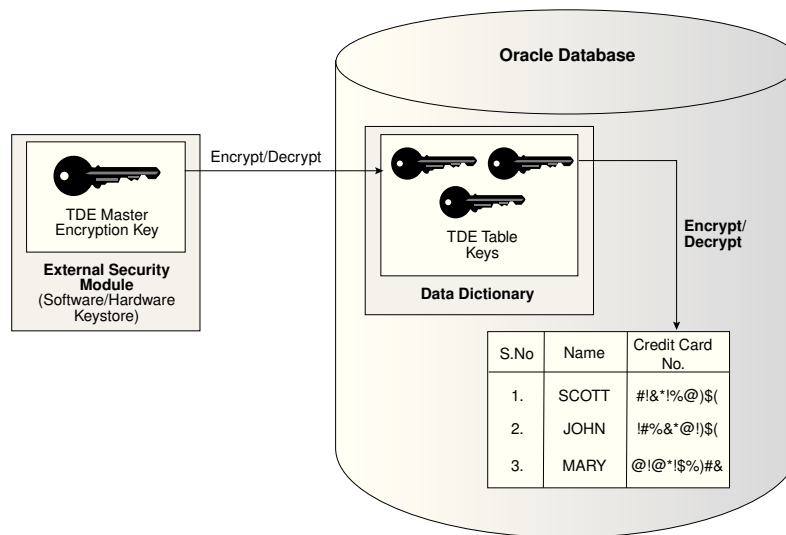
Both TDE column encryption and TDE tablespace encryption use a two-tiered key-based architecture. Unauthorized users, such as intruders who are attempting security attacks, cannot read the data from storage and back up media unless they have the TDE master encryption key to decrypt it.

2.4.2 How Transparent Data Encryption Column Encryption Works

Transparent Data Encryption (TDE) column encryption protects confidential data, such as credit card and Social Security numbers, that is stored in table columns.

TDE column encryption uses the two-tiered key-based architecture to transparently encrypt and decrypt sensitive table columns. The TDE master encryption key is stored in an external security module, which can be an Oracle software keystore or hardware keystore. This TDE master encryption key encrypts and decrypts the TDE table key, which in turn encrypts and decrypts data in the table column.

[Figure 2-1](#) (page 2-4) an overview of the TDE column encryption process.

Figure 2-1 TDE Column Encryption Overview

As shown in [Figure 2-1](#) (page 2-4), the TDE master encryption key is stored in an external security module that is outside of the database and accessible only to a user who was granted the appropriate privileges. For this external security module, Oracle Database uses an Oracle software keystore (wallet, in previous releases) or hardware security module (HSM) keystore. Storing the TDE master encryption key in this way prevents its unauthorized use.

Using an external security module separates ordinary program functions from encryption operations, making it possible to assign separate, distinct duties to database administrators and security administrators. Security is enhanced because the keystore password can be unknown to the database administrator, requiring the security administrator to provide the password.

When a table contains encrypted columns, TDE uses a single [TDE table key](#) regardless of the number of encrypted columns. Each TDE table key is individually encrypted with the TDE master encryption key. All of the TDE table keys are located together in the `colkey` column of the `ENC$` data dictionary table. No keys are stored in [plaintext](#).

2.4.3 How Transparent Data Encryption Tablespace Encryption Works

Transparent Data Encryption (TDE) tablespace encryption enables you to encrypt an entire tablespace.

All of the objects that are created in the encrypted tablespace are automatically encrypted. TDE tablespace encryption is useful if your tables contain sensitive data in multiple columns, or if you want to protect the entire table and not just individual columns. You do not need to perform a granular analysis of each table column to determine the columns that need encryption.

In addition, TDE tablespace encryption takes advantage of bulk encryption and caching to provide enhanced performance. The actual performance impact on applications can vary.

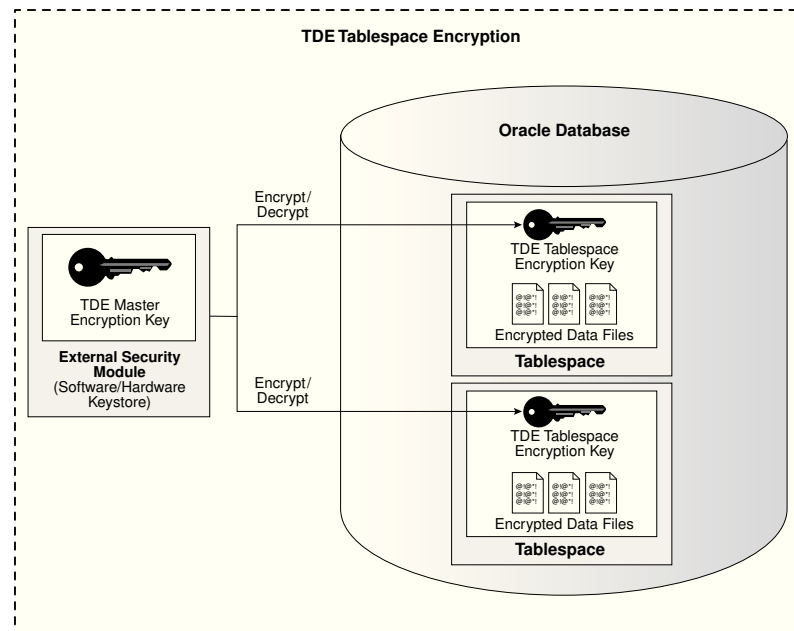
TDE tablespace encryption encrypts all of the data stored in an encrypted tablespace including its redo data. TDE tablespace encryption does not encrypt data that is stored outside of the tablespace. For example, `BFILE` data is not encrypted because it is stored outside the database. If you create a table with a `BFILE` column in an encrypted tablespace, then this particular column will not be encrypted.

All of the data in an encrypted tablespace is stored in encrypted format on the disk. Data is transparently decrypted for an authorized user having the necessary privileges to view or modify the data. A database user or application does not need to know if the data in a particular table is encrypted on the disk. In the event that the data files on a disk or backup media is stolen, the data is not compromised.

TDE tablespace encryption uses the two-tiered, key-based architecture to transparently encrypt (and decrypt) tablespaces. The TDE master encryption key is stored in an external security module (software or hardware keystore). This TDE master encryption key is used to encrypt the TDE [tablespace encryption key](#), which in turn is used to encrypt and decrypt data in the tablespace.

[Figure 2-2](#) (page 2-5) shows an overview of the TDE tablespace encryption process.

Figure 2-2 TDE Tablespace Encryption



Note:

The encrypted data is protected during operations such as JOIN and SORT. This means that the data is safe when it is moved to temporary tablespaces. Data in undo and redo logs is also protected.

TDE tablespace encryption also allows index range scans on data in encrypted tablespaces. This is not possible with TDE column encryption.

Oracle Database implements the following features to TDE tablespace encryption:

- It uses a unified TDE master encryption key for both TDE column encryption and TDE tablespace encryption.
- You can reset the unified TDE master encryption key. This provides enhanced security and helps meet security and compliance requirements.

2.4.4 How the Keystore for the Storage of TDE Master Encryption Keys Works

To control the encryption, you use a keystore and TDE master encryption key.

[About the Keystore Storage of TDE Master Encryption Keys](#) (page 2-6)

Oracle Database provides a key management framework for Transparent Data Encryption that stores and manages keys and credentials.

[Benefits of the Keystore Storage Framework](#) (page 2-6)

The key management framework provides several benefits for Transparent Data Encryption.

[Types of Keystores](#) (page 2-7)

Oracle Database supports software keystores and hardware (HSM-based) keystores.

2.4.4.1 About the Keystore Storage of TDE Master Encryption Keys

Oracle Database provides a key management framework for Transparent Data Encryption that stores and manages keys and credentials.

The key management framework includes the keystore to securely store the TDE master encryption keys and the management framework to securely and efficiently manage keystore and key operations for various database components.

The Oracle keystore stores a history of retired TDE master encryption keys, which enables you to change them and still be able to decrypt data that was encrypted under an earlier TDE master encryption key.

2.4.4.2 Benefits of the Keystore Storage Framework

The key management framework provides several benefits for Transparent Data Encryption.

- Enables separation of duty between the database administrator and the security administrator who manages the keys. You can grant the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege to users who are responsible for managing the keystore and key operations.
- Facilitates compliance, because it helps you to track encryption keys and implement requirements such as keystore password rotation and TDE master encryption key reset or rekey operations.
- Facilitates and helps enforce keystore backup requirements. A backup is a copy of the password-based software keystore that is created for all of the critical keystore operations.

You must make a backup of the keystore for all of the critical keystore operations. You must also make a backup of the TDE master encryption key before you reset or rekey this TDE master encryption key.

- Enables the keystore to be stored on an ASM file system. This is particularly useful for Oracle Real Application Clusters (Oracle RAC) environments where database instances share a unified file system view.
- Enables reverse migration from a hardware keystore to a file system-based software keystore. This option is useful if you must migrate back to a software keystore.

2.4.4.3 Types of Keystores

Oracle Database supports software keystores and hardware (HSM-based) keystores.

You can configure the following types of software keystores:

- **Password-based software keystores:** Password-based software keystores are protected by using a password that you create. You must open this type of keystore before the keys can be retrieved or used.
- **Auto-login software keystores:** Auto-login software keystores are protected by a system-generated password, and do not need to be explicitly opened by a security administrator. Auto-login software keystores are automatically opened when accessed. Auto-login software keystores can be used across different systems. If your environment does not require the extra security provided by a keystore that must be explicitly opened for use, then you can use an auto-login software keystore. Auto-login software keystores are ideal for unattended scenarios.
- **Local auto-login software keystores:** Local auto-login software keystores are auto-login software keystores that are local to the computer on which they are created. Local auto-login keystores cannot be opened on any computer other than the one on which they are created. This type of keystore is typically used for scenarios where additional security is required (that is, to limit the use of the auto-login for that computer) while supporting an unattended operation.

Software keystores can be stored on ASM disk groups or in a regular file system.

Hardware Security Modules are physical devices that provide secure storage for encryption keys, in hardware keystores. HSMs also provide secure computational space (memory) to perform encryption and decryption operations.

When using an HSM, all encryption and decryption operations that use the TDE master encryption key are performed inside the HSM. This means that the TDE master encryption key is never exposed in insecure memory.

2.4.5 Supported Encryption and Integrity Algorithms

By default, Transparent Data Encryption (TDE) Column encryption uses the Advanced Encryption Standard (AES).

The supported Advanced Encryption Standard cipher keys are the 192-bit length cipher key (AES192). Tablespace and database encryption use the 128-bit length cipher key (AES128)

In addition, `salt` is added by default to plaintext before encryption unless specified otherwise. You cannot add salt to indexed columns that you want to encrypt. For indexed columns, choose the `NO SALT` parameter for the `SQL ENCRYPT` clause.

For TDE tablespace encryption and database encryption, the default is to use the Advanced Encryption Standard with a 128-bit length cipher key (AES128). In addition, salt is always added to plaintext before encryption.

You can change encryption algorithms and encryption keys on existing encrypted columns by setting a different algorithm with the `SQL ENCRYPT` clause.

[Table 2-1](#) (page 2-8) lists the supported encryption algorithms.

Table 2-1 Supported Encryption Algorithms for Transparent Data Encryption

Algorithm	Key Size	Parameter Name
Advanced Encryption Standard (AES)	<ul style="list-style-type: none"> • 128 bits (default tablespace encryption setting) • 192 bits (default column level encryption) • 256 bits 	<ul style="list-style-type: none"> • AES192 for column-level encryption • AES128 for tablespace encryption • AES256
ARIA	<ul style="list-style-type: none"> • 128 bits • 192 bits • 256 bits 	<ul style="list-style-type: none"> • ARIA128 • ARIA192 • ARIA256
GOST	256 bits	GOST256
SEED	128 bits	SEED128
Triple Encryption Standard (DES)	168 bits	3DES168

For integrity protection of TDE column encryption, the SHA-1 hashing algorithm is used. If you have storage restrictions, then use the NOMAC option.

Configuring Transparent Data Encryption

You can configure software or hardware keystores, for use on both individual table columns or entire tablespaces.

[Configuring a Software Keystore](#) (page 3-1)

A software keystore is a container for the TDE master encryption key, and it resides in the software file system.

[Configuring a Hardware Keystore](#) (page 3-13)

A hardware keystore resides in a hardware security module (HSM), which is designed to store encryption keys.

[Encrypting Columns in Tables](#) (page 3-19)

You can use Transparent Data Encryption to encrypt individual columns in database tables.

[Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

You can perform encryption operations on both offline and online tablespaces and databases.

[Transparent Data Encryption Data Dynamic and Data Dictionary Views](#) (page 3-51)

You can query a set of dynamic and data dictionary views to find more information about Transparent Data Encryption (TDE) data.

3.1 Configuring a Software Keystore

A software keystore is a container for the TDE master encryption key, and it resides in the software file system.

[About Configuring a Software Keystore](#) (page 3-2)

A software keystore is a container that stores the Transparent Data Encryption master encryption key.

[Step 1: Set the Keystore Location in the sqlnet.ora File](#) (page 3-2)

The first step you must take to configure a software keystore is to designate a location for it in the `sqlnet.ora` file.

[Step 2: Create the Software Keystore](#) (page 3-6)

After you have specified a directory location for the software keystore, you can create the keystore.

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[Step 4: Set the Software TDE Master Encryption Key](#) (page 3-11)

Once the keystore is open, you can set a TDE master encryption key for it.

Step 5: Encrypt Your Data (page 3-13)

After you complete the software keystore configuration, you can begin to encrypt data.

3.1.1 About Configuring a Software Keystore

A software keystore is a container that stores the Transparent Data Encryption master encryption key.

Before you can configure the keystore, you first must define a location for it in the `sqlnet.ora` file. There is one keystore per database, and the database locates this keystore by checking the keystore location that you define in the `sqlnet.ora` file. You can create other keystores, such as copies of the keystore and export files that contain keys, depending on your needs. However, you must never remove or delete the keystore that you configured in the `sqlnet.ora` location, nor replace it with a different keystore.

After you configure the software keystore location in the `sqlnet.ora` file, you can log in to the database instance to create and open the keystore, and then set the TDE master encryption key. After you complete these steps, you can begin to encrypt data.

3.1.2 Step 1: Set the Keystore Location in the `sqlnet.ora` File

The first step you must take to configure a software keystore is to designate a location for it in the `sqlnet.ora` file.

About the Keystore Location in the `sqlnet.ora` File (page 3-3)

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

Configuring the `sqlnet.ora` File for a Software Keystore Location (page 3-3)

Use the `sqlnet.ora` file to configure the keystore location for a regular file system, for multiple database access, and for use with Oracle Automatic Storage Management (ASM).

Configuring an External Store for a Keystore Password (page 3-4)

An external store for a keystore password stores the software keystore password in a centrally accessed and managed location.

Example: Configuring a Software Keystore for a Regular File System
(page 3-5)

You can configure a software keystore for a regular file system.

Example: Configuring a Software Keystore When Multiple Databases Share the `sqlnet.ora` File (page 3-5)

You can configure multiple databases to share the `sqlnet.ora` file.

Example: Configuring a Software Keystore for Oracle Automatic Storage Management (page 3-5)

You can configure `sqlnet.ora` for an Automatic Storage Management (ASM) file system

Example: Configuring a Software Keystore for an Oracle Automatic Storage Management Disk Group (page 3-6)

You can configure `sqlnet.ora` for an Oracle Automatic Storage Management (ASM) disk group.

3.1.2.1 About the Keystore Location in the `sqlnet.ora` File

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

You must edit the `sqlnet.ora` file to define a directory location for the keystore that you plan to create. Ensure that this directory exists beforehand. Preferably, this directory should be empty.

Note the following behavior when you must edit the `sqlnet.ora` file in an Oracle Real Application Clusters (Oracle RAC) or a multitenant environment:

- **In an Oracle RAC environment:** If you are using the `srvctl` utility and if you want to include environment variables in the `sqlnet.ora` configuration file, then you must set these environment variables in both the operating system and the `srvctl` environment. Oracle recommends that you place the keystore on a shared file system, such as Oracle Automatic Storage Management (ASM) or NFS.
- **In a multitenant environment:** The keystore location is set for the entire multitenant container database (CDB), not for individual pluggable databases (PDBs).

In the `sqlnet.ora` file, you must set the `ENCRYPTION_WALLET_LOCATION` parameter to specify the keystore location. When determining which keystore to use, Oracle Database searches for the keystore location in the following places, in this order:

1. It attempts to use the keystore in the location specified by the parameter `ENCRYPTION_WALLET_LOCATION` in the `sqlnet.ora` file.
2. If the `ENCRYPTION_WALLET_LOCATION` parameter is not set, then it attempts to use the keystore in the location that is specified by the parameter `WALLET_LOCATION`.
3. If the `WALLET_LOCATION` parameter is also not set, then Oracle Database looks for a keystore at the default database location, which is `ORACLE_BASE/admin/DB_UNIQUE_NAME/wallet` or `ORACLE_HOME/admin/DB_UNIQUE_NAME/wallet`. (`DB_UNIQUE_NAME` is the unique name of the database specified in the initialization parameter file.) When the keystore location is not set in the `sqlnet.ora` file, then the `V$ENCRYPTION_WALLET` view displays the default location. You can check the location and status of the keystore in the `V$ENCRYPTION_WALLET` view.

By default, the `sqlnet.ora` file is located in the `ORACLE_HOMEdbs` directory or in the location set by the `TNS_ADMIN` environment variable. Ensure that you have properly set the `TNS_ADMIN` environment variable to point to the correct `sqlnet.ora` file.

See Also: *SQL*Plus User's Guide and Reference* for more information and examples of setting the `TNS_ADMIN` environment variable

3.1.2.2 Configuring the `sqlnet.ora` File for a Software Keystore Location

Use the `sqlnet.ora` file to configure the keystore location for a regular file system, for multiple database access, and for use with Oracle Automatic Storage Management (ASM).

- To create a software keystore on a regular file system, use the following format when you edit the `sqlnet.ora` file:

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=path_to_keystore)))
```

If the `path_to_keystore` will contain an environment variable, then set this variable in the environment where the database instance is started and before you start the database. If you are using the `srvctl` utility to start the database, then set the environment variable in the `srvctl` environment as well, using the following command:

```
srvctl setenv database -db database_name -env
"environment_variable_name=environment_variable_value"
```

3.1.2.3 Configuring an External Store for a Keystore Password

An external store for a keystore password stores the software keystore password in a centrally accessed and managed location.

An external store for a password is useful for situations in which you use automated tools to perform Transparent Data Encryption operations that require a password, when the scripts that run the automated tools include hard-coded passwords. To avoid hard-coding the password in a script, you can store this password in an external store on the database server. In a multitenant environment, different PDBs can make use of the external store.

You must complete the following steps before you use the `IDENTIFIED BY EXTERNAL STORE` clause in the `ADMINISTER KEY MANAGEMENT` statement.

1. Set the external keystore credential location by using one of the following methods:

- Run the `ALTER SYSTEM` statement for the `EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION` parameter. For example:

```
ALTER SYSTEM SET EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION = "/etc/ORACLE/
WALLETS/orcl/external_store" SCOPE = SPFILE;
```

- Edit the `init.ora` file for the database instance. For example:

```
EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION = "/etc/ORACLE/WALLETS/orcl/
external_store"
```

By default, the `init.ora` file is located in the `ORACLE_HOME/dbs` directory or in the location set by the `TNS_ADMIN` environment variable.

2. Log in as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege and who has the `ALTER SYSTEM` system privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as sysdba
Enter password: password
Connected.
```

3. Create an auto-login keystore that contains the keystore password, by including the `ADD SECRET` clause to the `ADMINISTER KEY MANAGEMENT` statement.

For example:


```
ADMINISTER KEY MANAGEMENT ADD SECRET 'password'
FOR CLIENT 'TDE_WALLET'
TO LOCAL AUTO_LOGIN KEYSTORE '/etc/ORACLE/WALLETS/orcl/external_store';
```

In this example, enter 'TDE_WALLET', in capital letters and enclosed in single quotation marks, for the *client_identifier* value set by the FOR CLIENT clause. This is a fixed value and must be entered as shown here for this application of the ADD SECRET clause. Otherwise, TDE will be unable to find this secret, and attempts to use the IDENTIFIED BY EXTERNAL STORE setting will generate an ORA-00988: missing or invalid password(s) error message.

4. Restart the database.

For example:

```
SHUTDOWN IMMEDIATE
STARTUP
```

Afterward, you must use the EXTERNAL STORE clause in the ADMINISTER KEY MANAGEMENT statement for the following operations: opening, closing, backing up the keystore; adding, updating, or deleting a secret keystore; creating, using, rekeying, tagging, importing, exporting encryption keys.

For example, to open a keystore in a multitenant environment:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY EXTERNAL STORE CONTAINER =
ALL;
```

3.1.2.4 Example: Configuring a Software Keystore for a Regular File System

You can configure a software keystore for a regular file system.

The following example shows how to configure a software keystore location in the `sqlnet.ora` file for a regular file system in which the database name is `orcl`.

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=/etc/ORACLE/WALLETS/orcl)))
```

3.1.2.5 Example: Configuring a Software Keystore When Multiple Databases Share the `sqlnet.ora` File

You can configure multiple databases to share the `sqlnet.ora` file.

The following example shows how to configure a software keystore location when multiple databases share the `sqlnet.ora` file.

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=/etc/ORACLE/WALLETS/$ORACLE_SID/)))
```

3.1.2.6 Example: Configuring a Software Keystore for Oracle Automatic Storage Management

You can configure `sqlnet.ora` for an Automatic Storage Management (ASM) file system

The following example shows how to configure a software keystore location in the `sqlnet.ora` file for an ASM file system:

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=+disk1/mydb/wallet)))
```

3.1.2.7 Example: Configuring a Software Keystore for an Oracle Automatic Storage Management Disk Group

You can configure `sqlnet.ora` for an Oracle Automatic Storage Management (ASM) disk group.

The following format shows how to configure a software keystore if you want to create a software keystore location on an ASM disk group:

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=+ASM_file_path_of_the_diskgroup)))
```

3.1.3 Step 2: Create the Software Keystore

After you have specified a directory location for the software keystore, you can create the keystore.

[About Creating Software Keystores](#) (page 3-6)

There are three different types of software keystores.

[Creating a Password-Based Software Keystore](#) (page 3-7)

A password-based software keystore requires a user password, which is used to protect the keys and credentials stored in the keystore.

[Creating an Auto-Login or a Local Auto-Login Software Keystore](#) (page 3-8)

As an alternative to password-based keystores, you can create either an auto-login or local auto-login software keystore.

3.1.3.1 About Creating Software Keystores

There are three different types of software keystores.

You can create password-based software keystores, auto-login software keystores, and local auto-login software keystores.

Be aware that executing the query `SELECT * FROM V$ENCRYPTION_WALLET` will automatically open an auto-login software keystore. For example, suppose you have a password-based keystore and an auto-login keystore. If the password-based keystore is open and you close the password-based keystore and then query the `V$ENCRYPTION_WALLET` view, then the output will indicate that a keystore is open. However, this is because `V$ENCRYPTION_WALLET` opened up the auto-login software keystore and then displayed the status of the auto-login keystore.

In a multitenant environment, you can create a secure external store for the software keystore. This feature enables you to hide the password from the operating system: it removes the need for storing clear-text keystore passwords in scripts or other tools that can access the database without user intervention, such as overnight batch scripts. The location for this keystore is set by the `EXTERNAL_KEYSTORE_CREDENTIAL_LOCATION` initialization parameter. In a

multitenant environment, different PDBs can access this external store location when you run the `ADMINISTER KEY MANAGEMENT` statement using the `IDENTIFIED BY EXTERNAL STORE` clause. This way, you can centrally locate the password and then update it only once in the external store.

Related Topics:

[Types of Keystores](#) (page 2-7)

Oracle Database supports software keystores and hardware (HSM-based) keystores.

3.1.3.2 Creating a Password-Based Software Keystore

A password-based software keystore requires a user password, which is used to protect the keys and credentials stored in the keystore.

Before you begin this procedure, ensure that you complete the procedure described in [Step 1: Set the Keystore Location in the sqlnet.ora File](#) (page 3-2).

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

If `SQL*Plus` is already open and you had modified the `sqlnet.ora` file during this time, then reconnect to `SQL*Plus`. The database session must be changed before the `sqlnet.ora` changes can take effect.

2. Run the `ADMINISTER KEY MANAGEMENT SQL` statement to create the keystore.

The syntax is as follows:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE 'keystore_location' IDENTIFIED BY
software_keystore_password;
```

In this specification:

- *keystore_location* is the path to the keystore directory location of the password-based keystore for which you want to create the auto-login keystore (for example, `/etc/ORACLE/WALLETS/orcl`). Enclose the *keystore_location* setting in single quotation marks (' '). To find this location, you can query the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view. (If the keystore was not created in the default location, then the `STATUS` column of the `V$ENCRYPTION_WALLET` view is `NOT_AVAILABLE`.)
- *software_keystore_password* is the password of the keystore that you, the security administrator, creates.

For example, to create the keystore in the `/etc/ORACLE/WALLETS/orcl` directory:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE '/etc/ORACLE/WALLETS/orcl'
IDENTIFIED BY password;
```

keystore altered.

After you run this statement, the `ewallet.p12` file, which is the keystore, appears in the keystore location.

3.1.3.3 Creating an Auto-Login or a Local Auto-Login Software Keystore

As an alternative to password-based keystores, you can create either an auto-login or local auto-login software keystore.

Both of these keystores have system-generated passwords. They are also PKCS#12-based files. The auto-login software keystore can be opened from different computers from the computer where this keystore resides, but the local auto-login software keystore can only be opened from the computer on which it was created. Both the auto-login and local auto-login keystores are created from the password-based software keystores.

Before you begin this procedure, ensure that you complete the procedure described in [Step 1: Set the Keystore Location in the `sqlnet.ora` File](#) (page 3-2).

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

If `SQL*Plus` is already open and you had modified the `sqlnet.ora` file during this time, then reconnect to `SQL*Plus`. The database session must be changed before the `sqlnet.ora` changes can take effect.

2. Create a password-based software keystore.

For example:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE '/etc/ORACLE/WALLETS/orcl' IDENTIFIED
BY keystore_password;
```

3. Run the `ADMINISTER KEY MANAGEMENT SQL` statement to create the keystore.

The syntax is as follows:

```
ADMINISTER KEY MANAGEMENT CREATE [LOCAL] AUTO_LOGIN KEYSTORE FROM KEYSTORE
'keystore_location' IDENTIFIED BY software_keystore_password;
```

In this specification:

- `LOCAL` enables you to create a local auto-login software keystore. Otherwise, omit this clause if you want the keystore to be accessible by other computers.
- `keystore_location` is the path to the directory location of the password-based keystore for which you want to create the auto-login keystore (for example, `/etc/ORACLE/WALLETS/orcl`). Enclose this setting in single quotation marks (' '). To find this location, query the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view.
- `software_keystore_password` is the password-based keystore for which you want to create the auto-login keystore.

For example, to create an auto-login software keystore of the password-based keystore that is located in the `/etc/ORACLE/WALLETS/orcl` directory:

```
ADMINISTER KEY MANAGEMENT CREATE AUTO_LOGIN KEYSTORE FROM KEYSTORE '/etc/
ORACLE/WALLETS/orcl' IDENTIFIED BY password;
```

keystore altered.

After you run this statement, the `cwallet.sso` file appears in the keystore location. The `ewallet.p12` file is the password-based wallet.

Note:

Do not remove the PKCS#12 wallet (`ewallet.p12` file) after you create the auto login keystore (`.sso` file). You must have the PKCS#12 wallet to regenerate or rekey the TDE master encryption key in the future. By default, this file is located in the `$ORACLE_HOME/admin/ORACLE_SID/wallet` directory.

Transparent Data Encryption uses an auto login keystore only if it is available at the correct location (`ENCRYPTION_WALLET_LOCATION`, `WALLET_LOCATION`, or the default keystore location), and the SQL statement to open an encrypted keystore has not already been executed. (Note that auto-login keystores are encrypted, because they have system-generated passwords.)

Related Topics:

[Creating a Password-Based Software Keystore](#) (page 3-7)

A password-based software keystore requires a user password, which is used to protect the keys and credentials stored in the keystore.

3.1.4 Step 3: Open the Software Keystore

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[About Opening Software Keystores](#) (page 3-9)

A password-based software keystore must be open before any TDE master encryption keys can be created or accessed in the keystore.

[Opening a Software Keystore](#) (page 3-10)

To open a software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

3.1.4.1 About Opening Software Keystores

A password-based software keystore must be open before any TDE master encryption keys can be created or accessed in the keystore.

You can either manually open a software keystore or, when you perform certain `ADMINISTER KEY MANAGEMENT` operations, have the keystore open temporarily during the course of the operation itself. You do not need to manually open auto-login or local auto-login software keystores. These keystore are automatically opened when it is required, that is, when an encryption operation must access the key. If necessary, you can explicitly close any of these types of keystores. Keystores can be in the following states: open, closed, open but with no master key, open but with an unknown master key, undefined, or not available (that is, not present in the `sqlnet.ora` location).

You can check the status of whether a keystore is open or not by querying the `STATUS` column of the `V$ENCRYPTION_WALLET` view.

After you open a keystore, it remains open until you manually close it. Each time you restart a database instance, you must manually open the password keystore to reenable encryption and decryption operations.

Related Topics:

[Performing Operations That Require a Keystore Password](#) (page 4-2)

Many `ADMINISTER KEY MANAGEMENT` operations require access to a keystore password, for both software and hardware keystores.

[How Keystore Open and Close Operations Work in a Multitenant Environment](#) (page 6-18)

You should be aware of how keystore open and close operations work in a multitenant environment.

3.1.4.2 Opening a Software Keystore

To open a software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

1. Ensure that you complete the procedure described in [Step 2: Create the Software Keystore](#) (page 3-6).
2. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, you must open the keystore first in the root before you can open it in a PDB. For example, to log in to the root:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

3. Run the `ADMINISTER KEY MANAGEMENT` statement.

Use the following syntax:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN [FORCE KEYSTORE]
IDENTIFIED BY [IDENTIFIED BY [EXTERNAL STORE | keystore_password]]
[CONTAINER = ALL | CURRENT];
```

In this specification:

- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `IDENTIFIED BY` permits the following authentication methods:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `keystore_password` is the same password that you used to create the software keystore in [Step 2: Create the Software Keystore](#) (page 3-6). In a multitenant environment, ensure that you enter the password that is specific to the PDB in which you are logged.

- CONTAINER is for use in a multitenant environment. Enter ALL to set the keystore in all of the PDBs in this CDB, or CURRENT for the current PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY password;
keystore altered.
```

If the auto-login keystore is open or if the password keystore is closed:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY EXTERNAL
STORE;
```

```
keystore altered.
```

Note that if the keystore is open but you have not created a TDE master encryption key yet (described next), the STATUS column of the V\$ENCRYPTION_WALLET view reminds you with an OPEN_NO_MASTER_KEY status.

3.1.5 Step 4: Set the Software TDE Master Encryption Key

Once the keystore is open, you can set a TDE master encryption key for it.

[About Setting the Software TDE Master Encryption Key](#) (page 3-11)

The TDE master encryption key is stored in the keystore.

[Setting the TDE Master Encryption Key in the Software Keystore](#) (page 3-11)

To set the TDE master encryption key in a software keystore, use the ADMINISTER KEY MANAGEMENT statement with the SET KEY clause.

3.1.5.1 About Setting the Software TDE Master Encryption Key

The TDE master encryption key is stored in the keystore.

This key protects the [TDE table keys](#) and [tablespace encryption keys](#). By default, the TDE master encryption key is a key that Transparent Data Encryption (TDE) generates. You can find if a keystore has no master key set or an unknown master key by querying the STATUS column of the V\$ENCRYPTION_WALLET view.

In a multitenant environment, you can create and manage the TDE master encryption key from either the root or the PDB.

You also can create TDE master encryption keys for use later on, and then manually activate them.

Related Topics:

[Creating TDE Master Encryption Keys for Later Use](#) (page 4-25)

You can create a TDE master encryption key that can be activated at a later date.

3.1.5.2 Setting the TDE Master Encryption Key in the Software Keystore

To set the TDE master encryption key in a software keystore, use the ADMINISTER KEY MANAGEMENT statement with the SET KEY clause.

1. For password software keystores, ensure that you complete the procedure described in [Step 3: Open the Software Keystore](#) (page 3-9) to open the key.

Auto-login or local auto-login software keys are opened automatically after you create them. Password-based software keystores must be open before you can set

the TDE master encryption key. If the auto-login software keystore is open, then you must close it and open the password-based software keystore. If both the password-based keystore and auto-login keystores are present in the configured location and the password-based keystore is open, then the TDE master encryption key is automatically written to the auto-login keystore as well.

2. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or to the PDB. For example, to log in to a PDB:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

3. Ensure that the database is open in `READ WRITE` mode.

You can set the TDE master encryption key if `OPEN_MODE` is set to `READ WRITE`. To find the status, for a non-multitenant environment, query the `OPEN_MODE` column of the `V$DATABASE` dynamic view. If you are using a multitenant environment, then query the `V$PDBS` view. (If you cannot access these views, then connect as `SYSDBA` and try the query again. In order to connect as `SYSKM` for this type of query, you must create a password file for it.

4. Connect using the `SYSKM` administrative privilege and then run the `ADMINISTER KEY MANAGEMENT SQL` statement to set the software management keystore.

```
ADMINISTER KEY MANAGEMENT SET KEY [USING TAG 'tag'] [FORCE KEYSTORE] IDENTIFIED
BY [EXTERNAL STORE] | keystore_password [WITH BACKUP [USING
'backup_identifier']] [CONTAINER = ALL | CURRENT];
```

In this specification:

- `tag` is the associated attributes and information that you define. Enclose this setting in single quotation marks (' ').
- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `keystore_password` is the mandatory keystore password that you created when you created the keystore in [Step 2: Create the Software Keystore](#) (page 3-6).
- `WITH BACKUP` creates a backup of the keystore. You must use this option for password-based keystores. Optionally, you can use the `USING` clause to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time_stamp_emp_key_backup.p12`, with `emp_key_backup` being the backup identifier). Follow the file naming conventions that your operating system uses.

- CONTAINER is for use in a multitenant environment. Enter ALL to set the key in all of the PDBs in this CDB, or CURRENT for the current PDB.

For example, if the password-based keystore is open:

```
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY keystore_password WITH BACKUP
USING 'emp_key_backup';
```

keystore altered.

If the auto-login keystore is open or if the keystore is closed:

```
ADMINISTER KEY MANAGEMENT SET KEY FORCE KEY IDENTIFIED BY keystore_password WITH
BACKUP USING 'emp_key_backup';
```

keystore altered.

Related Topics:

Oracle Database Administrator's Guide

3.1.6 Step 5: Encrypt Your Data

After you complete the software keystore configuration, you can begin to encrypt data.

You can encrypt data in individual table columns or in entire tablespaces or databases.

Related Topics:

[Encrypting Columns in Tables](#) (page 3-19)

You can use Transparent Data Encryption to encrypt individual columns in database tables.

[Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

You can perform encryption operations on both offline and online tablespaces and databases.

3.2 Configuring a Hardware Keystore

A hardware keystore resides in a hardware security module (HSM), which is designed to store encryption keys.

[About Configuring a Hardware \(External\) Keystore](#) (page 3-14)

A hardware keystore, also called an external keystore, is a separate server or device that provides security storage for encryption keys.

[Step 1: Set the Hardware Keystore Type in the sqlnet.ora File](#) (page 3-14)

Before you can configure a hardware keystore, you must modify the `sqlnet.ora` file.

[Step 2: Configure the Hardware Security Module](#) (page 3-14)

To configure a third-party hardware security module, you must copy the PKCS#11 library to the correct location and follow your vendor's instructions.

[Step 3: Open the Hardware Keystore](#) (page 3-15)

After you have configured the hardware security module, you must open the hardware keystore before it can be used.

Step 4: Set the Hardware Keystore TDE Master Encryption Key (page 3-17)

After you have opened the hardware keystore, you are ready to set the hardware keystore TDE master encryption key.

Step 5: Encrypt Your Data (page 3-19)

After you have completed the configuration for a hardware keystore or for an Oracle Key Vault keystore, you can begin to encrypt data.

3.2.1 About Configuring a Hardware (External) Keystore

A hardware keystore, also called an external keystore, is a separate server or device that provides security storage for encryption keys.

External keystores are external to an Oracle database. Oracle Database can interface with external keystores but cannot manipulate them outside of the Oracle interface. The Oracle database can request the external keystore to create a key but it cannot define how this key is stored in an external database. (Conversely, for software keystores that are created using TDE, Oracle Database has full control: that is, you can use SQL statements to manipulate this type of keystore.) Examples of external keystores are hardware security modules or Oracle Key Vault keystores. External keystores among multiple databases can be managed centrally, such as with Oracle Key Vault.

To configure a keystore for a hardware security module (hardware keystore), you must first include the keystore type in the `sqlnet.ora` file, configure and open the hardware keystore, and then set the hardware keystore TDE master encryption key. In short, there is one hardware keystore per database, and the database locates this keystore by checking the keystore type that you define in the `sqlnet.ora` file.

After you configure the hardware keystore, you are ready to begin encrypting your data.

3.2.2 Step 1: Set the Hardware Keystore Type in the `sqlnet.ora` File

Before you can configure a hardware keystore, you must modify the `sqlnet.ora` file.

By default, the `sqlnet.ora` file is located in the `ORACLE_HOME`\dbs directory or in the location set by the `TNS_ADMIN` environment variable.

- Use the following setting in the `sqlnet.ora` file to define the hardware keystore type, which is either `OKV` for Oracle Key Vault or `HSM` for hardware security module.

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=OKV))
```

3.2.3 Step 2: Configure the Hardware Security Module

To configure a third-party hardware security module, you must copy the PKCS#11 library to the correct location and follow your vendor's instructions.

If you are using Oracle Key Vault, then you can bypass this section.

1. Ensure that you complete the procedure described in [Step 1: Set the Hardware Keystore Type in the `sqlnet.ora` File](#) (page 3-14).
2. Copy the PKCS#11 library to its correct path.

Your hardware keystore vendor should provide you with an associated PKCS#11 library. Only one PKCS#11 library is supported at a time. If you want to use a hardware keystore from a new vendor, then you must replace the PKCS#11 library from the earlier vendor with the library from the new vendor.

Copy this library to the appropriate location to ensure that Oracle Database can find this library:

- **UNIX systems:** Use the following syntax to copy the library to this directory:

```
/opt/oracle/extapi/[32,64]/hsm/{VENDOR}/{VERSION}/libapiname.so
```

- **Windows systems:** Use the following syntax to copy the library to this directory:

```
%SYSTEM_DRIVE%\oracle\extapi\[32,64]\hsm\{VENDOR}\{VERSION}\libapiname.dll
```

In this specification:

- [32 , 64] specifies whether the supplied binary is 32 bits or 64 bits.
 - *VENDOR* stands for the name of the vendor supplying the library
 - *VERSION* refers to the version of the library. This should preferably be in the format, *number . number . number*
 - *apiname* requires no special format. However, the *apiname* must be prefixed with the word *lib*, as illustrated in the syntax.
3. Follow your vendor's instructions to set up the hardware keystore.
- Use your hardware keystore management interface and the instructions provided by your HSM vendor to set up the hardware keystore. Create the user account and password that must be used by the database to interact with the hardware keystore. This process creates and configures a hardware keystore that communicates with your Oracle database.

3.2.4 Step 3: Open the Hardware Keystore

After you have configured the hardware security module, you must open the hardware keystore before it can be used.

[About Opening Hardware Keystores](#) (page 3-15)

You must open the hardware keystore so that it is accessible to the database before you can perform any encryption or decryption.

[Opening a Hardware Keystore](#) (page 3-16)

To open a hardware keystore, use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

3.2.4.1 About Opening Hardware Keystores

You must open the hardware keystore so that it is accessible to the database before you can perform any encryption or decryption.

You can check the status of whether a keystore is open, closed, open but with no TDE master encryption key, or open but with an unknown master encryption key by querying the `STATUS` column of the `V$ENCRYPTION_WALLET` view.

You can either manually open a hardware keystore or, when you perform certain `ADMINISTER KEY MANAGEMENT` operations, have the keystore open temporarily

during the course of the operation itself. Keystores can be in the following states: open, closed, open but with no master key, open but with an unknown master key, undefined, or not available (that is, not present in the `sqlnet.ora` location).

Related Topics:

[How Keystore Open and Close Operations Work in a Multitenant Environment](#)
(page 6-18)

You should be aware of how keystore open and close operations work in a multitenant environment.

3.2.4.2 Opening a Hardware Keystore

To open a hardware keystore, use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

Before you begin this procedure, ensure that you complete the procedure described in [Step 2: Configure the Hardware Security Module](#) (page 3-14).

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, you must open the keystore first in the root before you can open it in a PDB. For example, to log in to the root:

```
sqlplus sec_admin as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

If SQL*Plus is already open and you had modified the `sqlnet.ora` file during this time, then reconnect to SQL*Plus. The database session must be changed before the `sqlnet.ora` changes can take effect.

2. Run the `ADMINISTER KEY MANAGEMENT SQL` statement using the following syntax:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN [FORCE KEYSTORE] IDENTIFIED BY
"user_id:password" [CONTAINER = ALL | CURRENT];
```

In this specification:

- `FORCE KEYSTORE` enables the keystore operation if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `user_id:password:user_id` is the user ID created for the hardware keystore; `password` is the password created for the hardware keystore. Enclose the `user_id:password` string in double quotation marks (" ") and separate `user_id` and `password` with a colon (:).
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to set the keystore in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
"psmith:password";
```

keystore altered.

Or, for a keystore that is stored externally:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
EXTERNAL STORE;
```

keystore altered.

3. Repeat this procedure each time you restart the database instance.

3.2.5 Step 4: Set the Hardware Keystore TDE Master Encryption Key

After you have opened the hardware keystore, you are ready to set the hardware keystore TDE master encryption key.

[About Setting the Hardware Keystore TDE Master Encryption Key](#) (page 3-17)

You must create a TDE master encryption key that is stored inside the hardware keystore.

[Setting a TDE Master Encryption Key if You Have Not Previously Configured One](#) (page 3-18)

You should complete this procedure if you have not previously configured a software keystore for Transparent Data Encryption.

[Migration of a Previously Configured TDE Master Encryption Key](#) (page 3-19)

You must migrate the previously configured TDE master encryption key if you previously configured a software keystore.

3.2.5.1 About Setting the Hardware Keystore TDE Master Encryption Key

You must create a TDE master encryption key that is stored inside the hardware keystore.

Oracle Database uses the TDE master encryption key to encrypt or decrypt [TDE table keys](#) or [tablespace encryption keys](#) inside the hardware security module.

If you have not previously configured a software keystore for Transparent Data Encryption, then you must set the master encryption key. If you have already configured a software keystore for TDE, then you must migrate it to the hardware security module.

Along with the current TDE master key, Oracle wallets maintain historical TDE master keys that are generated after every re-key operation that rotates the TDE master key. These historical TDE master keys help to restore Oracle database backups that were taken previously using one of the historical TDE master keys.

Related Topics:

[Setting a TDE Master Encryption Key if You Have Not Previously Configured One](#) (page 3-18)

You should complete this procedure if you have not previously configured a software keystore for Transparent Data Encryption.

[Migration of a Previously Configured TDE Master Encryption Key](#) (page 3-19)

You must migrate the previously configured TDE master encryption key if you previously configured a software keystore.

3.2.5.2 Setting a TDE Master Encryption Key if You Have Not Previously Configured One

You should complete this procedure if you have not previously configured a software keystore for Transparent Data Encryption.

In a multitenant environment, you can create and manage the TDE master encryption key from either the root or the PDB.

1. Ensure that you complete the procedure described in [Step 3: Open the Hardware Keystore](#) (page 3-15).
2. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or to the PDB. For example:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

3. Ensure that the database is open in `READ WRITE` mode.

You can set the TDE master encryption key if `OPEN_MODE` is set to `READ WRITE`. To find the status, for a non-multitenant environment, query the `OPEN_MODE` column of the `V$DATABASE` dynamic view. If you are in a multitenant environment, then query the `V$PDBS` view. (If you cannot access these views, then connect as `SYSDBA` and try the query again. In order to connect as `SYSKM` for this type of query, you must create a password file for it.

4. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET KEY [USING TAG 'tag'] [FORCE KEYSTORE] IDENTIFIED
BY [EXTERNAL STORE | "user_id:password"] [CONTAINER = ALL | CURRENT];
```

In this specification:

- *tag* is the associated attributes and information that you define. Enclose this setting in single quotation marks (' ').
- `FORCE KEYSTORE` enables the keystore operation if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `user_id:password: user_id` is the user ID created for the hardware keystore; `password` is the password created for the hardware keystore. Enclose the `user_id:password` string in double quotation marks (" ") and separate `user_id` and `password` with a colon (:).
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to set the keystore in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY "psmith:password";

keystore altered.
```

See Also:

- [Creating a TDE Master Encryption Key for Later Use](#) (page 4-26)
- *Oracle Database Administrator's Guide*

3.2.5.3 Migration of a Previously Configured TDE Master Encryption Key

You must migrate the previously configured TDE master encryption key if you previously configured a software keystore.

Tools such as Oracle Data Pump and Oracle Recovery Manager require access to the old software keystore to perform decryption and encryption operations on data exported or backed up using the software keystore. You can migrate from the software to the hardware keystore by following the instructions in [Migrating Between a Software Password Keystore and a Hardware Keystore](#) (page 4-13).

Along with the current TDE master key, Oracle wallets maintain historical TDE master keys that are generated after every re-key operation that rotates the TDE master key. These historical TDE master keys help to restore Oracle database backups that were taken previously using one of the historical TDE master keys.

3.2.6 Step 5: Encrypt Your Data

After you have completed the configuration for a hardware keystore or for an Oracle Key Vault keystore, you can begin to encrypt data.

Oracle Key Vault Administrator's Guide describes how to configure Oracle Key Vault keystores.

You can encrypt individual columns in a table or entire tablespaces.

Related Topics:

[Encrypting Columns in Tables](#) (page 3-19)

You can use Transparent Data Encryption to encrypt individual columns in database tables.

[Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

You can perform encryption operations on both offline and online tablespaces and databases.

3.3 Encrypting Columns in Tables

You can use Transparent Data Encryption to encrypt individual columns in database tables.

[About Encrypting Columns in Tables](#) (page 3-20)

You can encrypt individual columns in tables.

[Data Types That Can Be Encrypted with TDE Column Encryption](#) (page 3-20)

Oracle Database supports a specific set of data types that can be used with TDE column encryption.

[Restrictions on Using Transparent Data Encryption Column Encryption](#)

(page 3-21)

TDE encrypts at the SQL layer. Oracle Database utilities that bypass the SQL layer cannot use the TDE column encryption services.

[Creating Tables with Encrypted Columns](#) (page 3-22)

Oracle Database provides a selection of different algorithms that you can use to define the encryption used in encrypted columns.

[Encrypting Columns in Existing Tables](#) (page 3-26)

You can encrypt columns in existing tables. As with new tables, you have a choice of different algorithms to use to definite the encryption.

[Creating an Index on an Encrypted Column](#) (page 3-27)

You can create an index on an encrypted column.

[Adding Salt to an Encrypted Column](#) (page 3-28)

Salt, which is a random string added to data before encryption, is a way to strengthen the security of encrypted data. .

[Removing Salt from an Encrypted Column](#) (page 3-28)

You can use the ALTER TABLE SQL statement to remove salt from an encrypted column.

[Changing the Encryption Key or Algorithm for Tables with Encrypted Columns](#)

(page 3-28)

You can use the ALTER TABLE SQL statement to change the encryption key or algorithm used in encrypted columns.

3.3.1 About Encrypting Columns in Tables

You can encrypt individual columns in tables.

Whether you choose to encrypt individual columns or entire tablespaces depends on the data types that the table has. There are also several features that do not support TDE column encryption.

Related Topics:

[Data Types That Can Be Encrypted with TDE Column Encryption](#) (page 3-20)

Oracle Database supports a specific set of data types that can be used with TDE column encryption.

[Restrictions on Using Transparent Data Encryption Column Encryption](#)

(page 3-21)

TDE encrypts at the SQL layer. Oracle Database utilities that bypass the SQL layer cannot use the TDE column encryption services.

3.3.2 Data Types That Can Be Encrypted with TDE Column Encryption

Oracle Database supports a specific set of data types that can be used with TDE column encryption.

You can encrypt data columns that use a variety of different data types.

Supported data types are as follows:

- BINARY_DOUBLE
- BINARY_FLOAT

- CHAR
- DATE
- INTERVAL DAY TO SECOND
- INTERVAL YEAR TO MONTH
- NCHAR
- NUMBER
- NVARCHAR2
- RAW (legacy or extended)
- TIMESTAMP (includes `TIMESTAMP WITH TIME ZONE` and `TIMESTAMP WITH LOCAL TIME ZONE`)
- VARCHAR2 (legacy or extended)

You cannot encrypt a column if the encrypted column size is greater than the size allowed by the data type of the column.

[Table 3-1](#) (page 3-21) shows the maximum allowable sizes for various data types.

Table 3-1 Maximum Allowable Size for Data Types

Data Type	Maximum Size
CHAR	1932 bytes
VARCHAR2 (legacy)	3932 bytes
VARCHAR2 (extended)	32,699 bytes
NVARCHAR2 (legacy)	1966 bytes
NVARCHAR2 (extended)	16,315 bytes
NCHAR	966 bytes
RAW (extended)	32,699 bytes

Note:

TDE tablespace encryption does not have these data type restrictions.

Related Topics:

[Restrictions on Using Transparent Data Encryption Tablespace Encryption](#)
(page 3-31)

You should be aware of restrictions on using Transparent Data Encryption when you encrypt a tablespace.

3.3.3 Restrictions on Using Transparent Data Encryption Column Encryption

TDE encrypts at the SQL layer. Oracle Database utilities that bypass the SQL layer cannot use the TDE column encryption services.

Do not use TDE column encryption with the following database features:

- Index types other than B-tree
- Range scan search through an index
- Synchronous change data capture
- Transportable tablespaces
- Columns that have been created as identity columns

In addition, you cannot use TDE column encryption to encrypt columns used in foreign key constraints.

Applications that must use these unsupported features can use the `DBMS_CCRYPTO` PL/SQL package for their encryption needs.

Transparent Data Encryption protects data stored on a disk or other media. It does not protect data in transit. Use the network encryption solutions discussed in *Oracle Database Security Guide* to encrypt data over the network.

Related Topics:

[How Transparent Data Encryption Works with Export and Import Operations](#)
(page 6-1)

Oracle Data Pump can export and import tables that contain encrypted columns, as well as encrypt entire dump sets.

[Data Types That Can Be Encrypted with TDE Column Encryption](#) (page 3-20)

Oracle Database supports a specific set of data types that can be used with TDE column encryption.

3.3.4 Creating Tables with Encrypted Columns

Oracle Database provides a selection of different algorithms that you can use to define the encryption used in encrypted columns.

[About Creating Tables with Encrypted Columns](#) (page 3-23)

You can use the `CREATE TABLE` SQL statement to create a table with an encrypted column.

[Creating a Table with an Encrypted Column Using the Default Algorithm](#)
(page 3-23)

By default, TDE uses the AES encryption algorithm with a 192-bit key length (AES192).

[Creating a Table with an Encrypted Column Using No Algorithm or a Non-Default Algorithm](#) (page 3-24)

You can use the `CREATE TABLE` SQL statement to create a table with an encrypted column.

[Using the NOMAC Parameter to Save Disk Space and Improve Performance](#)
(page 3-24)

You can bypass checks that TDE performs. This can save up to 20 bytes of disk space per encrypted value.

[Example: Using the NOMAC Parameter in a CREATE TABLE Statement](#)
(page 3-25)

You can use the `CREATE TABLE` SQL statement to encrypt a table column using the `NOMAC` parameter.

[Example: Changing the Integrity Algorithm for a Table](#) (page 3-25)

You can use the `ALTER TABLE` SQL statement to change the integrity algorithm for a database table.

[Creating an Encrypted Column in an External Table](#) (page 3-25)

The external table feature enables you to access data in external sources as if the data were in a database table.

3.3.4.1 About Creating Tables with Encrypted Columns

You can use the `CREATE TABLE` SQL statement to create a table with an encrypted column.

To create relational tables with encrypted columns, you can specify the `SQL ENCRYPT` clause when you define database columns with the `CREATE TABLE` SQL statement.

3.3.4.2 Creating a Table with an Encrypted Column Using the Default Algorithm

By default, TDE uses the AES encryption algorithm with a 192-bit key length (`AES192`).

If you encrypt a table column without specifying an algorithm, then the column is encrypted using the `AES192` algorithm.

TDE adds `salt` to plaintext before encrypting it. Adding salt makes it harder for attackers to steal data through a brute force attack. TDE also adds a Message Authentication Code (MAC) to the data for integrity checking. The `SHA-1` integrity algorithm is used by default.

- To create a table that encrypts a column, use the `CREATE TABLE` SQL statement with the `ENCRYPT` clause.

For example, to encrypt a table column using the default algorithm:

```
CREATE TABLE employee (
  first_name VARCHAR2(128),
  last_name VARCHAR2(128),
  empID NUMBER,
  salary NUMBER(6) ENCRYPT);
```

This example creates a new table with an encrypted column (`salary`). The column is encrypted using the default encryption algorithm (`AES192`). Salt and MAC are added by default. This example assumes that the wallet is open and a master key is set.

Note:

If there are multiple encrypted columns in a table, then all of these columns must use the same pair of encryption and integrity algorithms.

Salt is specified at the column level. This means that an encrypted column in a table can choose not to use salt irrespective of whether or not other encrypted columns in the table use salt.

3.3.4.3 Creating a Table with an Encrypted Column Using No Algorithm or a Non-Default Algorithm

You can use the `CREATE TABLE` SQL statement to create a table with an encrypted column.

By default, TDE adds `salt` to plaintext before encrypting it. Adding salt makes it harder for attackers to steal data through a brute force attack. However, if you plan to index the encrypted column, then you must use the `NO SALT` parameter.

- To create a table that uses an encrypted column that is a non-default algorithm or no algorithm, run the `CREATE TABLE` SQL statement as follows:
 - If you do not want to use any algorithm, then include the `ENCRYPT NO SALT` clause.
 - If you want to use a non-default algorithm, then use the `ENCRYPT USING` clause, followed by one of the following algorithms enclosed in single quotation marks:

- * `3DES168`
- * `AES128`
- * `AES192` (default)
- * `AES256`

The following example shows how to specify encryption settings for the `empID` and `salary` columns.

```
CREATE TABLE employee (
  first_name VARCHAR2(128),
  last_name VARCHAR2(128),
  empID NUMBER ENCRYPT NO SALT,
  salary NUMBER(6) ENCRYPT USING '3DES168');
```

In this example:

- The `empID` column is encrypted and does not use salt. Both the `empID` and `salary` columns will use the `3DES168` encryption algorithm, because all of the encrypted columns in a table must use the same encryption algorithm.
- The `salary` column is encrypted using the `3DES168` encryption algorithm. Note that the string that specifies the algorithm must be enclosed in single quotation marks ('). The `salary` column uses salt by default.

3.3.4.4 Using the NOMAC Parameter to Save Disk Space and Improve Performance

You can bypass checks that TDE performs. This can save up to 20 bytes of disk space per encrypted value.

If the number of rows and encrypted columns in the table is large, then bypassing TDE checks can add up to a significant amount of disk space. In addition, this saves processing cycles and reduces the performance overhead associated with TDE.

TDE uses the `SHA-1` integrity algorithm by default. All of the encrypted columns in a table must use the same integrity algorithm. If you already have a table column using the `SHA-1` algorithm, then you cannot use the `NOMAC` parameter to encrypt another column in the same table.

- To bypass the integrity check during encryption and decryption operations, use the `NOMAC` parameter in the `CREATE TABLE` and `ALTER TABLE` statements.

Related Topics:

[Performance and Storage Overhead of Transparent Data Encryption](#) (page 5-3)
The performance of Transparent Data Encryption can vary.

3.3.4.5 Example: Using the `NOMAC` Parameter in a `CREATE TABLE` Statement

You can use the `CREATE TABLE` SQL statement to encrypt a table column using the `NOMAC` parameter.

[Example 3-1](#) (page 3-25) creates a table with an encrypted column. The `empID` column is encrypted using the `NOMAC` parameter.

Example 3-1 Using the `NOMAC` parameter in a `CREATE TABLE` statement

```
CREATE TABLE employee (
    first_name VARCHAR2(128),
    last_name VARCHAR2(128),
    empID NUMBER ENCRYPT 'NOMAC' ,
    salary NUMBER(6));
```

3.3.4.6 Example: Changing the Integrity Algorithm for a Table

You can use the `ALTER TABLE` SQL statement to change the integrity algorithm for a database table.

[Example 3-2](#) (page 3-25) shows how to change the integrity algorithm for encrypted columns in a table. The encryption algorithm is set to `3DES168` and the integrity algorithm is set to `SHA-1`. The second `ALTER TABLE` statement sets the integrity algorithm to `NOMAC`.

Example 3-2 Changing the Integrity Algorithm for a Table

```
ALTER TABLE EMPLOYEE REKEY USING '3DES168' 'SHA-1';

ALTER TABLE EMPLOYEE REKEY USING '3DES168' 'NOMAC';
```

3.3.4.7 Creating an Encrypted Column in an External Table

The external table feature enables you to access data in external sources as if the data were in a database table.

External tables can be updated using the `ORACLE_DATAPUMP` access driver.

- To encrypt specific columns in an external table, use the `ENCRYPT` clause when you define those columns:

A system-generated key encrypts the columns. For example, the following `CREATE TABLE` SQL statement encrypts the `ssn` column using the `3DES168` algorithm:

```
CREATE TABLE emp_ext (
    first_name,
    ....
    ssn ENCRYPT USING '3DES168',
    ....
```

If you plan to move an external table to a new location, then you cannot use a randomly generated key to encrypt the columns. This is because the randomly generated key will not be available at the new location.

For such scenarios, you should specify a password while you encrypt the columns. After you move the data, you can use the same password to regenerate the key required to access the encrypted column data at the new location.

Table partition exchange also requires a password-based [TDE table key](#).

[Example 3-3](#) (page 3-26) creates an external table using a password to create the [TDE table key](#).

Example 3-3 Creating a New External Table with a Password-Generated TDE Table Key

```
CREATE TABLE emp_ext (  
    first_name,  
    last_name,  
    empID,  
    salary,  
    ssn ENCRYPT IDENTIFIED BY password  
) ORGANIZATION EXTERNAL  
    (  
    TYPE ORACLE_DATAPUMP  
    DEFAULT DIRECTORY "D_DIR"  
    LOCATION('emp_ext.dat')  
    )  
    REJECT LIMIT UNLIMITED  
AS SELECT * FROM EMPLOYEE;
```

3.3.5 Encrypting Columns in Existing Tables

You can encrypt columns in existing tables. As with new tables, you have a choice of different algorithms to use to definite the encryption.

[About Encrypting Columns in Existing Tables](#) (page 3-26)

The `ALTER TABLE` SQL statement enables you to encrypt columns in an existing table.

[Adding an Encrypted Column to an Existing Table](#) (page 3-27)

You can encrypt columns in existing tables, use a different algorithm, and use `NO SALT` to index the column.

[Encrypting an Unencrypted Column](#) (page 3-27)

You can use the `ALTER TABLE MODIFY` statement to encrypt an existing unencrypted column.

[Disabling Encryption on a Column](#) (page 3-27)

You may want to disable encryption for reasons of compatibility or performance.

3.3.5.1 About Encrypting Columns in Existing Tables

The `ALTER TABLE` SQL statement enables you to encrypt columns in an existing table.

To add an encrypted column to an existing table, or to encrypt or decrypt an existing column, you use the `ALTER TABLE` SQL statement with the `ADD` or `MODIFY` clause.

3.3.5.2 Adding an Encrypted Column to an Existing Table

You can encrypt columns in existing tables, use a different algorithm, and use `NO SALT` to index the column.

- To add an encrypted column to an existing table, use the `ALTER TABLE ADD` statement, specifying the new column with the `ENCRYPT` clause.

Example 3-4 (page 3-27) adds an encrypted column, `ssn`, to an existing table, called `employee`. The `ssn` column is encrypted with the default `AES192` algorithm. Salt and MAC are added by default.

Example 3-4 Adding an Encrypted Column to an Existing Table

```
ALTER TABLE employee ADD (ssn VARCHAR2(11) ENCRYPT);
```

3.3.5.3 Encrypting an Unencrypted Column

You can use the `ALTER TABLE MODIFY` statement to encrypt an existing unencrypted column.

- To encrypt an existing unencrypted column, use the `ALTER TABLE MODIFY` statement, specifying the unencrypted column with the `ENCRYPT` clause.

The following example encrypts the `first_name` column in the `employee` table. The `first_name` column is encrypted with the default `AES192` algorithm. Salt is added to the data, by default. You can encrypt the column using a different algorithm. If you want to index a column, then you must specify `NO SALT`. You can also bypass integrity checks by using the `NOMAC` parameter.

```
ALTER TABLE employee MODIFY (first_name ENCRYPT);
```

The following example encrypts the `first_name` column in the `employee` table using the `NOMAC` parameter.

```
ALTER TABLE employee MODIFY (first_name ENCRYPT 'NOMAC');
```

3.3.5.4 Disabling Encryption on a Column

You may want to disable encryption for reasons of compatibility or performance.

- To disable column encryption, use the `ALTER TABLE MODIFY` command with the `DECRYPT` clause.

Example 3-5 (page 3-27) decrypts the `first_name` column in the `employee` table.

Example 3-5 Turning Off Column Encryption

```
ALTER TABLE employee MODIFY (first_name DECRYPT);
```

3.3.6 Creating an Index on an Encrypted Column

You can create an index on an encrypted column.

The column being indexed must be encrypted without `salt`. If the column is encrypted with salt, then the `ORA-28338: cannot encrypt indexed column(s) with salt` error is raised.

- To create an index on an encrypted column, use the `CREATE INDEX` statement with the `ENCRYPT NO SALT` clause.

[Example 3-6](#) (page 3-28) shows how to create an index on a column that has been encrypted without salt.

Example 3-6 Creating Index on a Column Encrypted Without Salt

```
CREATE TABLE employee (
  first_name VARCHAR2(128),
  last_name VARCHAR2(128),
  empID NUMBER ENCRYPT NO SALT,
  salary NUMBER(6) ENCRYPT USING '3DES168');

CREATE INDEX employee_idx on employee (empID);
```

3.3.7 Adding Salt to an Encrypted Column

Salt, which is a random string added to data before encryption, is a way to strengthen the security of encrypted data. .

Salt ensures that the same plaintext data does not always translate to the same encrypted text. Salt removes the one common method that intruders use to steal data, namely, matching patterns of encrypted text. Adding [salt](#) requires an additional 16 bytes of storage per encrypted data value.

- To add or remove salt from encrypted columns, use the ALTER TABLE MODIFY SQL statement.

For example, suppose you want to encrypt the `first_name` column using salt. If the `first_name` column was encrypted without salt earlier, then the ALTER TABLE MODIFY statement reencrypts it using salt.

```
ALTER TABLE employee MODIFY (first_name ENCRYPT SALT);
```

3.3.8 Removing Salt from an Encrypted Column

You can use the ALTER TABLE SQL statement to remove salt from an encrypted column.

- To remove salt from an encrypted column, use the ENCRYPT NO SALT clause in the ALTER TABLE SQL statement.

For example, suppose you wanted to remove salt from the `first_name` column. If you must index a column that was encrypted using salt, then you can use this statement to remove the salt before indexing

```
ALTER TABLE employee MODIFY (first_name ENCRYPT NO SALT);
```

3.3.9 Changing the Encryption Key or Algorithm for Tables with Encrypted Columns

You can use the ALTER TABLE SQL statement to change the encryption key or algorithm used in encrypted columns.

Each table can have only one [TDE table key](#) for its columns. You can regenerate the TDE table key with the ALTER TABLE statement. This process generates a new key, decrypts the data in the table using the previous key, reencrypts the data using the new key, and then updates the table metadata with the new key information. You can also use a different encryption algorithm for the new TDE table key.

- To change the encryption key or algorithm for tables that contain encrypted columns, use the ALTER TABLE SQL statement with the REKEY or REKEY USING clause.

For example:

```
ALTER TABLE employee REKEY;
```

[Example 3-7](#) (page 3-29) regenerates the TDE table key for the employee table by using the 3DES168 algorithm.

Example 3-7 Changing an Encrypted Table Column Encryption Key and Algorithm

```
ALTER TABLE employee REKEY USING '3DES168';
```

3.4 Encryption Conversions for Tablespaces and Databases

You can perform encryption operations on both offline and online tablespaces and databases.

[About Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

The `CREATE TABLESPACE SQL` statement can be used to encrypt new tablespaces. `ALTER TABLESPACE` can encrypt existing tablespaces.

[Restrictions on Using Transparent Data Encryption Tablespace Encryption](#) (page 3-31)

You should be aware of restrictions on using Transparent Data Encryption when you encrypt a tablespace.

[Creating an Encrypted New Tablespace](#) (page 3-32)

When you create a new tablespace, you can configure its encryption settings during the creation process.

[Encrypting Future Tablespaces](#) (page 3-36)

You can configure Oracle Database to automatically encrypt future tablespaces that you will create.

[Encryption Conversions for Existing Offline Tablespaces](#) (page 3-37)

You can perform offline encryption conversions by using the `ALTER TABLESPACE SQL` statement `OFFLINE`, `ENCRYPT`, and `DECRYPT` clauses.

[Encryption Conversions for Existing Online Tablespaces](#) (page 3-41)

You can encrypt and decrypt an online existing tablespace by using the `ALTER TABLESPACE SQL` statement with the `OFFLINE` and `ENCRYPT` or `DECRYPT` clauses.

[Encryption Conversions for Existing Databases](#) (page 3-48)

You can encrypt both offline and online databases.

3.4.1 About Encryption Conversions for Tablespaces and Databases

The `CREATE TABLESPACE SQL` statement can be used to encrypt new tablespaces. `ALTER TABLESPACE` can encrypt existing tablespaces.

In addition to encrypting new and existing tablespaces, you can encrypt full databases, which entails the encryption of the Oracle-supplied `SYS`, `SYSAUX`, `TEMP`, and `UNDO` tablespaces. To encrypt a full database, you use the `ALTER TABLESPACE` statement, not `ALTER DATABASE`, to encrypt the Oracle-supplied tablespaces.

The following table compares the differences between an offline and an online encryption conversion of tablespaces and databases.

Table 3-2 Offline and Online Tablespace and Database Encryption Conversions

Functionality	Offline Conversion	Online Conversion
Release with minimum conversion capability	Oracle Database 11g release 1 (11.1)	Oracle Database 12c release 2 (12.2)
What can be backported?	The ability to encrypt or decrypt a data file with the AES128 algorithm (using <code>ALTER DATABASE DATAFILE <i>data_file</i> ENCRYPT</code> or <code>DECRYPT</code>) can be used in Oracle Database releases 12.1.0.2 and 11.2.0.4.	No
Algorithms supported	AES128 only	All symmetric encryption algorithm that TDE supports. See About Encryption Conversions for Existing Online Tablespaces (page 3-43) for a list of the supported algorithms.
When can the conversion be run?	When the tablespace is offline or the database is in the mount stage.	When the tablespace is online and database is open in read/write mode.
Is auxiliary space required for the conversion?	No	Yes. See Encrypting an Existing Tablespace with Online Conversion (page 3-41) for guidelines.
Oracle Data Guard conversion guidelines	Convert both the primary and standby manually. Convert the standby first and then switch over to minimum downtime	After you convert the primary, the standby conversion takes place automatically. You cannot perform an online conversion directly on the standby.
Encrypt the SYSTEM, SYSAUX, and UNDO tablespaces (database conversion)	Oracle Database 12c release 2 (12.2) only. You must set <code>COMPATIBILITY</code> to <code>12.2.0.0</code> .	Oracle Database 12c release 2 (12.2) only. You must set <code>COMPATIBILITY</code> to <code>12.2.0.0</code> .
Can an existing TEMP tablespace be converted?	No, but you can create an encrypted TEMP tablespace in Oracle Database 12c release 2 (12.2), make it the default temporary tablespace, and then drop the original TEMP tablespace.	No, but you can create an encrypted TEMP tablespace in Oracle Database 12c release 2 (12.2), make it the default temporary tablespace, and then drop the original TEMP tablespace.

Table 3-2 (Cont.) Offline and Online Tablespace and Database Encryption Conversions

Functionality	Offline Conversion	Online Conversion
Can an existing tablespace be decrypted?	You only can decrypt a tablespace or data file that was previously encrypted by an offline encrypt operation. Oracle does not recommend that you decrypt the UNDO tablespace once it is encrypted.	Yes, but Oracle does not recommend that you decrypt the UNDO tablespace once it is encrypted.
Can encryption keys be rekeyed or rotated?	No, but after the tablespace is encrypted, you can then use online conversion to rekey in Oracle Database 12c release 2 (12.2) compatibility.	Yes
Can encryption operations be run in parallel?	You can run parallel encryption conversions at the data file level with multiple user sessions running.	You can run parallel encryption conversions at the tablespace level with multiple user sessions running.
What to do if an encryption conversion SQL statement fails to complete?	Re-issue the encryption or decryption SQL statement to ensure that all the data files within the tablespace are consistently either encrypted or decrypted.	Rerun the SQL statement but use the FINISH clause.

3.4.2 Restrictions on Using Transparent Data Encryption Tablespace Encryption

You should be aware of restrictions on using Transparent Data Encryption when you encrypt a tablespace.

Note the following restrictions:

- Transparent Data Encryption (TDE) tablespace encryption encrypts or decrypts data during read and write operations, as opposed to TDE column encryption, which encrypts and decrypts data at the SQL layer. This means that most restrictions that apply to TDE column encryption, such as data type restrictions and index type restrictions, do not apply to TDE tablespace encryption.
- To perform import and export operations, use Oracle Data Pump.
- If you encrypt the SYSTEM, SYSAUX, TEMP, or UNDO tablespace, then never close the keystore manually, even if you later decrypt the tablespace by using the ALTER TABLESPACE SQL statement.

See Also: *Oracle Database Utilities* for more information about Oracle Data Pump

3.4.3 Creating an Encrypted New Tablespace

When you create a new tablespace, you can configure its encryption settings during the creation process.

[Step 1: Set the COMPATIBLE Initialization Parameter for Tablespace Encryption](#) (page 3-32)

You must set the COMPATIBLE initialization parameter before creating an encrypted tablespace.

[Step 2: Set the Tablespace TDE Master Encryption Key](#) (page 3-34)

You should ensure that you have configured the TDE master encryption key.

[Step 3: Create the Encrypted Tablespace](#) (page 3-34)

After you have set the COMPATIBLE initialization parameter, you are ready to create the encrypted tablespace.

3.4.3.1 Step 1: Set the COMPATIBLE Initialization Parameter for Tablespace Encryption

You must set the COMPATIBLE initialization parameter before creating an encrypted tablespace.

[About Setting the COMPATIBLE Initialization Parameter for Tablespace Encryption](#) (page 3-32)

A minimum COMPATIBLE initialization parameter setting of 11.2.0.0 enables the full set of tablespace encryption features.

[Setting the COMPATIBLE Initialization Parameter for Tablespace Encryption](#) (page 3-33)

To set the COMPATIBLE initialization parameter, you must edit the initialization parameter file for the database instance.

3.4.3.1.1 About Setting the COMPATIBLE Initialization Parameter for Tablespace Encryption

A minimum COMPATIBLE initialization parameter setting of 11.2.0.0 enables the full set of tablespace encryption features.

Setting the compatibility to 11.2.0.0 enables the following functionality:

- The 11.2.0.0 setting enables the database to use any of the four supported algorithms for data encryption (3DES168, AES128, AES192, and AES256).
- The 11.2.0.0 setting enables the migration of a key from a software keystore to a hardware keystore (ensure that the TDE master encryption key was configured for the hardware keystore)
- The 11.2.0.0 setting enables resetting and rotating the TDE master encryption key

Be aware that once you set the COMPATIBLE parameter to 11.2.0.0, the change is irreversible. To use tablespace encryption, ensure that the compatibility setting is at the minimum, which is 11.2.0.0.

See Also:

- *Oracle Database SQL Language Reference* for more information about the COMPATIBLE parameter
- *Oracle Database Administrator's Guide* for more information about initialization parameter files

3.4.3.1.2 Setting the COMPATIBLE Initialization Parameter for Tablespace Encryption

To set the COMPATIBLE initialization parameter, you must edit the initialization parameter file for the database instance.

1. Log in to the database instance.

In a multitenant environment, log in to the PDB. For example:

```
sqlplus sec_admin@hrpdb
Enter password: password
Connected.
```

To find the available PDBs, query the DBA_PDBS data dictionary view. To check the current PDB, run the `show con_name` command.

2. Check the current setting of the COMPATIBLE parameter.

For example:

```
SHOW PARAMETER COMPATIBLE
```

NAME	TYPE	VALUE
compatible	string	11.2.0.0
noncdbcompatible	BOOLEAN	FALSE

3. If you must change the COMPATIBLE parameter, then complete the remaining steps in this procedure.

The value should be 11.2.0.0 or higher.

4. From the command line, locate the initialization parameter file for the database instance.

- **UNIX systems:** This file is in the `ORACLE_HOME/dbs` directory and is named `initORACLE_SID.ora` (for example, `initmydb.ora`).
- **Windows systems:** This file is in the `ORACLE_HOME\database` directory and is named `initORACLE_SID.ora` (for example, `initmydb.ora`).

5. Edit the initialization parameter file to use the new COMPATIBLE setting.

For example:

```
compatible=12.2.0.0.0
```

6. In SQL*Plus, connect as a user who has the SYSDBA administrative privilege, and then restart the database.

For example:

```
CONNECT /AS SYSDBA
SHUTDOWN
STARTUP
```

If tablespace encryption is in use, then open the keystore at the database mount. The keystore must be open before you can access data in an encrypted tablespace.

For example:

```
STARTUP MOUNT;
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY keystore_password;
ALTER DATABASE OPEN;
```

3.4.3.2 Step 2: Set the Tablespace TDE Master Encryption Key

You should ensure that you have configured the TDE master encryption key.

- Set the TDE master encryption key as follows:
 - For software TDE master encryption keys, see [Step 4: Set the Software TDE Master Encryption Key](#) (page 3-11).
 - For hardware TDE master encryption keys, see [Step 4: Set the Hardware Keystore TDE Master Encryption Key](#) (page 3-17).

3.4.3.3 Step 3: Create the Encrypted Tablespace

After you have set the `COMPATIBLE` initialization parameter, you are ready to create the encrypted tablespace.

[About Creating Encrypted Tablespaces](#) (page 3-34)

To create an encrypted tablespace, you can use the `CREATE TABLESPACE SQL` statement.

[Creating an Encrypted Tablespace](#) (page 3-35)

To create an encrypted tablespace, you must use the `CREATE TABLESPACE` statement with the `ENCRYPTION USING` clause.

[Example: Creating an Encrypted Tablespace That Uses AES192](#) (page 3-35)

You can use the `CREATE TABLESPACE SQL` statement to create an encrypted tablespace.

[Example: Creating an Encrypted Tablespace That Uses the Default Algorithm](#) (page 3-36)

You can use the `CREATE TABLESPACE SQL` statement to create an encrypted tablespace that uses the default algorithm.

3.4.3.3.1 About Creating Encrypted Tablespaces

To create an encrypted tablespace, you can use the `CREATE TABLESPACE SQL` statement.

You must have the `CREATE TABLESPACE` system privilege to create an encrypted tablespace.

You can import data into an encrypted tablespace by using Oracle Data Pump. You can also use a SQL statement such as `CREATE TABLE . . . AS SELECT . . .` or `ALTER TABLE . . . MOVE . . .` to move data into an encrypted tablespace. The `CREATE TABLE . . . AS SELECT . . .` statement creates a table from an existing table. The `ALTER TABLE . . . MOVE . . .` statement moves a table into the encrypted tablespace.

For security reasons, you cannot encrypt a tablespace with the `NO SALT` option.

You can query the `ENCRYPTED` column of the `DBA_TABLESPACES` and `USER_TABLESPACES` data dictionary views to verify if a tablespace was encrypted.

See Also:

Oracle Database Reference for more information about the `DBA_TABLESPACES` and `USER_TABLESPACES` data dictionary views

3.4.3.3.2 Creating an Encrypted Tablespace

To create an encrypted tablespace, you must use the `CREATE TABLESPACE` statement with the `ENCRYPTION USING` clause.

1. Log in to the database instance as a user who has been granted the `CREATE TABLESPACE` system privilege.

In a multitenant environment, log in to the PDB. For example:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Run the `CREATE TABLESPACE` statement, using its encryption clauses.

For example:

```
CREATE TABLESPACE encrypt_ts
  DATAFILE '$ORACLE_HOME/dbs/encrypt_df.dbf' SIZE 1M
  ENCRYPTION USING 'AES256' ENCRYPT;
```

In this specification:

- `ENCRYPTION USING 'AES256' ENCRYPT` specifies the encryption algorithm and the key length for the encryption. The `ENCRYPT` clause encrypts the tablespace. Enclose this setting in single quotation marks (' '). The key lengths are included in the names of the algorithms. If you do not specify an encryption algorithm, then the default encryption algorithm, `AES128`, is used.

See Also:

- [Supported Encryption and Integrity Algorithms](#) (page 2-7)
 - *Oracle Database SQL Language Reference*
-
-

3.4.3.3.3 Example: Creating an Encrypted Tablespace That Uses AES192

You can use the `CREATE TABLESPACE SQL` statement to create an encrypted tablespace.

Example 3-8 (page 3-36) creates a tablespace called `seurespace_1` that is encrypted using the 3DES algorithm. The key length is 168 bits.

Example 3-8 Creating an Encrypted Tablespace That Uses AES192

```
CREATE TABLESPACE securespace_1
DATAFILE '/home/user/oradata/secure01.dbf'
SIZE 150M
ENCRYPTION USING 'AES192' ENCRYPT;
```

3.4.3.3.4 Example: Creating an Encrypted Tablespace That Uses the Default Algorithm

You can use the `CREATE TABLESPACE` SQL statement to create an encrypted tablespace that uses the default algorithm.

[Example 3-9](#) (page 3-36) creates a tablespace called `securespace_2`. Because no encryption algorithm is specified, the default encryption algorithm (AES128) is used. The key length is 128 bits.

You cannot encrypt an existing tablespace.

Example 3-9 Creating an Encrypted Tablespace That Uses the Default Algorithm

```
CREATE TABLESPACE securespace_2
DATAFILE '/home/user/oradata/secure01.dbf'
SIZE 150M
ENCRYPTION ENCRYPT;
```

3.4.4 Encrypting Future Tablespaces

You can configure Oracle Database to automatically encrypt future tablespaces that you will create.

About Encrypting Future Tablespaces (page 3-36)

The ability to encrypt future tablespaces can help prevent data breaches in Oracle Cloud environments.

Setting Future Tablespaces to be Encrypted (page 3-37)

You can set the `ENCRYPT_NEW_TABLESPACES` database initialization parameter to automatically encrypt future tablespaces that you create.

3.4.4.1 About Encrypting Future Tablespaces

The ability to encrypt future tablespaces can help prevent data breaches in Oracle Cloud environments.

The `ENCRYPT_NEW_TABLESPACES` database initialization parameter controls how future databases are encrypted.

You can create and run an Oracle database completely in Oracle Cloud. Because this configuration hosts the customer's data in the Cloud, Oracle recommends that you enable encryption as much as possible. A long-term goal is to encrypt all data in Oracle Cloud. Alternatively, you can have the database both in the Cloud and on premises.

In an Oracle Cloud environment, the following scenarios may occur when you create encrypted tablespaces in Oracle Cloud and on-premises environments:

- You create a test database in Oracle Cloud and the tablespaces were encrypted by using when the `ENCRYPT_NEW_TABLESPACE` parameter has been set to automatically create the Cloud database as encrypted. However, you may not have the intention or even an Advanced Security Option license to bring the encrypted database back on premises.

- You create a hybrid definer's rights environment where the primary database is on premises and the standby database is on Oracle Cloud. If a switchover operation takes place, then the new primary is on Oracle Cloud. If a new tablespace is transparently encrypted, then a similar scenario to the first item in this list may occur. For example, suppose you do not have an Advanced Security Option (ASO) license, and you have an automatically encrypted tablespace in the Oracle Cloud. The standby database on premises is also automatically encrypted. In this case, because you do not have an ASO license, you cannot use the standby database. To remedy this problem, set the `ENCRYPT_NEW_TABLESPACES` to `DDL`, which prevents the encryption of the tablespace in Oracle Cloud.

3.4.4.2 Setting Future Tablespaces to be Encrypted

You can set the `ENCRYPT_NEW_TABLESPACES` database initialization parameter to automatically encrypt future tablespaces that you create.

- In SQL*Plus, enter the following `ALTER SYSTEM` statement:

```
ALTER SYSTEM SET ENCRYPT_NEW_TABLESPACES = value;
```

In this specification, *value* can be:

- `CLOUD_ONLY` transparently encrypts the tablespace in the Cloud using the AES128 algorithm if you do not specify the `ENCRYPTION` clause of the `CREATE TABLESPACE` SQL statement. It applies only to an Oracle Cloud environment. If you create the tablespace on premise, then it will follow the `CREATE TABLESPACE` statement specification that you enter. For example, if you omit the `ENCRYPTION` clause, then the tablespace is created unencrypted. If you include this clause and use a different algorithm, then the tablespace will use that algorithm. `CLOUD_ONLY` is the default.
- `ALWAYS` automatically encrypts the tablespace using the AES128 algorithm if you omit the `ENCRYPTION` clause of `CREATE TABLESPACE`, for both the Cloud and premises scenarios.
If you do provide the `ENCRYPTION` clause, however, the algorithm that you specify takes precedence over AES128.
- `DDL` encrypts the tablespace using the specified setting of the `ENCRYPTION` clause of `CREATE TABLESPACE`, for both Oracle Cloud and on-premise environments.

Related Topics:

`ENCRYPT_NEW_TABLESPACES`

3.4.5 Encryption Conversions for Existing Offline Tablespaces

You can perform offline encryption conversions by using the `ALTER TABLESPACE` SQL statement `OFFLINE`, `ENCRYPT`, and `DECRYPT` clauses.

[About Encryption Conversions for Existing Offline Tablespaces](#) (page 3-38)

You can encrypt or decrypt an existing data file of a user tablespace when the tablespace is offline or when the database is not open.

[Encrypting an Existing User-Defined Tablespace with Offline Conversion](#) (page 3-39)

To encrypt an existing tablespace with offline conversion, you can use the ALTER TABLESPACE SQL statement with the OFFLINE and ENCRYPT clauses.

[Decrypting an Existing Tablespace with Offline Conversion](#) (page 3-40)

To decrypt an existing tablespace with offline conversion, you can use the ALTER TABLESPACE SQL statement with the OFFLINE and DECRYPT clauses.

3.4.5.1 About Encryption Conversions for Existing Offline Tablespaces

You can encrypt or decrypt an existing data file of a user tablespace when the tablespace is offline or when the database is not open.

Use the offline encryption method if you do not plan to change the compatibility of your databases from Oracle Database 11g release 2 (11.2) or Oracle Database 12c release 1 (12.1) to Release 12.2, which is irreversible. The offline encryption method is also useful if you want to quickly make use of Transparent Data Encryption before you upgrade this database to release 12.2. You can both encrypt and decrypt offline tablespaces.

Note the following:

- If you want to encrypt the Oracle Database-supplied tablespaces (SYSTEM, SYSAUX, and UNDO) using the offline conversion method, then you must use the method that is described in [Encrypting an Existing Database with Offline Conversion](#) (page 3-49).
- You can use the online method to rekey a tablespace that was previously encrypted with the offline method.
- If you have configured Oracle Data Guard, you can minimize downtime by encrypting the tablespaces on the standby first, switching over to the primary, and then encrypting the tablespaces on the primary.
- You cannot specify the encryption algorithm in an offline conversion. In an offline conversion, all data files and tablespaces are encrypted using the AES128 encryption key. You can check the encryption key by querying the ENCRYPTIONALG column in the V\$DATABASE_KEY_INFO view.
- You can convert offline tablespaces in parallel by using multiple foreground sessions to encrypt different data files.
- If you are using Oracle Data Guard, you can minimize the downtime by encrypting the tablespaces on the standby first, switching over, and then encrypting the tablespaces on the original primary next.
- For Oracle Database 11g release 2 (11.2.0.4) and Oracle Database 12c release 1 (12.1.0.2), you cannot perform an offline encryption of the SYSTEM and SYSAUX tablespaces. Also, Oracle does not recommend encrypting offline the UNDO tablespace in these releases. Doing so prevents the keystore from being closed, and this prevents the database from functioning. In addition, encrypting the UNDO tablespace while the database is offline is not necessary because all undo records that are associated with any encrypted tablespaces are already automatically encrypted in the UNDO tablespace. If you want to encrypt the TEMP tablespace, you must drop and then recreate it as encrypted.

3.4.5.2 Encrypting an Existing User-Defined Tablespace with Offline Conversion

To encrypt an existing tablespace with offline conversion, you can use the `ALTER TABLESPACE SQL` statement with the `OFFLINE` and `ENCRYPT` clauses.

The procedure that is described in this section applies to the case where you want to encrypt individual user-created tablespaces within a database. These tablespaces can be encrypted offline. However, the Oracle Database-supplied `SYSTEM` and `UNDO` tablespaces cannot be brought offline. If you want to encrypt the tablespaces offline, then you must use the method that is described in [Encrypting an Existing Database with Offline Conversion](#) (page 3-49).

1. Connect as a user who has the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the `SYSDBA` administrative privilege if you plan to encrypt the `SYSTEM` and `SYSAUX` tablespaces. Otherwise, connect with the `SYSKM` administrative privilege.

2. Bring the tablespace offline.

```
ALTER TABLESPACE users OFFLINE NORMAL;
```

3. Back up the tablespace.

The offline conversion method does not use auxiliary disk space or files, and it operates directly in-place to the data files. Therefore, you should perform a full backup of the user tablespace before converting it offline.

4. As a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege, open the software keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY
software_keystore_password;
```

5. Encrypt the tablespace.

For example, to encrypt an entire tablespace, include its data files:

```
ALTER TABLESPACE users ENCRYPTION OFFLINE ENCRYPT;
```

To encrypt individual data files within a tablespace, use the `ALTER DATABASE DATAFILE SQL` statement. For example, to encrypt the data files `user_01.dbf` and `user_02.dbf`:

```
ALTER DATABASE DATAFILE 'user_01.dbf' ENCRYPT;
ALTER DATABASE DATAFILE 'user_02.dbf' ENCRYPT;
```

In the same database session, these statements encrypt each of the data files in sequence, one after another. If you execute each statement in its own database session, then they will be executed in parallel.

If the encryption process is interrupted, then rerun the `ALTER TABLESPACE` statement. The kinds of errors that you can expect in an interruption are general errors, such as file system or storage file system errors. The data files within the

tablespace must be consistently encrypted. For example, suppose you offline a tablespace that has 10 files but for some reason, the encryption only completes for nine of the files, leaving one decrypted. Before you can bring the tablespace back online, you must encrypt all of the files, not just some of them.

6. Bring the tablespace back online or open the database.

- To bring the tablespace back online:

```
ALTER TABLESPACE users ONLINE;
```

- To open a database in a non-multitenant environment:

```
ALTER DATABASE OPEN
```

- In a multitenant environment, you can encrypt a data file or tablespace with the offline method if the root is open and the PDB is not open. For example, for a PDB named `hr_pdb`:

```
ALTER PLUGGABLE DATABASE hr_pdb OPEN
```

Related Topics:

[Opening a Software Keystore](#) (page 3-10)

To open a software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

3.4.5.3 Decrypting an Existing Tablespace with Offline Conversion

To decrypt an existing tablespace with offline conversion, you can use the `ALTER TABLESPACE SQL` statement with the `OFFLINE` and `DECRYPT` clauses.

1. Connect as a user who has the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the `SYSDBA` administrative privilege if you plan to decrypt the `SYSTEM` and `SYSAUX` tablespaces. Otherwise, connect with the `SYSKM` administrative privilege.

2. Bring the tablespace offline.

```
ALTER TABLESPACE users OFFLINE NORMAL;
```

3. As a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege, open the keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY
software_keystore_password;
```

4. Run the `ALTER TABLESPACE SQL` statement to perform the decryption.

For example, for a tablespace called `users`:

```
ALTER TABLESPACE users ENCRYPTION OFFLINE DECRYPT;
```

If the encryption process is interrupted, then rerun the `ALTER TABLESPACE` statement. The kinds of errors that you can expect in an interruption are general

errors, such as file system or storage file system errors. The data files within the tablespace must be consistently decrypted. For example, suppose you offline a tablespace that has 10 files but for some reason, the decryption only completes for 9 of the files, leaving one encrypted. Before you can bring the tablespace back online, you must decrypt all of the files, not just some of them.

5. Bring the tablespace online.

```
ALTER TABLESPACE users ONLINE;
```

Related Topics:

[Opening a Software Keystore](#) (page 3-10)

To open a software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

3.4.6 Encryption Conversions for Existing Online Tablespaces

You can encrypt and decrypt an online existing tablespace by using the `ALTER TABLESPACE SQL` statement with the `OFFLINE` and `ENCRYPT` or `DECRYPT` clauses.

[Encrypting an Existing Tablespace with Online Conversion](#) (page 3-41)

To encrypt an existing tablespace with online conversion, use `ALTER TABLESPACE` with the `ONLINE` and `ENCRYPT` clauses.

[About Encryption Conversions for Existing Online Tablespaces](#) (page 3-43)

You can encrypt, decrypt, or rekey existing user tablespaces, and the `SYSTEM`, `SYSAUX`, and `UNDO` tablespace when they are online.

[Rekeying an Existing Tablespace with Online Conversion](#) (page 3-44)

To rekey an existing tablespace that is online, you can use the `REKEY` clause of the `ALTER TABLESPACE SQL` statement.

[Decrypting an Existing Tablespace with Online Conversion](#) (page 3-46)

To decrypt an existing tablespace with online conversion, you can use the `ALTER TABLESPACE SQL` statement with `DECRYPT` clause.

[Finishing an Interrupted Online Encryption Conversion](#) (page 3-47)

If an online encryption process is interrupted, then you can complete the conversion by rerunning the `ALTER TABLESPACE` statement using the `FINISH` clause.

3.4.6.1 Encrypting an Existing Tablespace with Online Conversion

To encrypt an existing tablespace with online conversion, use `ALTER TABLESPACE` with the `ONLINE` and `ENCRYPT` clauses.

1. Connect as a user who has the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the `SYSDBA` administrative privilege if you plan to encrypt the `SYSTEM` and `SYSAUX` tablespaces. Otherwise, connect with the `SYSKM` administrative privilege.

2. Ensure that the `COMPATIBLE` initialization parameter is set to 12.2.0.0.

You can use the `SHOW PARAMETER` command to check the current setting of a parameter.

3. Ensure that the database is open in read-write mode.

You can query the `STATUS` column of the `V$INSTANCE` dynamic view to find if a database is open and the `OPEN_MODE` column of the `V$DATABASE` view to find if it is in read-write mode.

4. If necessary, open the database in read-write mode.

```
ALTER DATABASE OPEN READ WRITE;
```

5. Ensure that the auxiliary space is at least the same size as the largest data file of this tablespace.

This size requirement is because Oracle Database performs the conversion one file at a time. For example, if the largest data file of the tablespace is 32 GB, then ensure that you have 32 GB of auxiliary space. To find the space used by a data file, query the `BYTES` or `BLOCKS` column of the `V$DATAFILE` dynamic performance view.

6. Create and open a master encryption key.

For example:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE 'keystore_location' IDENTIFIED BY
software_keystore_password;
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY
software_keystore_password;
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY software_keystore_password WITH
BACKUP;
```

7. Run the `ALTER TABLESPACE` statement using the `ENCRYPTION` and `ENCRYPT` clauses to perform the encryption.

For example, for a non-Oracle managed files tablespace named `users`:

```
ALTER TABLESPACE users ENCRYPTION ONLINE USING 'AES192' ENCRYPT FILE_NAME_CONVERT
= ('users.dbf', 'users_enc.dbf');
```

In this example:

- `ENCRYPTION ONLINE USING 'AES192' ENCRYPT` sets the statement to encrypt the tablespace `users` while it is online and assigns it the AES192 encryption algorithm. If you omit the `USING algorithm` clause, then the default algorithm, AES128, is used. For the `SYSTEM` and `UNDO` tablespaces, you can use the `ENCRYPT` clause to encrypt the tablespace, but you cannot specify an encryption algorithm because they must be encrypted with the existing database key the first time. After encrypting the tablespace, use the `REKEY` clause to specify the algorithm.
- `FILE_NAME_CONVERT` specifies one or more pairs of data files that are associated with the tablespace. The first name in the pair is an existing data file, and the second name is for the encrypted version of this data file, which will be created after the `ALTER TABLESPACE` statement successfully executes. If the tablespace has more than one data file, then you must process them all in this statement. Note the following:
 - Separate each file name with a comma, including multiple pairs of files. For example:

```
FILE_NAME_CONVERT = ('users1.dbf', 'users1_enc.dbf', 'users2.dbf',
'users2_enc.dbf')
```

- You can specify directory paths in the `FILE_NAME_CONVERT` clause. For example, the following clause converts and moves the matching files of the tablespace from the `dbfs` directory to the `dbfs/enc` directory:

```
FILE_NAME_CONVERT = ('dbfs', 'dbfs/enc')
```

- The `FILE_NAME_CONVERT` clause recognizes patterns. The following example converts the data files `users_1.dbf` and `users_2.dbf` to `users_enc1.dbf` and `users_enc2.dbf`:

```
FILE_NAME_CONVERT = ('users', 'users_enc')
```

- In an Oracle Data Guard environment, include the name of the standby database data file in the `FILE_NAME_CONVERT` settings.
- You must use the `FILE_NAME_CONVERT` clause for non-Oracle managed files. (In an Oracle-managed files configuration, new data files are created automatically.)
- You can find the data files for a tablespace by querying the `V$DATAFILE` or `V$DATAFILE_HEADER` dynamic views.

By default, data files are in the `$ORACLE_HOME/dbfs` directory. If the data files are located there, then you do not have to specify a path.

After you complete the conversion, you can check the encryption status by querying the `STATUS` column of the `V$ENCRYPTED_TABLESPACES` dynamic view. The `ENCRYPTIONALG` column of this view shows the encryption algorithm that is used. If the conversion process was interrupted, then you can resume it by running `ALTER TABLESPACE` with the `FINISH` clause. For example, if the primary data file converts but the standby data file does not, then you can run `ALTER TABLESPACE . . . FINISH` on the standby database for the standby data files.

Related Topics:

[Setting the COMPATIBLE Initialization Parameter for Tablespace Encryption](#)
(page 3-33)

To set the `COMPATIBLE` initialization parameter, you must edit the initialization parameter file for the database instance.

[Finishing an Interrupted Online Encryption Conversion](#) (page 3-47)

If an online encryption process is interrupted, then you can complete the conversion by rerunning the `ALTER TABLESPACE` statement using the `FINISH` clause.

3.4.6.2 About Encryption Conversions for Existing Online Tablespaces

You can encrypt, decrypt, or rekey existing user tablespaces, and the `SYSTEM`, `SYSAUX`, and `UNDO` tablespace when they are online.

However, you cannot encrypt, decrypt, or rekey a temporary tablespace online.

An online tablespace can be created by using the `ONLINE` clause of the `CREATE TABLESPACE` SQL statement. When you encrypt or rekey a tablespace online, the tablespace will have its own independent encryption keys and algorithms.

Note the following:

- If an offline tablespace has been encrypted, then you can rekey it online to use a different algorithm.
- You can encrypt multiple tablespaces online in parallel by using multiple foreground sessions to encrypt different tablespaces. Within each tablespace, the data files are encrypted sequentially.
- If the conversion is interrupted, then you can resume the process by issuing the `FINISH` clause of the `ALTER TABLESPACE SQL` statement.
- A redo log is generated for each online tablespace conversion.
- Do not encrypt the `SYSTEM` and `UNDO` tablespaces concurrently with other tablespaces.
- You cannot use the transportable tablespace feature with Oracle Data Pump while you are encrypting a tablespace.
- You cannot run the `ALTER TABLESPACE` statement concurrently with the following features:
 - `ADMINISTER KEY MANAGEMENT SET KEY SQL` statement
 - `FLASHBACK DATABASE SQL` statement
- If you are using Oracle-managed files for the data files, then the encryption process rekeys the data files that are associated with the tablespace and then copies or moves them to the default Oracle-managed files location.
- You can add new files to the tablespace after you have encrypted it. Oracle Database reformats the new file with the new encryption key. Blocks will be encrypted using the new key.
- Previous operations that took place in the root or the PDB may require the control files to be cross-checked against the data dictionary before you can begin the online conversion process. An `ORA-241 operation disallowed: control file is not yet checked against data dictionary` error may occur. To resolve this problem, restart the root or PDB, and then try issuing the online conversion commands again.

Related Topics:

[Supported Encryption and Integrity Algorithms](#) (page 2-7)

By default, Transparent Data Encryption (TDE) Column encryption uses the Advanced Encryption Standard (AES).

3.4.6.3 Rekeying an Existing Tablespace with Online Conversion

To rekey an existing tablespace that is online, you can use the `REKEY` clause of the `ALTER TABLESPACE SQL` statement.

Before you perform a rekey operation, be aware of the following:

- You cannot rekey the `TEMP` tablespace. If you want to assign a different encryption algorithm to a `TEMP` tablespace, then drop `TEMP` and recreate it with the correct encryption algorithm.

- Do not perform an online tablespace rekey operation with a master key operation concurrently. To find if any tablespaces are currently being rekeyed, issue the following query to find the rekey status of encrypted tablespaces:

```
SELECT TS#, ENCRYPTIONALG, STATUS FROM V$ENCRYPTED_TABLESPACES;
```

A status of `REKEYING` means that the corresponding tablespace is still being rekeyed. Do not rekey the master key while this status is in effect.

To rekey an existing tablespace with online conversion:

1. Connect as a user who has the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the `SYSDBA` administrative privilege if you plan to rekey the `SYSTEM` and `SYSAUX` tablespaces. Otherwise, connect with the `SYSKM` administrative privilege.

2. Ensure that the following requirements are met:

- The `COMPATIBLE` initialization parameter is set to 12.2.0.0.
- The database is open and in read-write mode.
- A master encryption key has been created and is open.

3. Query the `KEY_VERSION` and `STATUS` columns of the `V$ENCRYPTED_TABLESPACES` dynamic view to find the current status of the encryption algorithm used by the master encryption key.

4. Perform the rekey operation, based on the status returned by the `V$ENCRYPTED_TABLESPACES` dynamic view:

- If the key version status of the tablespace is `NORMAL`, then specify the new algorithm of the online tablespace rekey.

For example:

```
ALTER TABLESPACE users ENCRYPTION USING 'AES192' REKEY FILE_NAME_CONVERT =
('users.dbf', 'users_enc.dbf');
```

- If the key version status is `ENCRYPTING`, `DECRYPTING`, or `REKEYING`, then use the `FINISH` clause.

For example:

```
ALTER TABLESPACE users ENCRYPTION ONLINE FINISH REKEY FILE_NAME_CONVERT =
('users.dbf', 'users_enc.dbf');
```

5. If the `ORA-00241 operation disallowed: control file inconsistent with data dictionary` error appears, then restart the database.

In a multitenant environment, restart the CDB root database and then retry Step 4 (page 3-45).

If the conversion process was interrupted, then you can resume it by running `ALTER TABLESPACE` with the `FINISH` clause.

Related Topics:[Encrypting an Existing Tablespace with Online Conversion](#) (page 3-41)

To encrypt an existing tablespace with online conversion, use `ALTER TABLESPACE` with the `ONLINE` and `ENCRYPT` clauses.

[About Encryption Conversions for Existing Online Tablespaces](#) (page 3-43)

You can encrypt, decrypt, or rekey existing user tablespaces, and the `SYSTEM`, `SYSAUX`, and `UNDO` tablespace when they are online.

[Finishing an Interrupted Online Encryption Conversion](#) (page 3-47)

If an online encryption process is interrupted, then you can complete the conversion by rerunning the `ALTER TABLESPACE` statement using the `FINISH` clause.

3.4.6.4 Decrypting an Existing Tablespace with Online Conversion

To decrypt an existing tablespace with online conversion, you can use the `ALTER TABLESPACE SQL` statement with `DECRYPT` clause.

1. Connect as a user who has the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the `SYSDBA` administrative privilege if you plan to decrypt the `SYSTEM` and `SYSAUX` tablespaces. Otherwise, connect with the `SYSKM` administrative privilege.

2. Ensure that the following requirements are met:
 - The `COMPATIBLE` initialization parameter is set to 12.2.0.0.
 - The database is open and in read-write mode.
 - A master encryption key has been created and is open.
 - There is enough auxiliary space to complete the decryption.
3. Run the `ALTER TABLESPACE SQL` statement with the `DECRYPT` clause.

For example:

```
ALTER TABLESPACE users ENCRYPTION ONLINE DECRYPT FILE_NAME_CONVERT =
('users_enc.dbf', 'users.dbf');
```

In this specification:

- When you specify the files to decrypt, enter them in the reverse order in which they were originally encrypted. That is, first enter the name of the encrypted file (`users_enc.dbf`), followed by the data file (`users.dbf`).
- Do not provide an algorithm key for the decryption.

If the conversion process was interrupted, then you can resume it by running `ALTER TABLESPACE` with the `FINISH` clause.

Related Topics:[Encrypting an Existing Tablespace with Online Conversion](#) (page 3-41)

To encrypt an existing tablespace with online conversion, use `ALTER TABLESPACE` with the `ONLINE` and `ENCRYPT` clauses.

[Finishing an Interrupted Online Encryption Conversion](#) (page 3-47)

If an online encryption process is interrupted, then you can complete the conversion by rerunning the `ALTER TABLESPACE` statement using the `FINISH` clause.

3.4.6.5 Finishing an Interrupted Online Encryption Conversion

If an online encryption process is interrupted, then you can complete the conversion by rerunning the `ALTER TABLESPACE` statement using the `FINISH` clause.

An interrupted encryption process (encryption, rekey, or decryption) can be, for example, an `ORA-28425: missing a valid FILE_NAME_CONVERT clause` error in the `FILE_NAME_CONVERT` clause of the `ALTER TABLESPACE` SQL statement. Other examples of interrupted processes are if the conversion skips a data file, which can happen if there is an error when an Oracle Data Base WRiter (DBWR) process offlines a data file, or if there is not enough space for the auxiliary file. The tablespace should be operational even if you do not rerun the `ALTER TABLESPACE` statement with the `FINISH` clause.

1. Query the `V$ENCRYPTED_TABLESPACES` to check the `STATUS` column for the tablespace.

If the `STATUS` column reports `ENCRYPTING`, `DECRYPTING`, or `REKEYING`, then rerun the `ALTER TABLESPACE` statement with the `FINISH` clause, as described in this procedure. If the `STATUS` reports `NORMAL`, then you can rerun `ALTER TABLESPACE` without the `FINISH` clause.

You can find the tablespace name that matches the `TS#` and `TABLESPACE_NAME` columns by querying the `V$DATAFILE_HEADER` view.

2. If necessary query the following additional views to find information about the tablespace whose online conversion was interrupted:

- `DBA_TABLESPACES` to find if the `STATUS` of the tablespace indicates if it is online or offline.
- `V$ENCRYPTED_TABLESPACES` to find if the `STATUS` of the tablespace indicates if it is encrypted, and what the `KEY_VERSION` of the encryption key is.
- `V$DATAFILE` and `V$DATAFILE_HEADER` to find the data files that are associated with a tablespace.

3. Run the `ALTER TABLESPACE` statement using the `FINISH` clause.

Examples are as follows:

- For an encryption operation:

```
ALTER TABLESPACE users ENCRYPTION ONLINE FINISH ENCRYPT FILE_NAME_CONVERT = ('users.dbf', 'users_enc.dbf');
```

- For a decryption operation:

```
ALTER TABLESPACE users ENCRYPTION ONLINE FINISH DECRYPT FILE_NAME_CONVERT =  
( 'users_enc.dbf', 'users.dbf' );
```

Note the order in which the files are specified: first, the name of the encrypted file, and then the name of the data file. (In the encryption operation, the name of the data file is specified first, followed by the name of the encrypted file.)

- For a rekey operation:

```
ALTER TABLESPACE users ENCRYPTION ONLINE FINISH REKEY FILE_NAME_CONVERT =  
( 'users.dbf', 'users_enc.dbf' );
```

You cannot specify an algorithm when you use the `FINISH` clause in an `ALTER TABLESPACE` statement.

4. To check the conversion, query the `STATUS` column of the `V$ENCRYPTED_TABLESPACES` view.

The status should be `NORMAL`. In an Oracle Data Guard environment, if the database does not have `NORMAL` as the `STATUS`, then run the `ALTER TABLESPACE . . . FINISH` statement on the primary or the standby data file that did not successfully convert.

3.4.7 Encryption Conversions for Existing Databases

You can encrypt both offline and online databases.

[About Encryption Conversions for Existing Databases](#) (page 3-48)

The encryption conversion of an entire database encrypts all tablespaces, including the Oracle-supplied `SYSTEM`, `SYSAUX`, `UNDO`, and `TEMP` tablespaces.

[Encrypting an Existing Database with Offline Conversion](#) (page 3-49)

When you encrypt an existing database with offline conversion, you do not specify an encryption algorithm.

[Encrypting an Existing Database with Online Conversion](#) (page 3-50)

When you encrypt an existing database with online conversion, you do not specify an encryption algorithm.

3.4.7.1 About Encryption Conversions for Existing Databases

The encryption conversion of an entire database encrypts all tablespaces, including the Oracle-supplied `SYSTEM`, `SYSAUX`, `UNDO`, and `TEMP` tablespaces.

Note the following:

- To perform the encryption, you can use the offline and online functionality of the tablespace encryption conversions.
- You can encrypt any or all of the Oracle-supplied tablespaces, and in any order. The encryption of the Oracle-supplied tablespaces has no impact on the encryption of user-created tablespaces.
- When you encrypt the Oracle-supplied tablespaces, Oracle Database prevents the keystore from being closed.
- You cannot encrypt an existing temporary tablespace, but you can drop the existing temporary tablespace and then recreate it as encrypted.

- The UNDO and TEMP metadata that is generated from sensitive data in an encrypted tablespace is already automatically encrypted. Therefore, encrypting UNDO and TEMP is optional.
- Oracle recommends that you encrypt the Oracle-supplied tablespaces by using the default tablespace encryption algorithm, AES128. However, you can rekey any of these tablespaces afterwards to use a different encryption algorithm if you want. (To find the current encryption key for the current database, you can query the V\$DATABASE_KEY_INFO dynamic view.)
- The performance effect of encrypting all the tablespaces in a database depends on the workload and platform. Many modern CPUs provide built-in hardware acceleration, which results in a minimal performance impact.
- In a multitenant environment, you can encrypt any tablespaces in any pluggable databases (PDBs), including the Oracle-supplied tablespaces. However, the keystore in the CDB root must be open at all times so that a PDB can open its keystore. You can check the status of whether a keystore is open by querying the STATUS column of the V\$ENCRYPTION_WALLET view

3.4.7.2 Encrypting an Existing Database with Offline Conversion

When you encrypt an existing database with offline conversion, you do not specify an encryption algorithm.

1. Connect as a user who has the SYSDBA administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

You must have the SYSDBA administrative privilege to encrypt the SYSTEM and SYSAUX tablespaces.

2. Mount the database.

```
STARTUP MOUNT
```

3. Open the keystore.

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY keystore_password;
```

4. Run the ALTER TABLESPACE SQL statement to encrypt the SYSTEM, SYSAUX, and UNDO tablespaces. Do not specify an algorithm, and do not encrypt the SYSTEM tablespace concurrently with the encryption of other tablespaces.

For example, to encrypt the SYSTEM tablespace:

```
ALTER TABLESPACE SYSTEM ENCRYPTION OFFLINE ENCRYPT;
```

5. Open the database.

For example, to open the database in read/write mode:

```
ALTER DATABASE OPEN READ WRITE;
```

6. For a temporary tablespace, drop it and then recreate it as encrypted. Do not specify an algorithm.

For example, for a user-created tablespace:

```
DROP TABLESPACE temp_01;
CREATE TEMPORARY TABLESPACE temp_01
TEMPFILE 'temp01.dbf' SIZE 5M AUTOEXTEND ON
ENCRYPTION ENCRYPT;
```

You cannot drop the default TEMP tablespace. You must first create a new tablespace and make it the default before you can drop TEMP.

For example:

```
CREATE TEMPORARY TABLESPACE temp_01
TEMPFILE 'temp01.dbf' SIZE 5M AUTOEXTEND ON
ENCRYPTION ENCRYPT;

ALTER DATABASE DEFAULT TEMPORARY TABLESPACE temp_01;

DROP TABLESPACE TEMP;
```

7. Run the ALTER TABLESPACE SQL statement to encrypt other user tablespaces. Alternatively, you can proceed to the next step and open the database first, and then perform the steps described in [Encrypting an Existing User-Defined Tablespace with Offline Conversion](#) (page 3-39).
8. Open the database.

```
ALTER DATABASE OPEN;
```

See [Rotating the TDE Master Encryption Key for a Tablespace](#) (page 4-39) if you want to change the encryption algorithm of the tablespace.

3.4.7.3 Encrypting an Existing Database with Online Conversion

When you encrypt an existing database with online conversion, you do not specify an encryption algorithm.

The reason that you do not need to specify an encryption algorithm the first time you perform the encryption is that the tablespaces that you must use to encrypt the database are automatically encrypted with the database key. If you want to change the algorithm, then you can issue the ALTER TABLESPACE ENCRYPTION REKEY SQL statement after the initial encryption.

1. Perform the following tasks, which are described in [Encrypting an Existing Tablespace with Online Conversion](#) (page 3-41):
 - a. Connect as a user who has been granted the SYSDBA administrative privilege.
 - b. Ensure that the COMPATIBLE parameter is set to 12.2.0.0.
 - c. Ensure that the database is open in read-write mode.
 - d. Ensure that you have enough auxiliary space to complete the encryption.
 - e. Back up the tablespaces that you must encrypt.
 - f. Open the keystore.
2. Run the ALTER TABLESPACE SQL statement to encrypt the SYSTEM, SYSAUX, and UNDO tablespaces. Do not specify an algorithm, and do not encrypt the SYSTEM tablespace concurrently with the encryption of other tablespaces.

For example, to encrypt the SYSTEM tablespace:

```
ALTER TABLESPACE SYSTEM ENCRYPTION ONLINE ENCRYPT
FILE_NAME_CONVERT=('system01.dbf', 'system01_enc.dbf');
```

3. For a temporary tablespace, drop it and then recreate it as encrypted. Do not specify an algorithm.

For example, for a user-created tablespace:

```
DROP TABLESPACE temp_01;
CREATE TEMPORARY TABLESPACE temp_01
TEMPFILE 'temp01.dbf' SIZE 5M AUTOEXTEND ON
ENCRYPTION ENCRYPT;
```

You cannot drop the default TEMP tablespace. You must first create a new tablespace and make it the default before you can drop TEMP.

For example:

```
CREATE TEMPORARY TABLESPACE temp_01
TEMPFILE 'temp01.dbf' SIZE 5M AUTOEXTEND ON
ENCRYPTION ENCRYPT;

ALTER DATABASE DEFAULT TEMPORARY TABLESPACE temp_01;

DROP TABLESPACE TEMP;
```

Related Topics:

[Rotating the TDE Master Encryption Key for a Tablespace](#) (page 4-39)

You can use the REKEY clause of the ALTER TABLESPACE statement to rotate an encryption key for an encrypted tablespace.

3.5 Transparent Data Encryption Data Dynamic and Data Dictionary Views

You can query a set of dynamic and data dictionary views to find more information about Transparent Data Encryption (TDE) data.

[Table 3-3](#) (page 3-51) describes these dynamic and data dictionary views.

Table 3-3 Transparent Data Encryption Related Views

View	Description
ALL_ENCRYPTED_COLUMNS	Displays encryption information about encrypted columns in the tables accessible to the current user
DBA_ENCRYPTED_COLUMNS	Displays encryption information for all of the encrypted columns in the database
USER_ENCRYPTED_COLUMNS	Displays encryption information for encrypted table columns in the current user's schema
DBA_TABLESPACE_USAGE_METRICS	Describes tablespace usage metrics for all types of tablespaces, including permanent, temporary, and undo tablespaces

Table 3-3 (Cont.) Transparent Data Encryption Related Views

View	Description
V\$CLIENT_SECRETS	Lists the properties of the strings (secrets) that were stored in the keystore for various features (clients). In a multitenant environment, when you query this view in a PDB, then it displays information about keys that were created or activated for the current PDB. If you query this view in the root, then it displays this information about keys for all of the PDBs.
V\$DATABASE_KEY_INFO	Displays information about the default encryption key that is used for the current database. The default is AES128.
V\$ENCRYPTED_TABLESPACES	Displays information about the tablespaces that are encrypted
V\$ENCRYPTION_KEYS	When used with keys that have been rotated with the ADMINISTER KEY MANAGEMENT statement, displays information about the TDE master encryption keys. In a multitenant environment, when you query this view in a PDB, it displays information about keys that were created or activated for the current PDB. If you query this view in the root, it displays this information about keys for all of the PDBs.
V\$ENCRYPTION_WALLET	Displays information on the status of the keystore and the keystore location for TDE
V\$WALLET	Displays metadata information for a PKI certificate, which can be used as a master encryption key for TDE

See Also:

Oracle Database Reference for detailed information about these views

Managing the Keystore and the TDE Master Encryption Key

You can modify settings for the keystore and TDE master encryption key, and store Oracle Database and store Oracle GoldenGate secrets in a keystore.

[Managing the Keystore](#) (page 4-1)

You can perform maintenance activities on keystores such as changing passwords, and backing up, merging, and moving keystores.

[Managing the TDE Master Encryption Key](#) (page 4-25)

You can manage the TDE master encryption key in several ways.

[Storing Secrets Used by Oracle Database](#) (page 4-45)

Secrets are data that support internal Oracle Database features that integrate external clients such as Oracle GoldenGate into the database.

[Storing Oracle GoldenGate Secrets in a Keystore](#) (page 4-53)

You can store Oracle GoldenGate secrets in Transparent Data Encryption keystores.

4.1 Managing the Keystore

You can perform maintenance activities on keystores such as changing passwords, and backing up, merging, and moving keystores.

[Performing Operations That Require a Keystore Password](#) (page 4-2)

Many `ADMINISTER KEY MANAGEMENT` operations require access to a keystore password, for both software and hardware keystores.

[Changing the Password of a Software Keystore](#) (page 4-3)

Oracle Database enables you to easily change password-based software keystore passwords.

[Changing the Password of a Hardware Keystore](#) (page 4-4)

To change the password of a hardware keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement.

[Backing Up Password-Based Software Keystores](#) (page 4-5)

When you back up a password-based software keystore, you can create a backup identifier string to describe the backup type.

[Backups of the Hardware Keystore](#) (page 4-7)

You cannot use Oracle Database to back up hardware keystores.

[Merging Software Keystores](#) (page 4-8)

You can merge software keystores in a variety of ways.

[Moving a Software Keystore to a New Location](#) (page 4-11)

To move a software keystore to a new location, you must back up and close the keystore, edit `sqlnet.ora`, and then physically move the keystore to the new location.

[Moving a Software Keystore Out of Automatic Storage Management](#) (page 4-12)

You can use the `ADMINISTER KEY MANAGEMENT` statement to move a software keystore out Automatic Storage Management.

[Migrating Between a Software Password Keystore and a Hardware Keystore](#) (page 4-13)

You can migrate between password-based software keystores and hardware keystores.

[Migration of Keystores to and from Oracle Key Vault](#) (page 4-20)

You can use Oracle Key Vault to migrate both software and hardware keystores to and from Oracle Key Vault.

[Closing a Keystore](#) (page 4-20)

You can manually close software and hardware keystores.

[Using a Software Keystore That Resides on ASM Volumes](#) (page 4-23)

You can store a software keystore on an Automatic Storage Management (ASM) disk group.

[Backup and Recovery of Encrypted Data](#) (page 4-23)

For software keystores, you cannot access encrypted data without the TDE master encryption key.

[Deletion of Keystores](#) (page 4-24)

Oracle strongly recommends that you do not delete keystores, particularly after you have configured Transparent Data Encryption and the keystore is in use.

4.1.1 Performing Operations That Require a Keystore Password

Many `ADMINISTER KEY MANAGEMENT` operations require access to a keystore password, for both software and hardware keystores.

In some cases, a software keystore depends on an auto-login keystore before the operation can succeed. The auto-login keystore must be closed and the password-based keystore must be opened before the password can be accessed. Auto-login keystores open automatically when they are configured and a key is requested. They are generally used for operations where the keystore could be closed but a database operation needs a key (for example, after the database is restarted). Because the auto-login keystore opens automatically, it can be retrieved to perform a database operation without manual intervention. However, some keystore operations that require the keystore password cannot be performed when the auto-login keystore is open. The auto-login keystore must be closed and the password-based keystore must be opened for the keystore operations that require a password.

In a multitenant environment, the re-opening of keystores affects other PDBs. For example, an auto-login keystore in the root must be accessible by the PDBs in the CDB for this root.

You can temporarily open the keystore by including the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement when you perform the following operations: rotating a keystore password; creating, using, rekeying, tagging,

importing, exporting, migrating, or reverse migrating encryption keys; opening or backing up keystores; adding, updating, or deleting secret keystores. In a multitenant environment, if no keystore is open in the root, then `FORCE KEYSTORE` opens the password-based keystore in the root. After the operation completes, the keystore returns to its previous state.

4.1.2 Changing the Password of a Software Keystore

Oracle Database enables you to easily change password-based software keystore passwords.

[About Changing the Password of a Password-Based Software Keystore](#) (page 4-3)

You can only change (rotate) the password for password-based software keystores.

[Changing the Password-Based Software Keystore Password](#) (page 4-3)

To change the password of a password-based software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement.

4.1.2.1 About Changing the Password of a Password-Based Software Keystore

You can only change (rotate) the password for password-based software keystores.

You can change this password at any time, as per the security policies, compliance guidelines, and other security requirements of your site. As part of the command to change the password, you will be forced to specify the `WITH BACKUP` clause, and thus forced to make a backup of the current keystore. During the password change operation, Transparent Data Encryption operations such as encryption and decryption will continue to work normally.

You can change this password at any time. You may want to change this password if you think it was compromised.

4.1.2.2 Changing the Password-Based Software Keystore Password

To change the password of a password-based software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD [FORCE KEYSTORE] IDENTIFIED BY
old_password SET new_password [WITH BACKUP [USING 'backup_identifier']];
```

In this specification:

- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `old_password` is the current keystore password that you want to change.
- `new_password` is the new password that you set for the keystore.

- `WITH BACKUP` creates a backup of the current keystore before the password is changed. You must include this clause.
- `backup_identifier` specifies an optional identifier string for the backup that is created. The `backup_identifier` is added to the name of the backup file. Enclose `backup_identifier` in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time_stamp_emp_key_pwd_change.p12`).

The following example backs up the current keystore and then changes the password for the keystore:

```
ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD IDENTIFIED BY
old_password SET new_password WITH BACKUP USING 'pwd_change';
```

keystore altered.

This example performs the same operation but uses the `FORCE KEYSTORE` clause in case the auto-login keystore is in use or the keystore is closed.

```
ADMINISTER KEY MANAGEMENT ALTER KEYSTORE FORCE KEYSTORE PASSWORD IDENTIFIED BY
old_password SET new_password WITH BACKUP USING 'pwd_change';
```

keystore altered.

4.1.3 Changing the Password of a Hardware Keystore

To change the password of a hardware keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Close the hardware keystore.

Examples are as follows:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY "psmith:password";
```

For a keystore whose password is stored externally:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY EXTERNAL STORE;
```

3. From the hardware security module management interface, create a new hardware security module password.

4. In SQL*Plus, open the hardware keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY "psmith:new_password";
```

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY EXTERNAL STORE;
```

Related Topics:[Closing a Hardware Keystore](#) (page 4-22)

To close a hardware keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE CLOSE` clause.

[Opening a Hardware Keystore](#) (page 3-16)

To open a hardware keystore, use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

4.1.4 Backing Up Password-Based Software Keystores

When you back up a password-based software keystore, you can create a backup identifier string to describe the backup type.

[About Backing Up Password-Based Software Keystores](#) (page 4-5)

You must back up password-based software keystores, as per the security policy and requirements of your site.

[Creating a Backup Identifier String for the Backup Keystore](#) (page 4-6)

The backup file name of a software password keystore is derived from the name of the password-based software keystore.

[How the V\\$ENCRYPTION_WALLET View Interprets Backup Operations](#) (page 4-6)

The `BACKUP` column of the `V$ENCRYPTION_WALLET` view indicates a how a copy of the keystore was created.

[Backing Up a Password-Based Software Keystore](#) (page 4-6)

The `BACKUP KEYSTORE` clause of the `ADMINISTER KEY MANAGEMENT` statement backs up a password-based software keystore.

4.1.4.1 About Backing Up Password-Based Software Keystores

You must back up password-based software keystores, as per the security policy and requirements of your site.

A backup of the keystore contains all of the keys contained in the original keystore. Oracle Database prefixes the backup keystore with the creation time stamp (UTC). If you provide an identifier string, then this string is inserted between the time stamp and keystore name.

After you complete the backup operation, the keys in the original keystore are marked as "backed up". You can check the status of keys querying the `V$ENCRYPTION_WALLET` data dictionary view.

You cannot back up auto-login or local auto-login software keystores. No new keys can be added to them directly through the `ADMINISTER KEY MANAGEMENT` statement operations. The information in these keystores is only read and hence there is no need for a backup.

If you have not yet backed up the keystore, then you can include the `BACKUP` clause in the `ADMINISTER KEY MANAGEMENT` statement when you create the TDE master encryption key. This both backs up the keystore and creates the TDE master encryption key. ([Step 4: Set the Software TDE Master Encryption Key](#) (page 3-11) shows an example of how to accomplish this.)

4.1.4.2 Creating a Backup Identifier String for the Backup Keystore

The backup file name of a software password keystore is derived from the name of the password-based software keystore.

Oracle Database prefixes the software keystore password file name with the file creation time stamp in UTC format. If you provide an identifier string, then this string is inserted between the time stamp and keystore name.

- To create a backup identifier string for a backup keystore, use the `ADMINISTER KEY MANAGEMENT SQL` statement with the `BACKUP KEYSTORE` clause, with the following syntax:

```
ewallet_creation-time-stamp-in-UTC_user-defined-string.p12
```

When you create the backup identifier (*user_defined_string*), use the operating system file naming convention. For example, in UNIX systems, you may want to ensure that this setting does not have spaces.

Example 4-1 (page 4-6) shows the creation of a backup keystore that uses a bug number as the user-identified string, and how the resultant keystore appears in the file system. This example includes the `FORCE KEYSTORE` clause in the event the auto-login keystore is in use or the keystore is closed

Example 4-1 Creating a Backup Identifier String for a Backup Keystore

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE USING 'BUG1296' FORCE KEYSTORE IDENTIFIED
BY keystore_password;
```

This version is for a scenario in which the password is stored in an external store:

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE USING 'BUG1296' FORCE KEYSTORE IDENTIFIED
BY EXTERNAL STORE;
```

Resultant keystore file:

```
ewallet_2013041513244657_BUG1296.p12
```

4.1.4.3 How the `V$ENCRYPTION_WALLET` View Interprets Backup Operations

The `BACKUP` column of the `V$ENCRYPTION_WALLET` view indicates how a copy of the keystore was created.

The column indicates if a copy of the keystore had been created with the `WITH BACKUP` clause of the `ADMINISTER KEY MANAGEMENT` statement or the `ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE` statement.

When you modify a key or a secret, the modifications that you make do not exist in the previously backed-up copy, because you make a copy and then modify the key itself. Because there is no copy of the modification in the previous keystores, the `BACKUP` column is set to `NO`, even if the `BACKUP` had been set to `YES` previously. Hence, if the `BACKUP` column is `YES`, then after you perform an operation that requires a backup, such as adding a custom attribute tag, the `BACKUP` column value changes to `NO`.

4.1.4.4 Backing Up a Password-Based Software Keystore

The `BACKUP KEYSTORE` clause of the `ADMINISTER KEY MANAGEMENT` statement backs up a password-based software keystore.

- Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE [USING 'backup_identifier']
[FORCE KEYSTORE] IDENTIFIED BY [EXTERNAL STORE | software_keystore_password]
[TO 'keystore_location'];
```

In this specification:

- USING *backup_identifier* is an optional string that you can provide to identify the backup. Enclose this identifier in single quotation marks (' '). This identifier is appended to the named keystore file (for example, *ewallet_time-stamp_emp_key_backup.p12*).
- FORCE KEYSTORE enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- IDENTIFIED BY can be one of the following settings:
 - EXTERNAL STORE uses the keystore password stored in the external store to perform the keystore operation.
 - *software_keystore_password* is the password for the keystore.
- *keystore_location* is the path at which the backup keystore is stored. If you do not specify the *keystore_location*, then the backup is created in the same directory as the original keystore. Enclose this location in single quotation marks (' ').

The following example backs up a software keystore in the same location as the source keystore. FORCE KEYSTORE is used in case the auto-login keystore is in use or the keystore is closed.

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE USING 'hr.emp_keystore'
FORCE KEYSTORE IDENTIFIED BY
software_keystore_password TO '/etc/ORACLE/KEYSTORE/DB1/';
```

keystore altered.

In this version, the password for the keystore is external, so the EXTERNAL STORE clause is used.

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE USING 'hr.emp_keystore'
FORCE KEYSTORE IDENTIFIED BY EXTERNAL STORE;
```

After you run this statement, an *ewallet_identifier.p12* file (for example, *ewallet_time-stamp_hr.emp_keystore.p12*) appears in the keystore location.

4.1.5 Backups of the Hardware Keystore

You cannot use Oracle Database to back up hardware keystores.

See your HSM vendor instructions for information about backing up keys for hardware keystores.

4.1.6 Merging Software Keystores

You can merge software keystores in a variety of ways.

[About Merging Software Keystores](#) (page 4-8)

You can merge any combination of software keystores, but the merged keystore must be password-based. It can have a password that is different from the constituent keystores.

[Merging Two Software Keystores into a Third New Keystore](#) (page 4-9)

You can merge two software keystores into a third new keystore, so that the two existing keystores are not changed.

[Merging One Software Keystore into an Existing Software Keystore](#) (page 4-10)

You can use the `ADMINISTER KEY MANAGEMENT` statement with the `MERGE KEYSTORE` clause to merge one software keystore into another existing software keystore.

[Merging an Auto-Login Software Keystore into an Existing Password-Based Software Keystore](#) (page 4-10)

You can merge an auto-login software keystore into an existing password-based software keystore.

[Reversing a Software Keystore Merge Operation](#) (page 4-11)

You cannot directly reverse a keystore merge operation.

4.1.6.1 About Merging Software Keystores

You can merge any combination of software keystores, but the merged keystore must be password-based. It can have a password that is different from the constituent keystores.

To use the merged keystore, you must explicitly open the merged keystore after you create it, even if one of the constituent keystores was already open before the merge.

Whether a common key from two source keystores is added or overwritten to a merged keystore depends on how you write the `ADMINISTER KEY MANAGEMENT` merge statement. For example, if you merge Keystore 1 and Keystore 2 to create Keystore 3, then the key in Keystore 1 is added to Keystore 3. If you merge Keystore 1 into Keystore 2, then the common key in Keystore 2 is not overwritten.

The `ADMINISTER KEY MANAGEMENT` merge statement has no bearing on the configured keystore that is in use. However, the merged keystore can be used as the new configured database keystore if you want. Remember that you must reopen the keystore if you are using the newly created keystore as the keystore for the database at the location configured by the `sqlnet.ora` file.

Related Topics:

[Migrating Between a Software Password Keystore and a Hardware Keystore](#) (page 4-13)

You can migrate between password-based software keystores and hardware keystores.

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

4.1.6.2 Merging Two Software Keystores into a Third New Keystore

You can merge two software keystores into a third new keystore, so that the two existing keystores are not changed.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE 'keystore1_location'
[IDENTIFIED BY software_keystore1_password] AND KEYSTORE 'keystore2_location'
[IDENTIFIED BY software_keystore2_password]
INTO NEW KEYSTORE 'keystore3_location'
IDENTIFIED BY software_keystore3_password;
```

In this specification:

- *keystore1_location* is the directory location of the first keystore, which will be left unchanged after the merge. Enclose this path in single quotation marks (').
- The `IDENTIFIED BY` clause is required for the first keystore if it is a password-based keystore. *software_keystore1_password* is the current password for the first keystore.
- *keystore2_location* is the directory location of the second keystore. Enclose this path in single quotation marks (').
- The `IDENTIFIED BY` clause is required for the second keystore if it is a password-based keystore. *software_keystore2_password* is the current password for the second keystore.
- *keystore3_location* specifies the directory location of the new, merged keystore. Enclose this path in single quotation marks ('). If there is already an existing keystore at this location, the command exits with an error.
- *software_keystore3_password* is the new password for the merged keystore.

The following example merges an auto-login software keystore with a password-based keystore to create a merged password-based keystore at a new location:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE '/etc/ORACLE/KEYSTORE/DB1'
AND KEYSTORE '/etc/ORACLE/KEYSTORE/DB2'
IDENTIFIED BY existing_password_for_keystore_2
INTO NEW KEYSTORE '/etc/ORACLE/KEYSTORE/DB3'
IDENTIFIED BY new_password_for_keystore_3;
```

keystore altered.

4.1.6.3 Merging One Software Keystore into an Existing Software Keystore

You can use the `ADMINISTER KEY MANAGEMENT` statement with the `MERGE KEYSTORE` clause to merge one software keystore into another existing software keystore.

- To perform this type of merge, follow the steps in [Merging Two Software Keystores into a Third New Keystore](#) (page 4-9) but use the following SQL statement:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE 'keystore1_location'
[IDENTIFIED BY software_keystore1_password]
INTO EXISTING KEYSTORE 'keystore2_location'
IDENTIFIED BY software_keystore2_password
[WITH BACKUP [USING 'backup_identifier]];
```

In this specification:

- *keystore1_location* is the directory location of the first keystore, which will be left unchanged after the merge. Enclose this path in single quotation marks (' ').
- The `IDENTIFIED BY` clause is required for the first keystore if it is a password-based keystore. *software_keystore1_password* is the password for the first keystore.
- *keystore2_location* is the directory location of the second keystore into which the first keystore is to be merged. Enclose this path in single quotation marks (' ').
- *software_keystore2_password* is the password for the second keystore.
- `WITH BACKUP` creates a backup of the software keystore. Optionally, you can use the `USING` clause to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, *ewallet_time-stamp_emp_key_backup.p12*, with *emp_key_backup* being the backup identifier). Follow the file naming conventions that your operating system uses.

The resultant keystore after the merge operation is always a password-based keystore.

4.1.6.4 Merging an Auto-Login Software Keystore into an Existing Password-Based Software Keystore

You can merge an auto-login software keystore into an existing password-based software keystore.

- Use the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` SQL statement to merge an auto-login software keystore into an existing password-based software keystore.

[Example 4-2](#) (page 4-11) shows how to merge an auto-login software keystore into a password-based software keystore. It also creates a backup of the second keystore before creating the merged keystore.

Example 4-2 Merging a Software Auto-Login Keystore into a Password Keystore

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE '/etc/ORACLE/KEYSTORE/DB1'
INTO EXISTING KEYSTORE '/etc/ORACLE/KEYSTORE/DB2'
IDENTIFIED BY keystore_password WITH BACKUP;
```

In this specification:

- `MERGE KEYSTORE` must specify the auto-login keystore.
- `EXISTING KEYSTORE` refers to the password keystore.

4.1.6.5 Reversing a Software Keystore Merge Operation

You cannot directly reverse a keystore merge operation.

When you merge a keystore into an existing keystore (rather than creating a new one), you must include the `WITH BACKUP` clause in the `ADMINISTER KEY MANAGEMENT` statement to create a backup of this existing keystore. Later on, if you decide that you must reverse the merge, you can replace the merged software keystore with the one that you backed up.

In other words, suppose you want merge Keystore A into Keystore B. By using the `WITH BACKUP` clause, you create a backup for Keystore B before the merge operation begins. (The original Keystore A is still intact.) To reverse the merge operation, revert to the backup that you made of Keystore B.

- Use the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` SQL statement to perform merge operations.

- For example, to perform a merge operation into an existing keystore:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE '/etc/ORACLE/KEYSTORE/DB1'
INTO EXISTING KEYSTORE '/etc/ORACLE/KEYSTORE/DB2'
IDENTIFIED BY password WITH BACKUP USING "merge1";
```

Replace the new keystore with the backup keystore, which in this case would be named `ewallet_time-stamp_merge1.p12`.

- To merge an auto-login keystore into a password-based keystore, use the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` SQL statement.

4.1.7 Moving a Software Keystore to a New Location

To move a software keystore to a new location, you must back up and close the keystore, edit `sqlnet.ora`, and then physically move the keystore to the new location.

If you are using Oracle Key Vault, then you can configure a TDE direct connection where Key Vault directly manages the TDE master keys. In this case, you will never need to manually move the keystore to a new location. See *Oracle Key Vault Administrator's Guide* for more information about using a TDE direct connection.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or to the pluggable database (PDB). For example, to log in to a PDB called `hrpdb`:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Make a backup copy of the software keystore.

For example:

```
ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE USING 'hr.emp_keystore'
FORCE KEYSTORE IDENTIFIED BY
software_keystore_password TO '/etc/ORACLE/KEYSTORE/DB1/';
```

See [Backing Up Password-Based Software Keystores](#) (page 4-5).

3. Close the software keystore.

Examples of ways that you can close the keystore are as follows.

For an auto-login software keystore:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE;
```

For a password-based software keystore:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY
software_keystore_password;
```

For a keystore in which the password is stored externally:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY EXTERNAL STORE;
```

4. Exit the database session.

For example, if you are logged in to SQL*Plus:

```
EXIT
```

5. Back up and then manually edit the `sqlnet.ora` file to point to the new location where you want to move the keystore.

See the [Step 1: Set the Keystore Location in the sqlnet.ora File](#) (page 3-2) for more information.

6. Use the operating system move command (such as `mv`) to move the keystore with all of its keys to the new directory location.

4.1.8 Moving a Software Keystore Out of Automatic Storage Management

You can use the `ADMINISTER KEY MANAGEMENT` statement to move a software keystore out Automatic Storage Management.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Initialize a target keystore on the file system by using the following syntax:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE targetKeystorePath IDENTIFIED BY
targetKeystorePassword;
```

In this specification:

- *targetKeystorePath* is the directory path to the target keystore on the file system.
- *targetKeystorePassword* is a password that you create for the keystore.

For example:

```
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE '/etc/ORACLE/KEYSTORE/DB1/' IDENTIFIED
BY "targetKeystorePassword";
```

3. Copy the keystore from ASM to the target keystore that you just created.

This step requires that you merge the keystore from ASM to the file system, as follows:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE srcKeystorePath IDENTIFIED BY
srcKeystorePassword INTO EXISTING KEYSTORE targetKeystorePath IDENTIFIED BY
targetKeystorePassword WITH BACKUP USING backupIdentifier;
```

In this specification:

- *srcKeystorePath* is the directory path to the source keystore.
- *srcKeystorePassword* is the source keystore password.
- *targetKeystorePath* is the path to the target keystore.
- *targetKeystorePassword* is the target keystore password.
- *backupIdentifier* is the backup identifier to be added to the backup file name.

For example:

```
ADMINISTER KEY MANAGEMENT MERGE KEYSTORE '+DATAFILE' IDENTIFIED BY "srcPassword"
INTO EXISTING KEYSTORE '/etc/ORACLE/KEYSTORE/DB1/' IDENTIFIED BY
"targetKeystorePassword" WITH BACKUP USING "bkup";
```

4.1.9 Migrating Between a Software Password Keystore and a Hardware Keystore

You can migrate between password-based software keystores and hardware keystores.

[Migrating from a Password-Based Software Keystore to a Hardware Keystore](#)
(page 4-14)

You can migrate from a password-based software keystore to a hardware keystore.

[Migrating from a Hardware Keystore to a Password-Based Software Keystore](#)
(page 4-16)

You can migrate a hardware keystore to a software keystore.

[Keystore Order After a Migration](#) (page 4-19)

After you perform a migration, keystores can be either primary or secondary in their order.

4.1.9.1 Migrating from a Password-Based Software Keystore to a Hardware Keystore

You can migrate from a password-based software keystore to a hardware keystore.

[Step 1: Convert the Software Keystore to Open with the Hardware Keystore](#) (page 4-14)

Some Oracle tools require access to the old software keystore to encrypt or decrypt data that was exported or backed up using the software keystore.

[Step 2: Configure sqlnet.ora for the Migration of the Password-Based Software Keystore](#) (page 4-15)

You must edit `sqlnet.ora` to enable access to keys created in the software keystore when required.

[Step 3: Perform the Hardware Keystore Migration](#) (page 4-16)

You can use the `ADMINISTER KEY MANAGEMENT SQL` statement to perform a hardware keystore migration.

4.1.9.1.1 Step 1: Convert the Software Keystore to Open with the Hardware Keystore

Some Oracle tools require access to the old software keystore to encrypt or decrypt data that was exported or backed up using the software keystore.

Examples of these tools are Oracle Data Pump and Oracle Recovery Manager.

- Use the `ADMINISTER KEY MANAGEMENT SQL` statement to convert a software keystore to a open with a hardware keystore.
 - To set the software keystore password as that of the hardware keystore, use the following syntax:

```
ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD
FORCE KEYSTORE
IDENTIFIED BY software_keystore_password
SET "user_id:password" WITH BACKUP
[USING 'backup_identifier'];
```

In this specification:

- * *software_keystore_password* is the same password that you used when creating the software keystore.
 - * *user_id:password* is the new software keystore password which is the same as the password of the HSM.
 - * `WITH BACKUP` creates a backup of the software keystore. Optionally, you can use the `USING` clause to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time-stamp_emp_key_backup.p12`, with *emp_key_backup* being the backup identifier). Follow the file naming conventions that your operating system uses.
- To create an auto-login keystore for a software keystore, use the following syntax:

```
ADMINISTER KEY MANAGEMENT CREATE [LOCAL] AUTO_LOGIN KEYSTORE
FROM KEYSTORE 'keystore_location'
IDENTIFIED BY software_keystore_password;
```

In this specification:

- * LOCAL enables you to create a local auto-login software keystore. Otherwise, omit this clause if you want the keystore to be accessible by other computers.
- * *keystore_location* is the path to the keystore directory location of the keystore that is configured in the `sqlnet.ora` file.
- * *software_keystore_password* is the existing password of the configured software keystore.

4.1.9.1.2 Step 2: Configure `sqlnet.ora` for the Migration of the Password-Based Software Keystore

You must edit `sqlnet.ora` to enable access to keys created in the software keystore when required.

For the software keystore to open with the hardware keystore, either the software keystore must have the same password as the hardware keystore, or alternatively, you can create an auto-login keystore for the software keystore.

If you are migrating from a software keystore to a hardware keystore, then you must edit the `sqlnet.ora` file to use the `METHOD=HSM` setting.

- Use the following format in the `sqlnet.ora` file:

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=(METHOD=HSM)(METHOD_DATA=
(DIRECTORY=path_to_keystore)))
```

path_to_software_keystore is the path to the previously configured software keystore. Having both HSM and the `DIRECTORY` location in the `ENCRYPTION_WALLET_LOCATION` parameter indicates that you switched between using the software keystore and the hardware keystore in the past, and it also enables you to switch back easily in the future.

Note:

If a `DIRECTORY` value is present in the `ENCRYPTION_WALLET_LOCATION` parameter setting, then ensure that you do not delete it.

Although hardware keystores do not require a `DIRECTORY` value, Oracle Database uses this value to locate your software keystore when you migrate to and from a hardware security module.

[Example 4-3](#) (page 4-15) shows how to edit the `sqlnet.ora` file to format a software keystore to hardware security module-based keystore or the reverse:

Example 4-3 Sample `ENCRYPTION_WALLET_LOCATION` Entries

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=(METHOD=HSM)(METHOD_DATA=
(DIRECTORY=/app/wallet)))
```

Related Topics:

[About the Keystore Location in the sqlnet.ora File](#) (page 3-3)

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

4.1.9.1.3 Step 3: Perform the Hardware Keystore Migration

You can use the `ADMINISTER KEY MANAGEMENT` SQL statement to perform a hardware keystore migration.

To migrate from the software keystore to hardware keystore, you must use the `MIGRATE USING keystore_password` clause in the `ADMINISTER KEY MANAGEMENT SET KEY` SQL statement to decrypt the existing [TDE table keys](#) and the [tablespace encryption keys](#) with the TDE master encryption key in the software keystore and then reencrypt them with the newly created TDE master encryption key in the hardware keystore.

After you complete the migration, you do not need to restart the database, nor do you need to manually re-open the hardware keystore. The migration process automatically reloads the keystore keys in memory.

- Use the following syntax when you run the `ADMINISTER KEY MANAGEMENT` SQL statement for migration:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY IDENTIFIED BY "user_id:password"  
MIGRATE USING software_keystore_password [WITH BACKUP [USING  
'backup_identifier']];
```

In this specification:

- `user_id:password` is the user ID and password that was created in [Step 3](#) (page 3-15) under [Step 2: Configure the Hardware Security Module](#) (page 3-14) (in [Configuring Transparent Data Encryption](#) (page 3-1)). Enclose this setting in double quotation marks (" ") and separate `user_id` and `password` with a colon (:).
- `software_keystore_password` is the same password that you used when creating the software keystore or that you have changed to in [Step 1: Convert the Software Keystore to Open with the Hardware Keystore](#) (page 4-14).
- `USING` enables you to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time-stamp_emp_key_backup.p12`, with `emp_key_backup` being the backup identifier). Follow the file naming conventions that your operating system uses.

Note:

If the database contains columns encrypted with a public key, then the columns are decrypted and reencrypted with an AES symmetric key generated by HSM-based Transparent Data Encryption.

4.1.9.2 Migrating from a Hardware Keystore to a Password-Based Software Keystore

You can migrate a hardware keystore to a software keystore.

[About Migrating Back from a Hardware Keystore](#) (page 4-17)

To switch from using a hardware keystore solution to a software keystore, you can use reverse migration of the keystore.

[Step 1: Configure sqlnet.ora for the Reverse Migration](#) (page 4-17)

You must edit the keystore directory location in the `sqlnet.ora` file.

[Step 2: Configure the Keystore for the Reverse for the Reverse Migration](#) (page 4-18)

The `ADMINISTER KEY MANAGEMENT` statement with the `SET ENCRYPTION KEY` and `REVERSE MIGRATE` clauses can be used to reverse the migration of a keystore.

[Step 3: Configure the Hardware Keystore to Open with the Software Keystore](#) (page 4-19)

After you complete the migration, the migration process automatically reloads the keystore keys in memory.

4.1.9.2.1 About Migrating Back from a Hardware Keystore

To switch from using a hardware keystore solution to a software keystore, you can use reverse migration of the keystore.

After you complete the switch, keep the hardware security module, in case earlier backup files rely on the TDE master encryption keys in the hardware security module.

If you had originally migrated from the software keystore to the hardware security module and reconfigured the software keystore as described in [Migration of a Previously Configured TDE Master Encryption Key](#) (page 3-19), then you already have an existing keystore with the same password as the HSM password. Reverse migration configures this keystore to act as the new software keystore with a new password. If your existing keystore is an auto-login software keystore and you have the password-based software keystore for this auto-login keystore, then use the password-based keystore. If the password-based keystore is not available, then merge the auto-login keystore into a newly created empty password-based keystore, and use the newly create password-based keystore.

If you do not have an existing keystore, then you must specify a keystore location in the `sqlnet.ora` file using the `ENCRYPTION_WALLET_LOCATION` parameter. When you perform the reverse migration, migrate to the previous keystore so that you do not lose the keys.

Related Topics:[Merging Software Keystores](#) (page 4-8)

You can merge software keystores in a variety of ways.

4.1.9.2.2 Step 1: Configure sqlnet.ora for the Reverse Migration

You must edit the keystore directory location in the `sqlnet.ora` file.

- Set the following configuration in the `sqlnet.ora` file:

```
ENCRYPTION_WALLET_LOCATION=
  (SOURCE=(METHOD=FILE)(METHOD_DATA=
    (DIRECTORY=path_to_keystore)))
```

Replace *path_to_keystore* with the directory location of the destination keystore.

Related Topics:

[About the Keystore Location in the sqlnet.ora File](#) (page 3-3)

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

4.1.9.2.3 Step 2: Configure the Keystore for the Reverse for the Reverse Migration

The `ADMINISTER KEY MANAGEMENT` statement with the `SET ENCRYPTION KEY` and `REVERSE MIGRATE` clauses can be used to reverse the migration of a keystore.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY IDENTIFIED BY
software_keystore_password REVERSE MIGRATE USING "user_id:password" [WITH BACKUP
[USING 'backup_identifier']];
```

In this specification:

- `software_keystore_password` is the password for the existing keystore or the new keystore.
- `user_id:password` is the user ID and password that was created in [Step 3](#) (page 3-15) in [Step 2: Configure the Hardware Security Module](#) (page 3-14) (in [Configuring Transparent Data Encryption](#) (page 3-1)). If the pre-hardware security module software keystore is the new keystore, then you must ensure that it has the same password as the `user_id:password` before issuing the reverse migration command. Enclose this setting in double quotation marks ("").
- `WITH BACKUP` creates a backup of the software keystore. Optionally, you can include the `USING` clause to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time-stamp_emp_key_backup.p12`, with `emp_key_backup` being the backup identifier). Follow the file naming conventions that your operating system uses.

For example:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY IDENTIFIED BY password REVERSE
MIGRATE USING "psmith:password" WITH BACKUP;
```

keystore altered.

3. Optionally, change the keystore password.

For example:

```
ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD IDENTIFIED BY
old_password SET new_password WITH BACKUP USING 'pwd_change';
```

Related Topics:[Changing the Password of a Software Keystore](#) (page 4-3)

Oracle Database enables you to easily change password-based software keystore passwords.

4.1.9.2.4 Step 3: Configure the Hardware Keystore to Open with the Software Keystore

After you complete the migration, the migration process automatically reloads the keystore keys in memory.

You do not need to restart the database, nor do you need to manually re-open the software keystore.

The hardware keystore may still be required after reverse migration because the old keys are likely to have been used for encrypted backups or by tools such as Oracle Data Pump and Oracle Recovery Manager. You should cache the hardware keystore credentials in the keystore so that the HSM can be opened with the software keystore.

Related Topics:[Configuring Auto-Login Hardware Security Modules](#) (page 4-52)

A hardware security module can be configured to use the auto-login capability.

4.1.9.3 Keystore Order After a Migration

After you perform a migration, keystores can be either primary or secondary in their order.

The `WALLET_ORDER` column of the `V$ENCRYPTION_WALLET` dynamic view describes whether a keystore is primary (that is, it holds the current TDE master encryption key) or if it is secondary (it holds the previous TDE master encryption key). The `WRL_TYPE` column describes the type of locator for the keystore (for example, `FILE` for the `sqlnet.ora` file). The `WALLET_ORDER` column shows `SINGLE` if two keystores are not configured together and no migration was ever performed previously.

[Table 4-1](#) (page 4-19) describes how the keystore order works after you perform a migration.

Table 4-1 Keystore Order After a Migration

Type of Migration Done	WRL_TYPE	WALLET_ORDER	Description
Migration of software keystore to HSM	HSM	PRIMARY	Both the HSM and software keystore are configured. The TDE master encryption key can be either in the HSM or the software keystore. The TDE master encryption key is first searched in the HSM. If the TDE master encryption key is not in the primary keystore (HSM), then it will be searched for in the software keystore. All of the new TDE master encryption keys will be created in the primary keystore (in this case, the HSM).
	FILE	SECONDARY	

Table 4-1 (Cont.) Keystore Order After a Migration

Type of Migration Done	WRL_TYPE	WALLET_ORDER	Description
Reverse migration of HSM to software keystore	FILE HSM	PRIMARY SECONDARY	<p>Both the HSM and software keystore are configured. The TDE master encryption key can be either in the HSM or the software keystore.</p> <p>The TDE master encryption key is first searched for in the software keystore.</p> <p>If the TDE master encryption key is not present in the primary (that is, software) keystore, then it will be searched for in the HSM.</p> <p>All of the new TDE master encryption keys will be created in the primary keystore (in this case, the software keystore).</p>

4.1.10 Migration of Keystores to and from Oracle Key Vault

You can use Oracle Key Vault to migrate both software and hardware keystores to and from Oracle Key Vault.

This enables you to manage the keystores centrally, and then share the keystores as necessary with other TDE-enabled databases in your enterprise.

Oracle Key Vault enables you to upload a keystore to a container called a virtual wallet, and then create a new virtual wallet from the contents of previously uploaded Oracle keystores. For example, suppose you previously uploaded a keystore that contains 5 keys. You can create a new virtual wallet that consists of only 3 of these keys. You then can download this keystore to another TDE-enabled database. This process does not modify the original keystore.

In addition to Oracle keystores, Oracle Key Vault enables you to securely share other security objects, such as credential files and Java keystores, across the enterprise. It prevents the loss of keys and keystores due to forgotten passwords or accidentally deleted keystores. You can use Oracle Key Vault with products other than TDE: Oracle Real Application Security, Oracle Active Data Guard, and Oracle GoldenGate. Oracle Key Vault facilitates the movement of encrypted data using Oracle Data Pump and Oracle Transportable Tablespaces.

See Also:

Oracle Key Vault Administrator's Guide

4.1.11 Closing a Keystore

You can manually close software and hardware keystores.

[About Closing Keystores](#) (page 4-21)

After you open a keystore, it remains open until you shut down the database instance.

[Closing a Software Keystore](#) (page 4-21)

You can manually close password-based software keystores, auto-login software keystores, and local auto-login software keystores.

[Closing a Hardware Keystore](#) (page 4-22)

To close a hardware keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE CLOSE` clause.

4.1.11.1 About Closing Keystores

After you open a keystore, it remains open until you shut down the database instance.

When you restart the database instance, then auto-login and local auto-login software keystores automatically open when required (that is, when the TDE master encryption key must be accessed). However, software password-based and hardware keystores do not automatically open. You must manually open them again before you can use them.

When you close a software or hardware keystore, you disable all of the encryption and decryption operations on the database. Hence, a database user or application cannot perform any operation involving encrypted data until the keystore is reopened.

When you re-open a keystore after closing it, the keystore contents are reloaded back into the database. Thus, if the contents had been modified (such as during a migration), the database will have the latest keystore contents.

You can check if a keystore is closed by querying the `STATUS` column of the `V$ENCRYPTION_WALLET` view.

The following data operations will fail if the keystore is not accessible:

- `SELECT` data from an encrypted column
- `INSERT` data into on an encrypted column
- `CREATE` a table with encrypted columns
- `CREATE` an encrypted tablespace

Related Topics:

[How Keystore Open and Close Operations Work in a Multitenant Environment](#) (page 6-18)

You should be aware of how keystore open and close operations work in a multitenant environment.

4.1.11.2 Closing a Software Keystore

You can manually close password-based software keystores, auto-login software keystores, and local auto-login software keystores.

In the case of an auto-login keystore, which opens automatically when it is accessed, manually close it if you moved it to a new location. You do this if you are changing your configuration from an auto-login keystore to a password-based keystore: you move out the auto-login keystore, and then close the auto-login keystore.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, you must close the keystore first in the root. Afterward, all keystores in the PDBs will close as well. For example, to log in to the root:

```
sqlplus sec_admin as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Run the `ADMINISTER KEY MANAGEMENT SQL` statement.

- For a password-based software keystore, use the following syntax:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE [IDENTIFIED BY [ EXTERNAL
STORE | software_keystore_password]] [CONTAINER = ALL | CURRENT];
```

In this specification:

- `IDENTIFIED BY` can be one of the following:
 - * `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - * `software_keystore_password` is the password of the user who created the keystore.
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to close the keystore in all of the PDBs in this multitenant container database (CDB), or `CURRENT` for the current PDB. If you run this `ADMINISTER KEY MANAGEMENT` statement in the root, then all of the keystores on all of the PDBs will close, irrespective of whether `CONTAINER` is set to `ALL` or to `CURRENT`.

- For an auto-login or local auto-login software keystore, use the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE;
```

You do not need to specify a password for this statement.

Closing a keystore disables all of the encryption and decryption operations. Any attempt to encrypt or decrypt data or access encrypted data results in an error.

4.1.11.3 Closing a Hardware Keystore

To close a hardware keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE CLOSE` clause.

1. Log into the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY [EXTERNAL STORE |
"user_id:password"] [CONTAINER = ALL | CURRENT];
```

In this specification:

- IDENTIFIED BY can be one of the following settings:
 - EXTERNAL STORE uses the keystore password stored in the external store to perform the keystore operation.
 - `user_id:password: user_id` is the user ID and password that was created in Step 3 (page 3-15) in [Step 2: Configure the Hardware Security Module](#) (page 3-14). Enclose this setting in double quotation marks (" ") and separate `user_id` and `password` with a colon (:).
- CONTAINER is for use in a multitenant environment. Enter ALL to close the keystore in all of the PDBs in this CDB, or CURRENT for the current PDB. If you run this ADMINISTER KEY MANAGEMENT statement in the root, then all of the keystores on all of the PDBs will close, irrespective of whether CONTAINER is set to ALL or to CURRENT.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY "psmith:password";
```

4.1.12 Using a Software Keystore That Resides on ASM Volumes

You can store a software keystore on an Automatic Storage Management (ASM) disk group.

- Edit the `sqlnet.ora` file to use the location of an ASM disk group specified using the ASM file naming convention when you configure the DIRECTORY setting in the ENCRYPTION_WALLET_LOCATION setting. That is, you must use the plus sign (+) notation for the ASM file name.

For example:

```
ENCRYPTION_WALLET_LOCATION=
  (SOURCE=(METHOD=FILE)(METHOD_DATA=
    (DIRECTORY=+disk1/mydb/wallet)))
```

If you must move or merge software keystores between a regular file system and an ASM file system, then you can use the same keystore merge statements described in [Merging Software Keystores](#) (page 4-8).

To manage keystores in an ASM environment, you can use the ASMCMD utility.

See Also:

- [Configuring the sqlnet.ora File for a Software Keystore Location](#) (page 3-3)
 - *Oracle Database Storage Administrator's Guide* for detailed information about using the ASMCMD utility
-

4.1.13 Backup and Recovery of Encrypted Data

For software keystores, you cannot access encrypted data without the TDE master encryption key.

Because the TDE master encryption key is stored in the keystore, you should periodically back up the software keystore in a secure location. You must back up a

copy of the keystore whenever you set a new TDE master encryption key or perform any operation that writes to the keystore.

Do not back up the software keystore in the same location as the encrypted data. Back up the software keystore separately. This is especially true when you use the auto-login keystore, which does not require a password to open. In case the backup tape is lost, a malicious user should not be able to get both the encrypted data and the keystore.

Oracle Recovery Manager (Oracle RMAN) does not back up the software keystore as part of the database backup. When using a media manager such as Oracle Secure Backup with Oracle RMAN, Oracle Secure Backup automatically excludes auto-open keystores (the `cwallet.sso` files). However, it does not automatically exclude encryption keystores (the `ewallet.p12` files). It is a good practice to add the following `exclude data set` statement to your Oracle Secure Backup configuration:

```
exclude name *.p12
```

This setting instructs Oracle Secure Backup to exclude the encryption keystore from the backup set.

If you lose the software keystore that stores the TDE master encryption key, then you can restore access to encrypted data by copying the backed-up version of the keystore to the appropriate location. If you archived the restored keystore after the last time that you reset the TDE master encryption key, then you do not need to take any additional action.

If the restored software keystore does not contain the most recent TDE master encryption key, then you can recover old data up to the point when the TDE master encryption key was reset by rolling back the state of the database to that point in time. All of the modifications to encrypted columns after the TDE master encryption key was reset are lost.

See Also:

Oracle Database Backup and Recovery User's Guide for information about recovering a database

4.1.14 Deletion of Keystores

Oracle strongly recommends that you do not delete keystores, particularly after you have configured Transparent Data Encryption and the keystore is in use.

You can find if a keystore is in use by querying the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view after you open the keystore.

The reason you should not delete a keystore is because the keystore contains a list of all of the keys that were used for the database. Deleting the keystore deletes these keys, and could result in the loss of encrypted data. The deletion of a keystore can even hamper the normal functioning of the Oracle database. Even if you decrypted all of the data in your database, you still should not delete the keystore, because the TDE master encryption key in the keystore is also used for other Oracle Database features, such as off-lined tablespaces, Oracle Recovery Manager, and Oracle Secure Backup.

Even after you have migrated your keystores to a hardware security module, you should not delete the original keystore. The keys in the original keystore will be needed at a later time, for example when recovering an offline encrypted tablespace. Even if there is no data online that are not encrypted, the key may still be in use.

The exception is in the case of software auto-login (or auto-login local) keystores. If you do not want to use this type of keystore, then ideally you should move it to a secure directory. Only delete an auto-login keystore if you are sure that it comes from a specific password-based software keystore and that this keystore is available. The keystore should be available and known.

4.2 Managing the TDE Master Encryption Key

You can manage the TDE master encryption key in several ways.

[Creating TDE Master Encryption Keys for Later Use](#) (page 4-25)

You can create a TDE master encryption key that can be activated at a later date.

[Activation of TDE Master Encryption Keys](#) (page 4-28)

After you activate a TDE master encryption key, it can be used.

[TDE Master Encryption Key Attribute Management](#) (page 4-31)

Master encryption key attributes store information about the TDE master encryption key.

[Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#) (page 4-32)

Custom TDE master encryption key attributes enable you to defined attributes that are specific to your needs.

[Setting or Rotating the TDE Master Encryption Key in the Keystore](#) (page 4-34)

You can set or rotate the TDE master encryption key for both software keystores and hardware keystores.

[Exporting and Importing the TDE Master Encryption Key](#) (page 4-39)

You can export and import the TDE master encryption key in different ways to satisfy the needs of features such as a multitenant environment.

[Management of TDE Master Encryption Keys Using Oracle Key Vault](#) (page 4-45)

You can use Oracle Key Vault to manage and share TDE master encryption keys across an enterprise.

4.2.1 Creating TDE Master Encryption Keys for Later Use

You can create a TDE master encryption key that can be activated at a later date.

[About Creating a TDE Master Encryption Key for Later Use](#) (page 4-26)

The `CREATE KEY` clause of the `ADMINISTER KEY MANAGEMENT` statement can create a TDE master encryption key to be activated at a later date.

[Creating a TDE Master Encryption Key for Later Use](#) (page 4-26)

A keystore must be opened before you can create a TDE master encryption key for use later on.

[Example: Creating a TDE Master Encryption Key in a Single Database](#) (page 4-27)

You can use the `ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG` statement to create a TDE master encryption key in a single database.

Example: Creating a TDE Master Encryption Key in All PDBs (page 4-28)

You can use the `ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG` statement to create a TDE master encryption key in all PDBs.

4.2.1.1 About Creating a TDE Master Encryption Key for Later Use

The `CREATE KEY` clause of the `ADMINISTER KEY MANAGEMENT` statement can create a TDE master encryption key to be activated at a later date.

You then can activate this key on the same database or export it to another database and activate it there.

This method of TDE master encryption key creation is useful in a multitenant environment when you must re-create the TDE master encryption keys. The `CREATE KEY` clause enables you to use a single SQL statement to generate a new TDE master encryption key for all of the PDBs within a multitenant environment. The creation time of the new TDE master encryption key is later than the activation of the TDE master encryption key that is currently in use. Hence, the creation time can serve as a reminder to all of the PDBs to activate the most recently created TDE master encryption key as soon as possible.

4.2.1.2 Creating a TDE Master Encryption Key for Later Use

A keystore must be opened before you can create a TDE master encryption key for use later on.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Ensure that the keystore is open.

You can query the `STATUS` column of the `V$ENCRYPTION_WALLET` view to find if the keystore is open. If you find that you must open the software keystore, then you can optionally include the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement when you create the TDE master encryption key. This clause enables you to open a software keystore during the creation of the TDE master encryption key without having to separately open the auto-login keystore or the password-based keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
keystore_password;
```

3. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT CREATE KEY [USING TAG 'tag']
[FORCE KEYSTORE]
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING 'backup_identifier']]
[CONTAINER = (ALL|CURRENT)];
```

In this specification:

- *tag* is the associated attribute and information that you define. Enclose this setting in single quotation marks (' ').
- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is open (and in use), or if the keystore is closed. This clause enables you to open the keystore without having to explicitly open the keystore or explicitly open the password-based keystore.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - *keystore_password* is the mandatory keystore password that you used when you created the original keystore. It is case sensitive.
- `WITH BACKUP` backs up the TDE master encryption key in the same location as the key, as identified by the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view. To find the key locations for all of the database instances, query the `GV$ENCRYPTION_WALLET` view.

You must back up password-based software keystores. You do not need to back up auto-login or local auto-login software keystores. Optionally, include the `USING backup_identifier` clause to add a description of the backup. Enclose *backup_identifier* in single quotation marks (' ').

- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to set the encryption key in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

4. If necessary, activate the TDE master encryption key.

For example:

```
ADMINISTER KEY MANAGEMENT USE KEY
'ARaHD762tUkkvyLgPzAi6hMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA'
USING TAG 'quarter:second;description:Activate Key on standby'
IDENTIFIED BY password WITH BACKUP;
```

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[Opening a Hardware Keystore](#) (page 3-16)

To open a hardware keystore, use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

[Activation of TDE Master Encryption Keys](#) (page 4-28)

After you activate a TDE master encryption key, it can be used.

4.2.1.3 Example: Creating a TDE Master Encryption Key in a Single Database

You can use the `ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG` statement to create a TDE master encryption key in a single database.

[Example 4-4](#) (page 4-28) shows how to create a TDE master encryption key in a single database. After you run this statement, a TDE master encryption key with the tag definition is created in the keystore for that database. You can query the `TAG` column of the `V$ENCRYPTION_KEYS` view for the identifier of the newly created key. You can

query the `CREATION_TIME` column to find the most recently created key, which would be the key that you created from this statement. You can export this key to another database if you want or activate it locally later on, as described in [Activation of TDE Master Encryption Keys](#) (page 4-28).

Example 4-4 Creating a TDE Master Encryption Key in a Single Database

```
ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG
'source:admin@source;target:dbl@target'
IDENTIFIED BY password WITH BACKUP;
```

keystore altered.

4.2.1.4 Example: Creating a TDE Master Encryption Key in All PDBs

You can use the `ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG` statement to create a TDE master encryption key in all PDBs.

[Example 4-5](#) (page 4-28) shows how to create a TDE master encryption key in all of the PDBs in a multitenant environment. It uses the `FORCE KEYSTORE` clause in the event that the auto-login keystore in the root is open. The password is stored externally, so the `EXTERNAL STORE` setting is used for the `IDENTIFIED BY` clause. After you run this statement, a TDE master encryption key is created in each PDB. You can find the identifiers for these keys as follows:

- Log in to the PDB and then query the `TAG` column of the `V$ENCRYPTION_KEYS` view.
- Log in to the root and then query the `INST_ID` and `TAG` columns of the `GV$ENCRYPTION_KEYS` view.

You also can check the `CREATION_TIME` column of these views to find the most recently created key, which would be the key that you created from this statement. After you create the keys, you can individually activate the keys in each of the PDBs.

Example 4-5 Creating a TDE Master Encryption Key in All of the PDBs

```
ADMINISTER KEY MANAGEMENT CREATE KEY USING TAG
'scope:all pdbs;description:Create Key for ALL PDBS'
FORCE KEYSTORE IDENTIFIED BY EXTERNAL STORE
WITH BACKUP CONTAINER=ALL;
```

keystore altered.

4.2.2 Activation of TDE Master Encryption Keys

After you activate a TDE master encryption key, it can be used.

[About Activating TDE Master Encryption Keys](#) (page 4-29)

You can activate a previously created or imported TDE master encryption key by using the `USE KEY` clause of `ADMINISTER KEY MANAGEMENT`.

[Activating a TDE Master Encryption Key](#) (page 4-29)

To activate a TDE master encryption key, you must open the keystore and use `ADMINISTER KEY MANAGEMENT` with the `USE KEY` clause.

[Example: Activating a TDE Master Encryption Key](#) (page 4-30)

You can use the `ADMINISTER KEY MANAGEMENT SQL` statement to activate a TDE master encryption key.

4.2.2.1 About Activating TDE Master Encryption Keys

You can activate a previously created or imported TDE master encryption key by using the `USE KEY` clause of `ADMINISTER KEY MANAGEMENT`.

After you activate the key, it is available for use. The key will be used to protect all of the column keys and all of the [tablespace encryption keys](#). If you have deployed a logical standby database, then you must export the TDE master encryption keys after recreating them, and then import them into the standby database. You can have the TDE master encryption key in use on both the primary and the standby databases. To do so, you must activate the TDE master encryption key after you import it to the logical standby database.

4.2.2.2 Activating a TDE Master Encryption Key

To activate a TDE master encryption key, you must open the keystore and use `ADMINISTER KEY MANAGEMENT` with the `USE KEY` clause.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Ensure that the keystore is open.

You can query the `STATUS` column of the `V$ENCRYPTION_WALLET` view to find if the keystore is open. If you find that you must open the software keystore, then you can optionally include the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement when you activate the TDE master encryption key. This clause enables you to open a software keystore during the activation of the TDE master encryption key without having to separately open the auto-login keystore or the password-based keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
keystore_password;
```

3. Query the `KEY_ID` column of the `V$ENCRYPTION_KEYS` view to find the key identifier.

For example:

```
SELECT KEY_ID FROM V$ENCRYPTION_KEYS;

KEY_ID
-----
ARaHD762tUkkvyLgPzAi6hMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

4. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT USE KEY 'key_identifier'
[USING TAG 'tag'] [FORCE KEYSTORE]
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING 'backup_identifier']];
```

In this specification:

- *key_identifier* is the key identifier that you find from querying the `KEY_ID` column of the `V$ENCRYPTION_KEYS` view. Enclose this setting in single quotation marks (' ').
- *tag* is the associated attributes and information that you define. Enclose this setting in single quotation marks (' ').
- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - *keystore_password* is the mandatory keystore password that you used when you created the original keystore.
- `WITH BACKUP` backs up the TDE master encryption key in the same location as the key, as identified by the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view. To find the key locations for all of the database instances, query the `GV$ENCRYPTION_WALLET` view.
 You must back up password-based software keystores. You do not need to back up auto-login or local auto-login software keystores. Optionally, include the `USING backup_identifier` clause to add a description of the backup. Enclose *backup_identifier* in single quotation marks (' ').
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to set the encryption key in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

Related Topics:

[Opening a Software Keystore](#) (page 3-10)

To open a software keystore, you must use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

[Opening a Hardware Keystore](#) (page 3-16)

To open a hardware keystore, use the `ADMINISTER KEY MANAGEMENT` statement with the `SET KEYSTORE OPEN` clause.

4.2.2.3 Example: Activating a TDE Master Encryption Key

You can use the `ADMINISTER KEY MANAGEMENT SQL` statement to activate a TDE master encryption key.

[Example 4-6](#) (page 4-30) shows how to activate a previously imported TDE master encryption key and then update its tag. This key is activated with the current database time stamp and time zone.

Example 4-6 Activating a TDE Master Encryption Key

```
ADMINISTER KEY MANAGEMENT USE KEY
'ARaHD762tUkkvyLgPzAi6hMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA'
USING TAG 'quarter:second:description:Activate Key on standby'
IDENTIFIED BY password WITH BACKUP;
```

keystore altered.

In this version of the same operation, the `FORCE KEYSTORE` clause is added in the event that the auto-login keystore is in use, or if the keystore is closed. The password of the keystore is stored externally, so the `EXTERNAL STORE` setting is used for the `IDENTIFIED BY` clause.

```
ADMINISTER KEY MANAGEMENT USE KEY
'ARaHD762tUkkvyLgPzAi6hMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA'
USING TAG 'quarter:second:description:Activate Key on standby'
FORCE KEYSTORE IDENTIFIED BY EXTERNAL STORE WITH BACKUP;
```

keystore altered.

4.2.3 TDE Master Encryption Key Attribute Management

Master encryption key attributes store information about the TDE master encryption key.

[TDE Master Encryption Key Attributes](#) (page 4-31)

Master encryption key attributes include detailed information about the TDE master encryption key.

[Finding the TDE Master Encryption Key That Is in Use](#) (page 4-32)

A TDE master encryption key that is in use is the key that was activated most recently for the database.

4.2.3.1 TDE Master Encryption Key Attributes

Master encryption key attributes include detailed information about the TDE master encryption key.

The information contains the following types:

- **Key time stamp information:** Internal security policies and compliance policies usually determine the key rotation frequency. You should expire keys when they reach the end of their lifetimes and then generate new keys. Time stamp attributes such as key creation time and activation time help you to determine the key age accurately, and automate key generation.

The `V$ENCRYPTION_KEYS` view includes columns such as `CREATION_TIME` and `ACTIVATION_TIME`. See *Oracle Database Reference* for a complete description of the `V$ENCRYPTION_KEYS` view.

- **Key owner information:** Key owner attributes help you to determine the user who created or activated the key. These attributes can be important for security, auditing, and tracking purposes. Key owner attributes also include key use information, such as whether the key is used for standalone TDE operations or used in a multitenant environment.

The `V$ENCRYPTION_KEYS` view includes columns such as `CREATOR`, `CREATOR_ID`, `USER`, `USER_ID`, and `KEY_USE`.

- **Key source information:** Keys often must be moved between databases for operations such as import-export operations and Data Guard-related operations. Key source attributes enable you to track the origin of each key. You can track whether a key was created locally or imported, and the database name and instance number of the database that created the key. In a multitenant environment, you can track the PDB where the key was created.

The V\$ENCRYPTION_KEYS view includes columns such as CREATOR_DBNAME, CREATOR_DBID, CREATOR_INSTANCE_NAME, CREATOR_INSTANCE_NUMBER, CREATOR_PDBNAME, and so on.

- **Key usage information:** Key usage information determines the database or PDB where the key is being used. It also helps determine whether a key is in active use or not.

The V\$ENCRYPTION_KEYS view includes columns such as ACTIVATING_DBNAME, ACTIVATING_DBID, ACTIVATING_INSTANCE_NAME, ACTIVATING_PDBNAME, and so on.

- **User-defined information and other information:** When creating a key, you can tag it with information using the TAG option. Each key contains important information such as whether or not it has been backed up.

The V\$ENCRYPTION_KEYS view includes columns such as KEY_ID, TAG, and other miscellaneous columns, for example BACKED_UP.

4.2.3.2 Finding the TDE Master Encryption Key That Is in Use

A TDE master encryption key that is in use is the key that was activated most recently for the database.

In a multitenant environment, the master key in use of the PDB is the one that was activated most recently for that PDB.

- To find the master key, query the V\$ENCRYPTION_KEYS dynamic view.
 - To find the master key in use in a non-CDB:

```
SELECT KEY_ID
FROM V$ENCRYPTION_KEYS
WHERE ACTIVATION_TIME = (SELECT MAX(ACTIVATION_TIME)
                        FROM V$ENCRYPTION_KEYS
                        WHERE ACTIVATING_DBID = (SELECT DBID FROM V
$DATABASE));
```

- To find the master key in use in a CDB:

```
SELECT KEY_ID
FROM V$ENCRYPTION_KEYS
WHERE ACTIVATION_TIME = (SELECT MAX(ACTIVATION_TIME)
                        FROM V$ENCRYPTION_KEYS
                        WHERE ACTIVATING_PDBID = SYS_CONTEXT('USERENV',
'CON_ID'));
```

4.2.4 Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes

Custom TDE master encryption key attributes enable you to defined attributes that are specific to your needs.

[About Creating Custom Attribute Tags](#) (page 4-33)

Attribute tags enable you to monitor specific activities users perform, such as accessing a particular terminal ID.

[Creating a Custom Attribute Tag](#) (page 4-33)

To create a custom attribute tag, you must use the SET TAG clause of the ADMINISTER KEY MANAGEMENT statement.

4.2.4.1 About Creating Custom Attribute Tags

Attribute tags enable you to monitor specific activities users perform, such as accessing a particular terminal ID.

By default, Oracle Database defines a set of attributes that describe various characteristics of the TDE master encryption keys that you create, such as the creation time, database in which the TDE master encryption key is used, and so on. These attributes are captured by the `V$ENCRYPTION_KEY` dynamic view.

You can create custom attributes that can be captured by the `TAG` column of the `V$ENCRYPTION_KEYS` dynamic view. This enables you to define behaviors that you may want to monitor, such as users who perform activities on encryption keys. The tag can encompass multiple attributes, such as session IDs from a specific terminal.

4.2.4.2 Creating a Custom Attribute Tag

To create a custom attribute tag, you must use the `SET TAG` clause of the `ADMINISTER KEY MANAGEMENT` statement.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. If necessary, query the `TAG` column of the `V$ENCRYPTION_KEY` dynamic view to find a listing of existing tags for the TDE master encryption keys.

When you create a new tag for a TDE master encryption key, it overwrites the existing tag for that TDE master encryption key.

3. Create the tag as follows:

```
ADMINISTER KEY MANAGEMENT SET TAG 'tag' FOR 'master_key_identifier'
[FORCE KEYSTORE] IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING 'backup_identifier']];
```

In this specification

- *tag* is the associated attributes or information that you define. Enclose this information in single quotation marks (' ').
- *master_key_identifier* identifies the TDE master encryption key for which the *tag* is set. To find a list of TDE master encryption key identifiers, query the `KEY_ID` column of the `V$ENCRYPTION_KEYS` dynamic view.
- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - *keystore_password* is the password that was used to create the keystore.

- *backup_identifier* defines the tag values. Enclose this setting in single quotation marks (' ') and separate each value with a colon.

For example, to create a tag that uses two values, one to capture a specific session ID and the second to capture a specific terminal ID:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY USING TAG
'sessionid=3205062574:terminal=xcvt'
IDENTIFIED BY keystore_password WITH BACKUP;
```

keystore altered.

Both the session ID (3205062574) and terminal ID (xcvt) can derive their values by using either the SYS_CONTEXT function with the USERENV namespace, or by using the USERENV function.

After you create the tag for a TDE master encryption key, its name should appear in the TAG column of the V\$ENCRYPTION_KEYS view for that TDE master encryption key. If you create a tag for the secret, then the tag appears in the SECRET_TAG column of the V\$CLIENT_SECRETS view. If you create a secret with a tag, then the tag appears in the SECRET_TAG column of the V\$CLIENT_SECRETS view.

See Also:

- [Storing Oracle GoldenGate Secrets in a Keystore](#) (page 4-53) for information about creating secrets
 - *Oracle Database SQL Language Reference* for a full list of predefined parameters in the USERENV namespace in the SYS_CONTEXT function
-
-

4.2.5 Setting or Rotating the TDE Master Encryption Key in the Keystore

You can set or rotate the TDE master encryption key for both software keystores and hardware keystores.

[About Setting or Rotating the TDE Master Encryption Key in the Keystore](#)
(page 4-35)

You can set or rotate the TDE master encryption key for both software password-based and hardware keystores.

[Creating and Backing Up a TDE Master Encryption Key and Applying a Tag to It](#)
(page 4-35)

The ADMINISTER KEY MANAGEMENT statement enables you to create and back up a TDE master encryption key and apply a tag to it.

[About Rotating the TDE Master Encryption Key](#) (page 4-36)

Oracle Database uses a unified master encryption key for both TDE column encryption and TDE tablespace encryption.

[Rotating the TDE Master Encryption Key](#) (page 4-37)

You can use the ADMINISTER KEY MANAGEMENT statement to rotate (also called rekeying) a TDE master encryption key.

[Rotating the TDE Master Encryption Key for a Tablespace](#) (page 4-39)

You can use the REKEY clause of the ALTER TABLESPACE statement to rotate an encryption key for an encrypted tablespace.

4.2.5.1 About Setting or Rotating the TDE Master Encryption Key in the Keystore

You can set or rotate the TDE master encryption key for both software password-based and hardware keystores.

The TDE master encryption key is stored in an external security module (keystore), and it is used to protect the [TDE table keys](#) and [tablespace encryption keys](#). By default, the TDE master encryption key is a system-generated random value created by Transparent Data Encryption (TDE).

Use the `ADMINISTER KEY MANAGEMENT` statement to set or reset (`REKEY`) the TDE master encryption key. When the master encryption key is set, then TDE is considered enabled and cannot be disabled.

Before you can encrypt or decrypt database columns or tablespaces, you must generate a TDE master encryption key. Oracle Database uses the same TDE master encryption key for both TDE column encryption and TDE tablespace encryption. The instructions for setting a software or hardware TDE master encryption key explain how to generate a TDE master encryption key.

Related Topics:

[Step 4: Set the Software TDE Master Encryption Key](#) (page 3-11)

Once the keystore is open, you can set a TDE master encryption key for it.

[Step 4: Set the Hardware Keystore TDE Master Encryption Key](#) (page 3-17)

After you have opened the hardware keystore, you are ready to set the hardware keystore TDE master encryption key.

4.2.5.2 Creating and Backing Up a TDE Master Encryption Key and Applying a Tag to It

The `ADMINISTER KEY MANAGEMENT` statement enables you to create and back up a TDE master encryption key and apply a tag to it.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or to the PDB. For example:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY [USING TAG 'tag']
[FORCE KEYSTORE] IDENTIFIED BY [EXTERNAL STORE | keystore_password]
WITH BACKUP [USING 'backup_identifier'] [CONTAINER = ALL | CURRENT];
```

In this specification:

- `tag` is the tag that you want to create. Enclose this tag in single quotation marks (`'`).

- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `keystore_password` is either `software_keystore_password` or `user_id:password`. The `user_id:password` setting is the hardware keystore user ID and password that was created in Step 3 (page 3-15) under [Step 2: Configure the Hardware Security Module](#) (page 3-14). As with software passwords, it is case sensitive. You must enclose the password string in double quotation marks (" "). Separate `user_id` and `password` with a colon (:).
- `WITH BACKUP` backs the TDE master encryption key up in the same location as the key, as identified by the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view. To find the `WRL_PARAMETER` values for all of the database instances, query the `GV$ENCRYPTION_WALLET` view.
 You must back up password-based software keystores. You do not need to use it for auto-login or local auto-login software keystores. Optionally, include the `USING backup_identifier` clause to add a description of the backup. Enclose this identifier in single quotation marks (' ').
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to set the encryption key in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY USING TAG 'backups"
IDENTIFIED BY password WITH BACKUP USING 'hr.emp_key_backup';
```

keystore altered.

Oracle Database uses the keystore in the keystore location specified by the `ENCRYPTION_WALLET_LOCATION` parameter in the `sqlnet.ora` file to store the TDE master encryption key.

Related Topics:

[Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#) (page 4-32)

Custom TDE master encryption key attributes enable you to defined attributes that are specific to your needs.

[About the Keystore Location in the sqlnet.ora File](#) (page 3-3)

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

4.2.5.3 About Rotating the TDE Master Encryption Key

Oracle Database uses a unified master encryption key for both TDE column encryption and TDE tablespace encryption.

When you rotate (also called rekeying) the TDE master encryption key for TDE column encryption, the master encryption key for TDE tablespace encryption also is

rotated. Rotate the master encryption key only if it was compromised or as per the security policies of the organization. This process deactivates the previous TDE master encryption key.

You cannot change the TDE master encryption key or rotate a TDE master encryption key for an auto-login keystore. Because auto-login keystores do not have a password, an administrator or a privileged user can change the keys without the knowledge of the security officer. However, if both the auto-login and the password-based keystores are present in the configured location (as set in the `sqlnet.ora` file), then when you rotate the TDE master encryption key, a TDE master encryption key is added to both the auto-login and password-based keystores. If the auto-login keystore is in use in a location that is different from that of the password-based keystore, then you must re-create the auto-login keystore.

Do not perform a rotation operation of the master key concurrently with an online tablespace rekey operation. You can find if an online tablespace is in the process of being rotated by issuing the following query:

```
SELECT TS#, ENCRYPTIONALG, STATUS FROM V$ENCRYPTED_TABLESPACES;
```

A status of `REKEYING` means that the corresponding tablespace is still being rotated.

Note:

You cannot add new information to auto-login keystores separately.

4.2.5.4 Rotating the TDE Master Encryption Key

You can use the `ADMINISTER KEY MANAGEMENT` statement to rotate (also called rekeying) a TDE master encryption key.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or to the PDB. For example, to log in to a PDB called `hrpdb`:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Ensure that the keystore is open.

Query the `STATUS` column of the `V$ENCRYPTION_WALLET` view to find if the keystore is open. If you find that you must open the software keystore, then you can optionally include the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement when you rotate the TDE master encryption key. This clause enables you to open a software keystore during the rotation of the TDE master encryption key without having to separately open the auto-login keystore or the password-based keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
keystore_password;
```

3. If you are rotating the TDE master encryption key for a keystore that has auto login enabled, then ensure that both the auto login keystore, identified by the `.sso` file, and the encryption keystore, identified by the `.p12` file, are present.

You can find the location of these files by querying the `WRL_PARAMETER` column of the `V$ENCRYPTION_WALLET` view. To find the `WRL_PARAMETER` values for all of the database instances, query the `GV$ENCRYPTION_WALLET` view.

4. Rotate the TDE master encryption key by using the following statement:

```
ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY [USING TAG 'tag']
[FORCE KEYSTORE] IDENTIFIED BY [EXTERNAL STORE | keystore_password]
WITH BACKUP [USING 'backup_identifier'] [CONTAINER = ALL | CURRENT];
```

In this specification:

- `tag` is the associated attributes and information that you define. Enclose this setting in single quotation marks (' ').
- `FORCE KEYSTORE` enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `keystore_password` is the mandatory keystore password that you created when you created the keystore in [Step 2: Create the Software Keystore](#) (page 3-6).
- `WITH BACKUP` creates a backup of the keystore. You must use this option for password-based and hardware keystores. Optionally, you can use the `USING` clause to add a brief description of the backup. Enclose this description in single quotation marks (' '). This identifier is appended to the named keystore file (for example, `ewallet_time-stamp_emp_key_backup.p12`). Follow the file naming conventions that your operating system uses.
- `CONTAINER` is for use in a multitenant environment. Enter `ALL` to open the keystore in all of the PDBs in this CDB, or `CURRENT` for the current PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY keystore_password WITH BACKUP
USING 'emp_key_backup';
```

keystore altered.

For better security and to meet compliance regulations, periodically rotate the TDE master encryption key. This process deactivates the previous TDE master encryption key, creates a new TDE master encryption key, and then activates it. You can check the keys that were created recently by querying the `CREATION_TIME` column in the `V$ENCRYPTION_KEYS` view. To find the keys that were activated recently, query the `ACTIVATION_TIME` column in the `V$ENCRYPTION_KEYS` view.

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

Step 3: Open the Hardware Keystore (page 3-15)

After you have configured the hardware security module, you must open the hardware keystore before it can be used.

4.2.5.5 Rotating the TDE Master Encryption Key for a Tablespace

You can use the `REKEY` clause of the `ALTER TABLESPACE` statement to rotate an encryption key for an encrypted tablespace.

1. Ensure that the tablespace open in read-write mode.

You can query the `STATUS` column of the `V$INSTANCE` dynamic view to find if a database is open and the `OPEN_MODE` column of the `V$DATABASE` view to find if it in read-write mode.

2. If necessary, open the database in read-write mode.

```
ALTER DATABASE OPEN READ WRITE;
```

3. Run the `ALTER TABLESPACE` SQL statement to encrypt the tablespace.

For example:

```
ALTER TABLESPACE encrypt_ts ENCRYPTION USING 'AES256' ENCRYPT;
```

Related Topics:[About Encryption Conversions for Existing Online Tablespaces](#) (page 3-43)

You can encrypt, decrypt, or rekey existing user tablespaces, and the `SYSTEM`, `SYSAUX`, and `UNDO` tablespace when they are online.

4.2.6 Exporting and Importing the TDE Master Encryption Key

You can export and import the TDE master encryption key in different ways to satisfy the needs of features such as a multitenant environment.

[About Exporting and Importing the TDE Master Encryption Key](#) (page 4-40)

Oracle Database features such as transportable tablespaces and Oracle Data Pump move data that is possibly encrypted between databases.

[About Exporting TDE Master Encryption Keys](#) (page 4-40)

You can use `ADMINISTER KEY MANAGEMENT EXPORT` to export TDE master encryption keys from a keystore, and then import them into another keystore.

[Exporting a TDE Master Encryption Key](#) (page 4-41)

The `ADMINISTER KEY MANAGEMENT` statement with the `EXPORT [ENCRYPTION] KEYS WITH SECRET` clause exports a TDE master encryption key.

[Example: Exporting a TDE Master Encryption Key by Using a Subquery](#) (page 4-42)

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement can export a TDE master encryption key by using a subquery.

[Example: Exporting a List of TDE Master Encryption Key Identifiers to a File](#) (page 4-42)

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET` statement can export a list of TDE master encryption key identifiers to a file.

[Example: Exporting All TDE Master Encryption Keys of the Database](#)
(page 4-42)

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` SQL statement can export all TDE master encryption keys of a database.

[About Importing TDE Master Encryption Keys](#) (page 4-43)

The `ADMINISTER KEY MANAGEMENT IMPORT` statement can import exported TDE master encryption keys from a key export file into a target keystore.

[Importing a TDE Master Encryption Key](#) (page 4-43)

The `ADMINISTER KEY MANAGEMENT` statement with the `IMPORT [ENCRYPTION] KEYS WITH SECRET` clause can import a TDE master encryption key.

[Example: Importing a TDE Master Encryption Key](#) (page 4-44)

You can use the `ADMINISTER KEY MANAGEMENT IMPORT KEYS SQL` statement to import a TDE master encryption key.

[How Keystore Merge Differs from TDE Master Encryption Key Export or Import](#)
(page 4-44)

The keystore merge operation differs from the TDE master encryption key export and import operations.

4.2.6.1 About Exporting and Importing the TDE Master Encryption Key

Oracle Database features such as transportable tablespaces and Oracle Data Pump move data that is possibly encrypted between databases.

In addition, CDBs contain PDBs that can be plugged in or unplugged. These are some common scenarios in which you can choose to export and import TDE master encryption keys to move them between source and target keystores. For Data Guard (Logical Standby), you must copy the keystore that is in the primary database to the standby database. Instead of merging the primary database keystore with the standby database, you can export the TDE master encryption key that is in use and then import it to the standby database. Moving transportable tablespaces that are encrypted between databases requires that you export the TDE master encryption key at the source database and then import it into the target database.

4.2.6.2 About Exporting TDE Master Encryption Keys

You can use `ADMINISTER KEY MANAGEMENT EXPORT` to export TDE master encryption keys from a keystore, and then import them into another keystore.

A TDE master encryption key is exported together with its key identifier and key attributes. The exported keys are protected with a password (secret) in the export file.

You can specify the TDE master encryption keys to be exported by using the `WITH IDENTIFIER` clause of the `ADMINISTER KEY MANAGEMENT EXPORT` statement. To export the TDE master encryption keys, you can either specify their key identifiers as a comma-separated list, or you can specify a query that enumerates their key identifiers. Be aware that Oracle Database executes the query determining the key identifiers within the current user's rights and not with definer's rights.

If you omit the `WITH IDENTIFIER` clause, then all of the TDE master encryption keys of the database are exported.

In a consolidated database, you can export the keys from within a PDB for a PDB to be unplugged. In this scenario, you can only use the `WITH IDENTIFIER` clause in the root and not in a PDB.

Related Topics:

[Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12)

The `EXPORT` and `IMPORT` clauses of `ADMINISTER KEY MANAGEMENT` `EXPORT` can export or import TDE master encryption keys for a PDB.

[Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12)

The `EXPORT` and `IMPORT` clauses of `ADMINISTER KEY MANAGEMENT` `EXPORT` can export or import TDE master encryption keys for a PDB.

4.2.6.3 Exporting a TDE Master Encryption Key

The `ADMINISTER KEY MANAGEMENT` statement with the `EXPORT [ENCRYPTION] KEYS WITH SECRET` clause exports a TDE master encryption key.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. If necessary, open the keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY password;
```

3. Run the following SQL statement to export a set of TDE master encryption keys:

```
ADMINISTER KEY MANAGEMENT EXPORT [ENCRYPTION] KEYS
WITH SECRET "export_secret"
TO 'file_path'
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH IDENTIFIER IN 'key_id1', 'key_id2', 'key_idn' | (SQL_query)];
```

In this specification:

- `export_secret` is a password that you can specify to encrypt the export the file that contains the exported keys. Enclose this secret in double quotation marks (" "), or you can omit the quotation marks if the secret has no spaces.
- `file_path` is the complete path and name of the file to which the keys must be exported. Enclose this path in single quotation marks (' ').
- `IDENTIFIED BY` can be one of the following settings:
 - `EXTERNAL STORE` uses the keystore password stored in the external store to perform the keystore operation.
 - `software_keystore_password` is the password of the keystore containing the keys.
- `key_id1, key_id2, key_idn` is a string of one or more TDE master encryption key identifiers for the TDE master encryption key being exported. Separate each key identifier with a comma and enclose each of these key identifiers in single quotation marks (' '). To find a list of TDE master encryption key identifiers, query the `KEY_ID` column of the `V $ENCRYPTION_KEYS` dynamic view.

- *SQL_query* is a query that fetches a list of the TDE master encryption key identifiers. It should return only one column which contains the TDE master encryption key identifiers. This query is executed with current user rights.

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

4.2.6.4 Example: Exporting a TDE Master Encryption Key by Using a Subquery

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement can export a TDE master encryption key by using a subquery.

[Example 4-8](#) (page 4-42) shows how to export TDE master encryption keys whose identifiers are fetched by a query to a file called `export.exp`. The TDE master encryption keys in the file are encrypted using the secret `my_secret`. The `SELECT` statement finds the identifiers for the TDE master encryption keys to be exported.

Be aware that in a multitenant environment, the `WITH IDENTIFIER` clause is not supported when you try to import or export keys inside a PDB. It is only permitted in the root. See [Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12) for information about exporting keys in a PDB.

Example 4-7 Exporting a List of TDE Master Encryption Key Identifiers to a File

```
ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET "my_secret"
TO '/TDE/export.exp' IDENTIFIED BY password
WITH IDENTIFIER IN 'AdoxnJ0uH08cv7xkz83ovwsAAAAAAAAAAAAAAAAAAAAAAAAAAAA',
'AW5z3CoyKE/yv3cNT5CWCXUAAAAAAAAAAAAAAAAAAAAAAAAAAAAA';
```

keystore altered.

4.2.6.5 Example: Exporting a List of TDE Master Encryption Key Identifiers to a File

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET` statement can export a list of TDE master encryption key identifiers to a file.

[Example 4-7](#) (page 4-42) shows how to export TDE master encryption keys by specifying their identifiers as a list, to a file called `export.exp`. Master encryption keys in the file are encrypted using the secret `my_secret`. The identifiers of the TDE master encryption key to be exported are provided as a comma-separated list.

Example 4-8 Exporting TDE Master Encryption Key Identifiers by Using a Subquery

```
ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET "my_secret"
TO '/etc/TDE/export.exp' IDENTIFIED BY password
WITH IDENTIFIER IN (SELECT KEY_ID FROM V$ENCRYPTION_KEYS WHERE ROWNUM <3);
```

keystore altered.

4.2.6.6 Example: Exporting All TDE Master Encryption Keys of the Database

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS SQL` statement can export all TDE master encryption keys of a database.

[Example 4-9](#) (page 4-43) shows how to export all of the TDE master encryption keys of the database to a file called `export.exp`. The TDE master encryption keys in the file are encrypted using the secret `my_secret`.

Example 4-9 Exporting All of the TDE Master Encryption Keys of the Database

```
ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET "my_secret" TO
'/etc/TDE/export.exp' IDENTIFIED BY password;
```

keystore altered.

4.2.6.7 About Importing TDE Master Encryption Keys

The `ADMINISTER KEY MANAGEMENT IMPORT` statement can import exported TDE master encryption keys from a key export file into a target keystore.

You cannot re-import TDE master encryption keys that have already been imported.

In a consolidated database, you can import the keys from within a PDB for a PDB to be plugged.

Related Topics:

[Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12)

The `EXPORT` and `IMPORT` clauses of `ADMINISTER KEY MANAGEMENT EXPORT` can export or import TDE master encryption keys for a PDB.

4.2.6.8 Importing a TDE Master Encryption Key

The `ADMINISTER KEY MANAGEMENT` statement with the `IMPORT [ENCRYPTION] KEYS WITH SECRET` clause can import a TDE master encryption key.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root. The following command logs user `c##sec_admin` into the root.

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. If necessary, open the keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY password;
```

3. Run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT IMPORT [ENCRYPTION] KEYS
WITH SECRET "import_secret" FROM 'file_name' | FROM 'file_name'
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING 'backup_identifier']];
```

In this specification:

- `import_secret` is the same password that was used to encrypt the keys during the export operation. Enclose this secret in double quotation marks (" "), or you can omit the quotation marks if the secret has no spaces.
- `file_name` is the complete path and name of the file from which the keys need to be imported. Enclose this setting in single quotation marks (' ').
- `IDENTIFIED BY` can be one of the following settings:

- `EXTERNAL_STORE` uses the keystore password stored in the external store to perform the keystore operation.
- `software_keystore_password` is the password of the software keystore where the keys are being imported.
- `WITH BACKUP` must be used in case the target keystore was not backed up before the import operation. `backup_identifier` is an optional string that you can provide to identify the keystore backup. Enclose this setting in single quotation marks (' ').

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

4.2.6.9 Example: Importing a TDE Master Encryption Key

You can use the `ADMINISTER KEY MANAGEMENT IMPORT KEYS` SQL statement to import a TDE master encryption key.

[Example 4-10](#) (page 4-44) shows how to import the TDE master encryption key identifiers that are stored in the file `export.exp` and encrypted with the secret `my_secret`.

Example 4-10 Importing TDE Master Encryption Key Identifiers from an Export File

```
ADMINISTER KEY MANAGEMENT IMPORT KEYS WITH SECRET "my_secret"
FROM '/etc/TDE/export.exp' IDENTIFIED BY password WITH BACKUP;
```

keystore altered.

4.2.6.10 How Keystore Merge Differs from TDE Master Encryption Key Export or Import

The keystore merge operation differs from the TDE master encryption key export and import operations.

Even though both the `ADMINISTER KEY MANAGEMENT MERGE` statement and the `ADMINISTER KEY MANAGEMENT EXPORT` and `IMPORT` statements eventually move the TDE master encryption keys from one keystore to the next, there are differences in how these two statements function.

- The `MERGE` statement merges two keystores whereas the `EXPORT` and `IMPORT` statements export the keys to a file or import the keys from a file. The keystore is different from the export file, and the two cannot be used interchangeably. The export file is not a keystore and cannot be configured to be used with a database as a keystore. Similarly, the `IMPORT` statement cannot extract the TDE master encryption keys from the keystore.
- The `MERGE` statement merges all of the TDE master encryption keys of the specified keystores whereas the `EXPORT` and `IMPORT` statements can be selective.
- The `EXPORT` and `IMPORT` statements require the user to provide both a location (`filepath`) and the file name of the export file, whereas the `MERGE` statement only takes in the location of the keystores.

- The file name of the keystores is fixed and is determined by the `MERGE` operation and can be either `ewallet.p12` or `cwallet.sso`. The file names for the export files used in the `EXPORT` the `IMPORT` statements are specified by the user.
- The keystores on Automatic Storage Management (ASM) disk groups or regular file systems can be merged with `MERGE` statements. The export files used in the `EXPORT` and the `IMPORT` statements can only be a regular operating system file and cannot be located on an ASM disk group.
- The keystores merged using the `MERGE` statement do not need to be configured or in use with the database. The `EXPORT` statement can only export the keys from a keystore that is configured and in use with the database and is also open when the export is done. The `IMPORT` statement can only import the keys into a keystore that is open, configured, and in use with the database.
- The `MERGE` statement never modifies the metadata associated with the TDE master encryption keys. The `EXPORT` and `IMPORT` operations can modify the metadata of the TDE master encryption keys when required, such as during a PDB plug operation.

4.2.7 Management of TDE Master Encryption Keys Using Oracle Key Vault

You can use Oracle Key Vault to manage and share TDE master encryption keys across an enterprise.

Oracle Key Vault securely stores the keys in a central repository, along with other security objects such as credential files and Java keystores, and enables you to share these objects with other TDE-enabled databases.

See Also:

- [Migration of Keystores to and from Oracle Key Vault](#) (page 4-20) for additional benefits of using Oracle Key Vault
 - *Oracle Key Vault Administrator's Guide*
-
-

4.3 Storing Secrets Used by Oracle Database

Secrets are data that support internal Oracle Database features that integrate external clients such as Oracle GoldenGate into the database.

[About Storing Oracle Database Secrets in a Keystore](#) (page 4-46)

Keystores can store secrets that support internal Oracle Database features and integrate external clients such as Oracle GoldenGate.

[Storage of Oracle Database Secrets in a Software Keystore](#) (page 4-47)

The `ADMINISTER KEY MANAGEMENT ADD SECRET|UPDATE SECRET|DELETE CLIENT` statements can add secrets, update secrets, and delete secrets from a keystore.

[Example: Adding an HSM Password to a Software Keystore](#) (page 4-48)

The `ADMINISTER KEY MANAGEMENT ADD SECRET` statement can add an HSM password to a software keystore.

[Example: Changing an HSM Password Stored as a Secret in a Software Keystore](#)
(page 4-49)

The `ADMINISTER KEY MANAGEMENT UPDATE SECRET` statement can change an HSM password that is stored as a secret in a software keystore.

[Example: Deleting an HSM Password Stored as a Secret in a Software Keystore](#)
(page 4-49)

The `ADMINISTER KEY MANAGEMENT DELETE SECRET` statement can delete HSM passwords that are stored as secrets in a software keystore.

[Storage of Oracle Database Secrets in a Hardware Keystore](#) (page 4-49)

The `ADMINISTER KEY MANAGEMENT ADD SECRET | UPDATE SECRET | DELETE CLIENT` statements can add, update, and delete secrets.

[Example: Adding an Oracle Database Secret to a Hardware Keystore](#)
(page 4-51)

The `ADMINISTER KEY MANAGEMENT ADD SECRET` statement can add an Oracle Database secret to a hardware keystore.

[Example: Changing an Oracle Database Secret in a Hardware Keystore](#)
(page 4-51)

The `ADMINISTER KEY MANAGEMENT MANAGEMENT UPDATE SECRET` statement can change an Oracle Database secret in a hardware keystore.

[Example: Deleting an Oracle Database Secret in a Hardware Keystore](#)
(page 4-51)

The `ADMINISTER KEY MANAGEMENT DELETE SECRET FOR CLIENT` statement can delete an Oracle Database secret that is in a hardware keystore.

[Configuring Auto-Login Hardware Security Modules](#) (page 4-52)

A hardware security module can be configured to use the auto-login capability.

4.3.1 About Storing Oracle Database Secrets in a Keystore

Keystores can store secrets that support internal Oracle Database features and integrate external clients such as Oracle GoldenGate.

The secret key must be a string adhering to Oracle identifier rules. You can add, update, or delete a client secret in an existing keystore. The Oracle GoldenGate Extract process must have data encryption keys to decrypt the data that is in data files and in REDO or UNDO logs. Keys are encrypted with shared secrets when you share the keys between an Oracle database and an Oracle GoldenGate client. The software keystore stores the shared secrets.

Depending on your site's requirements, you may require automated open keystore operations even when a hardware security module is configured. For this reason, the hardware security module password can be stored in a software auto-login keystore, which enables the auto-login capability for the hardware security module. The Oracle Database side can also store the credentials for the database to log in to an external storage server in the software keystore.

You can store Oracle Database secrets in both software keystores and hardware keystores:

- **Software keystores:** You can store secrets in software password-based, auto-login, and local auto-login software keystores. If you want to store secrets in an auto-login (or auto-login local) keystore, then note the following:
 - If the software auto-login keystore is in the same location as its corresponding password-based software keystore, then the secrets are added automatically.
 - If the software auto-login keystore is in a different location from its corresponding password-based software keystore, then you must create the auto-login keystore again from the password-based keystore, and keep the two keystores in synchronization.
- **Hardware keystores:** You can store secrets in standard hardware security modules.

Related Topics:

[Storage of Oracle Database Secrets in a Hardware Keystore](#) (page 4-49)

The `ADMINISTER KEY MANAGEMENT ADD SECRET|UPDATE SECRET|DELETE CLIENT` statements can add, update, and delete secrets.

[Configuring Auto-Login Hardware Security Modules](#) (page 4-52)

A hardware security module can be configured to use the auto-login capability.

4.3.2 Storage of Oracle Database Secrets in a Software Keystore

The `ADMINISTER KEY MANAGEMENT ADD SECRET|UPDATE SECRET|DELETE CLIENT` statements can add secrets, update secrets, and delete secrets from a keystore.

As with all of the `ADMINISTER KEY MANAGEMENT` statements, you must have the `ADMINISTER KEY MANAGEMENT` or the `SYSKM` administrative privilege. In a multitenant environment, to add, update or delete secrets from a keystore in-use, you must run the statement in the root.

- **Adding a secret:** Use the following syntax:

```
ADMINISTER KEY MANAGEMENT
ADD SECRET 'secret' FOR CLIENT 'client_identifier'
[USING TAG 'tag']
[TO [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location_in_sqlnet.ora
[IDENTIFIED BY keystore_password]
[WITH BACKUP [USING backup_id]]
| [FORCE KEYSTORE]
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING backup_id]]];
```

- **Updating a secret:** Use the following syntax:

```
ADMINISTER KEY MANAGEMENT
UPDATE SECRET 'secret' FOR CLIENT 'client_identifier' [USING TAG 'tag']
[TO [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location_in_sqlnet.ora
[IDENTIFIED BY keystore_password] [WITH BACKUP [USING backup_id]]
| [FORCE KEYSTORE]
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING backup_id]]];
```

- **Deleting a secret:** Use the following syntax:

```

ADMINISTER KEY MANAGEMENT
DELETE SECRET FOR CLIENT 'client_identifier'
[FROM [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location_in_sqlnet.ora
[IDENTIFIED BY <keystore_password>] [WITH BACKUP [USING backup_id]]
| [FORCE KEYSTORE]
IDENTIFIED BY [EXTERNAL STORE | keystore_password]
[WITH BACKUP [USING backup_id]]];

```

In all of these statements, the specification is as follows:

- *secret* is the client secret key to be stored, updated, or deleted. Enclose this setting in single quotation marks (' ') or omit the quotation marks if the secret has no spaces.
- *client_identifier* is an alphanumeric string used to identify the secret key. *client_identifier* does not have a default value. Enclose this setting in single quotation marks (' ').
- TO [[LOCAL] AUTOLOGIN] KEYSTORE refers to the location of an auto-login keystore, which is specified in the `sqlnet.ora` file.
- *tag* is an optional, user-defined description for the secret key to be stored. You can use *tag* with the ADD and UPDATE operations. Enclose this setting in single quotation marks (' '). This tag appears in the SECRET_TAG column of the V \$CLIENT_SECRETS view..

WITH BACKUP is required in case the keystore was not backed up before the ADD, UPDATE, or DELETE operation. *backup_identifier* is an optional user-defined description for the backup. Enclose *backup_identifier* in single quotation marks (' ').

- [TO | FROM [[LOCAL] AUTOLOGIN] KEYSTORE specifies the location of the auto-login or password keystore, which is set in the `sqlnet.ora` file.
- FORCE KEYSTORE enables the keystore operation if the auto-login keystore is in use, or if the keystore is closed.
- IDENTIFIED BY can be one of the following settings:
 - EXTERNAL STORE uses the keystore password stored in the external store to perform the keystore operation.
 - *keystore_password* is the password for the keystore.

Related Topics:

[Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#)
(page 4-32)

Custom TDE master encryption key attributes enable you to defined attributes that are specific to your needs.

4.3.3 Example: Adding an HSM Password to a Software Keystore

The ADMINISTER KEY MANAGEMENT ADD SECRET statement can add an HSM password to a software keystore.

[Example 4-11](#) (page 4-49) shows how to add a hardware security module (HSM) password as a secret to a software keystore.

Example 4-11 Adding an Oracle Database Secret to a Software Keystore

```
ADMINISTER KEY MANAGEMENT
ADD SECRET 'psmith:password' FOR CLIENT 'HSM_PASSWORD'
USING TAG 'HSM credentials' FORCE KEYSTORE
IDENTIFIED BY password WITH BACKUP;
```

In this version, the keystore password is in an external store, so the EXTERNAL_STORE setting is used for IDENTIFIED BY:

```
ADMINISTER KEY MANAGEMENT
ADD SECRET 'psmith:password' FOR CLIENT 'HSM_PASSWORD'
USING TAG 'HSM credentials' FORCE KEYSTORE
IDENTIFIED BY EXTERNAL_STORE WITH BACKUP;
```

4.3.4 Example: Changing an HSM Password Stored as a Secret in a Software Keystore

The ADMINISTER KEY MANAGEMENT UPDATE SECRET statement can change an HSM password that is stored as a secret in a software keystore.

[Example 4-12](#) (page 4-49) shows how to change an HSM password that is stored as a secret in a software keystore.

Example 4-12 Changing an Oracle Database Secret to a Software Keystore

```
ADMINISTER KEY MANAGEMENT
UPDATE SECRET admin_password FOR CLIENT 'admin@myhost'
USING TAG 'new_host_credentials' FORCE KEYSTORE
IDENTIFIED BY software_keystore_password;
```

In this version, the password for the keystore is in an external store:

```
ADMINISTER KEY MANAGEMENT
UPDATE SECRET admin_password FOR CLIENT 'admin@myhost'
USING TAG 'new_host_credentials' FORCE KEYSTORE
IDENTIFIED BY EXTERNAL_STORE;
```

4.3.5 Example: Deleting an HSM Password Stored as a Secret in a Software Keystore

The ADMINISTER KEY MANAGEMENT DELETE SECRET statement can delete HSM passwords that are stored as secrets in a software keystore.

[Example 4-13](#) (page 4-49) shows how to delete an HSM password that is stored as a secret in the software keystore.

Example 4-13 Deleting an Oracle Database Secret in a Software Keystore

```
ADMINISTER KEY MANAGEMENT
DELETE SECRET FOR CLIENT 'HSM_PASSWORD'
FORCE KEYSTORE
IDENTIFIED BY password WITH BACKUP;
```

In this version, the password for the keystore is in an external store:

```
ADMINISTER KEY MANAGEMENT
DELETE SECRET FOR CLIENT 'HSM_PASSWORD'
FORCE KEYSTORE
IDENTIFIED BY EXTERNAL_STORE WITH BACKUP;
```

4.3.6 Storage of Oracle Database Secrets in a Hardware Keystore

The ADMINISTER KEY MANAGEMENT ADD SECRET|UPDATE SECRET|DELETE CLIENT statements can add, update, and delete secrets.

As with all `ADMINISTER KEY MANAGEMENT` statements, you must have the `ADMINISTER KEY MANAGEMENT` or the `SYSKM` administrative privilege. In a multitenant environment, run the statement in the root. When you add, update, or delete secrets from a keystore that is in use or presently open, then you must run `ADMINISTER KEY MANAGEMENT` in the root.

Note:

Before you attempt to add a secret to a hardware security module, ensure that it has PDCS#11 data object support.

- **Adding a secret:** Use the following syntax:

```
ADMINISTER KEY MANAGEMENT ADD SECRET 'secret'  
FOR CLIENT 'client_identifier' [USING TAG 'tag']  
[TO [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location  
[IDENTIFIED BY "user_id:password"]  
[WITH BACKUP [USING backup_id]];
```

- **Updating a secret:** Use the following syntax:

```
ADMINISTER KEY MANAGEMENT UPDATE SECRET 'secret'  
FOR CLIENT 'client_identifier' [USING TAG 'tag']  
[TO [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location  
[IDENTIFIED BY "user_id:password"]  
[WITH BACKUP [USING backup_id]];
```

- **Deleting a secret:** Use the following syntax:

```
ADMINISTER KEY MANAGEMENT DELETE SECRET  
FOR CLIENT 'client_identifier'  
[FROM [[LOCAL] AUTOLOGIN] KEYSTORE keystore_location  
[IDENTIFIED BY "user_id:password"]  
[WITH BACKUP [USING backup_id]];
```

In all of these statements, the specification as follows:

- *secret* is the client secret key to be stored, updated, or deleted. Enclose this setting in double quotation marks (') or omit the quotation marks if the secret has no spaces.
- *client_identifier* is an alphanumeric string used to identify the secret key. *client_identifier* does not have a default value. Enclose this setting in single quotation marks ('').
- *tag* is an optional, user-defined description for the secret key to be stored. You can use *tag* with the `ADD` and `UPDATE` operations. Enclose this setting in single quotation marks (''). This tag appears in the `SECRET_TAG` column of the `V$CLIENT_SECRETS` view.
- `[TO | FROM [[LOCAL] AUTOLOGIN] KEYSTORE]` specifies the location of the keystore used for the HSM.
- `IDENTIFIED BY` refers to the

Related Topics:

[Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#)
(page 4-32)

Custom TDE master encryption key attributes enable you to define attributes that are specific to your needs.

4.3.7 Example: Adding an Oracle Database Secret to a Hardware Keystore

The `ADMINISTER KEY MANAGEMENT ADD SECRET` statement can add an Oracle Database secret to a hardware keystore.

[Example 4-14](#) (page 4-51) shows how to add a password for a user to a hardware keystore.

Example 4-14 Adding an Oracle Database Secret to a Hardware Keystore

```
ADMINISTER KEY MANAGEMENT ADD SECRET 'password'
FOR CLIENT 'admin@myhost' USING TAG 'myhost admin credentials'
IDENTIFIED BY "psmith:password";
```

In this version, the keystore password is in an external store, so the `EXTERNAL STORE` setting is used for `IDENTIFIED BY`:

```
ADMINISTER KEY MANAGEMENT ADD SECRET 'password'
FOR CLIENT 'admin@myhost' USING TAG 'myhost admin credentials'
IDENTIFIED BY EXTERNAL STORE;
```

4.3.8 Example: Changing an Oracle Database Secret in a Hardware Keystore

The `ADMINISTER KEY MANAGEMENT MANAGEMENT UPDATE SECRET` statement can change an Oracle Database secret in a hardware keystore.

[Example 4-15](#) (page 4-51) shows how to change a password that is stored as a secret in a hardware keystore.

Example 4-15 Changing an Oracle Database Secret in a Hardware Keystore

```
ADMINISTER KEY MANAGEMENT MANAGEMENT UPDATE SECRET 'password2'
FOR CLIENT 'admin@myhost' USING TAG 'New host credentials'
IDENTIFIED BY "psmith:password";
```

In this version, the password for the keystore is in an external store:

```
ADMINISTER KEY MANAGEMENT MANAGEMENT UPDATE SECRET 'password2'
FOR CLIENT 'admin@myhost' USING TAG 'New host credentials'
IDENTIFIED BY EXTERNAL STORE;
```

4.3.9 Example: Deleting an Oracle Database Secret in a Hardware Keystore

The `ADMINISTER KEY MANAGEMENT DELETE SECRET FOR CLIENT` statement can delete an Oracle Database secret that is in a hardware keystore.

[Example 4-16](#) (page 4-51) shows how to delete a hardware security module password that is stored as a secret in the hardware keystore.

Example 4-16 Deleting an Oracle Database Secret in a Hardware Keystore

```
ADMINISTER KEY MANAGEMENT DELETE SECRET FOR CLIENT 'admin@myhost'
IDENTIFIED BY "psmith:password";
```

In this version, the password for the keystore is in an external store:

```
ADMINISTER KEY MANAGEMENT DELETE SECRET FOR CLIENT 'admin@myhost'
IDENTIFIED BY EXTERNAL STORE;
```

4.3.10 Configuring Auto-Login Hardware Security Modules

A hardware security module can be configured to use the auto-login capability.

[About Configuring Auto-Login Hardware Security Modules](#) (page 4-52)

An auto-login hardware security module stores the hardware security module credentials in an auto-login keystore.

[Configuring an Auto-Login Hardware Security Module](#) (page 4-52)

The `ADMINISTER KEY MANAGEMENT` statement configures an auto-login hardware security module.

4.3.10.1 About Configuring Auto-Login Hardware Security Modules

An auto-login hardware security module stores the hardware security module credentials in an auto-login keystore.

This configuration reduces the security of the system as a whole. However, this configuration does support unmanned or automated operations and is useful in deployments where automatic re-login of the hardware security module is necessary.

Be aware that executing the query `SELECT * FROM V$ENCRYPTION_WALLET` will automatically open an auto-login hardware security module. For example, suppose you have an auto-login hardware security module configured. If you close the keystore and query the `V$ENCRYPTION_WALLET` view, then the output will indicate that a keystore is open. This is because `V$ENCRYPTION_WALLET` opened up the auto-login hardware and then displayed the status of the auto-login keystore.

To enable the auto-login capability for a hardware security module, you must store the hardware security module credentials in the hardware keystore.

4.3.10.2 Configuring an Auto-Login Hardware Security Module

The `ADMINISTER KEY MANAGEMENT` statement configures an auto-login hardware security module.

Before you begin this procedure, ensure that you have configured the TDE hardware keystore. See [Configuring a Hardware Keystore](#) (page 3-13).

1. Reconfigure the `sqlnet.ora` file to include the keystore location of the software keystore, if it is not already present.

The software keystore location may already be present if the you have previously migrated to using HSM.

For example:

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=(METHOD=FILE)(METHOD_DATA=
(DIRECTORY=/etc/ORACLE/WALLETS/orcl)))
```

2. Add or update the secret in the software keystore.

The secret is the hardware security module password and the client is the `HSM_PASSWORD`. `HSM_PASSWORD` is an Oracle-defined client name that is used to represent the HSM password as a secret in the software keystore.

For example:

```
ADMINISTER KEY MANAGEMENT ADD SECRET 'user_id:password'
FOR CLIENT 'HSM_PASSWORD'
TO LOCAL AUTOLOGIN KEYSTORE software_keystore_location
WITH BACKUP;
```

In this example, *software_keystore_location* is the location of the software keystore that you just defined.

At this stage, the next time that a TDE operation executes, the hardware security module auto-login keystore opens automatically.

Related Topics:

[About the Keystore Location in the sqlnet.ora File](#) (page 3-3)

Oracle Database checks the `sqlnet.ora` file for the directory location of the keystore, whether it is a software keystore, a hardware module security (HSM) keystore, or an Oracle Key Vault keystore.

4.4 Storing Oracle GoldenGate Secrets in a Keystore

You can store Oracle GoldenGate secrets in Transparent Data Encryption keystores.

[About Storing Oracle GoldenGate Secrets in Keystores](#) (page 4-53)

You can use a keystore to store secret keys for tools and external clients such as Oracle GoldenGate.

[Oracle GoldenGate Extract Classic Capture Mode TDE Requirements](#) (page 4-54)

Ensure that you meet the requirements for Oracle GoldenGate Extract to support Transparent Data Encryption capture.

[Configuring TDE Keystore Support for Oracle GoldenGate](#) (page 4-54)

You can configure Transparent Data Encryption keystore support for Oracle GoldenGate by using a shared secret for the keystore.

4.4.1 About Storing Oracle GoldenGate Secrets in Keystores

You can use a keystore to store secret keys for tools and external clients such as Oracle GoldenGate.

The secret key must be a string adhering to Oracle identifier rules. You can add, update, or delete a client secret in an existing keystore. This section describes how to capture Transparent Data Encryption encrypted data in the Oracle GoldenGate Extract (Extract) process using classic capture mode.

TDE support when Extract is in classic capture mode requires the exchange of the following keys:

- TDE support for Oracle GoldenGate in the classic capture mode of the Extract process requires that an Oracle database and the Extract process share the secret to encrypt sensitive information being exchanged. The shared secret is stored securely in the Oracle database and Oracle GoldenGate domains. The shared secret is stored in the software keystore or the HSM as the database secret.
- The decryption key is a password known as the shared secret that is stored securely in the Oracle database and Oracle GoldenGate domains. Only a party that has possession of the shared secret can decrypt the table and redo log keys.

After you configure the shared secret, Oracle GoldenGate Extract uses the shared secret to decrypt the data. Oracle GoldenGate Extract does not handle the TDE master

encryption key itself, nor is it aware of the keystore password. The TDE master encryption key and password remain within the Oracle database configuration.

Oracle GoldenGate Extract only writes the decrypted data to the Oracle GoldenGate trail file, which Oracle GoldenGate persists during transit. You can protect this file using your site's operating system standard security protocols, as well as the Oracle GoldenGate AES encryption options. Oracle GoldenGate does not write the encrypted data to a discard file (specified with the `DISCARDFILE` parameter). The word `ENCRYPTED` will be written to any discard file that is in use.

Oracle GoldenGate does require that the keystore be open when processing encrypted data. There is no performance effect of Oracle GoldenGate feature on the TDE operations.

4.4.2 Oracle GoldenGate Extract Classic Capture Mode TDE Requirements

Ensure that you meet the requirements for Oracle GoldenGate Extract to support Transparent Data Encryption capture.

The requirements are as follows:

- To maintain high security standards, ensure that the Oracle GoldenGate Extract process runs as part of the Oracle user (the user that runs the Oracle database). That way, the keys are protected in memory by the same privileges as the Oracle user.
- Run the Oracle GoldenGate Extract process on the same computer as the Oracle database installation.

4.4.3 Configuring TDE Keystore Support for Oracle GoldenGate

You can configure Transparent Data Encryption keystore support for Oracle GoldenGate by using a shared secret for the keystore.

[Step 1: Decide on a Shared Secret for the Keystore](#) (page 4-54)

A shared secret for a keystore is a password.

[Step 2: Configure Oracle Database for TDE Support for Oracle GoldenGate](#) (page 4-55)

The `DBMS_INTERNAL_CLKM` PL/SQL package enables you to configure TDE support for Oracle GoldenGate.

[Step 3: Store the TDE GoldenGate Shared Secret in the Keystore](#) (page 4-55)

The `ADMINISTER KEY MANAGEMENT` statement can store a TDE GoldenGate shared secret in a keystore.

[Step 4: Set the TDE Oracle GoldenGate Shared Secret in the Extract Process](#) (page 4-57)

The GoldenGate Software Command Interface (GGSCI) utility set the TDE Oracle GoldenGate shared secret in the extract process.

4.4.3.1 Step 1: Decide on a Shared Secret for the Keystore

A shared secret for a keystore is a password.

- Decide on a shared secret that meets or exceeds Oracle Database password standards.

Do not share this password with any user other than trusted administrators who are responsible for configuring Transparent Data Encryption to work with Oracle GoldenGate Extract.

See Also:

Oracle Database Security Guide for guidelines on creating secure passwords

4.4.3.2 Step 2: Configure Oracle Database for TDE Support for Oracle GoldenGate

The `DBMS_INTERNAL_CLKM` PL/SQL package enables you to configure TDE support for Oracle GoldenGate.

1. Log in to the database instance as user `SYS` with the `SYSDBA` administrative privilege.

For example

```
sqlplus sys as sysdba
Enter password: password
Connected.
```

2. In a multitenant environment, connect to the appropriate PDB.

For example:

```
CONNECT SYS@hrpdb AS SYSDBA
Enter password: password
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

3. Load the Oracle Database-supplied `DBMS_INTERNAL_CLKM` PL/SQL package.

For example:

```
@?/app/oracle/product/12.2/rdbms/admin/prvtclkm.plb
```

The `prvtclkm.plb` file also enables Oracle GoldenGate to extract encrypted data from an Oracle database.

4. Grant the `EXECUTE` privilege on the `DBMS_INTERNAL_CLKM` PL/SQL package to the Oracle GoldenGate Extract database user.

For example:

```
GRANT EXECUTE ON DBMS_INTERNAL_CLKM TO psmith;
```

This procedure enables the Oracle database and Oracle GoldenGate Extract to exchange information.

5. Exit SQL*Plus.

4.4.3.3 Step 3: Store the TDE GoldenGate Shared Secret in the Keystore

The `ADMINISTER KEY MANAGEMENT` statement can store a TDE GoldenGate shared secret in a keystore.

Before you begin this procedure, ensure that you have configured the TDE software or hardware keystore, based on [Configuring a Software Keystore](#) (page 3-1) and [Configuring a Hardware Keystore](#) (page 3-13).

1. Set the Oracle GoldenGate-Transparent Data Encryption key in the keystore.

The syntax is as follows:

```
ADMINISTER KEY MANAGEMENT ADD|UPDATE|DELETE SECRET 'secret'
FOR CLIENT 'secret_identifier' [USING TAG 'tag']
IDENTIFIED BY keystore_password [WITH BACKUP [USING 'backup_identifier']];
```

In this specification:

- *secret* is the client secret key to be stored, updated, or deleted. Enclose this setting in single quotation marks (' ').
- *secret_identifier* is an alphanumeric string used to identify the secret key. *secret_identifier* does not have a default value. Enclose this setting in single quotation marks (' ').
- *tag* is an optional, user-defined description for the secret key to be stored. *tag* can be used with the ADD and UPDATE operations. Enclose this setting in single quotation marks (' '). This tag appears in the SECRET_TAG column of the V\$CLIENT_SECRETS view. [Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#) (page 4-32)
- *keystore_password* is the password for the keystore that is configured.
- WITH BACKUP is required in case the keystore was not backed up before the ADD, UPDATE or DELETE operation. *backup_identifier* is an optional user-defined description for the backup. Enclose *backup_identifier* in single quotation marks (' ').

The following example adds a secret key to the keystore and creates a backup in the same directory as the keystore:

```
ADMINISTER KEY MANAGEMENT ADD SECRET 'some_secret'
FOR CLIENT 'ORACLE_GG' USING TAG 'GoldenGate Secret'
IDENTIFIED BY password WITH BACKUP USING 'GG backup';
```

2. Verify the entry that you just created.

For example:

```
SELECT CLIENT, SECRET_TAG FROM V$CLIENT_SECRETS WHERE CLIENT = 'ORACLEGG';
```

```
CLIENT    SECRET_TAG
-----
ORACLEGG  some_secret
```

3. Switch the log files.

```
CONNECT / AS SYSDBA

ALTER SYSTEM SWITCH LOGFILE;
```


See Also:

- [Creating Custom TDE Master Encryption Key Attributes for Reporting Purposes](#) (page 4-32) for more information about tags
- *Oracle Database Administrator's Guide* for more information about switching log files
- [How Transparent Data Encryption Works with Oracle Real Application Clusters](#) (page 6-5) if you are having problems using this procedure in an Oracle Real Application Clusters environment

4.4.3.4 Step 4: Set the TDE Oracle GoldenGate Shared Secret in the Extract Process

The GoldenGate Software Command Interface (GGSCI) utility set the TDE Oracle GoldenGate shared secret in the extract process.

1. Start the GGSCI utility.

For example:

```
ggsci
```

2. In the GGSCI utility, run the `ENCRYPT PASSWORD` command to encrypt the shared secret so that it is obfuscated within the Oracle GoldenGate Extract parameter file.

```
ENCRYPT PASSWORD shared_secret algorithm ENCRYPTKEY keyname
```

In this specification:

- *shared_secret* is the clear-text shared secret that you created when you decided on a shared secret for the keystore. This setting is case sensitive.
- *algorithm* is one of the following values to specify AES encryption:
 - AES128
 - AES192
 - AES256
- *keyname* is the logical name of the encryption key in the `ENCKEYS` lookup file. Oracle GoldenGate uses this name to look up the actual key in the `ENCKEYS` file.

For example:

```
ENCRYPT PASSWORD password AES256 ENCRYPTKEY mykey1
```

3. In the Oracle GoldenGate Extract parameter file, set the `DBOPTIONS` parameter with the `DECRYPTPASSWORD` option.

As input, supply the encrypted shared secret and the Oracle GoldenGate-generated or user-defined decryption key.

```
DBOPTIONS DECRYPTPASSWORD shared_secret algorithm ENCRYPTKEY keyname
```

In this specification:

- *shared_secret* is the clear-text shared secret that you created when you decided on a shared secret for the keystore. This setting is case sensitive.
- *algorithm* is one of the following values to specify AES encryption:
 - AES128
 - AES192
 - AES256
- *keyname* is the logical name of the encryption key in the ENCKEYS lookup file.

For example:

```
DBOPTIONS DECRYPTPASSWORD AACAAAAAAAAAAIALCKDZIRHOJBHOJUH AES256  
ENCRYPTKEY mykey1
```

Related Topics:

[Step 1: Decide on a Shared Secret for the Keystore](#) (page 4-54)
A shared secret for a keystore is a password.

General Considerations of Using Transparent Data Encryption

When you use Transparent Data Encryption, you should consider factors such as security, performance, and storage overheads.

[Compression and Data Deduplication of Encrypted Data](#) (page 5-1)

With tablespace encryption, Oracle Database compresses tables and indexes before encrypting the tablespace.

[Security Considerations for Transparent Data Encryption](#) (page 5-2)

As with all Oracle Database features, you should consider security when you create TDE policies.

[Performance and Storage Overhead of Transparent Data Encryption](#) (page 5-3)

The performance of Transparent Data Encryption can vary.

[Modifying Your Applications for Use with Transparent Data Encryption](#) (page 5-5)

You can modify your applications to use Transparent Data Encryption.

[How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT](#) (page 5-6)

Many of the clauses from the ALTER SYSTEM statement correspond to the ADMINISTER KEY MANAGEMENT statement.

[Using Transparent Data Encryption with PKI Encryption](#) (page 5-9)

PKI encryption is deprecated, but if you are still using it, then there are several issues you must consider.

[Data Loads from External Files to Tables with Encrypted Columns](#) (page 5-11)

You can use SQL*Loader to perform data loads from files to tables that have encrypted columns.

[Transparent Data Encryption and Database Close Operations](#) (page 5-11)

You should ensure that the software or hardware keystore is open before you close the database.

5.1 Compression and Data Deduplication of Encrypted Data

With tablespace encryption, Oracle Database compresses tables and indexes before encrypting the tablespace.

This ensures that you receive the maximum space and performance benefits from compression, while also receiving the security of encryption at rest. In the CREATE TABLESPACE SQL statement, include both the COMPRESS and ENCRYPT clauses.

With column encryption, Oracle Database compresses the data after it encrypts the column. This means that compression will have minimal effectiveness on encrypted

columns. There is one notable exception: if the column is a SecureFiles LOB, and the encryption is implemented with SecureFiles LOB Encryption, and the compression (and possibly deduplication) are implemented with SecureFiles LOB Compression & Deduplication, then compression is performed before encryption. Similar to the `CREATE TABLESPACE` statement for tablespace encryption, include both the `COMPRESS` and `ENCRYPT` clauses.

See Also:

- *Oracle Database Backup and Recovery User's Guide* for more information about the Advanced Compression Option
 - *Oracle Database SecureFiles and Large Objects Developer's Guide* for information about SecureFiles LOB storage
 - *Oracle Database SecureFiles and Large Objects Developer's Guide* for information about SecureFiles Compression
-

5.2 Security Considerations for Transparent Data Encryption

As with all Oracle Database features, you should consider security when you create TDE policies.

[Transparent Data Encryption General Security Advice](#) (page 5-2)

Security considerations for Transparent Data Encryption (TDE) operate within the broader area of total system security.

[Transparent Data Encryption Column Encryption-Specific Advice](#) (page 5-3)

Additional security considerations apply to normal database and network operations when using TDE.

[Managing Security for Plaintext Fragments](#) (page 5-3)

You should remove old plaintext fragments that can appear over time.

5.2.1 Transparent Data Encryption General Security Advice

Security considerations for Transparent Data Encryption (TDE) operate within the broader area of total system security.

Follow these general guidelines:

- Identify the degrees of sensitivity of data in your database, the protection that they need, and the levels of risk to be addressed. For example, highly sensitive data requiring stronger protection can be encrypted with the AES256 algorithm. A database that is not as sensitive can be protected with no `salt` or the `nomac` option to enable performance benefits.
- Evaluate the costs and benefits that are acceptable to data and keystore protection. Protection of keys determines the type of keystore to be used: auto-login software keystores, password-based software keystores, or hardware keystores.
- Consider having separate security administrators for TDE and for the database.
- Consider having a separate and exclusive keystore for TDE.
- Implement protected back-up procedures for your encrypted data.

5.2.2 Transparent Data Encryption Column Encryption-Specific Advice

Additional security considerations apply to normal database and network operations when using TDE.

Encrypted column data stays encrypted in the data files, undo logs, redo logs, and the buffer cache of the system global area (SGA). However, data is decrypted during expression evaluation, making it possible for decrypted data to appear in the swap file on the disk. Privileged operating system users can potentially view this data.

Column values encrypted using TDE are stored in the data files in encrypted form. However, these data files may still contain some [plaintext](#) fragments, called ghost copies, left over by past data operations on the table. This is similar to finding data on the disk after a file was deleted by the operating system.

5.2.3 Managing Security for Plaintext Fragments

You should remove old plaintext fragments that can appear over time.

Old [plaintext](#) fragments may be present for some time until the database overwrites the blocks containing such values. If privileged operating system users bypass the access controls of the database, then they might be able to directly access these values in the data file holding the tablespace.

To minimize this risk:

1. Create a new tablespace in a new data file.

You can use the `CREATE TABLESPACE` statement to create this tablespace.

2. Move the table containing encrypted columns to the new tablespace. You can use the `ALTER TABLE MOVE` statement.

Repeat this step for all of the objects in the original tablespace.

3. Drop the original tablespace.

You can use the `DROP TABLESPACE tablespace INCLUDING CONTENTS KEEP DATAFILES` statement. Oracle recommends that you securely delete data files using platform-specific utilities.

4. Use platform-specific and file system-specific utilities to securely delete the old data file. Examples of such utilities include `shred` (on Linux) and `sdelete` (on Windows).

5.3 Performance and Storage Overhead of Transparent Data Encryption

The performance of Transparent Data Encryption can vary.

[Performance Overhead of Transparent Data Encryption](#) (page 5-4)

Transparent Data Encryption tablespace encryption has small associated performance overhead.

[Storage Overhead of Transparent Data Encryption](#) (page 5-4)

TDE tablespace encryption has no storage overhead, but TDE column encryption has some associated storage overhead.

5.3.1 Performance Overhead of Transparent Data Encryption

Transparent Data Encryption tablespace encryption has small associated performance overhead.

The actual performance impact on applications can vary. TDE column encryption affects performance only when data is retrieved from or inserted into an encrypted column. No reduction in performance occurs for operations involving unencrypted columns, even if these columns are in a table containing encrypted columns. Accessing data in encrypted columns involves small performance overhead, and the exact overhead you observe can vary.

The total performance overhead depends on the number of encrypted columns and their frequency of access. The columns most appropriate for encryption are those containing the most sensitive data.

Enabling encryption on an existing table results in a full table update like any other `ALTER TABLE` operation that modifies table characteristics. Keep in mind the potential performance and redo log impact on the database server before enabling encryption on a large existing table.

A table can temporarily become inaccessible for write operations while encryption is being enabled, [TDE table keys](#) are being rekeyed, or the encryption algorithm is being changed. You can use online table redefinition to ensure that the table is available for write operations during such procedures.

If you enable TDE column encryption on a very large table, then you may need to increase the redo log size to accommodate the operation.

Encrypting an indexed column takes more time than encrypting a column without indexes. If you must encrypt a column that has an index built on it, you can try dropping the index, encrypting the column with `NO SALT`, and then re-creating the index.

If you index an encrypted column, then the index is created on the encrypted values. When you query for a value in the encrypted column, Oracle Database transparently encrypts the value used in the SQL query. It then performs an index lookup using the encrypted value.

Note:

If you must perform range scans over indexed, encrypted columns, then use TDE tablespace encryption in place of TDE column encryption.

See Also:

- [Creating an Encrypted Column in an External Table](#) (page 3-25)
 - *Oracle Database Administrator's Guide* for information about redefining tables online
-
-

5.3.2 Storage Overhead of Transparent Data Encryption

TDE tablespace encryption has no storage overhead, but TDE column encryption has some associated storage overhead.

Encrypted column data must have more storage space than [plaintext](#) data. In addition, TDE pads out encrypted values to multiples of 16 bytes. This means that if a credit card number requires nine bytes for storage, then an encrypted credit card value will require an additional seven bytes.

Each encrypted value is also associated with a 20-byte integrity check. This does not apply if you have encrypted columns using the `NOMAC` parameter. If data was encrypted with `salt`, then each encrypted value requires an additional 16 bytes of storage.

The maximum storage overhead for each encrypted value is from one to 52 bytes.

Related Topics:

[Creating an Encrypted Column in an External Table](#) (page 3-25)

The external table feature enables you to access data in external sources as if the data were in a database table.

5.4 Modifying Your Applications for Use with Transparent Data Encryption

You can modify your applications to use Transparent Data Encryption.

1. Configure the software or hardware keystore for TDE, and then set the master encryption key.

See the following sections for more information:

- [Configuring a Software Keystore](#) (page 3-1)
- [Configuring a Hardware Keystore](#) (page 3-13)

2. Verify that the master encryption key was created by querying the `KEY_ID` column of the `V$ENCRYPTION_KEYS` view.

3. Identify the sensitive columns (such as those containing credit card data) that require Transparent Data Encryption protection.

4. Decide whether to use TDE column encryption or TDE tablespace encryption.

See the following sections for more information:

- [How Transparent Data Encryption Column Encryption Works](#) (page 2-3)
- [How Transparent Data Encryption Tablespace Encryption Works](#) (page 2-4)

5. Open the keystore.

See the following sections for more information:

- [Step 3: Open the Software Keystore](#) (page 3-9)
- [Step 3: Open the Hardware Keystore](#) (page 3-15)

6. Encrypt the columns or tablespaces.

See the following sections for more information:

- [Encrypting Columns in Tables](#) (page 3-19)
- [Encryption Conversions for Tablespaces and Databases](#) (page 3-29)

5.5 How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT

Many of the clauses from the ALTER SYSTEM statement correspond to the ADMINISTER KEY MANAGEMENT statement.

Table 5-1 (page 5-6) compares the Transparent Data Encryption usage of the ALTER SYSTEM statement and the orapki utility from previous releases with the ADMINISTER KEY MANAGEMENT statement.

Table 5-1 How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT

Behavior	ALTER SYSTEM or orapki	ADMINISTER KEY MANAGEMENT
Creating a keystore	<p>For software keystores (called wallets in previous releases):</p> <pre>ALTER SYSTEM SET ENCRYPTION KEY ["certificate_ID"] IDENTIFIED BY keystore_password;</pre> <p>For hardware keystores, the keystore is available after you configure the hardware security module.</p>	<p>For software keystores:</p> <pre>ADMINISTER KEY MANAGEMENT CREATE KEYSTORE 'keystore_location' IDENTIFIED BY software_keystore_password</pre> <p>For hardware keystores, the keystore is available after you configure the hardware security module.</p>
Creating an auto-login keystore	<pre>orapki wallet create -wallet wallet_location -auto_login [-pwd password]</pre>	<p>For software keystores:</p> <pre>ADMINISTER KEY MANAGEMENT CREATE [LOCAL] AUTO_LOGIN KEYSTORE FROM KEYSTORE 'keystore_location' IDENTIFIED BY software_keystore_password;</pre> <p>This type of keystore applies to software keystores only.</p>
Opening a keystore	<pre>ALTER SYSTEM SET ENCRYPTION WALLET OPEN IDENTIFIED BY password;</pre>	<pre>ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY keystore_password [CONTAINER = ALL CURRENT];</pre>
Closing a keystore	<pre>ALTER SYSTEM SET ENCRYPTION WALLET CLOSE IDENTIFIED BY password;</pre>	<p>For both software and hardware keystores:</p> <pre>ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY keystore_password [CONTAINER = ALL CURRENT];</pre>
Migrating from a hardware keystore to a software keystore	Not available	<pre>ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY IDENTIFIED BY software_keystore_password REVERSE MIGRATE USING "user_id:password" [WITH BACKUP [USING 'backup_identifier']];</pre>

Table 5-1 (Cont.) How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT

Behavior	ALTER SYSTEM or orapki	ADMINISTER KEY MANAGEMENT
Migrating from a software keystore to a hardware keystore	ALTER SYSTEM SET ENCRYPTION KEY IDENTIFIED BY "user_id:password" MIGRATE USING wallet_password;	ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY IDENTIFIED BY "user_id:password" MIGRATE USING software_keystore_password;
Changing a keystore password	orapki wallet change_pwd -wallet wallet_location [-oldpwd password] [-newpwd password]	For password-based software keystores: ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD IDENTIFIED BY software_keystore_old_password SET software_keystore_new_password [WITH BACKUP [USING 'backup_identifier']]; For hardware keystores, you close the keystore, change it in the hardware security module interface, and then reopen the keystore.
Backing up a password-based software keystore	Not available	ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE [USING 'backup_identifier'] IDENTIFIED BY software_keystore_password [TO 'keystore_location'];
Merging two software keystores into a third new keystore	Not available	ADMINISTER KEY MANAGEMENT MERGE KEYSTORE 'keystore1_location' [IDENTIFIED BY software_keystore1_password] AND KEYSTORE 'keystore2_location' [IDENTIFIED BY software_keystore2_password] INTO NEW KEYSTORE 'keystore3_location' IDENTIFIED BY software_keystore3_password;
Merging one software keystore into another existing keystore	Not available	ADMINISTER KEY MANAGEMENT MERGE KEYSTORE 'keystore1_location' [IDENTIFIED BY software_keystore1_password] INTO EXISTING KEYSTORE 'keystore2_location' IDENTIFIED BY software_keystore2_password [WITH BACKUP [USING 'backup_identifier']];

Table 5-1 (Cont.) How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT

Behavior	ALTER SYSTEM or orapki	ADMINISTER KEY MANAGEMENT
Setting or rotating the master encryption key	<p>For software wallets:</p> <pre>ALTER SYSTEM SET ENCRYPTION KEY ["certificate_ID"] IDENTIFIED BY keystore_password;</pre> <p>For hardware security modules:</p> <pre>ALTER SYSTEM SET ENCRYPTION KEY IDENTIFIED BY "user_id:password"</pre> <p>Note: The ALTER SYSTEM SET ENCRYPTION KEY statement does not update the V \$ENCRYPTION_KEYS dynamic view after you rotate the encryption key.</p>	<pre>ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY [USING TAG 'tag'] IDENTIFIED BY keystore_password WITH BACKUP [USING 'backup_identifier'] [CONTAINER = ALL CURRENT];</pre> <p>After you rotate the encryption key, the V \$ENCRYPTION_KEYS dynamic view is updated.</p>
Creating a master encryption key for later user	Not available	<pre>ADMINISTER KEY MANAGEMENT CREATE KEY [USING TAG 'tag'] IDENTIFIED BY keystore_password [WITH BACKUP [USING 'backup_identifier']] [CONTAINER = (ALL CURRENT)];</pre>
Activating a master encryption key	Not available	<pre>ADMINISTER KEY MANAGEMENT USE KEY 'key_identifier' [USING TAG 'tag'] IDENTIFIED BY keystore_password [WITH BACKUP [USING 'backup_identifier']];</pre>
Creating custom tags for master encryption keys	Not available	<pre>ADMINISTER KEY MANAGEMENT SET TAG 'tag' FOR 'master_key_identifier' IDENTIFIED BY keystore_password [WITH BACKUP [USING 'backup_identifier']];</pre>
Exporting a master encryption key	Not available	<pre>ADMINISTER KEY MANAGEMENT EXPORT [ENCRYPTION] KEYS WITH SECRET "export_secret" TO 'file_path' IDENTIFIED BY software_keystore_password [WITH IDENTIFIER IN 'key_id1', 'key_id2', 'key_idn' (SQL_query)]</pre>

Table 5-1 (Cont.) How ALTER SYSTEM and orapki Map to ADMINISTER KEY MANAGEMENT

Behavior	ALTER SYSTEM or orapki	ADMINISTER KEY MANAGEMENT
Importing a master encryption key	Not available	<pre>ADMINISTER KEY MANAGEMENT IMPORT [ENCRYPTION] KEYS WITH SECRET "import_secret" FROM 'file_name' IDENTIFIED BY software_keystore_password [WITH BACKUP [USING 'backup_identifier'];</pre>
Storing Oracle Database secrets in a keystore	Not available	<p>For software keystores:</p> <pre>ADMINISTER KEY MANAGEMENT ADD SECRET UPDATE SECRET DELETE SECRET "secret" FOR CLIENT 'client_identifier' [USING TAG 'tag'] IDENTIFIED BY keystore_password [WITH BACKUP [USING 'backup_identifier'];</pre> <p>For hardware keystores:</p> <pre>ADMINISTER KEY MANAGEMENT ADD SECRET UPDATE SECRET DELETE SECRET "secret" FOR CLIENT 'client_identifier' [USING TAG 'tag'] IDENTIFIED BY "user_id:password" [WITH BACKUP [USING 'backup_identifier'];</pre>

5.6 Using Transparent Data Encryption with PKI Encryption

PKI encryption is deprecated, but if you are still using it, then there are several issues you must consider.

Note:

The use of PKI encryption with Transparent Data Encryption is deprecated. To configure Transparent Data Encryption, use the ADMINISTER KEY MANAGEMENT SQL statement.

[Software Master Encryption Key Use with PKI Key Pairs](#) (page 5-10)

A master encryption key can be an existing key pair from a PKI certificate designated for encryption.

[TDE Tablespace and Hardware Keystores with PKI Encryption](#) (page 5-10)

PKI encryption is a cryptographic system that uses two keys, a public key and a private key, to encrypt data.

[Backup and Recovery of a PKI Key Pair](#) (page 5-10)

For software keystores, Transparent Data Encryption supports the use of PKI asymmetric key pairs as master encryption keys for column encryption.

5.6.1 Software Master Encryption Key Use with PKI Key Pairs

A master encryption key can be an existing key pair from a PKI certificate designated for encryption.

Note the following:

- If you have already deployed PKI in your organization, then you can use PKI services such as key escrow and recovery. However, encryption using current PKI algorithms requires significantly more system resources than symmetric key encryption. Using a PKI key pair as a master encryption key may result in greater performance degradation when accessing encrypted columns in the database.
- For PKI-based keys, certificate revocation lists are not enforced because enforcing certificate revocation may lead to losing access to all of the encrypted information in the database. However, you cannot use the same certificate to create the master encryption key again.

5.6.2 TDE Tablespace and Hardware Keystores with PKI Encryption

PKI encryption is a cryptographic system that uses two keys, a public key and a private key, to encrypt data.

You cannot use PKI-based encryption with TDE tablespace encryption or with hardware keystores.

5.6.3 Backup and Recovery of a PKI Key Pair

For software keystores, Transparent Data Encryption supports the use of PKI asymmetric key pairs as master encryption keys for column encryption.

This enables the database to use existing key backup, escrow, and recovery facilities from leading certificate authority vendors.

In current key escrow or recovery systems, the certificate authority with key recovery capabilities typically stores a version of the private key, or a piece of information that helps recover the private key. If the private key is lost, then you can recover the original key and certificate by contacting the certificate authority and initiating a key recovery process.

Typically, the key recovery process is automated and requires the user to present certain authenticating credentials to the certificate authority. TDE puts no restrictions on the key recovery process other than that the recovered key and its associated certificate be a PKCS#12 file that can be imported into an keystore. This requirement is consistent with the key recovery mechanisms of leading certificate authorities.

After obtaining the PKCS#12 file with the original certificate and private key, you must create an empty keystore in the same location as the previous keystore. You can then import the PKCS#12 file into the new keystore by using the same utility. Choose a strong password to protect the keystore.

After you use the `ADMINISTER KEY MANAGEMENT` statements to create the keystore and import the correct encryption keys, log in to the database and run the following `ALTER SYSTEM` statement at the SQL prompt to complete the recovery process:

```
ALTER SYSTEM SET ENCRYPTION KEY "cert_id" IDENTIFIED BY keystore_password;
```

In this specification:

- *cert_id* is the certificate ID of the certificate to be used as the master encryption key.
- *keystore_password* is a password that you create.

Note:

You must use the ALTER SYSTEM statement to regenerate encryption keys for PKI key pairs only. This restriction does not apply to non-PKI encryption keys.

5.7 Data Loads from External Files to Tables with Encrypted Columns

You can use SQL*Loader to perform data loads from files to tables that have encrypted columns.

Be aware that with SQL*Loader, you cannot include the ENCRYPT clause in the column definition of an external table of the type ORACLE_LOADER, but you can include it in the column definitions of external tables of type ORACLE_DATAPUMP.

- External tables of type ORACLE_LOADER

The reason that you cannot include the ENCRYPT clause in the column definitions of external tables of the type ORACLE_LOADER is because the contents of an external table with the ORACLE_LOADER type must come from a user-specified plaintext "backing file," and such plaintext files cannot contain any TDE encrypted data.

If you use the ENCRYPT clause in the definition of an external table of type ORACLE_LOADER, then when you query the TDE-encrypted column in this external table, the query fails. This is because TDE expects the external data to have been encrypted, and automatically tries to decrypt it on load. This action fails because the "backing file" only contains plaintext.

- External tables of type ORACLE_DATAPUMP

You can use TDE column encryption with external tables of type ORACLE_DATAPUMP. This is because for external tables of ORACLE_DATAPUMP type, the "backing file" is always created by Oracle Database (during an unload operation) and thus does have support for being populated with encrypted data.

5.8 Transparent Data Encryption and Database Close Operations

You should ensure that the software or hardware keystore is open before you close the database.

The master keys may be required during the database close operation. The database close operation automatically closes the software or hardware keystore.

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[Step 3: Open the Hardware Keystore](#) (page 3-15)

After you have configured the hardware security module, you must open the hardware keystore before it can be used.

Using Transparent Data Encryption with Other Oracle Features

You can use Oracle Data Encryption with other Oracle features, such as Oracle Data Guard or Oracle Real Application Clusters.

[How Transparent Data Encryption Works with Export and Import Operations](#) (page 6-1)

Oracle Data Pump can export and import tables that contain encrypted columns, as well as encrypt entire dump sets.

[How Transparent Data Encryption Works with Oracle Data Guard](#) (page 6-4)

For both software keystores and hardware keystores, Oracle Data Guard supports Transparent Data Encryption (TDE).

[How Transparent Data Encryption Works with Oracle Real Application Clusters](#) (page 6-5)

Oracle Real Application Clusters nodes can share both software and keystores and hardware security modules.

[How Transparent Data Encryption Works with SecureFiles](#) (page 6-7)

SecureFiles, which stores LOBS, has three features: compression, deduplication, and encryption.

[How Transparent Data Encryption Works in a Multitenant Environment](#) (page 6-8)

In a multitenant environment, the TDE operations that you can perform depend on whether you are in the root or a PDB.

[How Transparent Data Encryption Works with Oracle Call Interface](#) (page 6-20)

Transparent Data Encryption does not have any effect on the operation of Oracle Call Interface (OCI).

[How Transparent Data Encryption Works with Editions](#) (page 6-20)

Transparent Data Encryption does not have any effect on the Editions feature of Oracle Database.

[Configuring Transparent Data Encryption to Work in a Multidatabase Environment](#) (page 6-21)

Each Oracle database on the same server (such as databases sharing the same Oracle binary but using different data files) must access its own TDE keystore.

6.1 How Transparent Data Encryption Works with Export and Import Operations

Oracle Data Pump can export and import tables that contain encrypted columns, as well as encrypt entire dump sets.

[About Exporting and Importing Encrypted Data](#) (page 6-2)

You can use Oracle Data Pump to export and import tables that have encrypted columns.

[Exporting and Importing Tables with Encrypted Columns](#) (page 6-2)

You can export and import tables with encrypted columns using the `ENCRYPTION=ENCRYPTED_COLUMNS_ONLY` setting.

[Using Oracle Data Pump to Encrypt Entire Dump Sets](#) (page 6-3)

Oracle Data Pump can encrypt entire dump sets, not just Transparent Data Encryption columns.

6.1.1 About Exporting and Importing Encrypted Data

You can use Oracle Data Pump to export and import tables that have encrypted columns.

For both software and hardware keystores, the following points are important when you must export tables containing encrypted columns:

- Sensitive data should remain unintelligible during transport.
- Authorized users should be able to decrypt the data after it is imported at the destination.

When you use Oracle Data Pump to export and import tables containing encrypted columns, it uses the `ENCRYPTION` parameter to enable encryption of data in dump file sets. The `ENCRYPTION` parameter allows the following values:

- `ENCRYPTED_COLUMNS_ONLY`: Writes encrypted columns to the dump file set in encrypted format
- `DATA_ONLY`: Writes all of the data to the dump file set in encrypted format
- `METADATA_ONLY`: Writes all of the metadata to the dump file set in encrypted format
- `ALL`: Writes all of the data and metadata to the dump file set in encrypted format
- `NONE`: Does not use encryption for dump file sets

6.1.2 Exporting and Importing Tables with Encrypted Columns

You can export and import tables with encrypted columns using the `ENCRYPTION=ENCRYPTED_COLUMNS_ONLY` setting.

1. Ensure that the keystore is open before you attempt to export tables containing encrypted columns.

In a multitenant environment, if you are exporting data in a pluggable database (PDB), then ensure that the wallet is open in the PDB. If you are exporting into the root, then ensure that the wallet is open in the root.

To find if the keystore is open, query the `STATUS` column of the `V_$ENCRYPTION_WALLET` view. If you must open the keystore, then run the following SQL statement:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN
IDENTIFIED BY software_keystore_password
[CONTAINER = ALL | CURRENT];
```


The `software_keystore_password` setting is the password for the keystore. The keystore must be open because the encrypted columns must be decrypted using the [TDE table keys](#), which requires access to the TDE master encryption key. The columns are reencrypted using a password, before they are exported.

2. Run the EXPDP command, using the `ENCRYPTION_PASSWORD` parameter to specify a password that is used to encrypt column data in the export dump file set.

The following example exports the `employee_data` table. The `ENCRYPTION_PWD_PROMPT = YES` setting enables you to prompt for the password interactively, which is a recommended security practice.

```
expdp hr TABLES=employee_data DIRECTORY=dpump_dir
DUMPFILE=dpcd2be1.dmp ENCRYPTION=ENCRYPTED_COLUMNS_ONLY
ENCRYPTION_PWD_PROMPT = YES
```

Password: `password_for_hr`

3. To import the exported data into the target database, ensure that you specify the same password that you used for the export operation, as set by the `ENCRYPTION_PASSWORD` parameter.

The password is used to decrypt the data. Data is reencrypted with the new TDE table keys generated in the target database. The target database must have the keystore open to access the TDE master encryption key. The following example imports the `employee_data` table:

```
impdp hr TABLES=employee_data DIRECTORY=dpump_dir
DUMPFILE=dpcd2be1.dmp
ENCRYPTION_PWD_PROMPT = YES
```

Password: `password_for_hr`

6.1.3 Using Oracle Data Pump to Encrypt Entire Dump Sets

Oracle Data Pump can encrypt entire dump sets, not just Transparent Data Encryption columns.

While importing, you can use either the password or the keystore TDE master encryption key to decrypt the data. If the password is not supplied, then the TDE master encryption key in the keystore is used to decrypt the data. The keystore must be present and open at the target database. The open keystore is also required to reencrypt column encryption data at the target database.

You can use the `ENCRYPTION_MODE=TRANSPARENT` setting to transparently encrypt the dump file set with the TDE master encryption key stored in the keystore. A password is not required in this case. The keystore must be present and open at the target database, and it must *contain* the TDE master encryption key from the *source* database for a successful decryption of column encryption metadata during an import operation.

The open keystore is also required to reencrypt column encryption metadata at the target database. If a keystore already exists on the target database, then you can export the current TDE master encryption key *from* the keystore of the source database and import it *into* the keystore of the target database.

- Use the `ENCRYPTION_MODE` parameter to specify the encryption mode. `ENCRYPTION_MODE=DUAL` encrypts the dump set using the TDE master encryption key stored in the keystore and the password provided.

For example, to use dual encryption mode to export encrypted data:

```
expdp hr DIRECTORY=dpump_dir1 DUMPFILE=hr_enc.dmp  
ENCRYPTION=all ENCRYPTION_PASSWORD=encryption_password  
ENCRYPTION_ALGORITHM=AES256 ENCRYPTION_MODE=dual
```

Password: *password_for_hr*

See Also:

- [Exporting and Importing the TDE Master Encryption Key](#) (page 4-39)
 - *Oracle Database Utilities* for details on using Oracle Data Pump and the associated encryption parameters
 - [Creating an Encrypted Column in an External Table](#) (page 3-25)
-
-

6.2 How Transparent Data Encryption Works with Oracle Data Guard

For both software keystores and hardware keystores, Oracle Data Guard supports Transparent Data Encryption (TDE).

If the primary database uses TDE, then each standby database in a Data Guard configuration must have a copy of the encryption keystore from the primary database. If the primary database uses TDE, then each standby database in a Data Guard configuration must have an encryption keystore with the keystore from the primary database merged into it. If you reset the TDE master encryption key in the primary database, then you must merge the keystore on the primary database that contains the TDE master encryption key to each standby database.

Note the following:

- Encrypted data in log files remains encrypted when data is transferred to the standby database. Encrypted data also stays encrypted during transit.
- TDE works with SQL*Loader direct path loads. The data loaded into encrypted columns is transparently encrypted during the direct path load.
- Materialized views work with TDE tablespace encryption. You can create both materialized views and materialized view logs in encrypted tablespaces. Materialized views also work with TDE column encryption.

See Also:

- [Merging Software Keystores](#) (page 4-8)
 - *Oracle Data Guard Concepts and Administration* more information about the use of TDE with logical standby databases
 - *Oracle Key Vault Administrator's Guide* for information about how to use TDE with Oracle Data Guard in an Oracle Key Vault environment
-
-

6.3 How Transparent Data Encryption Works with Oracle Real Application Clusters

Oracle Real Application Clusters nodes can share both software and keystores and hardware security modules.

[About Using Transparent Data Encryption with Oracle Real Application Clusters](#)
(page 6-5)

Oracle Database enables Oracle Real Application Clusters nodes to share a software keystore. Hardware security modules use a network connection for each database instance.

[Using a Non-Shared File System to Store a Software Keystore in Oracle RAC](#)
(page 6-5)

If you do not use a shared file system to store the software keystore, then you must copy the keystore to the associated nodes.

6.3.1 About Using Transparent Data Encryption with Oracle Real Application Clusters

Oracle Database enables Oracle Real Application Clusters nodes to share a software keystore. Hardware security modules use a network connection for each database instance.

This eliminates the need to manually copy and synchronize the software keystore across all of the nodes. Oracle recommends that you create the software keystore on a shared file system. This enables all of the instances to access the same shared software keystore. If you configure Oracle RAC to use Automatic Storage Management (ASM), then store the keystore on the ASM disk group.

For hardware security modules, use a network connection for each database instance. Thus, all database instances have access to the hardware security module.

Keystore operations that must be performed or synchronized on all of the instances, such as opening or closing the keystore or rekeying can be performed on any one Oracle RAC instance. The synchronization operation applies to all of the other Oracle RAC instances in the cluster. This means that when you open and close a keystore for one instance, then it opens and closes for all of the Oracle RAC instances. Similarly, a TDE master encryption key rekey operation that you perform on one database instance applies to all of the database instances. You can perform other keystore operations, such as exporting TDE master encryption keys, rotating the keystore password, merging keystores, or backing up keystores, from a single instance only.

When using a shared file system, ensure that the `ENCRYPTION_WALLET_LOCATION` or `WALLET_LOCATION` parameter setting in the `sqlnet.ora` file for all of the Oracle RAC instances point to the same shared software keystore location. You also must ensure security of the shared software keystore by assigning the appropriate directory permissions.

6.3.2 Using a Non-Shared File System to Store a Software Keystore in Oracle RAC

If you do not use a shared file system to store the software keystore, then you must copy the keystore to the associated nodes.

1. Log in to the database instance as a user who has been granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

In a multitenant environment, log in to the root or the appropriate PDB. For example:

```
sqlplus sec_admin@hrpdb as syskm
Enter password: password
Connected.
```

2. Reset the TDE master encryption key on the first Oracle Real Application Clusters (Oracle RAC) node.

For example, for column encryption:

```
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY keystore_password WITH BACKUP
USING 'emp_key_backup';
```

For tablespace encryption:

```
ALTER TABLESPACE encrypt_ts ENCRYPTION USING 'AES256' ENCRYPT;
```

3. Copy the keystore file with the new TDE master encryption key from the first node to all of the other nodes.

To find the keystore file location, query the `WRL_PARAMETER` column in the `V$ENCRYPTION_WALLET` view. To find the `WRL_PARAMETER` settings for all of the database instances, query the `GV$ENCRYPTION_WALLET` view.

4. Close and then reopen the keystore on any node. (If you are using a multitenant container database (CDB), then run these statements in the root.)

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE IDENTIFIED BY
software_keystore_password;
```

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY
software_keystore_password [CONTAINER = ALL | CURRENT];
```

Note:

Any keystore operation, such as opening or closing the keystore, performed on any one Oracle RAC instance applies to all other Oracle RAC instances. This is true even if you are not using a shared file system.

All of the Oracle RAC nodes are now configured to use the new TDE master encryption key.

Related Topics:

[Setting or Rotating the TDE Master Encryption Key in the Keystore](#) (page 4-34)

You can set or rotate the TDE master encryption key for both software keystores and hardware keystores.

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[Closing a Software Keystore](#) (page 4-21)

You can manually close password-based software keystores, auto-login software keystores, and local auto-login software keystores.

6.4 How Transparent Data Encryption Works with SecureFiles

SecureFiles, which stores LOBS, has three features: compression, deduplication, and encryption.

[About Transparent Data Encryption and SecureFiles](#) (page 6-7)

SecureFiles encryption uses TDE to provide the encryption facility for LOBs.

[Example: Creating a SecureFiles LOB with a Specific Encryption Algorithm](#) (page 6-7)

The `CREATE TABLE` statement can create a SecureFiles LOB with encryption specified.

[Example: Creating a SecureFiles LOB with a Column Password Specified](#) (page 6-7)

The `CREATE TABLE` statement can create a SecureFiles LOB with a column password.

6.4.1 About Transparent Data Encryption and SecureFiles

SecureFiles encryption uses TDE to provide the encryption facility for LOBs.

When you create or alter tables, you can specify the SecureFiles encryption or LOB columns that must use the SecureFiles storage. You can enable the encryption for a LOB column by either using the current Transparent Data Encryption (TDE) syntax or by using the `ENCRYPT` clause as part of the LOB parameters for the LOB column. The `DECRYPT` option in the current syntax or the LOB parameters turn off encryption.

6.4.2 Example: Creating a SecureFiles LOB with a Specific Encryption Algorithm

The `CREATE TABLE` statement can create a SecureFiles LOB with encryption specified.

[Example 6-1](#) (page 6-7) shows how to create a SecureFiles LOB in a `CREATE TABLE` statement.

Example 6-1 *Creating a SecureFiles LOB with a Specific Encryption Algorithm*

```
CREATE TABLE table1 ( a BLOB ENCRYPT USING 'AES256')
  LOB(a) STORE AS SECUREFILE (
    CACHE
  );
```

6.4.3 Example: Creating a SecureFiles LOB with a Column Password Specified

The `CREATE TABLE` statement can create a SecureFiles LOB with a column password.

[Example 6-2](#) (page 6-7) shows an example of creating a SecureFiles LOB that uses password protections for the encrypted column.

All of the LOBS in the LOB column are encrypted with the same encryption specification.

Example 6-2 *Creating a SecureFiles LOB with a Column Password Specified*

```
CREATE TABLE table1 (a VARCHAR2(20), b BLOB)
  LOB(b) STORE AS SECUREFILE (
    CACHE
    ENCRYPT USING 'AES192' IDENTIFIED BY password
  );
```

6.5 How Transparent Data Encryption Works in a Multitenant Environment

In a multitenant environment, the TDE operations that you can perform depend on whether you are in the root or a PDB.

[About Using Transparent Data Encryption in a Multitenant Environment](#) (page 6-8)

TDE can be used for both columns and tablespaces in a multitenant environment.

[Operations That Must Be Performed in Root](#) (page 6-9)

You must perform specific `ADMINISTER KEY MANAGEMENT` keystore operations only in the root.

[Operations That Can Be Performed in Root or in a PDB](#) (page 6-11)

You can perform the some keystore operations in either the root or a PDB.

[Moving PDBs from One CDB to Another](#) (page 6-11)

You can automatically move a PDB from one CDB to another (for example, for load balancing or adding new functionality).

[Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12)

The `EXPORT` and `IMPORT` clauses of `ADMINISTER KEY MANAGEMENT EXPORT` can export or import TDE master encryption keys for a PDB.

[Unplugging and Plugging a PDB with Encrypted Data in a CDB](#) (page 6-14)

You can add or remove PDBs that have encrypted data to and from a CDB.

[Managing Cloned PDBs with Encrypted Data](#) (page 6-17)

You can clone a PDB that has encrypted data in a CDB.

[How Keystore Open and Close Operations Work in a Multitenant Environment](#) (page 6-18)

You should be aware of how keystore open and close operations work in a multitenant environment.

[Finding the Keystore Status for All of the PDBs in a Multitenant Environment](#) (page 6-19)

You can create a convenience function that uses the `V` `$_ENCRYPTION_WALLET` view to find the status for keystores in all PDBs in a CDB.

6.5.1 About Using Transparent Data Encryption in a Multitenant Environment

TDE can be used for both columns and tablespaces in a multitenant environment.

Note the following:

- **The keystore that you create resides in the host multitenant environment, not within any particular PDB.** Multiple PDBs can access a single keystore while running on this host. Each PDB that uses encryption has a Transparent Data Encryption TDE master encryption key stored in this keystore.
- **Each PDB has its own TDE master encryption key.** You must manage the TDE master encryption key for each PDB from within the PDB only, using the PDB-specific key management `ADMINISTER KEY MANAGEMENT` statements. From the

root or a PDB, you can query the appropriate views to find information about the TDE master encryption keys of the PDBs in a CDB. For example, the `PDBID` column of the `V$ENCRYPTION_KEYS` view indicates the PDBs to which a TDE master encryption key belongs.

- **You can manage the Transparent Data Encryption TDE master encryption keys independently within the keystore for each PDB.** You can rotate the TDE master encryption keys for each PDB individually. See [Exporting and Importing the TDE Master Encryption Key](#) (page 4-39) for more information.
- **You perform most of the keystore operations from the root.** Keystore operations such as rotating a keystore password, merging keystores, and so on must be performed in the root. There are a few key management operations that you can perform within a PDB, such as opening, closing, resetting, and creating keys. The operations can also be performed for all of the PDBs from the root. Where applicable, the `ADMINISTER KEY MANAGEMENT` statement has the `CONTAINER` clause. Setting `CONTAINER=ALL` performs the action on all of the PDBs.
- **An auto-login keystore is open (for example, in the root) or if the keystore is closed.** In this scenario, include the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement for the following operations: rotating a keystore password; creating, using, rekeying, tagging, importing, exporting, migrating, or reverse migrating encryption keys; opening or backing up keystores; adding, updating, or deleting secret keystores.
- **You can use an external store for passwords in a multitenant environment.** If you have configured the password-based software keystore to use an external store, then include the `IDENTIFIED BY EXTERNAL STORE` clause in the `ADMINISTER KEY MANAGEMENT` statement for the following operations: backing up, opening, or closing keystores; adding, updating, or deleting secret keystores; creating, using, rekeying, tagging, importing, or exporting encryption keys.
- **If you plan to move a PDB that uses Transparent Data Encryption to a new host computer, then you must move its TDE master encryption key as well.** To move the TDE master encryption key from one host computer to another, use the procedures described in [Exporting and Importing the TDE Master Encryption Key](#) (page 4-39).

Related Topics:

[Operations That Must Be Performed in Root](#) (page 6-9)

You must perform specific `ADMINISTER KEY MANAGEMENT` keystore operations only in the root.

[Operations That Can Be Performed in Root or in a PDB](#) (page 6-11)

You can perform the some keystore operations in either the root or a PDB.

6.5.2 Operations That Must Be Performed in Root

You must perform specific `ADMINISTER KEY MANAGEMENT` keystore operations only in the root.

These operations are as follows:

- **Creating password-based software keystores**, using the `ADMINISTER KEY MANAGEMENT CREATE KEYSTORE` statement

- **Creating auto-login software keystores**, using the `ADMINISTER KEY MANAGEMENT CREATE [LOCAL] AUTO_LOGIN KEYSTORE FROM KEYSTORE` statement
- **Changing the software keystore password**, using the `ADMINISTER KEY MANAGEMENT ALTER KEYSTORE PASSWORD` statement
- **Merging software keystores**, using the `ADMINISTER KEY MANAGEMENT MERGE KEYSTORE` statement
- **Backing up software keystores**, using the `ADMINISTER KEY MANAGEMENT BACKUP KEYSTORE` keystore
- **Migrating from a software keystore to a hardware keystore**, using the `ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY... MIGRATE USING` statement
- **Reverse migrating from a hardware security module to a software keystore**, using the `ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY... REVERSE MIGRATE` statement
- **Adding, updating, and deleting secrets**, using the `ADMINISTER KEY MANAGEMENT ADD|UPDATE|DELETE SECRET` statement
- **Selectively exporting and importing keys**, based on a query or identifier list

How the CONTAINER=ALL Setting Works for Key and Keystore Operations

You can specify the `CONTAINER=ALL` setting for the key and keystore operations described in this section. Specifying the `CONTAINER=ALL` setting performs the same operation on all of the PDBs within the CDB. Remember that you can only use the `CONTAINER=ALL` setting in the root. The `CONTAINER` clause is optional. If you omit the `CONTAINER` clause, then the default is `CONTAINER = CURRENT`.

The permitted `CONTAINER=ALL` operations are as follows:

- **Opening a keystore.** If you open the keystore using the `CONTAINER=ALL` setting, then the keystores on all of the associated PDBs open.
- **Closing a keystore.** Closing a keystore using the `CONTAINER=ALL` setting closes the keystores on all of the associated PDBs.
- **Creating a TDE master encryption key.** Creating a TDE master encryption key using the `CONTAINER=ALL` setting creates the key on all of the PDBs that are open. You can check the keys that were created recently by querying the `CREATION_TIME` column in the `V$ENCRYPTION_KEYS` view. You can also specify a tag with `CONTAINER=ALL` operation, but be aware that this operation creates the keys in all of the PDBs with the same tag. You should have individual tags for each TDE master encryption key, because the tags can help identify PDBs on which the create key operation succeeded in case of an error. You can modify the tag by using the `ADMINISTER KEY MANAGEMENT SET TAG` statement later on.
- **Performing a rekey operation.** Performing a rekey operation with the `CONTAINER=ALL` setting creates and then activates the key on all of the PDBs that are open. You can check the keys that were created recently by querying the `CREATION_TIME` column in the `V$ENCRYPTION_KEYS` view. To find the keys that were activated recently, query the `ACTIVATION_TIME` column in the `V$ENCRYPTION_KEYS` view. You can also specify a tag with `CONTAINER=ALL`

operation, but be aware that this operation creates the keys in all of the PDBs with the same tag. The tag can also help identify PDBs on which the create key operation succeeded in case of an error. You can modify the tag by using the `ADMINISTER KEY MANAGEMENT SET TAG` statement later on.

6.5.3 Operations That Can Be Performed in Root or in a PDB

You can perform the some keystore operations in either the root or a PDB.

These operations are as follows:

- **Opening keystores**, using the `ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN` statement
- **Closing keystores**, using the `ADMINISTER KEY MANAGEMENT SET KEYSTORE CLOSE` statement

You can perform the following key management operations either in the root or a PDB:

- **Creating a tag for the TDE master encryption key**, using the `ADMINISTER KEY MANAGEMENT SET TAG` statement
- **Creating a TDE master encryption key**, using the `ADMINISTER KEY MANAGEMENT CREATE KEY` statement
- **Resetting or rotating the TDE master encryption key**, using the `ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY` statement
- **Activating a TDE master encryption key**, using the `ADMINISTER KEY MANAGEMENT USE KEY` statement
- **Exporting TDE master encryption keys**, using the `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement
- **Importing TDE master encryption keys**, using the `ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS` statement

6.5.4 Moving PDBs from One CDB to Another

You can automatically move a PDB from one CDB to another (for example, for load balancing or adding new functionality).

If the PDB has TDE-encrypted tables or tablespaces, then you can set the `ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE` dynamic system parameter to enable the PDB to include the TDE keys in the PDB move operation. This parameter prevents the need of having to manually provide a keystore password when you import the TDE keys into the PDB after it has moved to a different CDB. When `ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE` is set to `TRUE`, the database caches the keystore password in memory, obfuscated at the system level, and then uses it for the import operation. The default for `ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE` is `FALSE`.

1. Before you begin the PDB move operation, log in to the root as a user who has been granted the `SYSDBA` administrative privilege.

For example:

```
sqlplus sec_admin as sysdba
Enter password: password
```

2. Set the `ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE` dynamic initialization parameter `TRUE`.

For example:

```
ALTER SYSTEM SET ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE = TRUE;
```

See Also: *Oracle Database Reference* for more information about `ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE`

6.5.5 Exporting and Importing TDE Master Encryption Keys for a PDB

The `EXPORT` and `IMPORT` clauses of `ADMINISTER KEY MANAGEMENT EXPORT` can export or import TDE master encryption keys for a PDB.

[About Exporting and Importing TDE Master Encryption Keys for a PDB](#)
(page 6-12)

You can export and import TDE master encryption keys from the root in the same way that you export and import this key for a non-CDB database.

[Exporting or Importing a TDE Master Encryption Key for a PDB](#) (page 6-13)

The `ADMINISTER KEY MANAGEMENT` statement can export or import a TDE master encryption key for a PDB.

[Example: Exporting a TDE Master Encryption Key from a PDB](#) (page 6-14)

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement can export TDE master encryption keys for a PDB.

[Example: Importing a TDE Master Encryption Key into a PDB](#) (page 6-14)

The `ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS` statement can import a TDE master encryption key into a PDB.

6.5.5.1 About Exporting and Importing TDE Master Encryption Keys for a PDB

You can export and import TDE master encryption keys from the root in the same way that you export and import this key for a non-CDB database.

You can export and import all of the TDE master encryption keys that belong to the PDB by exporting and importing the TDE master encryption keys from within a PDB. Export and import of TDE master encryption keys in a PDB supports the PDB unplug and plug operations. During a PDB unplug and plug, all of the TDE master encryption keys that belong to a PDB, as well as the metadata, are involved. Therefore, the `WITH IDENTIFIER` clause of the `ADMINISTER KEY MANAGEMENT EXPORT` statement is not allowed when you export keys from within a PDB. The `WITH IDENTIFIER` clause is only permitted in the root.

You should include the `FORCE KEYSTORE` clause if the root has an auto-login keystore or if the keystore is closed. If the keystore has been configured to use an external store for the password, then use the `IDENTIFIED BY EXTERNAL STORE` clause. For example, to perform an export operation for this scenario:

```
ADMINISTER KEY MANAGEMENT EXPORT KEYS WITH SECRET "my_secret"
TO '/etc/TDE/export.exp'
FORCE KEYSTORE IDENTIFIED BY EXTERNAL STORE;
```

This `ADMINISTER KEY MANAGEMENT EXPORT` operation exports not only the keys but creates metadata that is necessary for PDB environments (as well as for cloning operations).

Inside a PDB, the export operation of TDE master encryption keys exports the keys that were created or activated by a PDB with the same GUID as the PDB where the keys are being exported. Essentially, all of the keys that belong to a PDB where the export is being performed will be exported.

The importing of TDE master encryption keys from an export file within a PDB takes place only if the TDE master encryption key was exported from another PDB with the same GUID. To support the plug-in of a PDB into a CDB, the import will also import the TDE master encryption keys from an export file that contains the TDE master encryption keys of a non-CDB exported without the `WITH IDENTIFIER` clause. Because the PDB-specific details, such as the PDB name and database ID, can change from one CDB to the next, the PDB-specific information is modified during the import to reflect the updated PDB information.

Note:

Within a PDB, you can only export the keys of a PDB as a whole. The ability to export them selectively based on a query or an identifier is restricted to the root.

6.5.5.2 Exporting or Importing a TDE Master Encryption Key for a PDB

The `ADMINISTER KEY MANAGEMENT` statement can export or import a TDE master encryption key for a PDB.

1. Log in to the PDB as a user who was granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

For example:

```
sqlplus sec_admin@hr_pdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

2. Ensure that the keystore is open.

You can query the `STATUS` column of the `V$ENCRYPTION_WALLET` view to find if the keystore is open. If you find that you must open the software keystore, then you can optionally include the `FORCE KEYSTORE` clause in the `ADMINISTER KEY MANAGEMENT` statement when you perform the export or import operation. This clause enables you to open a software keystore during the operation without having to separately open the auto-login keystore or the password-based keystore.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
keystore_password;
```

3. Perform the export or import operation.

Related Topics:

[Step 3: Open the Software Keystore](#) (page 3-9)

Depending on the type of keystore you create, you must manually open the keystore before you can use it.

[Example: Exporting a TDE Master Encryption Key from a PDB](#) (page 6-14)

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement can export TDE master encryption keys for a PDB.

6.5.5.3 Example: Exporting a TDE Master Encryption Key from a PDB

The `ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS` statement can export TDE master encryption keys for a PDB.

[Example 6-3](#) (page 6-14) shows how to export a TDE master encryption key from the PDB `hr_pdb1`. In this example, the `FORCE KEYSTORE` clause is included in case the auto-login keystore is in use, or if the keystore is closed.

Example 6-3 Exporting a TDE Master Encryption Key from a PDB

```
sqlplus sec_admin@hr_pdb1 as syskm
Enter password: password
Connected.
```

```
ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET "my_secret" TO '/
export.p12' FORCE KEYSTORE IDENTIFIED BY password_cdb1;
```

6.5.5.4 Example: Importing a TDE Master Encryption Key into a PDB

The `ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS` statement can import a TDE master encryption key into a PDB.

[Example 6-4](#) (page 6-14) shows how to import a TDE master encryption key into the PDB `hr_pdb2`.

Example 6-4 Importing a TDE Master Encryption Key into a PDB

```
sqlplus sec_admin@hr_pdb2 as syskm
Enter password: password
Connected.
```

```
ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS WITH SECRET "my_secret" FROM '/tmp/
export.p12' FORCE KEYSTORE IDENTIFIED BY password_cdb2 WITH BACKUP;
```

6.5.6 Unplugging and Plugging a PDB with Encrypted Data in a CDB

You can add or remove PDBs that have encrypted data to and from a CDB.

[Unplugging a PDB That Has Encrypted Data](#) (page 6-15)

You can unplug a PDB from one CDB and then plug it into another CDB.

[Plugging a PDB That Has Encrypted Data into a CDB](#) (page 6-15)

To plug a PDB that has encrypted data into a CDB, you first plug in the PDB and then configure its encryption key.

[Unplugging a PDB That Has Master Keys Stored in an HSM](#) (page 6-16)

You can unplug a PDB from one CDB that has been configured with an HSM and then plug it into another CDB also configured with an HSM.

[Plugging a PDB That Has Master Keys Stored in an HSM](#) (page 6-16)

The `ADMINISTER KEY MANAGEMENT` statement can import an HSM master key to a PDB that has been moved to another CDB.

6.5.6.1 Unplugging a PDB That Has Encrypted Data

You can unplug a PDB from one CDB and then plug it into another CDB.

The database that is unplugged contains data files and other associated files. Because each PDB can have its own unique keystore, you do not need to export the TDE master encryption key of the PDB that you want to unplug.

- Unplug the PDB as you normally unplug PDBs, as described in *Oracle Database Administrator's Guide*.

6.5.6.2 Plugging a PDB That Has Encrypted Data into a CDB

To plug a PDB that has encrypted data into a CDB, you first plug in the PDB and then configure its encryption key.

When you plug an unplugged PDB into another CDB, the key version is set to 0 because this operation invalidates the history of the previous keys. You can check the key version by querying the `KEY_VERSION` column of the `V$ENCRYPTED_TABLESPACES` dynamic view. Similarly, if a control file is lost and recreated, then the previous history of the keys is reset to 0.

1. Create the PDB by plugging the unplugged PDB into the CDB, as described in *Oracle Database Administrator's Guide*.

During the open operation of the PDB after the plug operation, Oracle Database determines if the PDB has encrypted data. If so, it opens the PDB in the `RESTRICTED` mode. See *Oracle Database Administrator's Guide* for more information about the Open Mode of a PDB.

If you want to create the PDB by cloning another PDB or from a non-CDB, and if the source database has encrypted data or a keystore set, then you must provide the keystore password by including the `keystore identified by keystore_password` clause in the `CREATE PLUGGABLE DATABASE ... FROM SQL` statement. You must provide this password even if the source database is using an auto-login software keystore. You can find if the source database has encrypted data or a keystore by querying the `DBA_ENCRYPTED_COLUMNS` data dictionary view.

2. Import the TDE master encryption key into the PDB.
See [Exporting and Importing TDE Master Encryption Keys for a PDB](#) (page 6-12).
3. Close the PDB and then re-open the PDB, as described in *Oracle Database Administrator's Guide*.
4. Open the keystore.
See the following sections:
 - [Step 3: Open the Software Keystore](#) (page 3-9)
 - [Step 3: Open the Hardware Keystore](#) (page 3-15)
5. Set the TDE master encryption key for the PDB.
See the following topics:

- [Step 4: Set the Software TDE Master Encryption Key](#) (page 3-11)
- [Step 4: Set the Hardware Keystore TDE Master Encryption Key](#) (page 3-17)
- [Creating TDE Master Encryption Keys for Later Use](#) (page 4-25)

6.5.6.3 Unplugging a PDB That Has Master Keys Stored in an HSM

You can unplug a PDB from one CDB that has been configured with an HSM and then plug it into another CDB also configured with an HSM.

1. Unplug the PDB.
2. Move the master keys of the unplugged PDB in the HSM that was used at the source CDB to the HSM that is in use at the destination CDB.

Refer to the documentation for the HSM for information about moving master keys between HSMs.

See Also: *Oracle Database Administrator's Guide* for information about unplugging PDBs

6.5.6.4 Plugging a PDB That Has Master Keys Stored in an HSM

The `ADMINISTER KEY MANAGEMENT` statement can import an HSM master key to a PDB that has been moved to another CDB.

1. Plug that unplugged PDB into the destination CDB that has been configured with the HSM.

After the plug-in operation, the PDB that has been plugged in will be in restricted mode.

2. Ensure that the master keys from the HSM that has been configured with the source CDB are available in the HSM of the destination CDB.
3. Log in to the plugged PDB as a user who was granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

For example:

```
sqlplus sec_admin@hr_pdb as syskm
Enter password: password
Connected.
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `show con_name` command.

4. Open the master encryption key of the plugged PDB.

For example, for a PDB called `PDB1`:

```
ALTER SESSION SET CONTAINER = PDB1;
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY "keystore_password";
```

5. Import the HSM master key into the PDB.

```
ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS WITH SECRET "HSM" FROM 'HSM'
IDENTIFIED BY "keystore_password";
```

6. Restart the PDB.

```
ALTER PLUGGABLE DATABASE PDB1 CLOSE;
ALTER PLUGGABLE DATABASE PDB1 OPEN;
```

See Also: *Oracle Database Administrator's Guide* for information about plugging PDBs

6.5.7 Managing Cloned PDBs with Encrypted Data

You can clone a PDB that has encrypted data in a CDB.

About Managing Cloned PDBs That Have Encrypted Data (page 6-17)

When you clone a PDB, you must make the master key of the source PDB available to cloned PDB.

Cloning a PDB with Encrypted Data in a CDB (page 6-17)

The `CREATE PLUGGABLE DATABASE` statement with the `KEYSTORE IDENTIFIED BY` clause can clone a PDB that has encrypted data.

6.5.7.1 About Managing Cloned PDBs That Have Encrypted Data

When you clone a PDB, you must make the master key of the source PDB available to cloned PDB.

This allows a cloned PDB to operate on the encrypted data. To perform the clone, you do not need to export and import the keys because Oracle Database transports the keys for you even if the cloned PDB is in a remote CDB. However, you will need to provide the keystore password of the CDB where you are creating the clone.

6.5.7.2 Cloning a PDB with Encrypted Data in a CDB

The `CREATE PLUGGABLE DATABASE` statement with the `KEYSTORE IDENTIFIED BY` clause can clone a PDB that has encrypted data.

1. Log in to the root as a user who was granted the `ADMINISTER KEY MANAGEMENT` or `SYSKM` privilege.

For example:

```
sqlplus c##sec_admin as syskm
Enter password: password
Connected.
```

2. Ensure that the software keystore of the PDB that you plan to clone is open.

You can query the `STATUS` column of the `V$ENCRYPTION_WALLET` view to find if the software keystore is open.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN FORCE KEYSTORE IDENTIFIED BY
keystore_password;
```

3. Use the `CREATE PLUGGABLE DATABASE` statement with the `KEYSTORE IDENTIFIED BY` clause to clone the PDB.

For example:

```
CREATE PLUGGABLE DATABASE cdb1_pdb3 FROM cdb1_pdb1
FILE_NAME_CONVERT=('cdb1_pdb1', 'pdb3/cdb1_pdb3') KEystore IDENTIFIED BY
"keystore_password";
```

Replace *keystore_password* with the password of the keystore of the CDB where the *cdb1_pdb3* clone is created.

After you create the cloned PDB, encrypted data is still accessible by the clone using the master encryption key of the original PDB. After a PDB is cloned, there may be user data in the encrypted tablespaces. This encrypted data is still accessible because the master key of the source PDB is copied over to the destination PDB. Because the clone is a copy of the source PDB but will eventually follow its own course and have its own data and security policies, you should rekey the master key of the cloned PDB.

4. Rekey (rotate) the master key of the cloned PDB.

For example:

```
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY keystore_password WITH BACKUP
USING 'emp_key_backup';
```

Before you rekey the master key of the cloned PDB, the clone can still use master encryption keys that belong to the original PDB. However, these master keys do not appear in the cloned PDB *v\$* dynamic views. Rekeying the master key ensures that the cloned PDB uses its own unique keys, which will be viewable in the *v\$* views.

See Also:

- [Step 3: Open the Software Keystore](#) (page 3-9)
 - *Oracle Database Administrator's Guide* for information about cloning a PDB
 - [Rotating the TDE Master Encryption Key](#) (page 4-37)
-
-

6.5.8 How Keystore Open and Close Operations Work in a Multitenant Environment

You should be aware of how keystore open and close operations work in a multitenant environment.

For each PDB in a multitenant environment, you must explicitly open the password-based software keystore or hardware keystore in the PDB to enable the Transparent Data Encryption operations to proceed. (Auto-login and local auto-login software keystores open automatically.) Closing a keystore on a PDB blocks all of the Transparent Data Encryption operations on that PDB.

In a CDB, the open and close keystore operations in a PDB depends on the open and close status of the keystore in the root.

Note the following:

- Before you can manually open a software password-based or hardware keystore in an individual PDB, you must open the keystore in the root.
- If an auto-login keystore is in use, or if the keystore is closed, then include the `FORCE KEystore` clause in the `ADMINISTER KEY MANAGEMENT` statement when you open or close the keystore.

- If the keystore is a password-based software keystore that uses an external store for passwords, then set the `IDENTIFIED BY` clause to `EXTERNAL STORE`.
- Before you can set a TDE master encryption key in an individual PDB, you must set the key in the root.
- Auto-login and local auto-login software keystores open automatically. You do not need to manually open these from the root first, or from the PDB.
- If you close a keystore in the root, then the keystores in the dependent PDBs also close. A keystore close operation in the root is the equivalent of performing a keystore close operation with the `CONTAINER` clause set to `ALL`.
- If you open a keystore in the root and set the `CONTAINER` clause to `ALL`, then the keystores in the dependent PDBs also open.

6.5.9 Finding the Keystore Status for All of the PDBs in a Multitenant Environment

You can create a convenience function that uses the `V$ENCRYPTION_WALLET` view to find the status for keystores in all PDBs in a CDB.

The `V$ENCRYPTION_WALLET` view displays the status of the keystore in a PDB, whether it is open, closed, uses a software or hardware keystore, and so on.

- To create a function that uses the `V$ENCRYPTION_WALLET` view to find the keystore status, use the `CREATE PROCEDURE PL/SQL` statement.

[Example 6-5](#) (page 6-19) shows how to create this function.

Example 6-5 *Function to Find the Keystore Status of All of the PDBs in a CDB*

```
CREATE OR REPLACE PROCEDURE all_pdb_v$encryption_wallet
IS
    err_occ          BOOLEAN;
    curr_pdb         VARCHAR2(30);
    pdb_name         VARCHAR2(30);
    wrl_type         VARCHAR2(20);
    status           VARCHAR2(30);
    wallet_type      VARCHAR2(20);
    wallet_order     VARCHAR2(12);
    fully_backed_up  VARCHAR2(15);
    wrl_parameter    VARCHAR2(4000);
    cursor sel_pdb  IS SELECT NAME FROM V$CONTAINERS
                       WHERE NAME <> 'PDB$SEED' order by con_id desc;
BEGIN
    -- Store the original PDB name
    SELECT sys_context('userenv', 'con_name') INTO curr_pdb FROM DUAL;
    IF curr_pdb <> 'CDB$ROOT' THEN
        dbms_output.put_line('Operation valid in ROOT only');
    END IF;

    err_occ := FALSE;
    dbms_output.put_line('---');
    dbms_output.put_line('PDB_NAME                                WRL_TYPE STATUS                                ');
    dbms_output.put_line('-----');
    dbms_output.put_line('WALLET_TYPE          WALLET_ORDER FULLY_BACKED_UP');
    dbms_output.put_line('-----');
    dbms_output.put_line('WRL_PARAMETER');

    dbms_output.put_line('-----');
```

```

FOR pdbinfo IN sel_pdbs LOOP

    pdb_name := DBMS_ASSERT.ENQUOTE_NAME(pdbinfo.name, FALSE);
    EXECUTE IMMEDIATE 'ALTER SESSION SET CONTAINER = ' || pdb_name;

    BEGIN
        pdb_name := rpad(substr(pdb_name,1,30), 30, ' ');
        EXECUTE IMMEDIATE 'SELECT wr1_type from V$ENCRYPTION_WALLET' into wr1_type;
        wr1_type := rpad(substr(wr1_type,1,8), 8, ' ');
        EXECUTE IMMEDIATE 'SELECT status from V$ENCRYPTION_WALLET' into status;
        status := rpad(substr(status,1,30), 30, ' ');
        EXECUTE IMMEDIATE 'SELECT wallet_type from V$ENCRYPTION_WALLET' into wallet_type;
        wallet_type := rpad(substr(wallet_type,1,20), 20, ' ');
        EXECUTE IMMEDIATE 'SELECT wallet_order from V$ENCRYPTION_WALLET' into wallet_order;
        wallet_order := rpad(substr(wallet_order,1,9), 12, ' ');
        EXECUTE IMMEDIATE 'SELECT fully_backed_up from V$ENCRYPTION_WALLET' into fully_backed_up;
        fully_backed_up := rpad(substr(fully_backed_up,1,9), 15, ' ');
        EXECUTE IMMEDIATE 'SELECT wr1_parameter from V$ENCRYPTION_WALLET' into wr1_parameter;
        wr1_parameter := rpad(substr(wr1_parameter,1,79), 79, ' ');
        dbms_output.put_line(pdb_name || ' ' || wr1_type || ' ' || status);
        dbms_output.put_line(wallet_type || ' ' || wallet_order || ' ' || fully_backed_up);
        dbms_output.put_line(wr1_parameter);

    EXCEPTION

        WHEN OTHERS THEN
            err_occ := TRUE;

    END;

END LOOP;

IF err_occ = TRUE THEN
    dbms_output.put_line('One or more PDB resulted in an error');
END IF;

END;

.
/
set serveroutput on
exec all_pdb_v$encryption_wallet;

```

6.6 How Transparent Data Encryption Works with Oracle Call Interface

Transparent Data Encryption does not have any effect on the operation of Oracle Call Interface (OCI).

For most practical purposes, TDE is transparent to OCI except for the row shipping feature. You cannot use the OCI row shipping feature with TDE because the key to make the row usable is not available at the receipt-point.

6.7 How Transparent Data Encryption Works with Editions

Transparent Data Encryption does not have any effect on the Editions feature of Oracle Database.

For most practical purposes, TDE is transparent to Editions. Tables are always noneditioned objects. TDE Column Encryption encrypts columns of the table. Editions are not affected by TDE tablespace encryption.

6.8 Configuring Transparent Data Encryption to Work in a Multidatabase Environment

Each Oracle database on the same server (such as databases sharing the same Oracle binary but using different data files) must access its own TDE keystore.

Keystores are not designed to be shared among databases. By design, there must be one keystore per database. You cannot use the same keystore for more than one database.

- To configure the `sqlnet.ora` file for a multidatabase environment, use one of the following options:

- **Option 1:** If the databases share the same Oracle home, then keep the `sqlnet.ora` file in the default location, which is in the `ORACLE_HOME/network/admin` directory.

In this case, it is ideal to use the default location. Ensure that the `sqlnet.ora` file has no `WALLET_LOCATION` or `ENCRYPTION_WALLET_LOCATION` entries. Transparent Data Encryption accesses the keystore from the default `sqlnet.ora` location if these two entries are not in the `sqlnet.ora` file.

- **Option 2:** If Option 1 is not feasible for your site, then you can specify the keystore location based on an environment variable setting, such as `ORACLE_SID`.

For example:

```
ENCRYPTION_WALLET_LOCATION =
(SOURCE =
(METHOD = FILE)
(METHOD_DATA =
(DIRECTORY = /home/oracle/wallet/$ORACLE_SID)
```

- **Option 3:** If Options 1 and 2 are not feasible, then use separate `sqlnet.ora` files, one for each database. Ensure that you correctly set the `TNS_ADMIN` environment variable to point to the correct database configuration.

Caution:

Using a keystore from another database can cause partial or complete data loss.

See Also: *SQL*Plus User's Guide and Reference* for more information and examples of setting the `TNS_ADMIN` variable

Frequently Asked Questions About Transparent Data Encryption

Users frequently have questions about transparency and performance issues with Transparent Data Encryption.

[Transparency Questions About Transparent Data Encryption](#) (page 7-1)

Transparent Data encryption handles transparency in data in a variety of ways.

[Performance Questions About Transparent Data Encryption](#) (page 7-4)

There are several performance issues to consider when using Transparent Data Encryption.

7.1 Transparency Questions About Transparent Data Encryption

Transparent Data encryption handles transparency in data in a variety of ways.

Security auditors occasionally ask detailed questions about the encryption used by Oracle Advanced Security Transparent Data Encryption (TDE). They request information about TDE keys, algorithms, lengths, and keystores and then directly compare to requirements of regulations such as PCI-DSS. This topic contains important details about TDE encryption and key management. This information is current as of Oracle Database 12c (12.1.0.2). It is intended to help TDE customers respond to auditor questions quickly and accurately.

1. Is Transparent Data Encryption compatible with my application software?

Transparent Data Encryption is compatible with applications by default because it does not alter the inbound SQL statements or the outbound SQL query results. Oracle executes internal testing and validation of certain Oracle and third-party application software to capture helpful deployment tips or scripts, and to evaluate performance profiles.

Be aware of the difference between Transparent Data Encryption and the DBMS_CRYPTO PL/SQL package. This package is intended for different customer use cases. It is an API and toolkit solution and as such, it is non-transparent.

2. Is Transparent Data Encryption compatible with other Oracle Database tools and technologies that I am using?

One of the chief benefits of Transparent Data Encryption is its integration with frequently used Oracle Database tools and technologies such as high-availability clusters, storage compression, backup compression, data movement, database backup and restore, and database replication. Specific Oracle technologies that are integrated directly with Transparent Data Encryption include Oracle Real Application Clusters (Oracle RAC), Oracle Recovery Manager (RMAN), Oracle Data Guard, Advanced Compression, Oracle Data Pump, and Oracle GoldenGate,

among others. Transparent Data Encryption also has special points of integration with Oracle Exadata that fully use unique features of Oracle-engineered systems.

Transparent Data Encryption also works easily with security features of the Oracle Database. With Transparent Data Encryption, privilege grants, roles, Oracle Database Vault realms, Virtual Private Database policies, and Oracle Label Security labels remain in effect. You can use these and other security features in tandem with Transparent Data Encryption encryption.

3. Are there any known Transparent Data Encryption limitations or incompatibilities?

- **TDE column encryption:** TDE column encryption encrypts and decrypts data transparently when data passes through the SQL layer. Some features of Oracle will bypass the SQL layer, and hence cannot benefit from TDE column encryption. The following are known database features that TDE column encryption does not support, and their relevant software version numbers:

- Materialized View Logs (not supported prior to Oracle Database 11g Release 2)
- Streams (not supported prior to Oracle Database 11g Release 1)
- Synchronous and asynchronous change data capture for data warehousing (CDC)
- Transportable Tablespaces
- LOBs

Note that Secure Files were introduced in Oracle Database 11g Release 1, so it is not supported with TDE column encryption prior to that release

- **TDE tablespace encryption:** TDE tablespace encryption encrypts all content that is stored in the tablespace at the block level in storage, and it generally does not conflict with other database features. TDE tablespace encryption does not have any of the limitations that TDE column encryption has. However, you should be aware of the following:
 - You can use full transportable tablespaces (TTS) with Oracle Data Pump compression and encryption when going from a TDE-encrypted source to a TDE-encrypted destination. You must have an Oracle Database release 12c database instance available so that you can use its key export or keystore (wallet) merge capabilities to get the correct TDE master key to the destination database host without having to overwrite the original Oracle wallet file. This process is subject to the standard TTS limitations, and you must remember to check for compatible endianness.
 - Do not attempt to encrypt database internal objects such as the `SYSTEM`, `SYS_AUX`, `UNDO`, or `TEMP` tablespaces using TDE tablespace encryption. You should focus TDE tablespace encryption on tablespaces that hold application data, not on these core components of the Oracle database.

4. What types of keys and algorithms does TDE use?

TDE relies on two distinct sets of encryption keys. The first set of encryption keys are data encryption keys (DEK), which are used to transparently encrypt and decrypt stored data. DEKs are generated automatically by the database, stored internally in the database in encrypted form, and managed mostly behind the

scenes. One place where end-users interact with DEKs is when selecting the encryption algorithm and key length that TDE will use, which can be 3DES168, AES128, AES192, or AES256. This selection is made independently for each table containing encrypted columns and for each encrypted tablespace. You may also hear DEKs referred to as table keys (column encryption) or tablespace keys (tablespace encryption). The table keys are used in cipher block chaining (CBC) operating mode, and the tablespace keys are used in cipher feedback (CFB) operating mode.

The second set of encryption keys consists of current and historical key encryption keys (KEK), also known as TDE master keys. The TDE master keys are generated automatically by the database, used automatically to encrypt and decrypt DEKs as needed, and stored externally in a protected keystore. Users may interact with the current TDE master key by periodically rotating it, modifying certain key attributes, and so forth. Typically, the keystore for TDE master keys is either an Oracle wallet (out-of-the-box solution) or Oracle Key Vault (a specialized key management product). Although the database uses only one TDE master key at a time, all rotated TDE master keys are retained in the keystore for long-term recovery of encrypted data backups. TDE master keys always are AES256. They encrypt and decrypt DEKs using CBC operating mode. For both DEKs and TDE master keys, the underlying key material is not directly exposed. End-users see only attributes of keys necessary to manage TDE.

5. How are Oracle wallets containing TDE master keys protected?

There are three different types of wallets to consider when you use an Oracle wallet as the keystore for TDE master keys: password-based wallet, auto-login wallet, and local auto-login wallet. All of these wallets externalize TDE master keys, so they are separate from TDE-encrypted data. Oracle recommends that you place wallet files in local or network directories that are protected by tight file permissions and other security measures.

The password-based wallet is an encrypted key storage file (`ewallet.p12`) that follows the PKCS #12 standard. It is encrypted by a password-derived key according to the PKCS #5 standard. A human user must enter a command containing the password for the database to open the wallet, decrypt its contents, and gain access to keys. The password-based wallet is the default keystore for TDE master keys. In the past, it was encrypted using the 3DES168 encryption algorithm and CBC operating mode. The `orapki` command `convert wallet` enables you to convert password-based wallets to AES256 and CBC operating mode. *Oracle Database Security Guide* provides details about using `orapki` to convert wallets.

Auto-login wallets (`cwallet.sso`) optionally are derived from standard password-based wallets for special cases where automatic startup of the database is required with no human interaction to enter a wallet password. When using auto-login wallet, the master password-based wallet must be preserved because it is needed to rotate the TDE master key. In addition to the best practice of storing auto-login wallet in a local or network directory that is protected by tight file permissions, the file contents are scrambled by the database using a proprietary method for added security. A slight variation on the auto-login wallet called local auto-login wallet has similar behavior. One notable difference with local auto-login wallet is that its contents are scrambled using additional factors taken from the host machine where the file was created. This renders the local auto-login wallet unusable on other host machines. Details of the host factors and scrambling technique are proprietary.

6. What is Oracle Key Vault and how does it manage TDE master keys?

Oracle Key Vault centrally manages TDE master keys, Oracle wallets, Java keystores, and more. It helps you to take control of proliferating keys and key storage files. It includes optimizations specifically for TDE and other components of the Oracle stack. For more information about using Oracle Key Vault with TDE, see the product pages on www.oracle.com and Oracle Technology Network and *Oracle Key Vault Administrator's Guide*.

7.2 Performance Questions About Transparent Data Encryption

There are several performance issues to consider when using Transparent Data Encryption.

1. What is the typical performance overhead from Transparent Data Encryption?

There are many different variables involved in the creation of an accurate Transparent Data Encryption performance test. The results can vary depending on the test environment, test case or workload, measurement metrics or methods, and so on. Oracle cannot guarantee a specific performance overhead percentage that can apply in all possible scenarios. In practice, the performance tests by many Transparent Data Encryption customers are often in the low single digits as a percentage, but that is not universally the case. Customer examples that cite 1 percent and 2 percent overhead respectively are published on Oracle Technology Network in the following URL:

http://streaming.oracle.com/ebn/podcasts/media/12740910_ColumbiaU_120312.mp3

If possible, use Oracle Real Application Testing (Oracle RAT) to capture a real production workload and then replay it against Transparent Data Encryption to get a true indication of the performance overhead that the you can expect within your environment.

See also:

- [Performance and Storage Overhead of Transparent Data Encryption](#) (page 5-3)
- *Oracle Database Testing Guide* for more information about the Oracle Real Application Testing option

2. How can I tune for optimal Transparent Data Encryption performance?

- **TDE column encryption:**
 - Limit the crypto processing by only encrypting the subset of columns that are strictly required to be protected. In addition, turn off the optional integrity checking feature.
 - After you apply column encryption, rebuild the column indexes.
- **TDE tablespace encryption:** TDE tablespace encryption improves performance by caching unencrypted data in memory in the SGA buffer cache. This feature reduces the number of crypto operations that must be performed when users run `SELECT` queries, which draw from the SGA instead of drawing from disk. (Drawing from disk forces the database to perform decrypt operations.) Ensure that the size of the SGA buffer cache is large enough to take full advantage of this performance optimization.

Another major performance boost comes from using hardware and software that supports CPU-based cryptographic acceleration available in Intel AES-NI and Oracle SPARC T4/T5. To take advantage of this feature, you must be running a recent version of the database, have a recent version of the operating system installed, and be using hardware that includes crypto acceleration circuitry within its CPUs/cores.

Database compression further speeds up Transparent Data Encryption performance because the crypto processing occurs on data that already is compressed, resulting in less total data to encrypt and decrypt.

- **In general:**

- Ensure that you have applied the latest patches, which you can download from My Oracle Support at <https://support.oracle.com>
- When you specify an encryption algorithm, remember that AES is slightly faster than 3DES. Use AES128 where possible. Be aware that the performance benefit is small.
- Use Exadata, which includes additional performance benefits. For more information about Oracle Exadata, see *Oracle Database Testing Guide*.

3. Are there specific issues that may slow down TDE performance, and if so, how do I avoid them?

TDE tablespace performance is slower if the database cannot use CPU-based hardware acceleration on the host machine due to factors such as older hardware, an older database version, or an older operating system.

Note the following with regard to specific database workloads:

- **Encrypting the whole data set at once (for example, while doing “Bulk Data Load” into an Oracle data warehouse):** Lower crypto performance has been observed during bulk load of new data into the database or data warehouse. New data cannot be cached in SGA, so TDE tablespace encryption performance optimizations are bypassed. Hence, Transparent Data Encryption has no bonus performance benefits in this type of operation.

Follow these guidelines:

- Ensure that the database is running on servers with CPU-based cryptographic acceleration. This accelerates not only decrypt operations, but also encrypt operations as well (for loading new data). Take the crypto processing out of band by pre-encrypting the data set and then using Transportable Tablespaces (TTS) to load into the database. Try to parallelize this procedure where possible. This requires the database instance to copy the required TDE key to the keystore on the destination database. The procedure may not be feasible when there is a fixed time window for encryption and loading, and these must be done serially.
 - Consider using TDE column encryption. Encrypt only the handful of sensitive regulated columns instead of encrypting an entire tablespace.
- **Decrypting an entire data set at once (for example, while performing a full table scan by reading directly from disk, with no reading from SGA):**

Lower crypto performance is observed when running full table scan queries where data is read directly from storage. Certain performance optimizations of TDE tablespace encryption are bypassed (no caching). Hence, Transparent Data Encryption has no bonus performance benefits in this type of operation.

Follow these guidelines:

- Ensure that the database is running on servers with CPU-based cryptographic acceleration.
- Retest the full table scan queries with a larger SGA size to measure performance when data is read from cache. Try setting the Oracle event number 10949 to disable direct path read.
- Partition the database so that less data is scanned by full table scan operations. Production databases often use partitioning for this kind of scenario (that is, to limit the total amount of data scanned).
- Consider using TDE column encryption. Encrypt only the handful of sensitive regulated columns instead of encrypting an entire tablespace.

Part II

Using Oracle Data Redaction

Part II describes how to use Oracle Data Redaction.

[Introduction to Oracle Data Redaction](#) (page 8-1)

Oracle Data Redaction is the ability to redact sensitive data in real time.

[Oracle Data Redaction Features and Capabilities](#) (page 9-1)

Oracle Data Redaction provides a variety of ways to redact different types of data.

[Configuring Oracle Data Redaction Policies](#) (page 10-1)

An Oracle Data Redaction policy defines how to redact data in a column based on the table column type and the type of redaction you want to use.

[Using Oracle Data Redaction in Oracle Enterprise Manager](#) (page 11-1)

Oracle Enterprise Manager Cloud Control (Cloud Control) can manage Oracle Data Redaction policies and formats.

[Oracle Data Redaction Use with Oracle Database Features](#) (page 12-1)

Oracle Data Redaction can be used with other Oracle features, but some Oracle features may have restrictions with regard to Oracle Data Redaction.

[Security Considerations for Oracle Data Redaction](#) (page 13-1)

Oracle provides guidelines for using Oracle Data Redaction.

Introduction to Oracle Data Redaction

Oracle Data Redaction is the ability to redact sensitive data in real time.

[What Is Oracle Data Redaction?](#) (page 8-1)

Oracle Data Redaction enables you to mask (redact) data that is returned from queries issued by applications.

[When to Use Oracle Data Redaction](#) (page 8-2)

Use Oracle Data Redaction when you must disguise sensitive data that your applications and application users must access.

[Benefits of Using Oracle Data Redaction](#) (page 8-2)

Oracle Data Redaction provides several benefits when you use it to protect your data.

[Target Use Cases for Oracle Data Redaction](#) (page 8-2)

Oracle Data Redaction fulfills common use case scenarios.

8.1 What Is Oracle Data Redaction?

Oracle Data Redaction enables you to mask (redact) data that is returned from queries issued by applications.

You can redact column data by using one of the following methods:

- **Full redaction.** You redact all of the contents of the column data. The redacted value returned to the querying application user depends on the data type of the column. For example, columns of the `NUMBER` data type are redacted with a zero (0), and character data types are redacted with a single space.
- **Partial redaction.** You redact a portion of the column data. For example, you can redact a Social Security number with asterisks (*), except for the last 4 digits.
- **Regular expressions.** You can use regular expressions to look for patterns of data to redact. For example, you can use regular expressions to redact email addresses, which can have varying character lengths. It is designed for use with character data only.
- **Random redaction.** The redacted data presented to the querying application user appears as randomly generated values each time it is displayed, depending on the data type of the column.
- **No redaction.** The None redaction type option enables you to test the internal operation of your redaction policies, with no effect on the results of queries against tables with policies defined on them. You can use this option to test the redaction policy definitions before applying them to a production environment.

Oracle Database applies the redaction at runtime, when users access the data (that is, at query-execution time). This solution works well in a production system. During the

time that the data is being redacted, all of the data processing is performed normally, and the back-end referential integrity constraints are preserved.

Data redaction can help you to comply with industry regulations such as Payment Card Industry Data Security Standard (PCI DSS) and the Sarbanes-Oxley Act.

See Also:

- *Oracle Database 2 Day + Security Guide* for a tutorial about creating Oracle Data Redaction policies
 - *Oracle Database Security Guide* for information about using Transparent Sensitive Data Protection policies with Oracle Data Redaction
-
-

8.2 When to Use Oracle Data Redaction

Use Oracle Data Redaction when you must disguise sensitive data that your applications and application users must access.

Data Redaction enables you to easily disguise the data using several different redaction styles.

Oracle Data Redaction is ideal for situations in which you must redact specific characters out of the result set of queries of Personally Identifiable Information (PII) returned to certain application users. For example, you may want to present a U.S. Social Security number that ends with the numbers 4320 as ***-**-4320.

Oracle Data Redaction is particularly suited for call center applications and other applications that are read-only. Take care when using Oracle Data Redaction with applications that perform updates back to the database, because redacted data can be written back to this database.

8.3 Benefits of Using Oracle Data Redaction

Oracle Data Redaction provides several benefits when you use it to protect your data.

These benefits are as follows:

- You have different styles of redaction from which to choose.
- Because the data is redacted at runtime, Data Redaction is well suited to environments in which data is constantly changing.
- You can create the Data Redaction policies in one central location and easily manage them from there.
- The Data Redaction policies enable you to create a wide variety of function conditions based on `SYS_CONTEXT` values, which can be used at runtime to decide when the Data Redaction policies will apply to the results of the application user's query.

8.4 Target Use Cases for Oracle Data Redaction

Oracle Data Redaction fulfills common use case scenarios.

[Oracle Data Redaction Use with Database Applications](#) (page 8-3)

Oracle Data Redaction protects sensitive data that is displayed in database applications.

[Oracle Data Redaction with Ad Hoc Database Queries Considerations](#) (page 8-3)

You may encounter situations where it is convenient to redact sensitive data for ad hoc queries that are performed by database users.

8.4.1 Oracle Data Redaction Use with Database Applications

Oracle Data Redaction protects sensitive data that is displayed in database applications.

Data Redaction is transparent to application users because it preserves the original data type and (optionally) the formatting. It is highly transparent to the database because the data remains the same in buffers, caches, and storage—only being changed at the last minute just before SQL query results are returned to the caller. The redaction is enforced consistently across all of the applications that use the same underlying database. You can specify which application users should see only redacted data by checking application user information that is passed into the database through the `SYS_CONTEXT` function; you can redact data based on attributes of the current database or application user; and you can implement multiple logical conditions within a given redaction policy. In addition, Data Redaction is implemented in a way that minimizes performance overhead. These characteristics make Oracle Data Redaction particularly well suited for usage by a range of applications, analytics tools, reporting tools, and monitoring tools that share common production databases. Although its primary target is redaction of production data for applications, Oracle Data Redaction also can be used in combination with Oracle Enterprise Manager Data Masking and Subsetting Pack for protecting sensitive data in testing and development environments.

See Also:

- *Oracle Data Masking and Subsetting Guide* for more information about data masking and subsetting
 - [Oracle Data Redaction and Data Masking and Subsetting Pack](#)
(page 12-8)
-
-

8.4.2 Oracle Data Redaction with Ad Hoc Database Queries Considerations

You may encounter situations where it is convenient to redact sensitive data for ad hoc queries that are performed by database users.

For example, in the course of supporting a production application, a user may need to run ad hoc database queries to troubleshoot and fix an urgent problem with the application. This is different from the application-based scenarios described in [Oracle Data Redaction Use with Database Applications](#) (page 8-3), which typically generate a bounded set of SQL queries, use defined database accounts, and have fixed privileges.

Even though Oracle Data Redaction is not designed to prevent data exposure to database users who run ad hoc queries directly against the database, it can provide an additional layer to reduce the chances of accidental data exposure. Because such users may have rights to change data, alter the database schema, and circumvent the SQL query interface entirely, it is possible for a malicious user to bypass Data Redaction policies in certain circumstances.

Remember that the Oracle Database security tools are designed to be used together to improve overall security. By deploying one or more of these tools as a complement to Oracle Data Redaction, you can securely increase your overall security posture.

Related Topics:

[Oracle Data Redaction General Usage Guidelines](#) (page 13-1)

It is important to understand general guidelines for using Oracle Data Redaction.

Oracle Data Redaction Features and Capabilities

Oracle Data Redaction provides a variety of ways to redact different types of data.

[Full Data Redaction to Redact All Data](#) (page 9-1)

Full data redaction redacts the entire contents of the specified table or view column.

[Partial Data Redaction to Redact Sections of Data](#) (page 9-2)

In partial data redaction, you redact portions of the displayed output.

[Regular Expressions to Redact Patterns of Data](#) (page 9-3)

Regular expressions redact specific data within a column data value, based on a pattern search.

[Redaction Using Null Values](#) (page 9-4)

You can create an Oracle Data Redaction policy that redacts column data by replacing it with null values.

[Random Data Redaction to Generate Random Values](#) (page 9-4)

In random data redaction, the entire value is redacted by replacing it with a random value.

[Comparison of Full, Partial, and Random Redaction Based on Data Types](#) (page 9-5)

The full, partial, and random data redaction styles affect the Oracle built-in, ANSI, user-defined, and Oracle supplied types in different ways.

[No Redaction for Testing Purposes](#) (page 9-9)

You can create a Data Redaction policy that does not perform redaction.

[Central Management of Named Data Redaction Policy Expressions](#) (page 9-9)

You can create a library of named policy expressions that can be used in the columns of multiple tables and views.

9.1 Full Data Redaction to Redact All Data

Full data redaction redacts the entire contents of the specified table or view column.

By default the output is displayed as follows:

- **Character data types:** The output text is a single space.
- **Number data types:** The output text is a zero (0).
- **Date-time data types:** The output text is set to the first day of January, 2001, which appears as 01-JAN-01.

Full redaction is the default and is used whenever a Data Redaction policy specifies the column but omits the `function_type` parameter setting. When you run the `DBMS_REDACT.ADD_POLICY` procedure, to set the `function_type` parameter setting for full redaction, you enter the following setting:

```
function_type => DBMS_REDACT.FULL
```

You can use the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure to change the full redaction output to different values.

Related Topics:

[Syntax for Creating a Full Redaction Policy](#) (page 10-25)

The `DBMS_REDACT.ADD_POLICY` procedure enables you to create a full redaction policy.

9.2 Partial Data Redaction to Redact Sections of Data

In partial data redaction, you redact portions of the displayed output.

You can set the position within the [actual data](#) at which to begin the redaction, the number of characters to redact starting from that position, and the redaction character to use. This type of redaction is useful for situations where you want it to be obvious to the person viewing the data that it was redacted in some way. Typically, you use this type of redaction for credit cards or ID numbers.

Be aware that partial data redaction requires that your data width remain fixed. If you want to redact columns containing string values of variable length, then you must use regular expressions.

To specify partial redaction, you must set the `DBMS_REDACT.ADD_POLICY` procedure `function_type` parameter to `DBMS_REDACT.PARTIAL` and use the `function_parameters` parameter to define the partial redaction behavior.

The displayed output for partial data redaction can be as follows:

- Character data types:** When partially redacted, a Social Security number (represented as a hyphenated string within a character data type) with value 987-65-4320 could be redacted so that it is displayed as shown in the following examples. The code on the right specifies how to redact the character data: it specifies the expected input format of the actual data, the format to use for the display of the redacted output, the start position at which to begin the redaction, the character to use for the redaction, and how many characters to redact. The first example uses a predefined format (in previous releases called a shortcut) for character data type Social Security numbers, and the second example replaces the first five numbers with an asterisk (*) while preserving the hyphens (-) in between the numbers.

```
XXX-XX-4320    function_parameters => DBMS_REDACT.REDACT_US_SSN_F5,
***-**-4320    function_parameters => 'VVVFVVVFVVV,VVV-VV-VVVV',*,1,5',
```

- Number data types:** The partially redacted NUMBER data type Social Security number 987654328 could appear as follows. Both redact the first five digits. The first example uses a predefined format that is designed for Social Security numbers in the NUMBER data type, and the second replaces the first five numbers with the number 9, starting from the first digit.

```
XXXXX4328    function_parameters => DBMS_REDACT.REDACT_NUM_US_SSN_F5,
```

```
999994328    function_parameters => '9,1,5',
```

- **Date-time data types:** Partially redacted datetime values can appear simply as different dates. For example, the date 29-AUG-11 10.20.50.000000 AM could appear as follows. In the first example, the day of the month is redacted to 02 (using the setting d02) and in the second example, the month is redacted to DEC (using m12). The uppercase values show the actual month (M), year (Y), hour (H), minute (M), and second (S).

```
02-AUG-11 10.20.50.000000 AM    function_parameters => 'Md02YHMS',
```

```
29-DEC-11 10.20.50.000000 AM    function_parameters => 'm12DYHMS',
```

Related Topics:

[Regular Expressions to Redact Patterns of Data](#) (page 9-3)

Regular expressions redact specific data within a column data value, based on a pattern search.

[Syntax for Creating a Partial Redaction Policy](#) (page 10-31)

The DBMS_REDACT.ADD_POLICY statement enables you to create policies that redact specific parts of the data returned to the application.

9.3 Regular Expressions to Redact Patterns of Data

Regular expressions redact specific data within a column data value, based on a pattern search.

For example, you can redact the user name of email addresses, so that only the domain shows (for example, replacing hpreston in the email address hpreston@example.com with [redacted] so that it appears as [redacted]@example.com). To perform the redaction, set the DBMS_REDACT.ADD_POLICY procedure function_type parameter to either DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH, and then use the following parameters to build the regular expression:

- A string search pattern (that is, the values to search for), such as:

```
regexp_pattern    => '(.)@(.\.[A-Za-z]{2,4})'
```

This setting looks for a pattern of the following form:

```
one_or_more_characters@one_or_more_characters.2-4_characters_in_range_A-Z_or_a-z
```

- A replacement string, which replaces the value matched by the regexp_pattern setting. The replacement string can include back references to sub-expressions of the main regular expression pattern. The following example replaces the data before the @ symbol (from the regexp_pattern setting) with the text [redacted]. The \2 setting refers to the second match group, which is (.\.[A-Za-z]{2,4}) from the regexp_pattern setting.

```
regexp_replace_string => '[redacted]@\2'
```

- The starting position for the string search string, such as the first character of the data, such as:

```
regexp_position    => DBMS_REDACT.RE_BEGINNING
```

- The kind of search and replace operation to perform, such as the first occurrence, every fifth occurrence, or all of the occurrences, such as:

```
regexp_occurrence    => DBMS_REDACT.RE_ALL
```

- The default matching behavior for the search and replace operation, such as whether the search is case-sensitive (i sets it to be not case-sensitive):

```
regexp_match_parameter => 'i
```

In addition to the default parameters, you can use a set of predefined formats that enable you to use commonly used regular expressions for telephone numbers, email addresses, and credit card numbers.

Related Topics:

[Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40)

The `regexp_*` parameters of the `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression-based redaction policy.

9.4 Redaction Using Null Values

You can create an Oracle Data Redaction policy that redacts column data by replacing it with null values.

This feature enables you to use the `DBMS_REDACT.NULLIFY` function to hide all of the sensitive data in a table or view column and replace it with null values. You can set this function by using the `function_type` parameter of the `DBMS_REDACT.ADD_POLICY` or `DBMS_REDACT.ALTER_POLICY` procedure.

For example:

```
function_type        => DBMS_REDACT.NULLIFY
```

Related Topics:

[Creating a DBMS_REDACT.NULLIFY Redaction Policy](#) (page 10-28)

You can create Oracle Data Redaction policies that return null values for the displayed value of the table or view column.

9.5 Random Data Redaction to Generate Random Values

In random data redaction, the entire value is redacted by replacing it with a random value.

The redacted values displayed in the result set of the query change randomly each time application users run the query.

This type of redaction is useful in cases where you do not want it to be obvious that the data was redacted. It works especially well for number and datetime data types, where it is difficult to distinguish between random and real data.

The displayed output for random values changes based on the data type of the redacted column, as follows:

- **Character data types:** The random output is a mixture of characters (for example, HTU[G{\pjkEWcK). It behaves differently for the `CHAR` and `VARCHAR2` data types, as follows:
 - **CHAR data type:** The redacted output is always in the same character set as the character set of the column. The byte length of the redacted output is

always the same as the column definition length (that is, the column length that was provided at the time of table creation). For example, if the column is `CHAR(20)`, then a string of 20 random characters is provided in the redacted output of the user's query.

- **VARCHAR2 data type:** For random redaction of a `VARCHAR` data type, the redacted output is always in the same character set as the character set of the column. The length of the redacted output is limited based on the length of the **actual data** in the column. No characters in excess of the length of the actual data are displayed. For example, if the column is `VARCHAR2(20)` and the row being redacted contains actual data with a length of 12, then a string of 12 random characters (not 20) is provided in the redacted output of the user's query for that row.
- **Number data types:** Each actual number value is redacted by replacing it with a random, non-negative number modulo the absolute value of the actual data. This redaction results in random numbers that do not exceed the precision of the actual data. For example, the number 987654321 can be redacted by replacing it with any of the numbers 12345678, 13579, 0, or 987654320, but not by replacing it with any of the numbers 987654321, 99987654321, or -1. The number -123 could be redacted by replacing it with the numbers 122, 0, or 83, but not by replacing it with any of the numbers 123, 1123, or -2.

The only exception to the above is when the actual value is an integer between -1 and 9. In this case, the actual data is redacted by replacing it with a random, non-negative integer modulo ten (10).

- **Date-time data types:** When values of the date data type are redacted using random Data Redaction, Oracle Database displays them with random dates that are always different from those of the actual data.

The setting for using random redaction is as follows:

```
function_type => DBMS_REDACT.RANDOM
```

Related Topics:

[Syntax for Creating a Random Redaction Policy](#) (page 10-48)

A random redaction policy presents the redacted data to the querying application user as randomly generated values, based on the column data type.

9.6 Comparison of Full, Partial, and Random Redaction Based on Data Types

The full, partial, and random data redaction styles affect the Oracle built-in, ANSI, user-defined, and Oracle supplied types in different ways.

[Oracle Built-in Data Types Redaction Capabilities](#) (page 9-6)

Oracle Data Redaction handles the Oracle built-in data types depending on the type of Data Redaction policies that are used.

[ANSI Data Types Redaction Capabilities](#) (page 9-6)

Oracle Data Redaction converts ANSI data types in specific ways, depending on the type of redaction that the Data Redaction policy has.

[Built-in and ANSI Data Types Full Redaction Capabilities](#) (page 9-7)

For full redaction, the default redacted value depends on whether the data type is Oracle built-in or ANSI.

[User-Defined Data Types or Oracle Supplied Types Redaction Capabilities](#) (page 9-9)

Several data types or types are not supported by Oracle Data Redaction.

9.6.1 Oracle Built-in Data Types Redaction Capabilities

Oracle Data Redaction handles the Oracle built-in data types depending on the type of Data Redaction policies that are used.

[Table 9-1](#) (page 9-6) describes the Oracle Data Redaction support for Oracle built-in data types.

Table 9-1 Redaction Support for Oracle Built-in Data Types

Column Data Type	Full	Partial	Regexp	Random
Character ¹	Yes	Yes	Yes	Yes
Number ²	Yes	Yes	Yes	Yes
Raw ³	No	No	No	No
Date-time ⁴	Yes	Yes	Yes	Yes
Interval ⁵	No	No	No	No
BFILE	No	No	No	No
BLOB	Yes	No	No	No
CLOB	Yes	No	Yes	No
NCLOB	Yes	No	Yes	No
ROWID	No	No	No	No
UROWID	No	No	No	No

¹ Includes CHAR, VARCHAR2 (including long VARCHAR2, for example, VARCHAR2(2000)), NCHAR, NVARCHAR2

² Includes NUMBER, FLOAT, BINARY_FLOAT, BINARY_DOUBLE

³ Includes LONG RAW, RAW

⁴ Includes DATE, TIMESTAMP, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITH LOCAL TIME ZONE

⁵ Includes INTERVAL YEAR TO MONTH, INTERVAL DAY TO SECOND

9.6.2 ANSI Data Types Redaction Capabilities

Oracle Data Redaction converts ANSI data types in specific ways, depending on the type of redaction that the Data Redaction policy has.

[Table 9-2](#) (page 9-7) compares how the full, partial, and random redaction styles work for ANSI data types, with regard to how they are converted and their support status.

Table 9-2 Redaction Support for the ANSI Data Types

Data Type	How Converted	Full Redaction	Partial Redaction	Regexp	NULL Redaction	Random Redaction
CHARACTER(n), CHAR(n)	Converted to CHAR(n)	Yes	Yes	Yes	Yes	Yes
CHARACTER VARYING(n), CHAR VARYING(n)	Converted to VARCHAR 2(n)	Yes	Yes	Yes	Yes	Yes
NATIONAL CHARACTER(n), NATIONAL CHAR(n), NCHAR(n)	Converted to NCHAR(n)	Yes	Yes	Yes	Yes	Yes
NATIONAL CHARACTER VARYING(n), NATIONAL CHAR VARYING(n), NCHAR VARYING(n)	Converted to NVARCHAR 2(n)	Yes	Yes	Yes	Yes	Yes
NUMERIC[(p, s)] DECIMAL[(p, s)]	Converted to NUMBER(p, s)	Yes	Yes	Yes	Yes	Yes
INTEGER, INT, SMALLINT	Converted to NUMBER(38)	Yes	Yes	Yes	Yes	Yes
FLOAT, DOUBLE PRECISION	Converted to FLOAT(126)	Yes	Yes	Yes	Yes	Yes
REAL	Converted to FLOAT(63)	Yes	Yes	Yes	Yes	Yes
GRAPHIC, LONG VARGRAPHIC, VARGRAPHIC, TIME	No conversion	No	No	No	No	No

9.6.3 Built-in and ANSI Data Types Full Redaction Capabilities

For full redaction, the default redacted value depends on whether the data type is Oracle built-in or ANSI.

[ANSI Data Types Redaction Capabilities](#) (page 9-6) shows the default settings for both Oracle built-in and ANSI data type columns that use full redaction.

Table 9-3 Default Settings and Categories for Columns That Use Full Redaction

Data Type	Default Redacted Value	Data Type Category
CHARACTER	Single space (" ")	Oracle built-in
CHARACTER(n), CHAR(n)	Single space (" ")	ANSI
CHARACTER VARYING(n), CHAR VARYING(n)	Single space (" ")	ANSI
NATIONAL CHARACTER(n), NATIONAL CHAR(n), NCHAR(n)	Single space (" ")	ANSI
NATIONAL CHARACTER VARYING(n), NATIONAL CHAR VARYING(n), NCHAR VARYING(n)	Single space (" ")	ANSI
NUMBER	Zero (0)	Oracle built-in
NUMERIC(p, s) DECIMAL(p, s)	Zero (0)	Oracle built-in
INTEGER, INT, SMALLINT	Zero (0)	ANSI
FLOAT, DOUBLE PRECISION	Zero (0)	ANSI
REAL	Zero (0)	ANSI
DATE-TIME	01-01-01 or 01-01-01 01:00:00	Oracle built-in
BLOB	Oracle's raw representation of [redacted] ¹	Oracle built-in
CLOB	[redacted]	Oracle built-in
NCLOB	[redacted]	Oracle built-in

¹ If you have changed the character set, then you may need to invoke the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure to set the value to the raw representation in the new character set, as follows:

```
DECLARE
  new_red_blob BLOB;
BEGIN
  DBMS_LOB.CREATETEMPORARY(new_red_blob, TRUE);
  DBMS_LOB.WRITE(new_red_blob, 10, 1, UTL_RAW.CAST_TO_RAW('[redacted]'));

```



```

dbms_redact.update_full_redaction_values(
    blob_val => new_red_blob);
DBMS_LOB.FREETEMPORARY(new_red_blob);
END;
/

```

After you run this procedure, restart the database.

See also [Altering the Default Full Data Redaction Value](#) (page 10-26) for more information about using the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure.

9.6.4 User-Defined Data Types or Oracle Supplied Types Redaction Capabilities

Several data types or types are not supported by Oracle Data Redaction.

[Table 9-4](#) (page 9-9) compares how the full, partial, regular expression, and random redaction styles work for user-defined and Oracle-supplied types.

Table 9-4 Redaction Support for the User-Defined Data Types or Oracle-Supplied Types

Data Type or Type	Full Redaction	Partial Redaction	Regexp	NULL Redaction	Random Redaction
User-defined data types	No	No	No	No	No
Oracle supplied types: Any types, XML types, Oracle Spatial types, Oracle Media types	No	No	No	No	No

9.7 No Redaction for Testing Purposes

You can create a Data Redaction policy that does not perform redaction.

This is useful for cases in which you have a redacted base table, yet you want a specific application user to have a view that always shows the [actual data](#). You can create a new view of the redacted table and then define a Data Redaction policy for this view. The policy still exists on the base table, but no redaction is performed when the application queries using the view as long as the `DBMS_REDACT.NONE` `function_type` setting was used to create a policy on the view.

9.8 Central Management of Named Data Redaction Policy Expressions

You can create a library of named policy expressions that can be used in the columns of multiple tables and views.

By having named policy expressions, you can centrally manage all of the policy expressions within a database.

When you modify the policy expression, the change is reflected in all table columns that use the expression. The policy expression takes precedence over the `expression` setting in the Data Redaction policy. To create the policy expression, you must use the `DBMS_REDACT.CREATE_POLICY_EXPRESSION` procedure, and to apply the policy expression to a column, you use `DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL`. This feature provides flexibility to redact different columns in a table or view, based on different runtime conditions.

For example, consider a use case that involves a customer care application. A customer calls the customer care center to request a return on a recent purchase. A level 1 support representative of the call center must first verify the order ID, customer name, and customer address before initiating the return. During the process, there is no need

for the level 1 support representative to view the customer's credit card number. So, the credit card column is redacted when the support representative queries the customer details in the call center application. When the return is initiated, a sales representative from the return department may need to view the credit card number to process the return. However, there is no need for the sales representative to view the expiration date of the credit card. So, when the sales representative queries the customer details in the same application, the credit card number is visible but the expiration date is redacted.

In this use case, different columns in the customer details table must be redacted in different ways, based on who the logged in user is. Oracle Data Redaction simplifies the implementation of this use case by using named Data Redaction policy expressions. This type of policy expression enables you to define and associate different policy expressions on different columns in the same table or view. Moreover, you can centrally manage named policy expressions within a database. Any updates that you make to a named policy expression are immediately propagated to all of the associated table or view columns.

Related Topics:

[Creating and Managing Multiple Named Policy Expressions](#) (page 10-13)

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

Configuring Oracle Data Redaction Policies

An Oracle Data Redaction policy defines how to redact data in a column based on the table column type and the type of redaction you want to use.

[About Oracle Data Redaction Policies](#) (page 10-2)

An Oracle Data Redaction policy defines the conditions in which redaction must occur for a table or view.

[Who Can Create Oracle Data Redaction Policies?](#) (page 10-3)

Because data redaction involves the protection of highly sensitive data, only trusted users should create Oracle Data Redaction policies.

[Planning an Oracle Data Redaction Policy](#) (page 10-4)

Before you create a Oracle Data Redaction policy, you should plan the data redaction policy that best suits your site's needs.

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4)

To create a Data Redaction policy, you must use the DBMS_REDACT.ADD_POLICY procedure.

[Using Expressions to Define Conditions for Data Redaction Policies](#) (page 10-6)

The `expression` parameter in the DBMS_REDACT.ADD_POLICY procedure sets the conditions to which the policy applies.

[Creating and Managing Multiple Named Policy Expressions](#) (page 10-13)

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

[Creating a Full Redaction Policy and Altering the Full Redaction Value](#) (page 10-24)

You can create a full redaction policy to redact all contents in a data column, and optionally, you can alter the default full redaction value.

[Creating a DBMS_REDACT.NULLIFY Redaction Policy](#) (page 10-28)

You can create Oracle Data Redaction policies that return null values for the displayed value of the table or view column.

[Creating a Partial Redaction Policy](#) (page 10-30)

In partial data redaction, you can redact portions of data, and for different kinds of data types.

[Creating a Regular Expression-Based Redaction Policy](#) (page 10-39)

A regular expression-based redaction policy enables you to redact data based on a search-and-replace model.

[Creating a Random Redaction Policy](#) (page 10-47)

A random redaction policy presents redacted data as randomly generated values, such as UkJjs132[[]]s.

[Creating a Policy That Uses No Redaction](#) (page 10-49)

You can create policies that use no redaction at all, for when you want to test the policy in a development environment.

[Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)

You can exempt users from having Oracle Data Redaction policies applied to the data they access.

[Altering an Oracle Data Redaction Policy](#) (page 10-51)

The `DBMS_REDACT.ALTER_POLICY` procedure enables you to modify Oracle Data Redaction policies.

[Redacting Multiple Columns](#) (page 10-56)

You can redact more than one column in a Data Redaction policy.

[Disabling and Enabling an Oracle Data Redaction Policy](#) (page 10-57)

You can disable and then reenable Oracle Data Redactions policies as necessary.

[Dropping an Oracle Data Redaction Policy](#) (page 10-59)

The `DBMS_REDACT.DROP_POLICY` procedure drops Oracle Data Redaction policies.

[Tutorial: SQL Expressions to Build Reports with Redacted Values](#) (page 10-60)

SQL expressions can be used to build reports based on columns that have Oracle Data Redaction policies defined on them.

[Oracle Data Redaction Policy Data Dictionary Views](#) (page 10-61)

Oracle Database provides data dictionary views that list information about Data Redaction policies.

10.1 About Oracle Data Redaction Policies

An Oracle Data Redaction policy defines the conditions in which redaction must occur for a table or view.

A Data Redaction policy has the following characteristics:

- The Data Redaction policy defines the following: What kind of redaction to perform, how the redaction should occur, and when the redaction takes place. Oracle Database performs the redaction at execution time, just before the data is returned to the application.
- A Data Redaction policy can fully redact values, partially redact values, or randomly redact values. In addition, you can define a Data Redaction policy to not redact any data at all, for when you want to test your policies in a test environment.
- A Data Redaction policy can be defined with a policy expression which allows for different application users to be presented with either redacted data or **actual data**, based on whether the policy expression returns `TRUE` or `FALSE`. Redaction takes place when the boolean result of evaluating the policy expression is `TRUE`. For security reasons, the functions and operators that can be used in the policy expression are limited to `SYS_CONTEXT` and a few others. User-created functions

are not allowed. Policy expressions can make use of the `SYS_SESSION_ROLES` namespace with the `SYS_CONTEXT` function to check for enabled roles.

- Different Data Redaction policy expressions can be created and then applied individually for different columns within the same table or view.

[Table 10-1](#) (page 10-3) lists the procedures in the `DBMS_REDACT` package.

Table 10-1 DBMS_REDACT Procedures

Procedure	Description
<code>DBMS_REDACT.ADD_POLICY</code>	Adds a Data Redaction policy to a table or view
<code>DBMS_REDACT.ALTER_POLICY</code>	Modifies a Data Redaction policy
<code>DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL</code>	Applies a Data Redaction policy expression to a table or view column
<code>DBMS_REDACT.CREATE_POLICY_EXPRESSION</code>	Creates a Data Redaction policy expression
<code>DBMS_REDACT.DISABLE_POLICY</code>	Disables a Data Redaction policy
<code>DBMS_REDACT.DROP_POLICY</code>	Drops a Data Redaction policy
<code>DBMS_REDACT.DROP_POLICY_EXPRESSION</code>	Drops a Data Redaction policy expression
<code>DBMS_REDACT.ENABLE_POLICY</code>	Enables a Data Redaction policy
<code>DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES</code>	Globally updates the full redaction value for a given data type. You must restart the database instance before the updated values can be used.
<code>DBMS_REDACT.UPDATE_POLICY_EXPRESSION</code>	Updates a Data Redaction policy expression

See Also:

- [Oracle Database PL/SQL Packages and Types Reference](#) for detailed information about the `DBMS_REDACT` PL/SQL package
 - [Using Oracle Data Redaction in Oracle Enterprise Manager](#) (page 11-1) for information about using Oracle Enterprise Manager Cloud Control to create and manage Oracle Data Redaction policies and formats
-

10.2 Who Can Create Oracle Data Redaction Policies?

Because data redaction involves the protection of highly sensitive data, only trusted users should create Oracle Data Redaction policies.

To create redaction policies, you must have the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package. To find the privileges that a user has been granted, you can query the `DBA_SYS_PRIVS` data dictionary view.

You do not need any privileges to access the underlying tables or views that will be protected by the policy.

10.3 Planning an Oracle Data Redaction Policy

Before you create a Oracle Data Redaction policy, you should plan the data redaction policy that best suits your site's needs.

1. Ensure that you have been granted the EXECUTE privilege on the DBMS_REDACT PL/SQL package.
2. Determine the data type of the table or view column that you want to redact.
3. Ensure that this column is not used in an Oracle Virtual Private Database (VPD) row filtering condition. That is, it must not be part of the VPD predicate generated by the VPD policy function.
4. Decide on the type of redaction that you want to perform: full, random, partial, regular expressions, or none.
5. Decide which users to apply the Data Redaction policy to.
6. Based on this information, create the Data Redaction policy by using the DBMS_REDACT.ADD_POLICY procedure.
7. Configure the policy to have additional columns to be redacted.

After you create the Data Redaction policy, it is automatically enabled and ready to redact data.

Related Topics:

[Redacting Multiple Columns](#) (page 10-56)

You can redact more than one column in a Data Redaction policy.

10.4 General Syntax of the DBMS_REDACT.ADD_POLICY Procedure

To create a Data Redaction policy, you must use the DBMS_REDACT.ADD_POLICY procedure.

The complete syntax for the DBMS_REDACT.ADD_POLICY procedure is as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema          IN VARCHAR2 := NULL,
  object_name            IN VARCHAR2 := NULL,
  policy_name            IN VARCHAR2,
  policy_description     IN VARCHAR2 := NULL,
  column_name            IN VARCHAR2 := NULL,
  column_description    IN VARCHAR2 := NULL,
  function_type          IN BINARY_INTEGER := DBMS_REDACT.FULL,
  function_parameters    IN VARCHAR2 := NULL,
  expression             IN VARCHAR2,
  enable                 IN BOOLEAN := TRUE,
  regexp_pattern         IN VARCHAR2 := NULL,
  regexp_replace_string  IN VARCHAR2 := NULL,
  regexp_position       IN BINARY_INTEGER :=1,
  regexp_occurrence     IN BINARY_INTEGER :=0,
  regexp_match_parameter IN VARCHAR2 := NULL);
```

In this specification:

- `object_schema`: Specifies the schema of the object on which the Data Redaction policy will be applied. If you omit this setting (or enter NULL), then Oracle

Database uses the current user's name. Be aware that the meaning of "current user" here can change, depending on where you invoke the DBMS_REDACT.ADD_POLICY procedure.

For example, suppose user `mpike` grants user `fbrown` the EXECUTE privilege on a definer's rights PL/SQL package called `mpike.protect_data` in `mpike`'s schema. From within this package, `mpike` has coded a procedure called `protect_cust_data`, which invokes the DBMS_REDACT.ADD_POLICY procedure. User `mpike` has set the `object_schema` parameter to NULL.

When `fbrown` invokes the `protect_cust_data` procedure in the `mpike.protect_data` package, Oracle Database attempts to define the Data Redaction policy around the object `cust_data` in the `mpike` schema, not the `cust_data` object in the schema that belongs to `fbrown`.

- `object_name`: Specifies the name of the table or view to which the Data Redaction policy applies.
- `policy_name`: Specifies the name of the policy to be created. Ensure that this name is unique in the database instance. You can find a list of existing Data Redaction policies by querying the POLICY_NAME column of the REDACTION_POLICIES data dictionary view.
- `policy_description`: Specifies a brief description of the purpose of the policy.
- `column_name`: Specifies the column whose data you want to redact. Note the following:
 - **You can apply the Data Redaction policy to multiple columns.** If you want to apply the Data Redaction policy to multiple columns, then after you use DBMS_REDACT.ADD_POLICY to create the policy, run the DBMS_REDACT.ALTER_POLICY procedure as many times as necessary to add each of the remaining required columns to the policy. See [Altering an Oracle Data Redaction Policy](#) (page 10-51).
 - **Only one policy can be defined on a table or view.** You can, however, create a new view on the table, and by defining a second redaction policy on this new view, you can choose to redact the columns in a different way when a query is issued against this new view. When deciding how to redact a given column, Oracle Database uses the policy of the earliest view in a view chain.
 - **If you do not specify a column (for example, by entering NULL), then no columns are redacted by the policy.** This enables you to create your policies so that they are in place, and then later on, you can add the column specification when you are ready.
 - **Do not use a column that is currently used in an Oracle Virtual Private Database (VPD) row filtering condition.** In other words, the column should not be part of the VPD predicate generated by the VPD policy function. (See [Oracle Data Redaction and Oracle Virtual Private Database](#) (page 12-4) for more information about using Data Redaction with VPD.)
 - **You cannot define a Data Redaction policy on a virtual column.** In addition, you cannot define a Data Redaction policy on a column that is involved in the SQL expression of any virtual column.
- `column_description`: Specifies a brief description of the column that you are redacting.

- `function_type`: Specifies a function that sets the type of redaction. See the following sections for more information:
 - [Syntax for Creating a Full Redaction Policy](#) (page 10-25)
 - [Syntax for Creating a Partial Redaction Policy](#) (page 10-31)
 - [Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40)
 - [Syntax for Creating a Random Redaction Policy](#) (page 10-48)
 - [Syntax for Creating a Policy with No Redaction](#) (page 10-49)

If you omit the `function_type` parameter, then the default redaction `function_type` setting is `DBMS_REDACT.FULL`.

- `function_parameters`: Specifies how the column redaction should appear for partial redaction. See [Syntax for Creating a Partial Redaction Policy](#) (page 10-31).
- `expression`: Specifies a Boolean SQL expression to determine how the policy is applied. Redaction takes place only if this policy expression evaluates to `TRUE`. See [Using Expressions to Define Conditions for Data Redaction Policies](#) (page 10-6).
- `enable`: When set to `TRUE`, enables the policy upon creation. When set to `FALSE`, it creates the policy as a disabled policy. The default is `TRUE`. After you create the policy, you can disable or enable it. See the following sections:
 - [Disabling an Oracle Data Redaction Policy](#) (page 10-57)
 - [Enabling an Oracle Data Redaction Policy](#) (page 10-58)
- `regexp_pattern`, `regexp_replace_string`, `regexp_position`, `regexp_position`, `regexp_occurrence`, `regexp_match_parameter`: Enable you to use regular expressions to redact data, either fully or partially. If the `regexp_pattern` does not match anything in the **actual data**, then full redaction will take place, so be careful when specifying the `regexp_pattern`. Ensure that all of the values in the column conform to the semantics of the regular expression you are using. See [Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40) for more information.

10.5 Using Expressions to Define Conditions for Data Redaction Policies

The `expression` parameter in the `DBMS_REDACT.ADD_POLICY` procedure sets the conditions to which the policy applies.

[About Using Expressions in Data Redaction Policies](#) (page 10-7)

The `DBMS_REDACT.ADD_POLICY` and `DBMS_REDACT.ALTER_POLICY` `expression` parameter defines a Boolean expression that must evaluate to `TRUE` to enable a redaction.

[Supported Functions for Data Redaction Expressions](#) (page 10-7)

You can create expressions that use functions to return specific types of data, such as `SYS_CONTEXT` namespaces.

[Applying the Redaction Policy Based on User Environment](#) (page 10-11)

You can apply a Data Redaction policy based on the user's environment, such as the session user name or a client identifier.

[Applying the Redaction Policy Based on Database Roles](#) (page 10-11)

You can apply a Data Redaction policy based on a database role, such as the DBA role.

[Applying the Redaction Policy Based on Oracle Label Security Label Dominance](#) (page 10-12)

You can set a condition on which to apply a Data Redaction policy based on the dominance of Oracle Label Security labels.

[Applying the Redaction Policy Based on Application Express Session States](#) (page 10-12)

You can apply a Data Redaction policy based on an Oracle Application Express (APEX) session state.

[Applying the Redaction Policy to All Users](#) (page 10-13)

You can apply the policy irrespective of the context to any user, with no filtering.

10.5.1 About Using Expressions in Data Redaction Policies

The `DBMS_REDACT.ADD_POLICY` and `DBMS_REDACT.ALTER_POLICY` expression parameter defines a Boolean expression that must evaluate to `TRUE` to enable a redaction.

The expression that is defined in the `expression` parameter is the default expression for the Oracle Data Redaction policy. If you apply a named policy expression for the columns that will be redacted by the Data Redaction policy, then the named policy expression takes precedence over the expression defined in the Data Redaction policy.

You can create expressions that make use of other Oracle Database features. For example, you can create expressions that are based on a user's environment (using the `SYS_CONTEXT` and `XS_SYS_CONTEXT` functions), character string functions, the Oracle Label Security label dominance functions, or Oracle Application Express functions.

Follow these guidelines when you write the expression:

- Use only the following operators: `AND`, `OR`, `IN`, `NOT IN`, `=`, `!=`, `<>`, `<`, `>`, `>=`, `<=`
- Because the expression must evaluate to `TRUE` for redaction, be careful when making comparisons with `NULL`. Remember that in SQL the value `NULL` is undefined, so comparisons with `NULL` tend to return `FALSE`.
- Do not use user-created functions in the `expression` parameter; this is not permitted.
- Remember that for user `SYS` and users who have the `EXEMPT REDACTION POLICY` privilege, all of the Data Redaction policies are bypassed, so the results of their queries are not redacted. See the following topics for more information about users who are exempted from Data Redaction policies:
 - [Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)
 - [Oracle Data Pump Security Model for Oracle Data Redaction](#) (page 12-5)

10.5.2 Supported Functions for Data Redaction Expressions

You can create expressions that use functions to return specific types of data, such as `SYS_CONTEXT` namespaces.

Expressions Using Namespace Functions (page 10-8)

You can use the `SYS_CONTEXT` and `XS_SYS_CONTEXT` namespace functions in Data Redaction expressions.

Expressions Using the SUBSTR Function (page 10-9)

You can use the `SUBSTR` function, which returns portion (such as characters 1–3) of the character string specified, in Data Redaction expressions. The first parameter must be a constant string or a call to the `SYS_CONTEXT` function or the `XS_SYS_CONTEXT` function.

Expressions Using Length of Character String Functions (page 10-9)

You can use the following functions, which return the length of character strings, in Data Redaction expressions. Oracle Database also checks that the arguments to each of these operators is either a constant string or a call to the `SYS_CONTEXT` or `XS_SYS_CONTEXT` function.

Expressions Using Oracle Application Express Functions (page 10-10)

You can use Oracle Application Express functions in Data Redaction expressions.

Expressions Using Oracle Label Security Functions (page 10-10)

You can use Oracle Label Security functions with Data Redaction expressions.

10.5.2.1 Expressions Using Namespace Functions

You can use the `SYS_CONTEXT` and `XS_SYS_CONTEXT` namespace functions in Data Redaction expressions.

Table 10-2 Expressions Using Namespace Functions

Namespace Function	Description
<code>SYS_CONTEXT</code>	Returns the value associated with a namespace. The following namespace functions are valid: <ul style="list-style-type: none"> <code>USERENV</code> (default namespace), which includes values such as <code>SESSION_USER</code> and <code>CLIENT_IDENTIFIER</code>. <code>SYS_SESSION_ROLES</code>, which contains attributes for each role <code>XS\$SESSION</code>, which contains attributes for the user session. User-defined namespaces, but these must exist in the <code>DBA_CONTEXT</code> catalog view before the policy expression is created.
<code>XS_SYS_CONTEXT</code>	Similar to <code>SYS_CONTEXT</code> but designed for an Oracle Real Application Security environment. <code>XS_SYS_CONTEXT</code> supports the same namespaces that <code>SYS_CONTEXT</code> supports.

See Also:

- *Oracle Database SQL Language Reference* for more information about `SYS_CONTEXT`
- *Oracle Database Real Application Security Administrator's and Developer's Guide* for more information about `XS_SYS_CONTEXT`

10.5.2.2 Expressions Using the SUBSTR Function

You can use the SUBSTR function, which returns portion (such as characters 1–3) of the character string specified, in Data Redaction expressions. The first parameter must be a constant string or a call to the SYS_CONTEXT function or the XS_SYS_CONTEXT function.

Table 10-3 Expressions Using SUBSTR String Functions

SUBSTR String Function	Description
SUBSTR	Returns a portion of the input char value, beginning at <i>character position</i> , <i>substring_length</i> characters long. SUBSTR calculates length using characters as defined by the input character set.
SUBSTRB	Returns the specified portion of the input value in bytes
SUBSTRC	Returns the specified portion of the input value in Unicode complete characters
SUBSTR2	Returns the specified portion of the input value in UCS2 code points
SUBSTR4	Returns the specified portion of the input value in UCS4 code points

See Also: *Oracle Database SQL Language Reference* for more information about the SUBSTR functions

10.5.2.3 Expressions Using Length of Character String Functions

You can use the following functions, which return the length of character strings, in Data Redaction expressions. Oracle Database also checks that the arguments to each of these operators is either a constant string or a call to the SYS_CONTEXT or XS_SYS_CONTEXT function.

Table 10-4 Expressions Using Character String Functions

Character String Function	Description
LENGTH	Returns the length of the input char value. LENGTH calculates length using characters as defined by the input character set.
LENGTHB	Returns the length of the input value in bytes
LENGTHC	Returns the length of the input value in Unicode complete characters
LENGTH2	Returns the length of the input value in UCS2 code points
LENGTH4	Returns the length of the input value in UCS4 code points

See Also: *Oracle Database SQL Language Reference* for more information about the `LENGTH` functions

10.5.2.4 Expressions Using Oracle Application Express Functions

You can use Oracle Application Express functions in Data Redaction expressions.

Table 10-5 Oracle Application Express Functions

Oracle Application Express Function	Description
V	Returns the session state for an item. It is a wrapper for the <code>APEX_UTIL.GET_SESSION_STATE</code> function
NV	Returns the numeric value for a numeric item. It is a wrapper for the <code>APEX_UTIL.GET_NUMERIC_SESSION_STATE</code> function

See Also: *Oracle Application Express API Reference* for more information about the Oracle Application Express functions

10.5.2.5 Expressions Using Oracle Label Security Functions

You can use Oracle Label Security functions with Data Redaction expressions.

For the functions in the bold font, Oracle Data Redaction checks that their parameters are either constants or calls to only one of the `SA_UTL.NUMERIC_LABEL`, `CHAR_TO_LABEL`, and `SA_SESSION.LABEL` functions, and that the arguments to those functions are constant.

Table 10-6 Oracle Label Security Functions

Oracle Label Security Function	Description
<code>LBACSYS.OLS_LABEL_DOMINATES</code>	Checks if the session label of an Oracle Label Security policy dominates or is equal to another OLS label
DOMINATES	Checks if one OLS label is dominant to a second OLS label. Deprecated in Oracle Database 12c release 1 (12.1); use the <code>OLS_DOMINATES</code> or <code>OLS_DOM</code> function instead.
<code>OLS_DOMINATES</code>	Checks if one OLS label is dominant to a second OLS label
<code>OLS_DOM</code>	Checks if one OLS label is dominant to a second OLS label
<code>DOM</code>	Checks if one OLS label is dominant to a second OLS label
<code>OLS_STRICTLY_DOMINATES</code>	Checks if one OLS label is dominant to a second OLS label and is not equal to it

Table 10-6 (Cont.) Oracle Label Security Functions

Oracle Label Security Function	Description
<code>STRICTLY_DOMINATES</code>	Checks if one OLS label is dominant to a second OLS label and is not equal to it
<code>S_DOM</code>	Checks if one OLS label is dominant to a second OLS label and is not equal to it. Deprecated in Oracle Database 12c release 1 (12.1); use the <code>OLS_DOMINATES</code> or <code>OLS_DOM</code> function instead.
<code>SA_UTL.DOMINATES</code>	Checks if one OLS label dominates a second OLS label or if the session label for a given OLS policy dominates an OLS label
<code>SA_UTL.CHECK_READ</code>	Checks if a user can read a policy-protected row
<code>SA_UTL.NUMERIC_LABEL</code>	Returns the current session OLS label
<code>CHAR_TO_LABEL</code>	Converts a character string to an OLS label tag
<code>SA_SESSION.LABEL</code>	Returns the label that is associated with the specified OLS policy

See Also: *Oracle Application Express API Reference* for more information about the Oracle Label Security functions

10.5.3 Applying the Redaction Policy Based on User Environment

You can apply a Data Redaction policy based on the user's environment, such as the session user name or a client identifier.

- Use the `USERENV` namespace of the `SYS_CONTEXT` function in the `DBMS_REDACT.ADD_POLICY` expression parameter to apply the policy based on a user's environment.

For example, to apply the policy only to the session user name `psmith`:

```
expression => 'SYS_CONTEXT(''USERENV'', 'SESSION_USER') = 'PSMITH''
```

See Also:

Oracle Database SQL Language Reference for information about more namespaces that you can use with the `SYS_CONTEXT` function

10.5.4 Applying the Redaction Policy Based on Database Roles

You can apply a Data Redaction policy based on a database role, such as the DBA role.

- Use the `SYS_SESSION_ROLES` namespace in the `SYS_CONTEXT` function to apply the policy based on a user role.

This namespace contains attributes for each role. The value of the attribute is `TRUE` if the specified role is enabled for the querying application user; the value is `FALSE` if the role is not enabled.

For example, suppose you wanted only supervisors to be allowed to see the [actual data](#). The following example shows how to use the `DBMS_REDACT.ADD_POLICY` expression parameter to set the policy to show the actual data to any application user who has the `supervisor` role enabled, but redact the data for all of the other application users.

```
expression => 'SYS_CONTEXT(''SYS_SESSION_ROLES'', 'SUPERVISOR') = 'FALSE'''
```

10.5.5 Applying the Redaction Policy Based on Oracle Label Security Label Dominance

You can set a condition on which to apply a Data Redaction policy based on the dominance of Oracle Label Security labels.

- Use the public standalone function `OLS_LABEL_DOMINATES` to check the dominance of a session label. This function returns 1 (`TRUE`) if the session label of the specified `policy_name` value dominates or is equal to the label that is specified by the `label` parameter; otherwise, it returns 0 (`FALSE`).

For example, to apply a Data Redaction policy only in cases where the session label for the policy `hr_ols_pol` does not dominate nor is equal to label `hs`:

```
expression => 'OLS_LABEL_DOMINATES (''hr_ols_pol'', 'hs') = 0'
```

10.5.6 Applying the Redaction Policy Based on Application Express Session States

You can apply a Data Redaction policy based on an Oracle Application Express (APEX) session state.

- Use either of the following public Application Express APIs in the `DBMS_REDACT.ADD_POLICY` expression parameter to apply the policy on an Oracle Application Express session state:
 - `V`, which is a synonym for the `APEX_UTIL.GET_SESSION_STATE` function
 - `NV`, which is a synonym for the `APEX_UTIL.GET_NUMERIC_SESSION_STATE` function

For example, to set the `DBMS_REDACT.ADD_POLICY` expression parameter if you wanted redaction to take place when the application item called `G_JOB` has the value `CLERK`:

```
expression => 'V(''APP_USER'') != 'mavis@example.com' or V(''APP_USER'') is null'
```

You can, for example, use these functions to redact data based on a job or a privilege role that is stored in a session state in an APEX application.

If you want redaction to take place when the querying user is *not* within the context of an APEX application (when the query is issued from outside the APEX framework, for example directly through SQL*Plus), then use an `IS NULL` clause as follows. This policy expression causes actual data to be shown to user `mavis` only when her query comes from within an APEX application. Otherwise, the query result is redacted.

See Also:

Oracle Application Express API Reference

10.5.7 Applying the Redaction Policy to All Users

You can apply the policy irrespective of the context to any user, with no filtering.

However, be aware that user `SYS` and users who have the `EXEMPT REDACTION POLICY` privilege are always except from Oracle Data Redaction policies.

- To apply the policy to users who are not `SYS` or have been granted the `EXEMPT REDACTION POLICY` privilege, write the `DBMS_REDACT.ADD_POLICY` expression parameter to evaluate to `TRUE`.

For example:

```
expression => '1=1'
```

See Also:

[Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)

10.6 Creating and Managing Multiple Named Policy Expressions

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

[About Data Redaction Policy Expressions to Define Conditions](#) (page 10-13)

A named Oracle Data Redaction policy expression is designed to work as an alternative to the policy expression that is used in existing Data Redaction policies.

[Creating and Applying a Named Data Redaction Policy Expression](#) (page 10-15)

The `DBMS_REDACT.CREATE_POLICY_EXPRESSION` and `DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL` enable you to create and apply a named Data Redaction policy expression.

[Updating a Named Data Redaction Policy Expression](#) (page 10-16)

You can use the `DBMS_REDACT.UPDATE_POLICY_EXPRESSION` procedure to update a Data Redaction policy expression. The update takes place immediately and is reflected in all columns that use the policy expression.

[Dropping a Named Data Redaction Expression Policy](#) (page 10-16)

You can use the `DBMS_REDACT.DROP_POLICY_EXPRESSION` procedure to drop a Data Redaction expression policy.

[Tutorial: Creating and Sharing a Named Data Redaction Policy Expression](#) (page 10-17)

This tutorial shows how to create an Oracle Data Redaction policy expression, apply it to multiple tables, and centrally manage the policy expression.

10.6.1 About Data Redaction Policy Expressions to Define Conditions

A named Oracle Data Redaction policy expression is designed to work as an alternative to the policy expression that is used in existing Data Redaction policies.

A named policy expression enables you to redact data based on runtime conditions. This type of policy can only affect whether or not redaction takes place on columns of

the table or view on which the redaction policy is defined. By default, a Data Redaction policy expression applies to all the columns of a table or view. Alternatively, you can choose to create and associate a policy expression for individual columns of a table or view. These column level expressions are called as named policy expressions; in other words, a policy expression with a name. A named policy expressions has the following properties:

You can use Data Redaction policy expressions in the following ways.:

- A single Data Redaction policy expression can be shared by more than one Data Redaction policy by applying it to columns that are a part of separate Data Redaction policies.
- Each named policy expression can be associated with multiple columns of the same or different tables or views.
- Each named policy expression can be associated with columns within the same or different Data Redaction policies.
- The named policy expression overrides the default policy expression of the associated columns. The default policy expression still applies to redaction columns that have no named policy expressions applied to them.
- Any updates made to a named policy expression apply to all of the column associations of the expression.
- You cannot associate multiple named policy expressions for the same column.
- In a multitenant environment, you cannot associate named policy expressions with columns in a different pluggable database (PDB).

The column to which you apply a named policy expression must already be redacted by a Data Redaction policy. After the named policy expression is applied, the result of its evaluation takes precedence over that of the default policy expression when deciding whether or not to redact the column. When you modify a named policy expression, the changes are applied to all the tables and views that use it. In a multitenant environment, as with Data Redaction policies, a named policy expression is valid only in the PDB in which it was created, and can only be applied to columns of objects within the PDB in which it was created.

[Table 10-7](#) (page 10-14) describes the DBMS_REDACT PL/SQL procedures that you can use to create and manage named policy expressions. To find information about policy expressions, query the REDACTION_EXPRESSIONS data dictionary view.

Table 10-7 DBMS_REDACT Policy Expression Procedures

Procedure	Description
DBMS_REDACT.CREATE_POLICY_EXPRESSION	Creates a Data Redaction policy expression
DBMS_REDACT.UPDATE_POLICY_EXPRESSION	Updates a Data Redaction policy expression
DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL	Applies a Data Redaction policy expression to a table or a view column
DBMS_REDACT.DROP_POLICY_EXPRESSION	Drops a Data Redaction policy expression

Related Topics:

[Managing Named Data Redaction Policy Expressions Using Enterprise Manager](#)
(page 11-17)

You can manage Oracle Data Redaction policy expressions in Enterprise Manager Cloud Control.

10.6.2 Creating and Applying a Named Data Redaction Policy Expression

The `DBMS_REDACT.CREATE_POLICY_EXPRESSION` and `DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL` enable you to create and apply a named Data Redaction policy expression.

1. Ensure that the `COMPATIBLE` initialization parameter is set to `12.2.0.0`.

To find the current setting, use the `SHOW PARAMETER` command.

2. To create the policy expression, run the `DBMS_REDACT.CREATE_POLICY_EXPRESSION` procedure.

For example:

```
BEGIN
  DBMS_REDACT.CREATE_POLICY_EXPRESSION (
    policy_expression_name => 'redact_pol',
    expression              => '1=1',
    description             => 'Determines whether the column will be redacted');
END;
/
```

3. Run the `DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL` procedure to apply the policy expression to a table or view column.

For example, assume that you have already created a Data Redaction policy on the `SALARY` column of the `HR.EMPLOYEES` table, as follows:

```
BEGIN
  DBMS_REDACT.ADD_POLICY (
    object_schema    => 'hr',
    object_name      => 'employees',
    policy_name      => 'overall_policy',
    expression       => '1=0');
END;
/

BEGIN
  DBMS_REDACT.ALTER_POLICY (
    object_schema    => 'hr',
    object_name      => 'employees',
    policy_name      => 'overall_policy',
    function_type    => DBMS_REDACT.FULL,
    action           => DBMS_REDACT.ADD_COLUMN,
    column_name      => 'SALARY ');
END;
/
```

Then you can apply the policy expression to the `SALARY` table as follows:

```
BEGIN
  DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL (
    object_schema    => 'hr',
    object_name      => 'employees',
```

```
column_name          => 'salary',
policy_expression_name => 'redact_pol');
END;
/
```

In this specification:

- `object_schema`: Specifies the schema of the object on which the policy expression will be used. If you omit this setting (or enter `NULL`), then Oracle Database uses the name of the current schema.
- `object_name`: Specifies the name of the table or view to be used for the policy expression.
- `column_name`: Specifies the column to which you want to apply the policy expression.
- `policy_expression_name`: Specifies the name of the policy expression.

After you create an Oracle Data Redaction policy expression, you can apply it to a column of a table or view which is part of an existing Data Redaction policy.

10.6.3 Updating a Named Data Redaction Policy Expression

You can use the `DBMS_REDACT.UPDATE_POLICY_EXPRESSION` procedure to update a Data Redaction policy expression. The update takes place immediately and is reflected in all columns that use the policy expression.

You can query the `REDACTION_EXPRESSIONS` data dictionary view to find existing Data Redaction policy expressions.

1. Ensure that the `COMPATIBLE` initialization parameter is set to `12.2.0.0`.

To find the current setting, use the `SHOW PARAMETER` command.

2. Run the `DBMS_REDACT.UPDATE_POLICY_EXPRESSION` procedure to perform the update.

For example:

```
BEGIN
DBMS_REDACT.UPDATE_POLICY_EXPRESSION(
  policy_expression_name => 'redact_pol',
  expression              => '1=0');
END;
/
```

10.6.4 Dropping a Named Data Redaction Expression Policy

You can use the `DBMS_REDACT.DROP_POLICY_EXPRESSION` procedure to drop a Data Redaction expression policy.

You can query the `REDACTION_EXPRESSIONS` data dictionary view to find existing Data Redaction policy expressions.

1. Ensure that the `COMPATIBLE` initialization parameter is set to `12.2.0.0`.

To find the current setting, use the `SHOW PARAMETER` command.

2. Remove the named policy expression's association with any table or view column.

You cannot drop a policy expression if it is associated with an existing table or view column. To remove a given column's association with a named policy expression (to revert to redacting that column based on the evaluation result of the default policy expression), you must set the `policy_expression_name` parameter of the `DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL` procedure to `NULL`.

For example:

```
BEGIN
DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
  object_schema      => 'hr',
  object_name        => 'employees',
  column_name        => 'salary',
  policy_expression_name => null);
END;
/
```

3. Run `DBMS_REDACT.DROP_POLICY_EXPRESSION` to drop the policy expression.

For example:

```
BEGIN
DBMS_REDACT.DROP_POLICY_EXPRESSION(
  policy_expression_name => 'redact_pol');
END;
/
```

10.6.5 Tutorial: Creating and Sharing a Named Data Redaction Policy Expression

This tutorial shows how to create an Oracle Data Redaction policy expression, apply it to multiple tables, and centrally manage the policy expression.

[Step 1: Create Users for This Tutorial](#) (page 10-18)

You must create two users for this tutorial: `dr_admin`, who will create the Oracle Data Redaction policies, and `hr_clerk`, who will test them.

[Step 2: Create an Oracle Data Redaction Policy](#) (page 10-18)

User `dr_admin` is ready to create an Oracle Data Redaction policy to protect the `HR.EMPLOYEES` and `HR.JOBS` tables.

[Step 3: Test the Oracle Data Redaction Policy](#) (page 10-19)

User `hr_clerk` is ready to query the tables that have redacted data.

[Step 4: Create and Apply a Policy Expression to the Redacted Table Columns](#) (page 10-20)

Next, user `dr_admin` is ready to create a Data Redaction policy expression and apply it to two of the three redacted table columns.

[Step 5: Test the Data Redaction Policy Expression](#) (page 10-21)

User `hr_clerk` is now ready to test the `hr_redact_pol` policy expression.

[Step 6: Modify the Data Redaction Policy Expression](#) (page 10-21)

User `dr_admin` decides to modify the Data Redaction policy expression so that user `HR` will have access to the redacted data, not user `hr_clerk`.

[Step 7: Test the Modified Policy Expression](#) (page 10-22)

Users `HR` and `hr_clerk` are ready to test the modified Data Redaction policy expression.

Step 8: Remove the Components of This Tutorial (page 10-23)

If you do not need the components of this tutorial, then you can remove them.

10.6.5.1 Step 1: Create Users for This Tutorial

You must create two users for this tutorial: `dr_admin`, who will create the Oracle Data Redaction policies, and `hr_clerk`, who will test them.

Before you begin this tutorial, ensure that the `COMPATIBLE` initialization parameter is set to `12.2.0.0`. You can check this setting by using the `SHOW PARAMETER` command.

1. Log in to SQL*Plus as user `SYS` with the `SYSDBA` administrative privilege.

```
sqlplus sys as sysdba
Enter password: password
```

2. In a multitenant environment, connect to the appropriate PDB.

For example:

```
CONNECT SYS@my_pdb AS SYSDBA
Enter password: password
```

To find the available PDBs, query the `DBA_PDBS` data dictionary view. To check the current PDB, run the `SHOW CON_NAME` command.

3. Create the `dr_admin` and `hr_clerk` user accounts.

```
GRANT CREATE SESSION TO dr_admin IDENTIFIED BY password;
GRANT CREATE SESSION TO hr_clerk IDENTIFIED BY password;
```

4. Grant the `EXECUTE` privilege to the `dr_admin` user.

```
GRANT EXECUTE ON DBMS_REDACT TO dr_admin;
```

5. Connect as user `HR`.

```
CONNECT HR --Or, for a PDB, CONNECT hr@my_pdb
Enter password: password
```

6. Grant `hr_clerk` the `SELECT` privilege on the `EMPLOYEES` and `JOBS` tables.

```
GRANT SELECT on EMPLOYEES to hr_clerk;
GRANT SELECT on JOBS to hr_clerk;
```

10.6.5.2 Step 2: Create an Oracle Data Redaction Policy

User `dr_admin` is ready to create an Oracle Data Redaction policy to protect the `HR.EMPLOYEES` and `HR.JOBS` tables.

1. Connect as user `dr_admin`.

```
CONNECT dr_admin --Or, for a PDB, CONNECT dr_admin@my_pdb
Enter password: password
```

2. Create the `hr_emp_redact_comp_pol` policy, which will perform full redaction of the `HR.EMPLOYEES.SALARY` column.

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'hr',
```

```

object_name      => 'employees',
column_name      => 'salary',
policy_name      => 'hr_emp_redact_comp_pol',
function_type    => DBMS_REDACT.FULL,
expression       => '1=1');
END;
/

```

3. Alter the `hr_redact_comp_pol` policy to also redact the `COMMISSION_PCT` column of the `HR.EMPLOYEES` table.

```

BEGIN
DBMS_REDACT.ALTER_POLICY(
  object_schema  => 'hr',
  object_name    => 'employees',
  policy_name    => 'hr_emp_redact_comp_pol',
  action         => DBMS_REDACT.ADD_COLUMN,
  column_name    => 'commission_pct',
  function_type  => DBMS_REDACT.FULL,
  expression     => '1=1');
END;
/

```

4. Create the `hr_jobs_redact_comp_pol` policy for the `max_salary` column of the `HR.JOBS` table.

```

BEGIN
DBMS_REDACT.ADD_POLICY(
  object_schema  => 'hr',
  object_name    => 'jobs',
  column_name    => 'max_salary',
  policy_name    => 'hr_jobs_redact_comp_pol',
  function_type  => DBMS_REDACT.FULL,
  expression     => '1=1');
END;
/

```

At this stage, the data in the `HR.EMPLOYEES.SALARY`, `HR.EMPLOYEES.COMMISSION_PCT`, and `HR.JOBS.MAX_SALARY` columns are redacted.

10.6.5.3 Step 3: Test the Oracle Data Redaction Policy

User `hr_clerk` is ready to query the tables that have redacted data.

1. Connect as user `hr_clerk`.

```

CONNECT hr_clerk --Or, for a PDB, CONNECT hr_clerk@my_pdb
Enter password: password

```

2. Query the `HR.EMPLOYEES` table.

```

SELECT SALARY, COMMISSION_PCT FROM HR.EMPLOYEES WHERE SALARY > 15000;

```

The output should be as follows:

```

      SALARY COMMISSION_PCT
-----
          0
          0
          0

```

3. Query the HR.JOBS table.

```
SELECT MAX_SALARY FROM HR.JOBS WHERE MAX_SALARY > 15000;
```

The output should be as follows:

```
MAX_SALARY
-----
          0
          0
          0
          0
          0
```

10.6.5.4 Step 4: Create and Apply a Policy Expression to the Redacted Table Columns

Next, user dr_admin is ready to create a Data Redaction policy expression and apply it to two of the three redacted table columns.

This policy expression will enable user hr_clerk to view the redacted data.

1. Connect as user dr_admin.

```
CONNECT dr_admin --Or, for a PDB, CONNECT dr_admin@my_pdb
Enter password: password
```

2. Create the policy expression.

```
BEGIN
  DBMS_REDACT.CREATE_POLICY_EXPRESSION(
    policy_expression_name => 'hr_redact_pol',
    expression              => 'SYS_CONTEXT(''USERENV'', ''SESSION_USER'') !=
''HR_CLERK''');
END;
/
```

This expression returns FALSE for the hr_clerk user, which enables the hr_clerk user to view actual data in the HR.EMPLOYEES and HR.JOBS tables that are subject to the Data Redaction policies.

3. Apply the hr_redact_pol policy expression to the HR.EMPLOYEES.SALARY column.

```
BEGIN
  DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
    object_schema      => 'hr',
    object_name        => 'employees',
    column_name        => 'salary',
    policy_expression_name => 'hr_redact_pol');
END;
/
```

4. Apply the hr_redact_pol policy expression to the HR.JOBS.MAX_SALARY column.

```
BEGIN
  DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
    object_schema      => 'hr',
    object_name        => 'jobs',
    column_name        => 'max_salary',
    policy_expression_name => 'hr_redact_pol');
```

```
END;
/
```

User `hr_clerk` can view data in the `HR.EMPLOYEES.SALARY` and `HR.JOBS.MAX_SALARY`, but the data in the `HR.EMPLOYEES.COMMISSION_PCT` column will still be redacted for this user.

10.6.5.5 Step 5: Test the Data Redaction Policy Expression

User `hr_clerk` is now ready to test the `hr_redact_pol` policy expression.

1. Connect as user `hr_clerk`.

```
CONNECT hr_clerk --Or, for a PDB, CONNECT hr_clerk@my_pdb
Enter password: password
```

2. Query the `HR.EMPLOYEES` table.

```
SELECT SALARY, COMMISSION_PCT FROM HR.EMPLOYEES WHERE SALARY > 15000;
```

The output should be as follows:

```

SALARY COMMISSION_PCT
-----
24000
17000
17000
```

User `hr_clerk` now can view the `SALARY` column data, but still has not access to the `COMMISSION_PCT` column data.

3. Query the `HR.JOBS` table.

```
SELECT MAX_SALARY FROM HR.JOBS WHERE MAX_SALARY > 15000;
```

The output should be as follows:

```

MAX_SALARY
-----
40000
30000
16000
16000
20080
```

User `hr_clerk` now can view the `MAX_SALARY` column data.

10.6.5.6 Step 6: Modify the Data Redaction Policy Expression

User `dr_admin` decides to modify the Data Redaction policy expression so that user `HR` will have access to the redacted data, not user `hr_clerk`.

1. Connect as user `dr_admin`.

```
CONNECT dr_admin --Or, for a PDB, CONNECT dr_admin@my_pdb
Enter password: password
```

2. Modify the `hr_redact_pol` policy as follows:

```
BEGIN
DBMS_REDACT.UPDATE_POLICY_EXPRESSION(
  policy_expression_name => 'hr_redact_pol',
```

```

        expression          => 'SYS_CONTEXT(''USERENV'', 'SESSION_USER') !=
        ''HR'');
    END;
/

```

10.6.5.7 Step 7: Test the Modified Policy Expression

Users HR and hr_clerk are ready to test the modified Data Redaction policy expression.

1. Connect as user HR.

```

CONNECT HR --Or, for a PDB, CONNECT HR@my_pdb
Enter password: password

```

2. Query the HR.EMPLOYEES table.

```

SELECT SALARY, COMMISSION_PCT FROM HR.EMPLOYEES WHERE SALARY > 15000;

```

The output should be as follows:

```

      SALARY COMMISSION_PCT
-----
      24000
      17000
      17000

```

User HR now has access to the redacted data. A query by HR on the HR.JOBS.MAX_SALARY column will produce similar results.

```

SELECT MAX_SALARY FROM HR.JOBS WHERE MAX_SALARY > 15000;

```

```

MAX_SALARY
-----
      40000
      30000
      16000
      16000
      20080

```

3. Connect as user hr_clerk.

```

CONNECT hr_clerk --Or, for a PDB, CONNECT hr_clerk@my_pdb
Enter password: password

```

4. Query the HR.EMPLOYEES and HR.JOBS tables and then observe the results.

```

SELECT SALARY, COMMISSION_PCT FROM HR.EMPLOYEES WHERE SALARY > 15000;

```

```

      SALARY COMMISSION_PCT
-----
           0
           0
           0

```

```

SELECT MAX_SALARY FROM HR.JOBS WHERE MAX_SALARY > 15000;

```

```

MAX_SALARY
-----
           0
           0
           0

```



```
0
0
```

10.6.5.8 Step 8: Remove the Components of This Tutorial

If you do not need the components of this tutorial, then you can remove them.

1. Connect as user `dr_admin`.

```
CONNECT dr_admin --Or, for a PDB, CONNECT dr_admin@my_pdb
Enter password: password
```

2. Modify the policy expression so that it is no longer associated with the table columns that are associated with the expression.

To do so, you must set the `policy_expression_name` parameter to `NULL`.

```
BEGIN
DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
  object_schema => 'hr',
  object_name   => 'employees',
  column_name   => 'salary',
  policy_expression_name => null);
END;
/
```

```
BEGIN
DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
  object_schema => 'hr',
  object_name   => 'jobs',
  column_name   => 'max_salary',
  policy_expression_name => null);
END;
/
```

3. Drop the policy expressions.

```
BEGIN
DBMS_REDACT.DROP_POLICY_EXPRESSION(
  policy_expression_name => 'hr_redact_pol');
END;
/
```

4. Drop the `hr_emp_redact_comp_pol` and `hr_jobs_redact_comp_pol` Data Redaction policies.

```
BEGIN
DBMS_REDACT.DROP_POLICY (
  object_schema => 'hr',
  object_name   => 'employees',
  policy_name   => 'hr_emp_redact_comp_pol');
END;
/
```

```
BEGIN
DBMS_REDACT.DROP_POLICY (
  object_schema => 'hr',
  object_name   => 'jobs',
  policy_name   => 'hr_jobs_redact_comp_pol');
END;
/
```

5. Connect as the `SYSTEM` user or a user who has privileges to drop user accounts.

For example:

```
CONNECT SYSTEM
Enter password: password
```

6. Drop the `dr_admin` and `hr_clerk` user accounts.

```
DROP USER dr_admin;
DROP USER hr_clerk;
```

10.7 Creating a Full Redaction Policy and Altering the Full Redaction Value

You can create a full redaction policy to redact all contents in a data column, and optionally, you can alter the default full redaction value.

[Creating a Full Redaction Policy](#) (page 10-24)

A full data redaction policy redacts all the contents of a data column.

[Altering the Default Full Data Redaction Value](#) (page 10-26)

The `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure alters the default full data redaction value.

10.7.1 Creating a Full Redaction Policy

A full data redaction policy redacts all the contents of a data column.

[About Creating Full Data Redaction Policies](#) (page 10-24)

To set a redaction policy to redact all data in the column, you must set the `function_type` parameter to `DBMS_REDACT.FULL`.

[Syntax for Creating a Full Redaction Policy](#) (page 10-25)

The `DBMS_REDACT.ADD_POLICY` procedure enables you to create a full redaction policy.

[Example: Full Redaction Policy](#) (page 10-25)

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure to create a full redaction policy.

[Example: Fully Redacted Character Values](#) (page 10-26)

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure to create a policy that fully redacts character values.

10.7.1.1 About Creating Full Data Redaction Policies

To set a redaction policy to redact all data in the column, you must set the `function_type` parameter to `DBMS_REDACT.FULL`.

By default, `NUMBER` data type columns are replaced with zero (0) and character data type columns are replaced with a single space (). You can modify this default by using the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure.

Related Topics:

[Altering the Default Full Data Redaction Value](#) (page 10-26)

The `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure alters the default full data redaction value.

10.7.1.2 Syntax for Creating a Full Redaction Policy

The `DBMS_REDACT.ADD_POLICY` procedure enables you to create a full redaction policy.

The `DBMS_REDACT.ADD_POLICY` fields for creating a full data redaction policy are as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type      IN BINARY_INTEGER := NULL,
  expression         IN VARCHAR2,
  enable            IN BOOLEAN := TRUE);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4).
- `function_type`: Specifies the function used to set the type of redaction. Enter `DBMS_REDACT.FULL`.

If you omit the `function_type` parameter, then the default redaction `function_type` setting is `DBMS_REDACT.FULL`.

Remember that the data type of the column determines which `function_type` settings that you are permitted to use. See [Comparison of Full, Partial, and Random Redaction Based on Data Types](#) (page 9-5).

10.7.1.3 Example: Full Redaction Policy

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure to create a full redaction policy.

[Example 10-1](#) (page 10-25) shows how to use full redaction for all the values in the `HR.EMPLOYEES` table `COMMISSION_PCT` column. The `expression` parameter applies the policy to any user querying the table, except for users who have been granted the `EXEMPT REDACTION POLICY` system privilege.

Example 10-1 Full Data Redaction Policy

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'hr',
    object_name   => 'employees',
    column_name   => 'commission_pct',
    policy_name   => 'redact_com_pct',
    function_type => DBMS_REDACT.FULL,
    expression    => '1=1');
END;
/
```

Query and redacted result:

```
SELECT COMMISSION_PCT FROM HR.EMPLOYEES;

COMMISSION_PCT
-----
```

0
0
0

Related Topics:

[Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)

You can exempt users from having Oracle Data Redaction policies applied to the data they access.

10.7.1.4 Example: Fully Redacted Character Values

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure to create a policy that fully redacts character values.

[Example 10-2](#) (page 10-26) shows how to redact fully the user IDs of the `user_id` column in the `mavis.cust_info` table. The `user_id` column is of the `VARCHAR2` data type. The output is a blank string. The `expression` setting enables users who have the `MGR` role to view the user IDs.

Example 10-2 Fully Redacted Character Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    column_name   => 'user_id',
    policy_name   => 'redact_cust_user_ids',
    function_type => DBMS_REDACT.FULL,
    expression    => 'SYS_CONTEXT(''SYS_SESSION_ROLES'', 'MGR') = ''FALSE''');
END;
/
```

Query and redacted result:

```
SELECT user_id FROM mavis.cust_info;
```

```
USER_ID
-----
0
0
0
```

10.7.2 Altering the Default Full Data Redaction Value

The `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure alters the default full data redaction value.

[About Altering the Default Full Data Redaction Value](#) (page 10-27)

You can alter the default displayed values for full Data Redaction policies.

[Syntax for the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES Procedure](#) (page 10-27)

The `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure accommodates the standard supported Oracle Database data types.

[Modifying the Default Full Data Redaction Value](#) (page 10-28)

To modify the default full data redaction value, use the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure.

10.7.2.1 About Altering the Default Full Data Redaction Value

You can alter the default displayed values for full Data Redaction policies.

By default, 0 is the redacted value when Oracle Database performs full redaction (DBMS_REDACT.FULL) on a column of the NUMBER data type. If you want to change it to another value (for example, 7), then you can run the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure to modify this value. The modification applies to all of the Data Redaction policies in the current database instance. After you modify a value, you must restart the database for it to take effect. You can find the current values by querying the REDACTION_VALUES_FOR_TYPE_FULL data dictionary view.

Be aware that this change affects all Data Redaction policies in the database that use full data redaction. Before you alter the default full data redaction value, examine the affect that this change would have on existing full Data Redaction policies.

10.7.2.2 Syntax for the DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES Procedure

The DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES procedure accommodates the standard supported Oracle Database data types.

The syntax is as follows:

```
DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES (
  number_val      IN NUMBER                NULL,
  binfloat_val   IN BINARY_FLOAT          NULL,
  bindouble_val  IN BINARY_DOUBLE         NULL,
  char_val       IN CHAR                   NULL,
  varchar_val    IN VARCHAR2              NULL,
  nchar_val      IN NCHAR                  NULL,
  nvarchar_val   IN NVARCHAR2             NULL,
  date_val       IN DATE                   NULL,
  ts_val         IN TIMESTAMP              NULL,
  tswtz_val      IN TIMESTAMP WITH TIME ZONE NULL,
  blob_val       IN BLOB                   NULL,
  clob_val       IN CLOB                   NULL,
  nclob_val      IN NCLOB                  NULL);
```

In this specification:

- `number_val` modifies the default value for columns of the NUMBER data type.
- `binfloat_val` modifies the default value for columns of the BINARY_FLOAT data type.
- `bindouble_val` modifies the default value for columns of the BINARY_DOUBLE data type.
- `char_val` modifies the default value for columns of the CHAR data type.
- `varchar_val` modifies the default value for columns of the VARCHAR2 data type.
- `nchar_val` modifies the default value for columns of the NCHAR data type.
- `nvarchar_val` modifies the default value for columns of the NVARCHAR2 data type.
- `date_val` modifies the default value for columns of the DATE data type.

- `ts_val` modifies the default value for columns of the `TIMESTAMP` data type.
- `tswtz_val` modifies the default value for columns of the `TIMESTAMP WITH TIME ZONE` data type.
- `blob_val` modifies the default value for columns of the `BLOB` data type.
- `clob_val` modifies the default value for columns of the `CLOB` data type.
- `nclob` modifies the default value for columns of the `NCLOB` data type.

10.7.2.3 Modifying the Default Full Data Redaction Value

To modify the default full data redaction value, use the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure.

1. Log in to the database instance as a user who has been granted the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package and who has administrative privileges, such as users who have been granted the `DBA` role.
2. Check the value that you want to change.

For example, to check the current value for columns that use the `NUMBER` data type:

```
SELECT NUMBER_VALUE FROM REDACTION_VALUES_FOR_TYPE_FULL;

NUMBER_VALUE
-----
           0
```

3. Run the `DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES` procedure to modify the value.

For example:

```
EXEC DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES (number_val => 7);
```

4. Restart the database instance.

For example:

```
SHUTDOWN IMMEDIATE

STARTUP
```

10.8 Creating a DBMS_REDACT.NULLIFY Redaction Policy

You can create Oracle Data Redaction policies that return null values for the displayed value of the table or view column.

[About Creating a Policy That Returns Null Values](#) (page 10-29)

The `DBMS_REDACT.NULLIFY` `function_type` parameter redacts all the data in a column and replace it with null values.

[Syntax for Creating a Policy That Returns Null Values](#) (page 10-29)

The `DBMS_REDACT.ADD_POLICY` procedure can create a redaction policy that performs a full redaction and displays null values for the redacted columns.

Example: Redaction Policy That Returns Null Values (page 10-29)

The `DBMS_REDACT.ADD_POLICY` procedure will return null values for the `COMMISSION_PCT` column of the `HR.EMPLOYEES` table.

10.8.1 About Creating a Policy That Returns Null Values

The `DBMS_REDACT.NULLIFY` `function_type` parameter redacts all the data in a column and replace it with null values.

You can use this function type on all supported column types that the `DBMS_REDACT.FULL` function type supports. It also supports the `CLOB` and `NCLOB` data types. To use the `DBMS_REDACT.NULLIFY` function, you must first ensure that the `COMPATIBLE` parameter is set at a minimum to `12.2.0.0.0`.

10.8.2 Syntax for Creating a Policy That Returns Null Values

The `DBMS_REDACT.ADD_POLICY` procedure can create a redaction policy that performs a full redaction and displays null values for the redacted columns.

The syntax for using `DBMS_REDACT.ADD_POLICY` to return null values is as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type       IN BINARY_INTEGER := NULL,
  expression         IN VARCHAR2,
  enable             IN BOOLEAN := TRUE);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4).
- `function_type`: Specifies the function used to set the type of redaction. Enter `DBMS_REDACT.NULLIFY`.

If you omit the `function_type` parameter, then the default setting is `DBMS_REDACT.FULL`.

Remember that the data type of the column determines which `function_type` settings that you are permitted to use. See [Comparison of Full, Partial, and Random Redaction Based on Data Types](#) (page 9-5).

10.8.3 Example: Redaction Policy That Returns Null Values

The `DBMS_REDACT.ADD_POLICY` procedure will return null values for the `COMMISSION_PCT` column of the `HR.EMPLOYEES` table.

The `expression` parameter applies the policy to any user who queries the table, except for users who have been granted the `EXEMPT REDACTION POLICY` system privilege.

[Example 10-3](#) (page 10-29) shows how to create the Oracle Data Redaction policy.

Example 10-3 Redaction Policy That Returns Null Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
```

```
object_schema => 'hr',
object_name   => 'employees',
column_name   => 'commission_pct',
policy_name   => 'nullify_com_pct',
function_type => DBMS_REDACT.NULLIFY,
expression    => '1=1';
END;
/
```

Query and redacted result:

```
SELECT COMMISSION_PCT FROM HR.EMPLOYEES;

COMMISSION_PCT
-----
```

Related Topics:

[Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)

You can exempt users from having Oracle Data Redaction policies applied to the data they access.

10.9 Creating a Partial Redaction Policy

In partial data redaction, you can redact portions of data, and for different kinds of data types.

[About Creating Partial Redaction Policies](#) (page 10-30)

In partial data redaction, only a portion of the data, such as the first five digits of an identification number, are redacted.

[Syntax for Creating a Partial Redaction Policy](#) (page 10-31)

The `DBMS_REDACT.ADD_POLICY` statement enables you to create policies that redact specific parts of the data returned to the application.

[Creating Partial Redaction Policies Using Fixed Character Formats](#) (page 10-31)

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter enables you to use fixed character formats.

[Creating Partial Redaction Policies Using Character Data Types](#) (page 10-35)

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter enables you to redact character data types.

[Creating Partial Redaction Policies Using Number Data Types](#) (page 10-36)

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter can redact number data types.

[Creating Partial Redaction Policies Using Date-Time Data Types](#) (page 10-38)

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter can redact date-time data types.

10.9.1 About Creating Partial Redaction Policies

In partial data redaction, only a portion of the data, such as the first five digits of an identification number, are redacted.

For example, you can redact most of a credit card number with asterisks (*), except for the last 4 digits. You can create policies for columns that use character, number, or

date-time data types. For policies that redact character data types, you can use fixed character redaction formats. If you have the Enterprise Manager for Oracle Database 12.1.0.7 plug-in deployed on your system, then you can also create and save custom redaction formats.

Note:

In previous releases, the term shortcut was used for the term format.

10.9.2 Syntax for Creating a Partial Redaction Policy

The `DBMS_REDACT.ADD_POLICY` statement enables you to create policies that redact specific parts of the data returned to the application.

The `DBMS_REDACT.ADD_POLICY` fields for creating a partial redaction policy are as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type       IN BINARY_INTEGER := NULL,
  function_parameters IN VARCHAR2 := NULL,
  expression         IN VARCHAR2,
  enable             IN BOOLEAN := TRUE);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4)
- `function_type`: Specifies the function used to set the type of redaction. Enter `DBMS_REDACT.PARTIAL`.
- `function_parameters`: The parameters that you set here depend on the data type of the column specified for the `column_name` parameter. See the following sections for details:
 - [Creating Partial Redaction Policies Using Fixed Character Formats](#) (page 10-31)
 - [Creating Partial Redaction Policies Using Character Data Types](#) (page 10-35)
 - [Creating Partial Redaction Policies Using Number Data Types](#) (page 10-36)
 - [Creating Partial Redaction Policies Using Date-Time Data Types](#) (page 10-38)

10.9.3 Creating Partial Redaction Policies Using Fixed Character Formats

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter enables you to use fixed character formats.

[Settings for Fixed Character Formats](#) (page 10-32)

Oracle Data Redaction provides special predefined formats to configure policies that use fixed characters.

[Example: Partial Redaction Policy Using a Fixed Character Format](#) (page 10-34)
 You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure to create a partial redaction policy that uses a fixed character format.

10.9.3.1 Settings for Fixed Character Formats

Oracle Data Redaction provides special predefined formats to configure policies that use fixed characters.

[Table 10-8](#) (page 10-32) describes `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter formats that you can use for commonly redacted identity numbers (such as Social Security numbers or Canadian Social Insurance Numbers), postal codes, and credit cards that use either the `VARCHAR2` or `NUMBER` data types for their columns.

Table 10-8 Partial Fixed Character Redaction Formats

Format	Description
<code>DBMS_REDACT.REDACT_US_SSN_F5</code>	Redacts the first 5 numbers of Social Security numbers when the column is a <code>VARCHAR2</code> data type. For example, the number 987-65-4320 becomes XXX-XX-4320.
<code>DBMS_REDACT.REDACT_US_SSN_L4</code>	Redacts the last 4 numbers of Social Security numbers when the column is a <code>VARCHAR2</code> data type. For example, the number 987-65-4320 becomes 987-65-XXXX.
<code>DBMS_REDACT.REDACT_US_SSN_ENTIRE</code>	Redacts the entire Social Security number when the column is a <code>VARCHAR2</code> data type. For example, the number 987-65-4320 becomes XXX-XX-XXXX.
<code>DBMS_REDACT.REDACT_NUM_US_SSN_F5</code>	Redacts the first 5 numbers of Social Security numbers when the column is a <code>NUMBER</code> data type. For example, the number 987654320 becomes XXXXX4320.
<code>DBMS_REDACT.REDACT_NUM_US_SSN_L4</code>	Redacts the last 4 numbers of Social Security numbers when the column is a <code>NUMBER</code> data type. For example, the number 987654320 becomes 98765XXXX.
<code>DBMS_REDACT.REDACT_NUM_US_SSN_ENTIRE</code>	Redacts the entire Social Security number when the column is a <code>NUMBER</code> data type. For example, the number 987654320 becomes XXXXXXXXX.
<code>DBMS_REDACT.REDACT_SIN_NUMBER</code>	Redacts the Canadian Social Insurance number by replacing the first 6 digits by 9 (number). For example, 123456789 is redacted to 999999789.
<code>DBMS_REDACT.REDACT_SIN_UNFORMATTED</code>	Redacts the Canadian Social Insurance number by replacing the first 6 digits by X (string). For example, 123456789 is redacted to XXXXXX789.

Table 10-8 (Cont.) Partial Fixed Character Redaction Formats

Format	Description
DBMS_REDACT.REDACT_SIN_FORMATTED	Redacts the Canadian Social Insurance Number by replacing the first 6 digits by X (string). For example, 123-456-789 is redacted to XXX-XXX-789.
DBMS_REDACT.REDACT_UK_NIN_FORMATTED	Redacts the UK National Insurance number by replacing the first 6 digits by X (string) but leaving the alphabetic characters as is. For example, ET 27 02 23 D is redacted to ET XX XX XX D.
DBMS_REDACT.REDACT_UK_NIN_UNFORMATTED	Redacts the UK National Insurance number by replacing the first 6 digits by X (string) and leaving the alphabetic characters as is. For example, ET270223D is redacted to ETXXXXXXD.
DBMS_REDACT.REDACT_CCN_FORMATTED	Redacts the credit card number (other than American Express) by replacing everything but the last 4 digits by *. For example, the credit card number 5105-1051-0510-5100 is redacted to ****-****-****-5100.
DBMS_REDACT.REDACT_CCN_NUMBER	Redacts the credit card number (other than American Express) by replacing everything but the last 4 digits by 0. For example, the credit card number 5105105105105100 is redacted to *****5100.
DBMS_REDACT.REDACT_CCN16_F12	Redacts a 16-digit credit card number (other than American Express), leaving the last 4 digits displayed. For example, 5105 1051 0510 5100 becomes ****-****-****-5100.
DBMS_REDACT.REDACT_AMEX_CCN_FORMATTED	Redacts the American Express credit card number by replacing the digits with * except the last 5 digits. For example, the credit card number 3782 822463 10005 is redacted to **** * 10005.
DBMS_REDACT.REDACT_AMEX_CCN_NUMBER	Redacts the American Express Credit Card Number by replacing the digits with 0 except the last 5 digits. For example, the credit card number 3782 822463 10005 is redacted to 0000 000000 10005.
DBMS_REDACT.REDACT_ZIP_CODE	Redacts a 5-digit postal code when the column is a VARCHAR2 data type. For example, 95476 becomes XXXXX.
DBMS_REDACT.REDACT_NUM_ZIP_CODE	Redacts a 5-digit postal code when the column is a NUMBER data type. For example, 95476 becomes XXXXX.
DBMS_REDACT.REDACT_DATE_EPOCH	Redacts all dates to 01-JAN-70.

Table 10-8 (Cont.) Partial Fixed Character Redaction Formats

Format	Description
DBMS_REDACT.REDACT_NA_PHONE_FORMATTED	Redacts the North American phone number by leaving the area code, but replacing everything else with X. For example, 650-555-0100 is redacted to 650-XXX-XXXX.
DBMS_REDACT.REDACT_NA_PHONE_NUMBER	Redacts the North American phone number by leaving the area code, but replacing everything else with 0. For example, 6505550100 gets redacted to 650000000.
DBMS_REDACT.REDACT_NA_PHONE_UNFORMATTED	Redacts the North American phone number by leaving the area code, but replacing everything else with X. For example, 6505550100 is redacted to 650XXXXXXXX.
DBMS_REDACT.REDACT_DATE_MILLENNIUM	Redacts dates that are in the DD-MON-YY format to 01-JAN-00 (January 1, 2000).

See Also:

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4) for information about other DBMS_REDACT.ADD_POLICY parameters

10.9.3.2 Example: Partial Redaction Policy Using a Fixed Character Format

You can use the DBMS_REDACT.ADD_POLICY PL/SQL procedure to create a partial redaction policy that uses a fixed character format.

[Example 10-4](#) (page 10-34) shows how Social Security numbers in a VARCHAR2 data type column and can be redacted using the REDACT_US_SSN_F5 format.

Example 10-4 Partially Redacted Character Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    column_name  => 'ssn',
    policy_name   => 'redact_cust_ssns3',
    function_type => DBMS_REDACT.PARTIAL,
    function_parameters => DBMS_REDACT.REDACT_US_SSN_F5,
    expression   => 'l=1',
    policy_description => 'Partially redacts 1st 5 digits in SS numbers',
    column_description => 'ssn contains Social Security numbers');
END;
/
```

Query and redacted result:

```
SELECT ssn FROM mavis.cust_info;

SSN
-----
XXX-XX-4320
XXX-XX-4323
```

XXX-XX-4325
XXX-XX-4329

10.9.4 Creating Partial Redaction Policies Using Character Data Types

The `DBMS_REDACT.ADD_POLICY function_parameters` parameter enables you to redact character data types.

[Settings for Character Data Types](#) (page 10-35)

Oracle Data Redaction provides special settings to configure policies that use character data types.

[Example: Partial Redaction Policy Using a Character Data Type](#) (page 10-36)

The `DBMS_REDACT.ADD_POLICY PL/SQL` procedure can create a partial redaction policy that uses a character data type.

10.9.4.1 Settings for Character Data Types

Oracle Data Redaction provides special settings to configure policies that use character data types.

When you set the `DBMS_REDACT.ADD_POLICY function_parameters` parameter to define partial redaction of character data types, enter values for the following settings in the order shown. Separate each value with a comma

Note:

Be aware that you must use a fixed width character set for the partial redaction. In other words, each character redacted must be replaced by another of equal byte length. If you want to use a variable-length character set (for example, UTF-8), then you must use a regular expression-based redaction. See [Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40) for more information.

The settings are as follows:

1. **Input format:** Defines how the data is currently formatted. Enter `V` for each character that potentially can be redacted, such as all of the digits in a credit card number. Enter `F` for each character that you want to format using a formatting character, such as hyphens or blank spaces in the credit card number. Ensure that each character has a corresponding `V` or `F` value. (The input format values are not case-sensitive.)
2. **Output format:** Defines how the displayed data should be formatted. Enter `V` for each character to be potentially redacted. Replace each `F` character in the input format with the character that you want to use for the displayed output, such as a hyphen. (The output format values are not case-sensitive.)
3. **Mask character:** Specifies the character to be used for the redaction. Enter a single character to use for the redaction, such as an asterisk (*).
4. **Starting digit position:** Specifies the starting `V` digit position for the redaction.
5. **Ending digit position:** Specifies the ending `V` digit position for the redaction. Do not include the `F` positions when you decide on the ending position value.

For example, the following setting redacts the first 12 V digits of the credit card number 5105 1051 0510 5100, and replaces the F positions (which are blank spaces) with hyphens to format it in a style normally used for credit card numbers, resulting in ****-****-****-4320.

```
function_parameters => 'VVVVFVVVVFVVVVFVVV,VVVV-VVVV-VVVV-VVVV,*,1,12',
```

See Also:

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4) for information about other DBMS_REDACT.ADD_POLICY parameters

10.9.4.2 Example: Partial Redaction Policy Using a Character Data Type

The DBMS_REDACT.ADD_POLICY PL/SQL procedure can create a partial redaction policy that uses a character data type.

[Example 10-5](#) (page 10-36) shows how to redact Social Security numbers that are in a VARCHAR2 data type column and to preserve the character hyphens in the Social Security number.

Example 10-5 Partially Redacted Character Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema    => 'mavis',
    object_name      => 'cust_info',
    column_name      => 'ssn',
    policy_name      => 'redact_cust_ssns2',
    function_type     => DBMS_REDACT.PARTIAL,
    function_parameters => 'VVVVFVVVVFVVV,VVV-VV-VVVV,*,1,5',
    expression       => '1=1',
    policy_description => 'Partially redacts Social Security numbers',
    column_description => 'ssn contains character Social Security numbers');
END;
```

Query and redacted result:

```
SELECT ssn FROM mavis.cust_info;
```

```
SSN
-----
***-**-4320
***-**-4323
***-**-4325
***-**-4329
```

10.9.5 Creating Partial Redaction Policies Using Number Data Types

The DBMS_REDACT.ADD_POLICY function_parameters parameter can redact number data types.

[Settings for Number Data Types](#) (page 10-37)

When you set values for the number data type, you must specify a mask character, a starting digit position, and ending digit position.

Example: Partial Redaction Policy Using a Number Data Type (page 10-37)
 The `DBMS_REDACT.ADD_POLICY` procedure can create a partial redaction policy that uses a number data type.

10.9.5.1 Settings for Number Data Types

When you set values for the number data type, you must specify a mask character, a starting digit position, and ending digit position.

For partial redaction of number data types, you can enter values for the following settings for the `function_parameters` parameter of the `DBMS_REDACT.ADD_POLICY` procedure, in the order shown.

1. **Mask character:** Specifies the character to display. Enter a number from 0 to 9.
2. **Starting digit position:** Specifies the starting digit position for the redaction, such as 1 for the first digit.
3. **Ending digit position:** Specifies the ending digit position for the redaction.

For example, the following setting redacts the first five digits of the Social Security number 987654321, resulting in 999994321.

```
function_parameters => '9,1,5',
```

See Also:

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4)
 for information about other `DBMS_REDACT.ADD_POLICY` parameters

10.9.5.2 Example: Partial Redaction Policy Using a Number Data Type

The `DBMS_REDACT.ADD_POLICY` procedure can create a partial redaction policy that uses a number data type.

Example 10-6 (page 10-37) shows how to partially redact a set of Social Security numbers in the `mavis.cust_info` table, for any application user who logs in. (Hence, the `expression` parameter evaluates to `TRUE`.)

This type of redaction is useful when the application is expecting a formatted number and not a string. In this scenario, the Social Security numbers are in a column of the data type `NUMBER`. In other words, the `ssn` column contains numbers only, not other characters such as hyphens or blank spaces.

Example 10-6 Partially Redacted Data Redaction Numeric Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema      => 'mavis',
    object_name        => 'cust_info',
    column_name        => 'ssn',
    policy_name        => 'redact_cust_ssns1',
    function_type       => DBMS_REDACT.PARTIAL,
    function_parameters => '7,1,5',
    expression         => '1=1',
    policy_description  => 'Partially redacts Social Security numbers',
    column_description => 'ssn contains numeric Social Security numbers');
END;
/
```

Query and redacted result:

```
SELECT ssn FROM mavis.cust_info;
```

```
SSN  
-----  
777774320  
777774323  
777774325  
777774329
```

10.9.6 Creating Partial Redaction Policies Using Date-Time Data Types

The `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter can redact date-time data types.

[Settings for Date-Time Data Types](#) (page 10-38)

Oracle Data Redaction provides special settings for configuring date-time data types.

[Example: Partial Redaction Policy Using Date-Time Data Type](#) (page 10-39)

The `DBMS_REDACT.ADD_POLICY` procedure can create a partial redaction policy that uses the date-time data type.

10.9.6.1 Settings for Date-Time Data Types

Oracle Data Redaction provides special settings for configuring date-time data types.

For partial redaction of date-time data types, enter values for the following `DBMS_REDACT.ADD_POLICY` `function_parameters` parameter settings.

Enter these values in the order shown:

1. `m`: Redacts the month. To redact with a month name, append 1–12 to lowercase `m`. For example, `m5` displays as `MAY`. To omit redaction, enter an uppercase `M`.
2. `d`: Redacts the day of the month. To redact with a day of the month, append 1–31 to a lowercase `d`. For example, `d7` displays as `07`. If you enter a higher number than the days of the month (for example, 31 for the month of February), then the last day of the month is displayed (for example, 28). To omit redaction, enter an uppercase `D`.
3. `y`: Redacts the year. To redact with a year, append 1–9999 to a lowercase `y`. For example, `y1984` displays as `84`. To omit redaction, enter an uppercase `Y`.
4. `h`: Redacts the hour. To redact with an hour, append 0–23 to a lowercase `h`. For example, `h20` displays as `20`. To omit redaction, enter an uppercase `H`.
5. `m`: Redacts the minute. To redact with a minute, append 0–59 to a lowercase `m`. For example, `m30` displays as `30`. To omit redaction, enter an uppercase `M`.
6. `s`: Redacts the second. To redact with a second, append 0–59 to a lowercase `s`. For example, `s45` displays as `45`. To omit redaction, enter an uppercase `S`.

See Also:

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4)
for information about other `DBMS_REDACT.ADD_POLICY` parameters

10.9.6.2 Example: Partial Redaction Policy Using Date-Time Data Type

The `DBMS_REDACT.ADD_POLICY` procedure can create a partial redaction policy that uses the date-time data type.

[Example 10-7](#) (page 10-39) shows how to partially redact a date. This example redacts the birth year of customers; replacing it with 13, but retaining the remaining values.

Example 10-7 Partially Redacted Data Redaction Using Date-Time Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    column_name   => 'birth_date',
    policy_name   => 'redact_cust_bdate',
    function_type => DBMS_REDACT.PARTIAL,
    function_parameters => 'mdy2013HMS',
    expression    => '1=1',
    policy_description => 'Replaces birth year with 2013',
    column_description => 'birth_date contains customer's birthdate');
END;
/
```

Query and redacted result:

```
SELECT birth_date FROM mavis.cust_info;
```

```
BIRTH_DATE
07-DEC-13 09.45.40.000000 AM
12-OCT-13 04.23.29.000000 AM
```

10.10 Creating a Regular Expression-Based Redaction Policy

A regular expression-based redaction policy enables you to redact data based on a search-and-replace model.

[About Creating Regular Expression-Based Redaction Policies](#) (page 10-39)

Regular expression-based redaction enables you to search for patterns of data to redact.

[Syntax for Creating a Regular Expression-Based Redaction Policy](#) (page 10-40)

The `regexp_*` parameters of the `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression-based redaction policy.

[Regular Expression-Based Redaction Policies Using Formats](#) (page 10-42)

The `DBMS_REDACT.ADD_POLICY` procedure `regexp_pattern` and `regexp_replace_string` parameters both support formats.

[Custom Regular Expression Redaction Policies](#) (page 10-46)

You can customize regular expressions in Data Redaction policies.

10.10.1 About Creating Regular Expression-Based Redaction Policies

Regular expression-based redaction enables you to search for patterns of data to redact.

For example, you can use regular expressions to redact email addresses, which can have varying character lengths. It is designed for use with character data only. You

can use formats for the search and replace operation, or you can create custom pattern formats.

You cannot use regular expressions to redact a subset of the values in a column. The `REGEXP_PATTERN` (regular expression pattern) must match *all* of the values in order for the `REGEXP_REPLACE_STRING` setting to take effect, and the `REGEXP_REPLACE_STRING` must change the value.

For rows where the `REGEXP_PATTERN` fails to match, Data Redaction performs `DBMS_REDACT.FULL` redaction. This mitigates the risk of a mistake in the `REGEXP_PATTERN` which causes the regular expression to fail to match all of the values in the column, from showing the actual data for those rows which it failed to match.

In addition, if no change to the value occurs as a result of the `REGEXP_REPLACE_STRING` setting during regular expression replacement operation, Data Redaction performs `DBMS_REDACT.FULL` redaction.

10.10.2 Syntax for Creating a Regular Expression-Based Redaction Policy

The `regexp_*` parameters of the `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression-based redaction policy.

The `DBMS_REDACT.ADD_POLICY` fields for creating a regular expression-based data redaction policy are as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type       IN BINARY_INTEGER := NULL,
  expression         IN VARCHAR2,
  enable             IN BOOLEAN := TRUE,
  regexp_pattern     IN VARCHAR2 := NULL,
  regexp_replace_string IN VARCHAR2 := NULL,
  regexp_position    IN BINARY_INTEGER := 1,
  regexp_occurrence  IN BINARY_INTEGER := 0,
  regexp_match_parameter IN VARCHAR2 := NULL);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4).
- `function_type`: Specifies the type of redaction. For regular expression based redaction, use either `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`.

If you use the `DBMS_REDACT.REGEXP` redaction type, then no truncation occurs. This applies even if the redacted value is wider than the column width, and if the Oracle Call Interface width attribute (`OCI_ATTR_CHAR_SIZE`) of the column is *not* preserved. (It becomes 4000, just as it does when the `REGEXP_REPLACE` SQL operator is used on a column.)

Using the `DBMS_REDACT.REGEXP_WIDTH` redaction type truncates any redacted value that exceeds the width of the column, and ensures that the OCI width attribute of the column (`OCI_ATTR_CHAR_SIZE`) remains unchanged.

Note the following:

- Use the `DBMS_REDACT.REGEXP_WIDTH` function type if your applications depend on the value of the `OCI_ATTR_CHAR_SIZE` attribute. For example, applications that are built using the Oracle OLE DB Provider interface are sensitive to the value of the `OCI_ATTR_CHAR_SIZE` attribute. If you use `DBMS_REDACT.REGEXP` as the redaction type, then the `OCI_ATTR_CHAR_SIZE` always becomes 4000. This setting makes it unsuitable as the redaction type of policies on tables that are used by Oracle OLE DB based applications. See *Oracle Call Interface Programmer's Guide* for more information about Oracle Call Interface parameter attributes.
- When you set the `function_type` parameter to `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`, omit the `function_parameters` parameter from the `DBMS_REDACT.ADD_POLICY` procedure.
- Specify the regular expression parameters in much the same way that you specify the `pattern`, `replace`, `position`, `occurrence`, and `match_parameter` arguments to the `REGEXP_REPLACE` SQL function. See *Oracle Database SQL Language Reference* for information about the `REGEXP_REPLACE` SQL function.
- `regexp_pattern`: Describes the search pattern for data that must be matched. If it finds a match, then Oracle Database replaces the data as specified by the `regexp_replace_string` setting. See the following sections for more information:
 - [Regular Expression-Based Redaction Policies Using Formats](#) (page 10-42)
 - [Custom Regular Expression Redaction Policies](#) (page 10-46)
- `regexp_replace_string`: Specifies how you want to replace the data to be redacted. See the following sections for more information:
 - [Regular Expression-Based Redaction Policies Using Formats](#) (page 10-42)
 - [Custom Regular Expression Redaction Policies](#) (page 10-46)
- `regexp_position`: Specifies the starting position for the string search. The value that you enter must be a positive integer indicating the character of the `column_name` data where Oracle Database should begin the search. The default is 1 or the `DBMS_REDACT.RE_BEGINNING` format, meaning that Oracle Database begins the search at the first character of the `column_name` data.
- `regexp_occurrence`: Specifies how to perform the search and replace operation. The value that you enter must be a nonnegative integer indicating the occurrence of the replace operation:
 - If you specify 0 or the `DBMS_REDACT.RE_ALL` format, then Oracle Database replaces all the occurrences of the match.
 - If you specify the `DBMS_REDACT.RE_FIRST` format, then Oracle Database replaces the first occurrence of the match.
 - If you specify a positive integer *n*, then Oracle Database replaces the *n*th occurrence of the match.

If the occurrence is greater than 1, then the database searches for the second occurrence beginning with the first character following the first occurrence of pattern, and so forth.

- `regexp_match_parameter`: Specifies a text literal that lets you change the default matching behavior of the function. The behavior of this parameter is the same for this function as for the `REGEXP_REPLACE` SQL function. See *Oracle Database SQL Language Reference* for detailed information.

To filter the search so that it is not case sensitive, specify the `RE_MATCH_CASE_INSENSITIVE` format.

10.10.3 Regular Expression-Based Redaction Policies Using Formats

The `DBMS_REDACT.ADD_POLICY` procedure `regexp_pattern` and `regexp_replace_string` parameters both support formats.

Regular Expression Formats (page 10-42)

The regular expression formats represent commonly used expressions, such as the replacement of digits within a credit card number.

Example: Regular Expression Redaction Policy Using Formats (page 10-45)

The `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression redaction policy that uses formats.

10.10.3.1 Regular Expression Formats

The regular expression formats represent commonly used expressions, such as the replacement of digits within a credit card number.

[Table 10-9](#) (page 10-42) describes the formats that you can use with the `regexp_pattern` parameter in the `DBMS_REDACT.ADD_POLICY` procedure.

Table 10-9 *Formats for the `regexp_pattern` Parameter*

Format	Description
<code>DBMS_REDACT.RE_PATTERN_ANY_DIGIT</code>	<p>Searches for any digit. Replaces the identified pattern with the characters specified by the <code>regexp_replace_string</code> parameter. The <code>DBMS_REDACT.RE_PATTERN_ANY_DIGIT</code> is commonly used with the following values of the <code>regexp_replace_string</code> parameter:</p> <pre> regexp_replace_string => DBMS_REDACT.RE_REDACT_WITH_SINGLE_X, </pre> <p>This setting replaces any matched digit with the X character.</p> <p>The following setting replaces any matched digit with the 1 character.</p> <pre> regexp_replace_string => DBMS_REDACT.RE_REDACT_WITH_SINGLE_1, </pre>

Table 10-9 (Cont.) Formats for the `regexp_pattern` Parameter

Format	Description
<code>DBMS_REDACT.RE_PATTERN_CC_L6_T4</code>	Searches for the middle digits of any credit card (other than American Express) that has 6 leading digits and 4 trailing digits. Replaces the identified pattern with the characters specified by the <code>regexp_replace_string</code> parameter. The appropriate <code>regexp_replace_string</code> setting to use with this format is <code>DBMS_REDACT.RE_REDACT_CC_MIDDLE_DIGITS</code> , which finds any credit card that could have 6 leading and 4 trailing digits left as actual data . It then redacts the middle digits.
<code>DBMS_REDACT.RE_PATTERN_CCN</code>	Matches credit card numbers other than American Express credit card numbers. The appropriate <code>regexp_replace_string</code> setting to use with this format is <code>DBMS_REDACT.RE_REDACT_CCN</code> . The end result is a redaction of all the digits except the last 4.
<code>DBMS_REDACT.RE_PATTERN_AMEX_CCN</code>	Matches American Express credit card numbers. The appropriate <code>regexp_replace_string</code> setting to use with this format is <code>DBMS_REDACT.RE_REDACT_AMEX_CCN</code> . The end result is a redaction of all the digits except the last 5.
<code>DBMS_REDACT.RE_PATTERN_US_PHONE</code>	Searches for any U.S. telephone number. Replaces the identified pattern with the characters specified by the <code>regexp_replace_string</code> parameter. The appropriate <code>regexp_replace_string</code> setting to use with this format is <code>DBMS_REDACT.RE_REDACT_US_PHONE_L7</code> , which finds United States phone numbers and then redacts the last 7 digits.
<code>DBMS_REDACT.RE_PATTERN_EMAIL_ADDRESS</code>	Searches for any email address. Replaces the identified pattern with the characters specified by the <code>regexp_replace_string</code> parameter. The appropriate <code>regexp_replace_string</code> settings that you can use with this format are as follows: <code>RE_REDACT_EMAIL_NAME</code> , which finds any email address and redacts the email user name <code>RE_REDACT_EMAIL_DOMAIN</code> , which finds any email address and redacts the email domain <code>RE_REDACT_EMAIL_ENTIRE</code> , which finds any email address and redacts the entire email address

Table 10-9 (Cont.) Formats for the regexp_pattern Parameter

Format	Description
DBMS_REDACT.RE_PATTERN_IP_ADDRESS	Searches for an IP address. Replaces the identified pattern with the characters specified by the <code>regexp_replace_string</code> parameter. The appropriate <code>regexp_replace_string</code> setting to use with this format is <code>DBMS_REDACT.RE_REDACT_IP_L3</code> , which replaces the last section of the dotted decimal string representation of an IP address with the characters 999 to indicate that it was redacted.

Table 10-10 (page 10-44) describes formats that you can use with the `regexp_replace_string` parameter in the `DBMS_REDACT.ADD_POLICY` procedure.

Table 10-10 Formats for the regexp_replace_string Parameter

Format	Description
DBMS_REDACT.RE_REDACT_WITH_SINGLE_X	Replaces the data with a single X character for each of the actual data characters. For example, the credit card number 5105 1051 0510 5100 could be replaced with XXXX XXXX XXXX XXXX.
DBMS_REDACT.RE_REDACT_WITH_SINGLE_1	Replaces the data with a single 1 digit for each of the actual data digits. For example, the credit card number 5105 1051 0510 5100 could be replaced with 1111 1111 1111 1111.
DBMS_REDACT.RE_REDACT_CC_MIDDLE_DIGITS	Redacts the middle digits in credit card numbers, as specified by setting the <code>regexp_pattern</code> parameter with the <code>RE_PATTERN_CC_L6_T4</code> format. The redaction replaces each redacted character with an X. For example, the credit card number 5105 1051 0510 5100 could be replaced with 5105 10XX XXXX 5100.
DBMS_REDACT.RE_REDACT_CCN	Redacts the first 12 digits of a credit card number other than an American Express card number. For example, 4012888888881881 is redacted to *****1881.
DBMS_REDACT.RE_REDACT_AAMEX_CCN	Redacts the first 10 digits of an American Express number. For example, 378282246310005 is redacted to *****10005.

Table 10-10 (Cont.) Formats for the `regexp_replace_string` Parameter

Format	Description
<code>DBMS_REDACT.RE_REDACT_PHONE_L7</code>	Redacts the last 7 digits of U.S. telephone numbers, as specified by setting the <code>regexp_pattern</code> parameter with the <code>RE_PATTERN_US_PHONE</code> format. The redaction replaces each redacted character with an X. This setting only applies to hyphenated phone numbers, not phone numbers with spaces. For example, the telephone number 415-555-0100 could be replaced with 415-XXX-XXXX.
<code>DBMS_REDACT.RE_REDACT_EMAIL_NAME</code>	Redacts the email name as specified by setting the <code>regexp_pattern</code> parameter with the <code>RE_PATTERN_EMAIL_ADDRESS</code> format. The redaction replaces the email user name with four x characters. For example, the email address <code>psmith@example.com</code> could be replaced with <code>xxxx@example.com</code> .
<code>DBMS_REDACT.RE_REDACT_EMAIL_DOMAIN</code>	Redacts the email domain name as specified by setting the <code>regexp_pattern</code> parameter with the <code>RE_PATTERN_EMAIL_ADDRESS</code> format. The redaction replaces the domain with five x characters. For example, the email address <code>psmith@example.com</code> could be replaced with <code>psmith@xxxxx.com</code> .
<code>DBMS_REDACT.RE_REDACT_IP_L3</code>	Redacts the last three digits of the IP address as specified by setting the <code>regexp_pattern</code> parameter with the <code>RE_PATTERN_IP_ADDRESS</code> format. For example, the IP address 192.0.2.254 could be replaced with 192.0.2.999, which is an invalid IP address.

See Also:

[General Syntax of the `DBMS_REDACT.ADD_POLICY` Procedure](#) (page 10-4) for information about other `DBMS_REDACT.ADD_POLICY` parameters

10.10.3.2 Example: Regular Expression Redaction Policy Using Formats

The `DBMS_REDACT.ADD_POLICY` procedure can create a regular expression redaction policy that uses formats.

[Example 10-8](#) (page 10-45) shows how to use regular expression formats to redact credit card numbers.

Example 10-8 Regular Expression Data Redaction Character Values

```
BEGIN
DBMS_REDACT.ADD_POLICY(
  object_schema => 'mavis',
  object_name   => 'cust_info',
  column_name   => 'cc_num',
```

```

policy_name          => 'redact_cust_cc_nums',
function_type        => DBMS_REDACT.REGEXP,
function_parameters  => NULL,
expression           => '1=1',
regexp_pattern       => DBMS_REDACT.RE_PATTERN_CC_L6_T4,
regexp_replace_string => DBMS_REDACT.RE_REDACT_CC_MIDDLE_DIGITS,
regexp_position      => DBMS_REDACT.RE_BEGINNING,
regexp_occurrence    => DBMS_REDACT.RE_FIRST,
regexp_match_parameter => DBMS_REDACT.RE_MATCH_CASE_INSENSITIVE,
policy_description   => 'Regular expressions to redact credit card numbers',
column_description   => 'cc_num contains customer credit card numbers');
END;
/

```

Query and redacted result:

```
SELECT cc_num FROM mavis.cust_info;
```

```

CC_NUM
-----
401288XXXXXX1881
411111XXXXXX1111
555555XXXXXX1111
511111XXXXXX1118

```

10.10.4 Custom Regular Expression Redaction Policies

You can customize regular expressions in Data Redaction policies.

[Settings for Custom Regular Expressions](#) (page 10-46)

Oracle Data Redaction provides special settings to configure policies that use regular expressions.

[Example: Custom Regular Expression Redaction Policy](#) (page 10-47)

The `DBMS_REDACT.ADD_POLICY` procedure `regexp*` parameters can create a custom regular expression redaction policy.

10.10.4.1 Settings for Custom Regular Expressions

Oracle Data Redaction provides special settings to configure policies that use regular expressions.

To create custom regular expression redaction policies, you use the following parameters in the `DBMS_REDACT.ADD_POLICY` procedure:

- `regexp_pattern`: This pattern is usually a text literal and can be of any of the data types `CHAR`, `VARCHAR2`, `NCHAR`, or `NVARCHAR2`. The pattern can contain up to 512 bytes. For further information about writing the regular expression for the `regexp_pattern` parameter, see the description of the `pattern` argument of the `REGEXP_REPLACE` SQL function in *Oracle Database SQL Language Reference*, because the support that Data Redaction provides for regular expression matching is similar to that of the `REGEXP_REPLACE` SQL function.
- `regexp_replace_string`: This data can be of any of the data types `CHAR`, `VARCHAR2`, `NCHAR`, or `NVARCHAR2`. The `regexp_replace_string` can contain up to 500 back references to subexpressions in the form `\n`, where `n` is a number from 1 to 9. If you want to include a backslash (`\`) in the `regexp_replace_string` setting, then you must precede it with the escape character, which is also a backslash. For example, to literally replace the matched pattern with `\2` (rather than replace it with the second matched subexpression of

the matched pattern), you enter `\\2` in the `regexp_replace_string` setting. For more information, see *Oracle Database SQL Language Reference*.

See Also:

[General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4) for information about other `DBMS_REDACT.ADD_POLICY` parameters

10.10.4.2 Example: Custom Regular Expression Redaction Policy

The `DBMS_REDACT.ADD_POLICY` procedure `regexp*` parameters can create a custom regular expression redaction policy.

[Example 10-9](#) (page 10-47) shows how to use regular expressions to redact the `emp_id` column data. In this example, taken together, the `regexp_pattern` and `regexp_replace_string` parameters do the following: first, find the pattern of 9 digits. For reference, break them into three groups that contain the first 3, the next 2, and then the last 4 digits. Then, replace all 9 digits with `XXXXX` concatenated with the third group (the last 4 digits) as found in the original pattern.

Query and redacted result:

```
SELECT emp_id FROM mavis.cust_info;
```

```
EMP_ID
-----
XXXXX1234
XXXXX5678
```

Example 10-9 Partially Redacted Data Redaction Using Regular Expressions

```
BEGIN
DBMS_REDACT.ADD_POLICY(
  object_schema      => 'mavis',
  object_name        => 'cust_info',
  column_name        => 'emp_id',
  policy_name        => 'redact_cust_ids',
  function_type      => DBMS_REDACT.REGEXP,
  expression         => '1=1',
  regexp_pattern     => '(\d\d\d)(\d\d)(\d\d\d\d)',
  regexp_replace_string => 'XXXXX\3',
  regexp_position    => 1,
  regexp_occurrence  => 0,
  regexp_match_parameter => 'i',
  policy_description => 'Redacts customer IDs using regular expression',
  column_description => 'emp_id contains employee ID numbers');
END;
```

10.11 Creating a Random Redaction Policy

A random redaction policy presents redacted data as randomly generated values, such as `Ukjs132[[]]]s`.

[Syntax for Creating a Random Redaction Policy](#) (page 10-48)

A random redaction policy presents the redacted data to the querying application user as randomly generated values, based on the column data type.

Example: Random Redaction Policy (page 10-48)

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure create a random redaction policy.

10.11.1 Syntax for Creating a Random Redaction Policy

A random redaction policy presents the redacted data to the querying application user as randomly generated values, based on the column data type.

Be aware that LOB columns are not supported.

The `DBMS_REDACT.ADD_POLICY` fields for creating a random redaction policy are as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type      IN BINARY_INTEGER := NULL,
  expression         IN VARCHAR2,
  enable             IN BOOLEAN := TRUE);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4).
- `function_type`: Specifies the function used to set the type of redaction. Enter `DBMS_REDACT.RANDOM`.

If you omit the `function_type` parameter, then the default redaction `function_type` setting is `DBMS_REDACT.FULL`.

Remember that the data type of the column determines which `function_type` settings that you are permitted to use. See [Comparison of Full, Partial, and Random Redaction Based on Data Types](#) (page 9-5).

10.11.2 Example: Random Redaction Policy

You can use the `DBMS_REDACT.ADD_POLICY` PL/SQL procedure create a random redaction policy.

[Example 10-10](#) (page 10-48) shows how to generate random values. Each time you run the `SELECT` statement, the output will be different.

Example 10-10 Randomly Redacted Data Redaction Values

```
BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    column_name   => 'login_username',
    policy_name   => 'redact_cust_rand_username',
    function_type => DBMS_REDACT.RANDOM,
    expression    => 'SYS_CONTEXT(''USERENV'', ''SESSION_USER'') = ''APP_USER''';
END;
/
```

Query and redacted result:

```
SELECT login_username FROM mavis.cust_info;

LOGIN_USERNAME
-----
N[CG{\pTVcK
```

10.12 Creating a Policy That Uses No Redaction

You can create policies that use no redaction at all, for when you want to test the policy in a development environment.

[Syntax for Creating a Policy with No Redaction](#) (page 10-49)

The None redaction type option can be used to test the internal operation of redaction policies.

[Example: Performing No Redaction](#) (page 10-49)

The `DBMS_REDACT.ADD_POLICY` procedure can create a policy that performs no redaction.

10.12.1 Syntax for Creating a Policy with No Redaction

The None redaction type option can be used to test the internal operation of redaction policies.

The None redaction type has no effect on the query results against tables that have policies defined on them. You can use this option to test the redaction policy definitions before applying them to a production environment. Be aware that LOB columns are not supported.

The `DBMS_REDACT.ADD_POLICY` fields for creating a policy with no redaction are as follows:

```
DBMS_REDACT.ADD_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2,
  column_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  function_type       IN BINARY_INTEGER := NULL,
  expression          IN VARCHAR2,
  enable             IN BOOLEAN := TRUE);
```

In this specification:

- `object_schema`, `object_name`, `column_name`, `policy_name`, `expression`, `enable`: See [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4).
- `function_type`: Specifies the functions used to set the type of data redaction. Enter `DBMS_REDACT.NONE`.

If you omit the `function_type` parameter, then the default redaction `function_type` setting is `DBMS_REDACT.FULL`.

10.12.2 Example: Performing No Redaction

The `DBMS_REDACT.ADD_POLICY` procedure can create a policy that performs no redaction.

[Example 10-11](#) (page 10-50) shows how to create a Data Redaction policy that does not redact any of the displayed values.

Example 10-11 No Redacted Data Redaction Values

```

BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    column_name  => 'user_name',
    policy_name  => 'redact_cust_no_vals',
    function_type => DBMS_REDACT.NONE,
    expression   => '1=1');
END;
/

```

Query and redacted result:

```
SELECT user_name FROM mavis.cust_info;
```

```

USER_NAME
-----
IDA NEAU

```

10.13 Exemption of Users from Oracle Data Redaction Policies

You can exempt users from having Oracle Data Redaction policies applied to the data they access.

To do so, you should grant the users the `EXEMPT REDACTION POLICY` system privilege. Grant this privilege to trusted users only.

In addition to users who were granted this privilege, user `SYS` is also exempt from all Data Redaction policies. The person who creates the Data Redaction policy is by default not exempt from it, unless this person is user `SYS` or has the `EXEMPT REDACTION POLICY` system privilege.

Note the following:

- Users who have the `INSERT` privilege on a table can insert values into a redacted column, regardless of whether a Data Redaction policy exists on the table. Data Redaction only affects SQL `SELECT` statements (that is, queries) issued by a user, and has no effect on any other SQL issued by a user, including `INSERT`, `UPDATE`, or `DELETE` statements. (See the next bullet for exceptions to this rule.)
- Users cannot perform a `CREATE TABLE AS SELECT` where any of the columns being selected (source columns) is protected by a Data Redaction policy (and similarly, any DML operation where the source is a redacted column), unless the user was granted the `EXEMPT REDACTION POLICY` system privilege.
- The `EXEMPT REDACTION POLICY` system privilege is included in the `DBA` role, but this privilege must be granted explicitly to users because it is not included in the `WITH ADMIN OPTION` for `DBA` role grants. Users who were granted the `DBA` role are exempt from redaction policies because the `DBA` role contains the `EXP_FULL_DATABASE` role, which is granted the `EXEMPT REDACTION POLICY` system privilege.

Related Topics:

[Restriction of Administrative Access to Oracle Data Redaction Policies](#)
(page 13-2)

You can restrict the list of users who can create, view and edit Data Redaction policies.

[Oracle Data Pump Security Model for Oracle Data Redaction](#) (page 12-5)

The `DATAPUMP_EXP_FULL_DATABASE` role includes the powerful `EXEMPT REDACTION POLICY` system privilege.

10.14 Altering an Oracle Data Redaction Policy

The `DBMS_REDACT.ALTER_POLICY` procedure enables you to modify Oracle Data Redaction policies.

[About Altering Oracle Data Redaction Policies](#) (page 10-51)

The `DBMS_REDACT.ALTER_POLICY` procedure alters a Data Redaction policy.

[Syntax for the DBMS_REDACT.ALTER_POLICY Procedure](#) (page 10-51)

The `DBMS_REDACT.ALTER_POLICY` procedure syntax can be used to alter all types of the Data Redaction policies.

[Parameters Required for DBMS_REDACT.ALTER_POLICY Actions](#) (page 10-52)

The `DBMS_REDACT.ALTER_POLICY` procedure provides parameters that can perform various actions, such as adding or modifying a column.

[Tutorial: Altering an Oracle Data Redaction Policy](#) (page 10-53)

You can redact multiple columns in a table or view, with each column having its own redaction setting.

10.14.1 About Altering Oracle Data Redaction Policies

The `DBMS_REDACT.ALTER_POLICY` procedure alters a Data Redaction policy.

If the policy is already enabled, then you do not need to disable it first, and after you alter the policy, it remains enabled.

You can find the names of existing Data Redaction policies by querying the `POLICY_NAME` column of the `REDACTION_POLICIES` data dictionary view, and information about the columns, functions, and parameters specified in a policy by querying the `REDACTION_COLUMNS` view. To find the current value for policies that use full data redaction, you can query the `REDACTION_VALUES_FOR_TYPE_FULL` data dictionary view.

The `action` parameter specifies the type of modification that you want to perform. At a minimum, you must include the `object_name` and `policy_name` parameters when you run this procedure.

10.14.2 Syntax for the DBMS_REDACT.ALTER_POLICY Procedure

The `DBMS_REDACT.ALTER_POLICY` procedure syntax can be used to alter all types of the Data Redaction policies.

The syntax for the `DBMS_REDACT.ALTER_POLICY` procedure is as follows:

```
DBMS_REDACT.ALTER_POLICY (
  object_schema      IN VARCHAR2 := NULL,
  object_name        IN VARCHAR2 := NULL,
  policy_name        IN VARCHAR2,
  action              IN BINARY_INTEGER := DBMS_REDACT.ADD_COLUMN,
  column_name        IN VARCHAR2 := NULL,
  function_type      IN BINARY_INTEGER := DBMS_REDACT.FULL,
  function_parameters IN VARCHAR2 := NULL,
```

```
expression          IN VARCHAR2 := NULL,  
regexp_pattern      IN VARCHAR2 := NULL,  
regexp_replace_string IN VARCHAR2 := NULL,  
regexp_position     IN BINARY_INTEGER := NULL,  
regexp_occurrence   IN BINARY_INTEGER := NULL,  
regexp_match_parameter IN VARCHAR2 := NULL,  
policy_description  IN VARCHAR2 := NULL,  
column_description  IN VARCHAR2 := NULL);
```

In this specification:

- `action`: Enter one of the following values to define the kind of action to use:
 - `DBMS_REDACT.MODIFY_COLUMN` if you plan to change the `column_name` value.
 - `DBMS_REDACT.ADD_COLUMN` if you plan to add a new column (in addition to columns that are already protected by the policy) for redaction. This setting is the default for the `action` parameter.
 - `DBMS_REDACT.DROP_COLUMN` if you want to remove redaction from a column.
 - `DBMS_REDACT.MODIFY_EXPRESSION` if you plan to change the `expression` value. Each policy can have only one policy expression. In other words, when you modify the policy expression, you are replacing the existing policy expression with a new policy expression.
 - `DBMS_REDACT.SET_POLICY_DESCRIPTION` if you want to change the description of the policy.
 - `DBMS_REDACT.SET_COLUMN_DESCRIPTION` if you want to change the description of the column.

See Also:

- [Parameters Required for DBMS_REDACT.ALTER_POLICY Actions](#) (page 10-52)
 - [General Syntax of the DBMS_REDACT.ADD_POLICY Procedure](#) (page 10-4) for information about the remaining parameters
-
-

10.14.3 Parameters Required for DBMS_REDACT.ALTER_POLICY Actions

The `DBMS_REDACT.ALTER_POLICY` procedure provides parameters that can perform various actions, such as adding or modifying a column.

[Table 10-11](#) (page 10-53) shows the combinations of these parameters.

Table 10-11 Parameters Required for Various DBMS_REDACT.ALTER_POLICY Actions

Desired Alteration	Parameters to Set
Add or modify a column	<ul style="list-style-type: none"> action (DBMS_REDACT.MODIFY_COLUMN) column_name function_type function_parameters (if necessary) regexp* (if necessary)
Change the policy expression	<ul style="list-style-type: none"> action (DBMS_REDACT.MODIFY_EXPRESSION) expression
Change the description of the policy	<ul style="list-style-type: none"> action (DBMS_REDACT.SET_POLICY_DESCRIPTION) policy_description
Change the description of the column	<ul style="list-style-type: none"> action (DBMS_REDACT.SET_COLUMN_DESCRIPTION) column_description
Drop a column	<ul style="list-style-type: none"> action (DBMS_REDACT.DROP_COLUMN) column_name

10.14.4 Tutorial: Altering an Oracle Data Redaction Policy

You can redact multiple columns in a table or view, with each column having its own redaction setting.

The exercise in this section shows how to modify a Data Redaction policy so that multiple columns are redacted. It also shows how to change the `expression` setting for the policy. To accomplish this, you must run the `DBMS_REDACT.ALTER_POLICY` procedure in stages.

1. Connect as a user who has privileges to create users and grant them privileges.

For example:

```
CONNECT sec_admin
Enter password: password
```

2. Create the following users:

```
GRANT CREATE SESSION TO dr_admin IDENTIFIED BY password;
GRANT CREATE SESSION TO sales_rep IDENTIFIED BY password;
GRANT CREATE SESSION TO support_rep IDENTIFIED BY password;
```

3. Grant EXECUTE on the DBMS_REDACT PL/SQL package to user dr_admin:

```
GRANT EXECUTE ON DBMS_REDACT TO dr_admin;
```

4. Connect as user OE.

```
CONNECT OE
Enter password: password
```

5. Create and populate a table that contains customer credit card information.

```
CREATE TABLE cust_order_info(
  first_name varchar2(20),
```

```

last_name varchar2(20),
address varchar2(30),
city varchar2(30),
state varchar2(3),
zip varchar2(5),
cc_num varchar(19),
cc_exp varchar2(7));

```

```

INSERT INTO cust_order_info VALUES ('Jane','Dough','39 Mockingbird Lane', 'San
Francisco', 'CA', 94114, '5105 1051 0510 5100', '10/2018');
INSERT INTO cust_order_info VALUES ('Mary','Hightower','2319 Maple Street',
'Sonoma', 'CA', 95476, '5111 1111 1111 1118', '03/2019');
INSERT INTO cust_order_info VALUES ('Herbert','Donahue','292 Winsome Way', 'San
Francisco', 'CA', 94117, '5454 5454 5454 5454', '08/2018');

```

- Grant the `SELECT` privilege on the `cust_order_info` table to the `sales_rep` and `support_rep` users.

```
GRANT SELECT ON cust_order_info TO sales_rep, support_rep;
```

- Connect as user `dr_admin`.

```
CONNECT dr_admin
Enter password: password
```

- Create and enable policy to redact the credit card number.

```

BEGIN DBMS_REDACT.ADD_POLICY(
  object_schema      => 'oe',
  object_name        => 'cust_order_info',
  column_name        => 'cc_num',
  policy_name        => 'redact_cust_cc_info',
  function_type      => DBMS_REDACT.REGEXP,
  function_parameters => NULL,
  expression         => '1=1',
  regexp_pattern     => DBMS_REDACT.RE_PATTERN_CCN,
  regexp_replace_string => DBMS_REDACT.RE_REDACT_CCN,
  regexp_position    => NULL,
  regexp_occurrence  => NULL,
  regexp_match_parameter => NULL,
  policy_description => 'Partially redacts credit card info',
  column_description => 'cc_num_number lists credit card numbers');
END;
/

```

- Modify the policy to include redaction of the expiration date.

```

BEGIN DBMS_REDACT.ALTER_POLICY(
  object_schema      => 'oe',
  object_name        => 'cust_order_info',
  policy_name        => 'redact_cust_cc_info',
  action             => DBMS_REDACT.ADD_COLUMN,
  column_name        => 'cc_exp',
  function_type      => DBMS_REDACT.RANDOM,
  expression         => '1-1');
END;
/

```

- Modify the policy again, to use a condition so that the `sales_rep` user views the redacted values and the `support_rep` user views the [actual data](#).


```

BEGIN
  DBMS_REDACT.ALTER_POLICY(
    object_schema => 'oe',
    object_name   => 'cust_order_info',
    policy_name   => 'redact_cust_cc_info',
    action        => DBMS_REDACT.MODIFY_EXPRESSION,
    expression    => 'SYS_CONTEXT(''USERENV'', ''SESSION_USER'') =
''SALES_REP''');
END;
/

```

11. To test the policy, have the two users query the `cust_order_info` table.

```

CONNECT suport_rep
Enter password: password

SELECT cc_num, cc_exp FROM OE.cust_order_info;

CC_NUM          CC_EXP
-----
5105 1051 0510 5100  10/2018
5111 1111 1111 1118  03/2019
5454 5454 5454 5454  08/2018

```

User `support_rep` can view the actual data.

```

CONNECT sales_rep
Enter password: password

SELECT cc_num, cc_exp FROM OE.cust_order_info;

CC_NUM          CC_EXP
-----
*****5100     1ST=033
*****1119     OZA.w4C
*****5454     B(9+;01

```

The actual data is redacted using for user `sales_rep`.

12. Alter the `cust_order_info` to include a condition so that only `support_rep` can see the redacted data but `sales_rep` cannot.

```

CONNECT dr_admin
Enter password: password

BEGIN
  DBMS_REDACT.ALTER_POLICY(
    object_schema => 'oe',
    object_name   => 'cust_order_info',
    policy_name   => 'redact_cust_cc_info',
    action        => DBMS_REDACT.MODIFY_EXPRESSION,
    expression    => 'SYS_CONTEXT(''USERENV'', ''SESSION_USER'') =
''SUPPORT_REP''');
END;
/

```

13. Have the users test the policy again.

```

CONNECT support_rep
Enter password: password

SELECT cc_num, cc_exp FROM OE.cust_order_info;

```

```

CC_NUM          CC_EXP
-----
*****5100     1^XMF~`
*****1119     qz+9=#S
*****5454     *KCaUkm

```

User `support_rep` can no longer view the actual data; it is now redacted.

```

CONNECT sales_rep
Enter password: password

SELECT cc_num, cc_exp FROM OE.cust_order_info;

```

```

CC_NUM          CC_EXP
-----
5105 1051 0510 5100 10/2018
5111 1111 1111 1118 03/2019
5454 5454 5454 5454 08/2018

```

User `sales_rep` now can view the actual data.

14. If you do not need the components of this tutorial, then you can remove them as follows:

```

CONNECT dr_admin
Enter password: password

BEGIN
  DBMS_REDACT.DROP_POLICY (
    object_schema => 'oe',
    object_name   => 'cust_order_info',
    policy_name   => 'redact_cust_cc_info');
END;
/

CONNECT sec_admin
Enter password: password

DROP USER dr_admin;
DROP USER sales_rep;
DROP USER support_rep;

CONNECT OE
Enter password: password

DROP TABLE cust_order_info;

```

10.15 Redacting Multiple Columns

You can redact more than one column in a Data Redaction policy.

[Adding Columns to a Data Redaction Policy for a Single Table or View](#)
(page 10-57)

You can redact columns of different data types, using different redaction types, for one table or view.

[Example: Redacting Multiple Columns](#) (page 10-57)

The `DBMS_REDACT.ALTER_POLICY` procedure can redact multiple columns.

10.15.1 Adding Columns to a Data Redaction Policy for a Single Table or View

You can redact columns of different data types, using different redaction types, for one table or view.

1. Create the policy for the first column that you want to redact.
2. Use the `DBMS_REDACT.ALTER_POLICY` procedure to add the next column to the policy.

As necessary, set the `action`, `column_name`, `function_type`, and `function_parameters` (or the parameters that begin with `regexp_`) parameters to define the redaction for the new column, but do not change the `object_schema`, `object_name`, `policy_name`, or `expression` parameters. Each redacted column continues to have the same redaction parameters that were used to create it.

10.15.2 Example: Redacting Multiple Columns

The `DBMS_REDACT.ALTER_POLICY` procedure can redact multiple columns.

[Example 10-12](#) (page 10-57) shows how to add a column to an existing Data Redaction policy. In this example, the `action` parameter specifies that a new column must be added, using `DBMS_REDACT.ADD_COLUMN`. The name of the new column, `card_num`, is set by the `column_name` parameter.

Example 10-12 Adding a Column to a Data Redaction Policy

```
BEGIN
  DBMS_REDACT.ALTER_POLICY(
    object_schema => 'mavis',
    object_name   => 'cust_info',
    policy_name   => 'redact_cust_user_ids',
    action        => DBMS_REDACT.ADD_COLUMN,
    column_name  => 'card_num',
    function_type => DBMS_REDACT.FULL,
    function_parameters => '',
    expression    => 'SYS_CONTEXT(''SYS_SESSION_ROLES'', 'ADM') = 'TRUE''');
END;
```

10.16 Disabling and Enabling an Oracle Data Redaction Policy

You can disable and then reenable Oracle Data Redactions policies as necessary.

[Disabling an Oracle Data Redaction Policy](#) (page 10-57)

The `DBMS_REDACT.DISABLE_POLICY` procedure disables Oracle Data Redaction policies.

[Enabling an Oracle Data Redaction Policy](#) (page 10-58)

The `DBMS_REDACT.ENABLE_POLICY` procedure enables Oracle Data Redaction policies.

10.16.1 Disabling an Oracle Data Redaction Policy

The `DBMS_REDACT.DISABLE_POLICY` procedure disables Oracle Data Redaction policies.

You can find the names of existing Data Redaction policies and whether they are enabled by querying the `POLICY_NAME` and `ENABLE` columns of the `REDACTION_POLICIES` view. However, as long as the policy still exists, you cannot create another policy for that table or view, even if the original policy is disabled. In other words, if you want to create a different policy on the same table column, then you must drop the first policy before you can create and use the new policy.

- To disable a Data Redaction policy, run the `DBMS_REDACT.DISABLE_POLICY` procedure, using the following syntax:

```
DBMS_REDACT.DISABLE_POLICY (
  object_schema      IN VARCHAR2 DEFAULT NULL,
  object_name        IN VARCHAR2,
  policy_name        IN VARCHAR2);
```

In this specification:

- `object_schema`: Specifies the schema of the object on which the Data Redaction policy will be applied. If you omit this setting (or enter `NULL`), then Oracle Database uses the name of the current schema.
- `object_name`: Specifies the name of the table or view to be used for the Data Redaction policy.
- `policy_name`: Specifies the name of the policy to be disabled.

[Example 10-13](#) (page 10-58) shows how to disable a Data Redaction policy.

Example 10-13 Disabling a Data Redaction Policy

```
BEGIN
  DBMS_REDACT.DISABLE_POLICY (
    object_schema => 'mavis',
    object_name   => 'cust_info',
    policy_name   => 'redact_cust_user_ids');
END;
/
```

10.16.2 Enabling an Oracle Data Redaction Policy

The `DBMS_REDACT.ENABLE_POLICY` procedure enables Oracle Data Redaction policies.

Immediately after you create a new policy, you do not need to enable it; the creation process handles that for you. To find the names of existing Data Redaction policies and whether they are enabled, you can query the `POLICY_NAME` and `ENABLE` columns of the `REDACTION_POLICIES` view. After you run the procedure to enable the policy, the enablement takes effect immediately.

- To enable a Data Redaction policy, run the `DBMS_REDACT.ENABLE_POLICY` procedure, using the following syntax.

```
DBMS_REDACT.ENABLE_POLICY (
  object_schema      IN VARCHAR2 DEFAULT NULL,
  object_name        IN VARCHAR2,
  policy_name        IN VARCHAR2);
```

In this specification:

- `object_schema`: Specifies the schema of the object on which the Data Redaction policy will be applied. If you omit this setting (or enter `NULL`), then Oracle Database uses the name of the current schema.
- `object_name`: Specifies the name of the table or view to be used for the Data Redaction policy.
- `policy_name`: Specifies the name of the policy to be enabled.

[Example 10-14](#) (page 10-59) shows how to enable a Data Redaction policy.

Example 10-14 Enabling a Data Redaction Policy

```
BEGIN
  DBMS_REDACT.ENABLE_POLICY (
    object_schema => 'mavis',
    object_name   => 'cust_info',
    policy_name   => 'redact_cust_user_ids');
END;
/
```

10.17 Dropping an Oracle Data Redaction Policy

The `DBMS_REDACT.DROP_POLICY` procedure drops Oracle Data Redaction policies.

You can drop an Oracle Data Redaction policy whether it is enabled or disabled. You can find the names of existing Data Redaction policies, by querying the `POLICY_NAME` column of the `REDACTION_POLICIES` view. When you drop a table or view that is associated with an Oracle Data Redaction policy, the policy is automatically dropped. As a best practice, drop the policy first, and then drop the table or view afterward. See [Dropped Oracle Data Redaction Policies When the Recycle Bin Is Enabled](#) (page 13-3) for more information.

- To drop a Data Redaction policy, run the `DBMS_REDACT.DROP_POLICY` procedure.

Use the following syntax:

```
DBMS_REDACT.DROP_POLICY (
  object_schema      IN VARCHAR2 DEFAULT NULL,
  object_name        IN VARCHAR2,
  policy_name        IN VARCHAR2);
```

In this specification:

- `object_schema`: Specifies the schema of the object to which the Data Redaction policy applies. If you omit this setting (or enter `NULL`), then Oracle Database uses the name of the current schema.
- `object_name`: Specifies the name of the table or view to be used for the Data Redaction policy.
- `policy_name`: Specifies the name of the policy to be dropped.

After you run the `DBMS_REDACT.DROP_POLICY` procedure, the drop takes effect immediately.

[Example 10-15](#) (page 10-59) shows how to drop a Data Redaction policy.

Example 10-15 Dropping a Data Redaction Policy

```
BEGIN
  DBMS_REDACT.DROP_POLICY (
```

```

object_schema => 'mavis',
object_name   => 'cust_info',
policy_name   => 'redact_cust_user_ids');
END;
/

```

10.18 Tutorial: SQL Expressions to Build Reports with Redacted Values

SQL expressions can be used to build reports based on columns that have Oracle Data Redaction policies defined on them.

The values used in the SQL expression will be redacted. This redaction occurs in such a way that the redaction takes place before the SQL expression is evaluated: the result value that is displayed in the report is the end result of the evaluated SQL expression over the redacted values, rather than the redacted result of the SQL expression as a whole.

1. Create the following Data Redaction policy for the `HR.EMPLOYEES` table.

This policy will replace the first 4 digits of the value from the `SALARY` column with the number 9 and the first digit of the value from the `COMMISSION_PCT` column with a 9.

```

BEGIN
  DBMS_REDACT.ADD_POLICY(
    object_schema      => 'HR',
    object_name        => 'EMPLOYEES',
    column_name        => 'SALARY',
    column_description => 'emp_sal_comm shows employee salary and commission',
    policy_name        => 'redact_emp_sal_comm',
    policy_description => 'Partially redacts the emp_sal_comm column',
    function_type      => DBMS_REDACT.PARTIAL,
    function_parameters => '9,1,4',
    expression         => '1=1');
END;
/
BEGIN
  DBMS_REDACT.ALTER_POLICY(
    object_schema      => 'HR',
    object_name        => 'EMPLOYEES',
    policy_name        => 'redact_emp_sal_comm',
    action             => DBMS_REDACT.ADD_COLUMN,
    column_name        => 'COMMISSION_PCT',
    function_type      => DBMS_REDACT.PARTIAL,
    function_parameters => '9,1,1',
    expression         => '1=1');
END;
/

```

2. Log in to the `HR` schema and then run the following report.

This report will use the SQL expression `(SALARY + COMMISSION_PCT)` to combine the employees' salaries and commissions.

```

SELECT (SALARY + COMMISSION_PCT) total_emp_compensation
FROM HR.EMPLOYEES
WHERE DEPARTMENT_ID = 80;

```

```

TOTAL_EMP_COMPENSATION
-----
                9999.9
                9999.95

```

99990.95

...

3. Use SQL expressions for the report, including concatenation.

For example:

```
SELECT 'Employee ID '           || EMPLOYEE_ID ||
       ' has a salary of '      || SALARY ||
       ' and a commission of '  || COMMISSION_PCT || '.' detailed_emp_compensation
FROM HR.EMPLOYEES
WHERE DEPARTMENT_ID = 80
ORDER BY EMPLOYEE_ID;
```

DETAILED_EMP_COMPENSATION

```
-----
Employee ID 150 has a salary of 99990 and a commission of .9.
Employee ID 151 has a salary of 9999 and a commission of .95.
Employee ID 152 has a salary of 9999 and a commission of .95.
```

...

4. Connect the user who created the `redact_emp_sal_comm` Data Redaction policy and then run the following statement to drop the policy.

```
BEGIN
  DBMS_REDACT.DROP_POLICY (
    object_schema => 'HR',
    object_name   => 'EMPLOYEES',
    policy_name   => 'redact_emp_sal_comm');
END;
/
```

10.19 Oracle Data Redaction Policy Data Dictionary Views

Oracle Database provides data dictionary views that list information about Data Redaction policies.

Before you can query these views, you must be granted the `SELECT_CATALOG_ROLE` role.

[Table 10-12](#) (page 10-61) lists the Data Redaction data dictionary views.

Table 10-12 Data Redaction Views

View	Description
<code>REDACTION_COLUMNS</code>	Describes all of the redacted columns in the database, providing the the owner of the table or view within which the column resides, the object name, the column name, the type of redaction function, the parameters to the redaction function (if any), and a description of the redaction policy. If a policy expression has been created, displays the default object-wide policy expression's SQL expression.
<code>REDACTION_EXPRESSIONS</code>	Displays the names of existing policy expressions and their SQL expressions

Table 10-12 (Cont.) Data Redaction Views

View	Description
REDACTION_POLICIES	Describes all of the data redaction policies in the database. It includes information about the object owner, object name, policy name, policy expression, whether the policy is enabled, and a description of the Data Redaction policy.
REDACTION_VALUES_FOR_TYPE _FULL	Shows the current redaction values for Data Redaction policies that use full redaction

Using Oracle Data Redaction in Oracle Enterprise Manager

Oracle Enterprise Manager Cloud Control (Cloud Control) can manage Oracle Data Redaction policies and formats.

[About Using Oracle Data Redaction in Oracle Enterprise Manager](#) (page 11-1)

Oracle Enterprise Manager Cloud Control provides an unified user interface for creating and managing Oracle Data Redaction policies.

[Oracle Data Redaction Workflow](#) (page 11-2)

First, you should create sensitive column types and formats if necessary, and then create the Oracle Data Redaction policy afterward.

[Management of Sensitive Column Types in Enterprise Manager](#) (page 11-2)

A sensitive column type categorizes table column sensitive information into a sensitive information type, such as credit card numbers.

[Managing Oracle Data Redaction Formats Using Enterprise Manager](#) (page 11-4)

Oracle Data Redaction provides redaction formats to be used directly within a redaction policy to redact data.

[Managing Oracle Data Redaction Policies Using Enterprise Manager](#) (page 11-9)

You can create, edit, view, and delete Oracle Data Redaction policies in Enterprise Manager Cloud Control.

[Managing Named Data Redaction Policy Expressions Using Enterprise Manager](#) (page 11-17)

You can manage Oracle Data Redaction policy expressions in Enterprise Manager Cloud Control.

11.1 About Using Oracle Data Redaction in Oracle Enterprise Manager

Oracle Enterprise Manager Cloud Control provides an unified user interface for creating and managing Oracle Data Redaction policies.

You can do the following:

- Create and manage custom Oracle Data Redaction formats, which were previously known as Data Redaction shortcuts. (This functionality is not available from the command line.)
- Create and manage sensitive column types directly from the Oracle Data Redaction pages. While you create a Data Redaction policy, Cloud Control uses sensitive column types to obtain the Oracle Data Redaction formats that are relevant to the column that you are redacting.

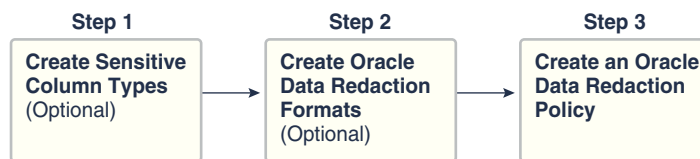
Note:

Ensure that you have the latest plug-in for Oracle Enterprise Manager. For information about how to deploy a plug-in, see *Enterprise Manager Cloud Control Administrator's Guide*. If you have the Oracle Database plug-in release 13.1.1.0.0, then you can create named Data Redaction policy expressions in Oracle Enterprise Manager.

11.2 Oracle Data Redaction Workflow

First, you should create sensitive column types and formats if necessary, and then create the Oracle Data Redaction policy afterward.

The following figure illustrates this process:



1. (Optional) If you want to map the database columns (that contain the data that you want to redact) to new sensitive column types, then create the required sensitive column types as described in [Management of Sensitive Column Types in Enterprise Manager](#) (page 11-2).
2. (Optional) If you want to redact the data (present in a particular database column) using a custom redaction format, then create the required redaction format as described in [Creating a Custom Oracle Data Redaction Format Using Enterprise Manager](#) (page 11-5).
3. Create an Oracle Data Redaction policy for the required database, as described in [Creating an Oracle Data Redaction Policy Using Enterprise Manager](#) (page 11-11).

Note:

When you create an Oracle Data Redaction policy, it is enabled by default. For information on how to disable an enabled redaction policy, see [Enabling or Disabling an Oracle Data Redaction Policy in Enterprise Manager](#) (page 11-16).

11.3 Management of Sensitive Column Types in Enterprise Manager

A sensitive column type categorizes table column sensitive information into a sensitive information type, such as credit card numbers.

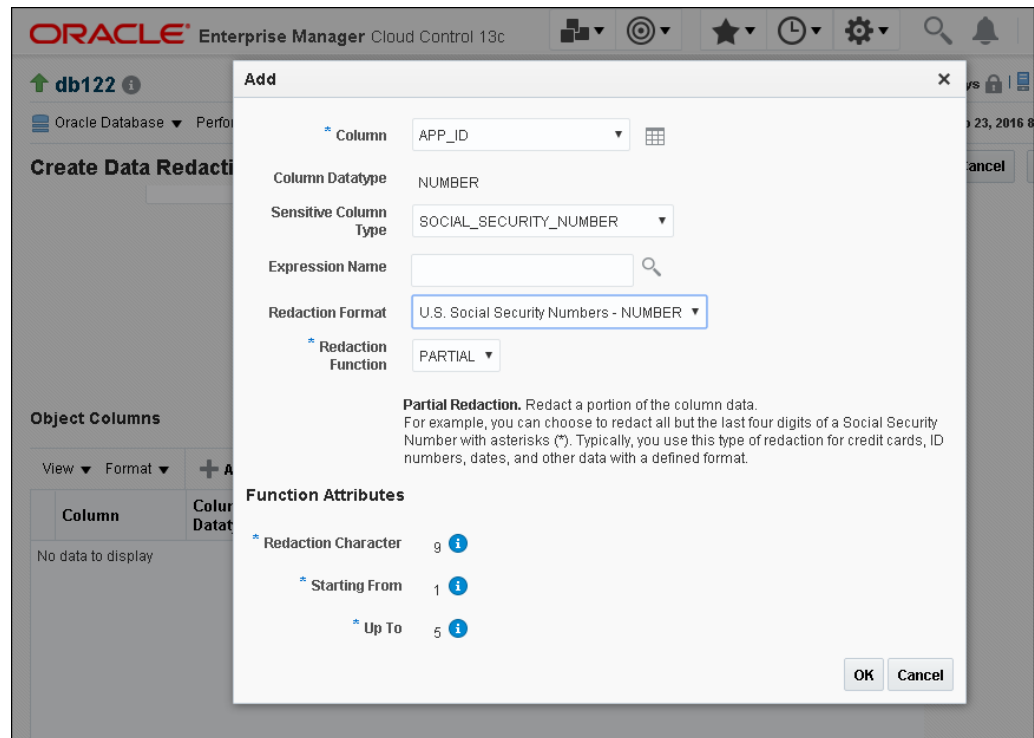
Sensitive column types use a combination of the column name, column comments, and the data pattern defined using a regular expression to tag a column to a particular sensitive information type.

While you create Oracle Data Redaction policies, redaction formats are filtered on the basis of the chosen sensitive column type, thus saving time and effort. For example, if the database table column that you want to redact contains U.S. Social Security numbers, and you select the `SOCIAL_SECURITY_NUMBER` sensitive column type for

the column while adding it to the Oracle Data Redaction policy, the default redaction formats that you can use to redact the column data are filtered, and only the relevant redaction formats are displayed.

Figure 11-1 (page 11-3) illustrates the filtering of Oracle Data Redaction formats based on sensitive column types.

Figure 11-1 Oracle Data Redaction Formats Filtered on the Basis of Sensitive Column Types



Note:

This functionality is available only if you have the Enterprise Manager for Oracle Database plug-in 12.1.0.7 or later deployed in your system.

For information on how to verify the plug-ins deployed in your environment, see *Enterprise Manager Cloud Control Administrator's Guide*.

As part of the Application Data Modeling feature, Oracle provides a number of default sensitive column types that a database column can be mapped to.

Figure 11-2 (page 11-4) displays some of the default sensitive column types. To access this page, click **Manage Sensitive Column Types** on the Data Redaction Formats page.

Figure 11-2 Default Sensitive Column Types

Name	Description	Author
CREDIT_CARD_NUMBER	Identifies credit card number columns. Samples: 5199-1234-1234-123...	Oracle
EMAIL_ID	Identifies email address columns. Samples: jsmith@comgmt.com, Ja...	Oracle
IP_ADDRESS	Identifies IP address columns. Samples: 7.7.7.1, 78.78.78.12, 789.789...	Oracle
ISBN_10	Identifies 10 digit International Standard Book Number columns. Samp...	Oracle
ISBN_13	Identifies 13 digit International Standard Book Number columns. Samp...	Oracle
NATIONAL_INSURANCE_NUMBER	Identifies National Insurance number (UK) columns. Samples: ZR 50 1...	Oracle
PHONE_NUMBER	Identifies phone number columns. Samples: 555-1212, (123)555-121...	Oracle
SOCIAL_INSURANCE_NUMBER	Identifies Social Insurance Number (Canada) columns. Samples: 884-...	Oracle
SOCIAL_SECURITY_NUMBER	Identifies Social Security number columns. Samples: 123-45-6789, 12...	Oracle
UNDEFINED	Sensitive column type not defined.	Oracle
UNIVERSAL_PRODUCT_CODE	Identifies Universal Product Code columns. Samples: 1-23456-78901-...	Oracle

If none of the default sensitive column types are suitable for the database column that contains the data that you want to redact, you can create a new sensitive column type, or create a sensitive column type that is based on an existing sensitive column type, as described in *Oracle Database Testing Guide*.

11.4 Managing Oracle Data Redaction Formats Using Enterprise Manager

Oracle Data Redaction provides redaction formats to be used directly within a redaction policy to redact data.

[About Managing Oracle Data Redaction Formats Using Enterprise Manager](#) (page 11-5)

The Oracle Data Redaction formats are used for commonly redacted data, such as ID numbers, credit cards, or phone numbers.

[Creating a Custom Oracle Data Redaction Format Using Enterprise Manager](#) (page 11-5)

You can create and save custom Oracle Data Redaction formats using Enterprise Manager Cloud Control.

[Editing a Custom Oracle Data Redaction Format Using Enterprise Manager](#) (page 11-7)

You can edit custom Oracle Data Redaction formats using Enterprise Manager Cloud Control, but not in SQL*Plus.

[Viewing Oracle Data Redaction Formats Using Enterprise Manager](#) (page 11-8)

Enterprise Manager Cloud Control displays the details of the Oracle-supplied and custom Oracle Data Redaction formats.

[Deleting a Custom Oracle Data Redaction Format Using Enterprise Manager](#) (page 11-9)

You can delete custom Oracle Data Redaction formats using Enterprise Manager Cloud Control.

11.4.1 About Managing Oracle Data Redaction Formats Using Enterprise Manager

The Oracle Data Redaction formats are used for commonly redacted data, such as ID numbers, credit cards, or phone numbers.

You can use several default Oracle Data Redaction formats (previously known as Oracle Data Redaction templates). As an example of the Oracle Data Redaction formats, a set of Social Security number formats enable you to quickly designate ways to redact Social Security numbers, such as redacting the first five numbers of the Social Security number.

Figure 11-3 (page 11-5) displays the default Oracle Data Redaction formats.

Figure 11-3 Default Oracle Data Redaction Formats

Format Name	Sensitive Column Type	Function Type	Description
American Express Credit Card Numbers - Formatted	CREDIT_CARD_NUMBER	PARTIAL	Redact the American Express Credit Card Number by replacing every
American Express Credit Card Numbers - NUMBER	CREDIT_CARD_NUMBER	PARTIAL	Redact the American Express Credit Card Number by replacing every
American Express Credit Card Numbers - Partially Redacted	CREDIT_CARD_NUMBER	REGEX	Redact the American Express Credit Card Number by replacing every
American Express Credit Card Numbers - Random	CREDIT_CARD_NUMBER	RANDOM	Redact the American Express Credit Card Number by replacing every
Canadian Social Insurance Numbers - Formatted	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance Number by replacing every
Canadian Social Insurance Numbers - NUMBER	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance number by replacing every
Canadian Social Insurance Numbers - Random	SOCIAL_INSURANCE_NUMBER	RANDOM	Redact the Canadian Social Insurance Number by replacing every
Canadian Social Insurance Numbers - VARCHAR	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance number by replacing every
Credit Card Numbers - Formatted	CREDIT_CARD_NUMBER	PARTIAL	Redact the Credit Card Number by replacing every
Credit Card Numbers - NUMBER	CREDIT_CARD_NUMBER	PARTIAL	Redact the Credit Card Number by replacing every
Credit Card Numbers - Partially Redacted	CREDIT_CARD_NUMBER	REGEX	Redact the Credit Card Number by replacing every

Each default Oracle Data Redaction format consists of a specific redaction function that determines the redacted output when the redaction format is used in an Oracle Data Redaction policy. For example, the `Credit Card Numbers - NUMBER` default redaction format replaces the first twelve digits of the column data with the digit 0, when it is used in an Oracle Data Redaction policy. That is, if the column data is 5555555555554444, the redacted output will be 0000000000004444.

If you have deployed the Enterprise Manager for Oracle Database plug-in 12.1.0.7 or higher on your system, then you can also create and save custom redaction formats, which you can then use in your redaction policies.

11.4.2 Creating a Custom Oracle Data Redaction Format Using Enterprise Manager

You can create and save custom Oracle Data Redaction formats using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target.

4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.
 Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.
6. Select the **Formats** tab.
7. Do one of the following:
 - To create a new redaction format, click **Create**.
 - To create a redaction format that is based on a default format, select the format and then click **Create Like**.

If you select **Create**, then the following dialog box appears:

Create
✕

* **Format Name**

* **Description**

Sensitive Column Type

* **Redaction Function**

Regular Expression Based Redaction. Specifies a regular expression that represents the column data that will be redacted.

Function Attributes

* **Pattern**

Specifies the regular expression pattern to be searched. Example: 'd\d\d\d\d\d678' for number like '012345678'

* **Replace String**

Example: Use 'XXXXXXXX3' (replace string) to redact '012345678' (actual value) which matches '(d\d\d\d) (d\d\d\d) (d\d\d\d)' (regexp pattern) to 'XXXXXXXX678' (redacted value). Note that the '3' in the replace string preserves the actual data in the third set of parentheses in the pattern.

* **Position**

Specifies the starting position of the string search. The default is 1, meaning it begins the search from the first character of column data.

* **Occurrence**

Specifies how to perform the search and replace operation. Zero means it replaces all occurrences. Positive integer 'n' would replace nth occurrence of the string.

Match Parameter

Specifies the matching parameters for the REGEX redaction function.

8. Provide a name and a description for the redaction format that you want to create.

If you want to map the redaction format to a particular sensitive column type (such that the created redaction format can be used to redact the data of a column that is associated with the sensitive column type), then select a value for **Sensitive Column Type**.

Select the function that the format should use to redact the column data. For **Redaction Function**, select as follows:

- **FULL** if the format should redact the entire column data.
- **PARTIAL** if the format should redact only a part of the column data. Ensure that you provide the function attributes, as well as the data type that you want to use the redaction format for.
- **REGEX** if the format should redact data based on a regular expression. Ensure that you provide the function attributes.
- **RANDOM** if the format should redact data in a random manner, using randomly generated values
- **NONE** if the format will be used to only test the definition of a redaction policy, and not redact any column data

9. Click **OK** to create and save the custom redaction format.

This format now can be used to create a redaction policy.

Related Topics:

[Oracle Data Redaction Features and Capabilities](#) (page 9-1)

Oracle Data Redaction provides a variety of ways to redact different types of data.

[Creating an Oracle Data Redaction Policy Using Enterprise Manager](#) (page 11-11)

You can create an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

11.4.3 Editing a Custom Oracle Data Redaction Format Using Enterprise Manager

You can edit custom Oracle Data Redaction formats using Enterprise Manager Cloud Control, but not in SQL*Plus.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.

5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. Select the **Formats** tab.
7. Select the custom redaction format that you want to edit, then click **Edit**.

A dialog box similar to the following appears:

Edit [X]

* **Format Name** American Express Credit Cards - FULL

* **Description** Redact the American E:

Sensitive Column Type CREDIT_CARD_NUMBER ▼

* **Redaction Function** FULL ▼

Full Redaction. Redact all the contents of the column data. The redacted value returned to the querying user depends on the data type of the column. For example, columns of the NUMBER data type are redacted with a zero (0) and character data types are redacted with a blank space. These default values can be changed if necessary.

OK Cancel

8. (Optional) Choose to edit the format description, sensitive column type, redaction function, and the redaction function attributes.
9. Click **OK** to save the edited format.

11.4.4 Viewing Oracle Data Redaction Formats Using Enterprise Manager

Enterprise Manager Cloud Control displays the details of the Oracle-supplied and custom Oracle Data Redaction formats.

1. Log into Oracle Enterprise Manager Cloud Control as either user SYSTEM or SYSMAN.

The URL is as follows:

`https://host:port/em`

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. Select the **Formats** tab.

The Data Redaction Formats page appears, similar to the following page.

Format Name	Sensitive Column Type	Function Type	Description
American Express Credit Card Numbers - Formatted	CREDIT_CARD_NUMBER	PARTIAL	Redact the American Express Credit Card Number
American Express Credit Card Numbers - NUMBER	CREDIT_CARD_NUMBER	PARTIAL	Redact the American Express Credit Card Number
American Express Credit Card Numbers - Partially Redacted	CREDIT_CARD_NUMBER	REGEX	Redact the American Express Credit Card Number
American Express Credit Card Numbers - Random	CREDIT_CARD_NUMBER	RANDOM	Redact the American Express Credit Card Number
Canadian Social Insurance Numbers - Formatted	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance Number
Canadian Social Insurance Numbers - NUMBER	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance number
Canadian Social Insurance Numbers - Random	SOCIAL_INSURANCE_NUMBER	RANDOM	Redact the Canadian Social Insurance Number
Canadian Social Insurance Numbers - VARCHAR	SOCIAL_INSURANCE_NUMBER	PARTIAL	Redact the Canadian Social Insurance number
Credit Card Numbers - Formatted	CREDIT_CARD_NUMBER	PARTIAL	Redact the Credit Card Number by replacing ever
Credit Card Numbers - NUMBER	CREDIT_CARD_NUMBER	PARTIAL	Redact the Credit Card Number by replacing ever
Credit Card Numbers - Partially Redacted	CREDIT_CARD_NUMBER	REGEX	Redact the Credit Card Number by replacing ever

7. Select the required redaction format, then click **View**.

11.4.5 Deleting a Custom Oracle Data Redaction Format Using Enterprise Manager

You can delete custom Oracle Data Redaction formats using Enterprise Manager Cloud Control.

You can only delete custom Oracle Data Redaction formats, and not the redaction formats that are provided by Oracle.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package.

6. Select the **Formats** tab.
7. Select the custom redaction format that you want to delete, and then click **Delete**.
8. In the Confirmation dialog box, click **Yes** or **No**.

11.5 Managing Oracle Data Redaction Policies Using Enterprise Manager

You can create, edit, view, and delete Oracle Data Redaction policies in Enterprise Manager Cloud Control.

[About Managing Oracle Data Redaction Policies Using Enterprise Manager](#)
(page 11-10)

Use the Data Redaction page in Cloud Control to manage Oracle Data Redaction policies.

[Creating an Oracle Data Redaction Policy Using Enterprise Manager](#)
(page 11-11)

You can create an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

[Editing an Oracle Data Redaction Policy Using Enterprise Manager](#) (page 11-14)
You can edit an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

[Viewing Oracle Data Redaction Policy Details Using Enterprise Manager](#)
(page 11-16)

You can find Oracle Data Redaction policy details such as whether the policy is enabled by using Enterprise Manager Cloud Control.

[Enabling or Disabling an Oracle Data Redaction Policy in Enterprise Manager](#)
(page 11-16)

An Oracle Data Redaction policy is executed at run time only if it is enabled. When you create an Oracle Data Redaction policy, it is enabled by default.

[Deleting an Oracle Data Redaction Policy Using Enterprise Manager](#)
(page 11-17)

You can delete an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

11.5.1 About Managing Oracle Data Redaction Policies Using Enterprise Manager

Use the Data Redaction page in Cloud Control to manage Oracle Data Redaction policies.

To redact the data present in a particular database table or view column, you must create an Oracle Data Redaction policy. Data is redacted using a redaction format that is specified by the Oracle Data Redaction policy. To redact data, you can use any of the Oracle-supplied redaction formats, or create and use a custom redaction format. If the table or view column that contains the data that you want to redact is mapped to a sensitive column type, Oracle uses the mapping to recommend suitable redaction formats for the data. Thus, Oracle Data Redaction policies encapsulate database schemas, database table and view columns, sensitive column types, and Oracle Data Redaction formats.

[Figure 11-4](#) (page 11-11) shows the Data Redaction page, which enables you to create and manage Oracle Data Redaction policies in Cloud Control.

Figure 11-4 Oracle Data Redaction Policies Page

Data Redaction
Oracle Data Redaction provides an easy way to quickly redact sensitive information that is displayed in applications without altering the underlying database blocks on disk or in cache. Data is redacted in real-time according to flexible multi-factor policies. Data Redaction is licensed as part of Oracle Advanced Security.

Policies Expressions Formats

▲ Search Data Redaction Policies

Schema %
Table/View %
Policy Name %
Go

Policies

View Format Create Edit View Enable Disable Delete Freeze Wrap

Schema	Table/View	Policy Name	Enabled	Redacted Columns
HR	EMPLOYEES	test	✓	1
HR	JOBS	dr_redact	✓	0

11.5.2 Creating an Oracle Data Redaction Policy Using Enterprise Manager

You can create an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target for which you want to create an Oracle Data Redaction policy.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package.

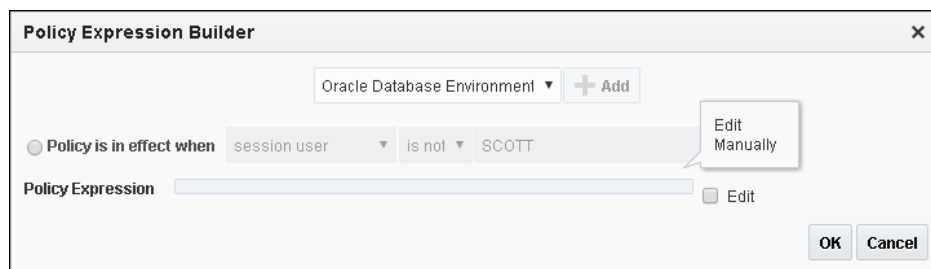
6. In the Policies section of the Policies tab, select **Create**.

If this is the first time that you are creating a Data Redaction policy, then the Data Redaction: Set up for enabling column sensitive type discovery dialog box appears. This feature enables the use of column sensitive type discovery for Data Redaction policies.

To accomplish this, Enterprise Manager creates the `GET_COL_DATA_SENSITIVE_TYPES` procedure in the `DBSNMP` schema. To perform a sensitive type discovery for a selected column while creating an Oracle Data Redaction policy, a user must have the `EXECUTE` privilege on the `DBSNMP.GET_COL_DATA_SENSITIVE_TYPES` procedure. If the database is

protected by Oracle Database Vault, then ensure that any users who must create Data Redaction policies are participants to realms that protect the DBSNMP schema.

7. If the Data Redaction: Set up for enabling column sensitive type discovery dialog box appears and if the current login user does not have the correct requirements, select a credential of a user who has the EXECUTE privilege on DBSNMP.GET_COL_DATA_SENSITIVE_TYPES. Then click OK.
8. On the Create Data Redaction Policy page, enter the following information:
 - **Schema:** Enter (or search for) the name of the schema that contains the data you want to redact.
 - **Table/View:** Enter (or search for) the table or field that contains the column you want to redact.
 - **Policy Name:** Enter a for the policy, such as emp_wages_pol.
 - **Default Expression:** Enter the default expression. The default setting is 1=1, which means that the policy always will be enforced. If you are not familiar with the components of a policy expression, then click the pencil icon beside the **Policy Expression** field to use Policy Expression Builder. Select **Policy is in effect when**, select the required conditions, then click **Add**. Click **Edit** if you want to edit the policy expression manually. After building the required policy expression, click **OK**. The Policy Expression Builder appears as follows:



9. In the Object Columns section, click **Add** to add a table or view column to the redaction policy.

A dialog box similar to the following appears:

The redaction policy is applied only on the table or view columns that are added to it.

10. From the **Column** menu, select the table or view column to which you want to apply the redaction policy.

To the right of the **Column** menu is an icon that you can click to view the contents of the selected column.

If the column contains sensitive data and has been mapped to a sensitive column type, then from the **Sensitive Column Type** menu, select the sensitive column type that it has been mapped to. If the search pattern in the **Sensitive Column Type** menu matches, then the sensitive column type is selected by default. For example, for a column listing credit card numbers, if there is a match, then the menu will list **Undefined** and **CREDIT_CARD_TYPE**. If there is no sensitive column type created, then the default **Sensitive Column Type** menu listing is only **Undefined**.

11. From the **Redaction Format** menu, select the redaction format that you want to use.

The drop-down list is populated with the Oracle Database-supplied redaction formats, as well as the custom redaction formats that you have created and saved.

If you do not want to use a pre-defined redaction format (that is, an Oracle-Database supplied redaction format, or a custom redaction format that you have

created), and instead want to specify the redaction details while creating the redaction policy, select **CUSTOM** for **Redaction Format**.

The Add dialog box adjusts to accommodate the type of redaction format and function that you select. For example, if you select the **CUSTOM** redaction format and the **REGEX** redaction function, then the Function Attributes region appears in the dialog box.

12. From the **Redaction Function** menu, select the function that you want to use to redact the column data.

Select **FULL** if you want to redact the entire column data, **PARTIAL** if you want to redact only a part of the column data, **REGEX** if you want to redact the column data based on a regular expression, **RANDOM** if you want to redact the column data in a random manner, using randomly generated values, or **NONE** if you only want to test the definition of the redaction policy, and not redact any column data. Note that all the redaction functions may not be applicable for a particular redaction format. The drop-down list displays only the redaction functions that are applicable for the selected redaction format.

If you selected **CUSTOM** for **Redaction Format** in the previous step, and **PARTIAL** or **REGEX** for **Redaction Function**, ensure that you specify the function attributes.

13. Click **OK**.
14. Repeat these steps starting with Step 8 for all the columns that you want to add to the redaction policy.
15. On the Create Data Redaction Policy page, click **OK** to create the data redaction policy.

When you create an Oracle Data Redaction policy, it is enabled by default.

Related Topics:

[Creating a Custom Oracle Data Redaction Format Using Enterprise Manager](#)
(page 11-5)

You can create and save custom Oracle Data Redaction formats using Enterprise Manager Cloud Control.

[Oracle Data Redaction Features and Capabilities](#) (page 9-1)

Oracle Data Redaction provides a variety of ways to redact different types of data.

[Enabling or Disabling an Oracle Data Redaction Policy in Enterprise Manager](#)
(page 11-16)

An Oracle Data Redaction policy is executed at run time only if it is enabled. When you create an Oracle Data Redaction policy, it is enabled by default.

11.5.3 Editing an Oracle Data Redaction Policy Using Enterprise Manager

You can edit an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

`https://host:port/em`

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then search for and click the name of the database target for which the Oracle Data Redaction policy that you want to edit was created.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. In the Policies section of the **Policies** tab, select the redaction policy that you want to edit, then click **Edit**.

7. On the Edit Data Redaction Policy page, choose to edit the policy expression, add new columns to the redaction policy, modify the redaction details of a column that is a part of the policy, or delete a column from the redaction policy.

You can do the following:

- To add a new column to the redaction policy, in the Object Columns section, click **Add**, select the table or view column that you want to add, then specify the redaction details.
 - To modify the redaction details of a column that is a part of the policy, select the column, click **Modify**, then edit the redaction details.
 - To delete a column from the redaction policy, select the column, then click **Delete**.
8. On the Edit Data Redaction Policy page, after editing the required fields, click **OK** to save and enable the edited redaction policy.

11.5.4 Viewing Oracle Data Redaction Policy Details Using Enterprise Manager

You can find Oracle Data Redaction policy details such as whether the policy is enabled by using Enterprise Manager Cloud Control.

You can disable an enabled redaction policy, or enable a disabled redaction policy using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then search for and click the name of the database target for which the Oracle Data Redaction policy that you want to view was created.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT PL/SQL` package.

6. In the Policies section of the Policies tab, do one of the following:
 - Select the name of the policy in the table.
 - Select the required redaction policy, then click **View**.
7. To exit, click **OK**.

11.5.5 Enabling or Disabling an Oracle Data Redaction Policy in Enterprise Manager

An Oracle Data Redaction policy is executed at run time only if it is enabled. When you create an Oracle Data Redaction policy, it is enabled by default.

You can disable an enabled redaction policy, or enable a disabled redaction policy using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then search for and click the name of the database target for which the Oracle Data Redaction policy that you want to enable or disable was created.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.

5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. In the Policies section of the Policies tab, select the redaction policy that you want to enable or disable, and then click **Enable** or **Disable**.

Schema	Table/View	Policy Name	Enabled	Redacted Columns
HR	EMPLOYEES	emp_comm_pol	✓	2

7. In the Confirmation dialog box, click **Yes** or **No**.

11.5.6 Deleting an Oracle Data Redaction Policy Using Enterprise Manager

You can delete an Oracle Data Redaction policy using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user SYSTEM or SYSMAN.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then search for and click the name of the database target for which the Oracle Data Redaction policy that you want to delete was created.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. In the Policies section of the Policies tab, select the redaction policy that you want to delete, and then click **Delete**.
7. In the Confirmation dialog box, click **Yes** or **No**.

11.6 Managing Named Data Redaction Policy Expressions Using Enterprise Manager

You can manage Oracle Data Redaction policy expressions in Enterprise Manager Cloud Control.

[About Named Data Redaction Policy Expressions in Enterprise Manager](#)
(page 11-18)

You can create and apply named Oracle Data Redaction policy expression to multiple columns in tables and views in Oracle Enterprise Manager Cloud Control.

[Creating a Named Data Redaction Policy Expression in Enterprise Manager](#)
(page 11-18)

You can create and apply a named Oracle Data Redaction policy expression using Enterprise Manager Cloud Control.

[Editing a Named Data Redaction Policy Expression in Enterprise Manager](#)
(page 11-19)

You can edit a named Oracle Data Redaction policy expression using Enterprise Manager Cloud Control.

[Viewing Named Data Redaction Policy Expressions in Enterprise Manager](#)
(page 11-20)

You can view named Oracle Data Redaction policy expressions using Enterprise Manager Cloud Control.

[Deleting a Named Data Redaction Policy Expression in Enterprise Manager](#)
(page 11-21)

You can delete named Oracle Data Redaction policy expressions using Enterprise Manager Cloud Control.

11.6.1 About Named Data Redaction Policy Expressions in Enterprise Manager

You can create and apply named Oracle Data Redaction policy expression to multiple columns in tables and views in Oracle Enterprise Manager Cloud Control.

When you modify the policy expression, the change is reflected in all redacted columns in the database instance that use the policy expression. Cloud Control enables you to create, edit, view, apply to columns, and delete policy expressions. Before you can create and use named Data Redaction policy expressions, ensure that the `COMPATIBLE` initialization parameter is set to `12.2.0.0`.

Related Topics:

[Creating and Managing Multiple Named Policy Expressions](#) (page 10-13)

A named, centrally managed Oracle Data Redaction policy expression can be used in multiple redaction policies and applied to multiple tables or views.

11.6.2 Creating a Named Data Redaction Policy Expression in Enterprise Manager

You can create and apply a named Oracle Data Redaction policy expression using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target for which you want to create an Oracle Data Redaction policy.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT PL/SQL` package.

6. On the Oracle Data Redaction page, select the **Expressions** tab.
7. Click **Create**.

The Create dialog box appears.

8. In the Create dialog box, enter the following information:
 - **Expression Name:** Enter a name for the policy expression. Existing policy expressions are listed on the Data Redaction page.
 - **Description:** Enter a brief description of the policy.
 - **Expression:** Enter the expression. For more complex expressions, such as applying or exempting the policy from specific users, click the Policy Expression Builder icon at the right of the **Expression** field. Click **OK** in the Policy Expression Builder to create the expression.
9. Click **OK** in the Create dialog box.

After you create the policy expression, it is listed in the Data Redaction page and ready to be associated with a Data Redaction policy.
10. In the Data Redaction page, select the **Policies** tab.
11. Under Policies, select the row for the policy that redacts the column to which you want to apply the policy expression, and then click **Edit**.
12. Under Object Columns, select the column that you want and then click the **Modify** button.
13. In the Modify dialog box, select the expression from the **Expression Name** list.
14. Click **OK**, and then click **OK** again in the Edit Data Redaction Policy dialog box.

11.6.3 Editing a Named Data Redaction Policy Expression in Enterprise Manager

You can edit a named Oracle Data Redaction policy expression using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

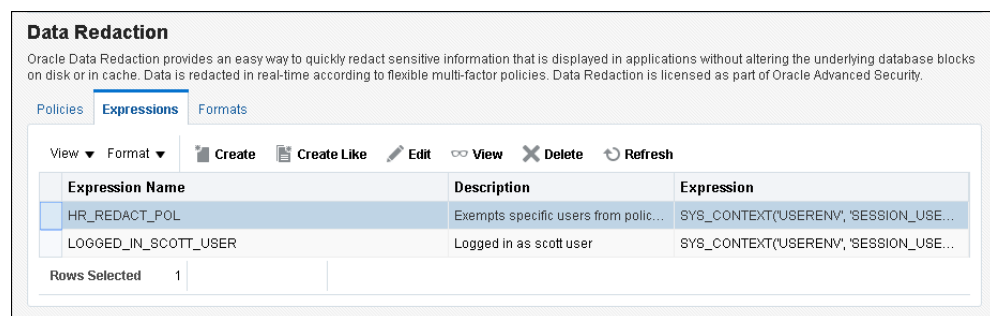
The URL is as follows:

`https://host:port/em`

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target for which you want to create an Oracle Data Redaction policy.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the EXECUTE privilege on the DBMS_REDACT PL/SQL package.

6. On the Oracle Data Redaction page, select the **Expressions** tab.
7. Select the policy expression that you want to edit and then click **Edit**.



8. In the Edit dialog box, modify the **Description** and **Expression** fields as necessary. For more complex expressions, click the Policy Expression Builder icon, and then click **OK** after you have recreated the expression.
9. Click **OK** in the Edit dialog box.

11.6.4 Viewing Named Data Redaction Policy Expressions in Enterprise Manager

You can view named Oracle Data Redaction policy expressions using Enterprise Manager Cloud Control.

1. Log into Oracle Enterprise Manager Cloud Control as either user SYSTEM or SYSMAN.

The URL is as follows:

`https://host:port/em`

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target for which you want to create an Oracle Data Redaction policy.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package.

6. On the Oracle Data Redaction page, select the **Expressions** tab.
7. Select the policy expression that you want to view and then click **View**.
The View dialog box appears, showing the definition of the policy expression.
8. Click **OK** to exit the View dialog box.

11.6.5 Deleting a Named Data Redaction Policy Expression in Enterprise Manager

You can delete named Oracle Data Redaction policy expressions using Enterprise Manager Cloud Control.

The deletion process first dissociates the policy expression from all columns to which it is applied.

1. Log into Oracle Enterprise Manager Cloud Control as either user `SYSTEM` or `SYSMAN`.

The URL is as follows:

```
https://host:port/em
```

2. From the **Targets** menu, select **Databases**.
3. Select **Search List**, then click the name of a database target for which you want to create an Oracle Data Redaction policy.
4. On the home page of the database target, from the **Security** menu, select **Data Redaction**.
5. Log in to the database, if you are prompted to do so.

Ensure that you log in to the database as a user that has the `EXECUTE` privilege on the `DBMS_REDACT` PL/SQL package.

6. On the Oracle Data Redaction page, select the **Expressions** tab.
7. Select the policy expression that you want to delete, and then click **Delete**.
The Delete Expressions confirmation dialog box appears.
8. Click **OK**.

Oracle Data Redaction Use with Oracle Database Features

Oracle Data Redaction can be used with other Oracle features, but some Oracle features may have restrictions with regard to Oracle Data Redaction.

[Oracle Data Redaction and DML and DDL Operations](#) (page 12-2)

Oracle Data Redaction affects DML and DDL operations, especially for users who issue ad-hoc SQL against tables with redacted columns.

[Oracle Data Redaction and Nested Functions, Inline Views, and the WHERE Clause](#) (page 12-2)

You can use Oracle Data Redaction with nested functions, inline views, and the WHERE clause in SELECT statements.

[Oracle Data Redaction and Database Links](#) (page 12-3)

Do not create Oracle Data Redaction policies on database views that reference database links.

[Oracle Data Redaction and Aggregate Functions](#) (page 12-3)

Aggregate functions can affect performance overhead on Oracle Data Redaction policies.

[Oracle Data Redaction and Object Types](#) (page 12-3)

You can use object types to model real-world entities such as customer accounts.

[Oracle Data Redaction and XML Generation](#) (page 12-3)

You cannot use XML generation functions on columns that have Oracle Data Redaction policies defined on them.

[Oracle Data Redaction and Editions](#) (page 12-3)

You cannot redact editioned views.

[Oracle Data Redaction in a Multitenant Environment](#) (page 12-3)

In a multitenant environment, Oracle Data Redaction policies apply only to the objects within the current pluggable database (PDB).

[Oracle Data Redaction and Oracle Virtual Private Database](#) (page 12-4)

Oracle Data Redaction does not affect Oracle Virtual Private Database policies because the VPD inline view, which contains the VPD predicate, acts on actual values.

[Oracle Data Redaction and Oracle Database Real Application Security](#) (page 12-4)

Oracle Data Redaction differs from Oracle Database Real Application Security because of how security is implemented for applications.

[Oracle Data Redaction and Oracle Database Vault](#) (page 12-4)

You can use Oracle Data Redaction in an Oracle Database Vault environment.

[Oracle Data Redaction and Oracle Data Pump](#) (page 12-5)

Oracle Data Pump export operations can affect objects that have Oracle Data Redaction policies.

[Oracle Data Redaction and Data Masking and Subsetting Pack](#) (page 12-8)

Oracle Enterprise Manager Data Masking and Subsetting Pack can be used to create a development or test copy of a production database.

[Oracle Data Redaction and JSON](#) (page 12-8)

JavaScript Object Notation (JSON) can be used to create `is json` constraints on table columns.

12.1 Oracle Data Redaction and DML and DDL Operations

Oracle Data Redaction affects DML and DDL operations, especially for users who issue ad-hoc SQL against tables with redacted columns.

Note the following:

- Oracle Data Redaction treats the `RETURNING INTO` clause of a DML statement as a query, even though the result is not displayed. The result that is sent to the buffer is what would have been displayed had the `RETURNING INTO` clause been run as an ordinary SQL query, rather than as part of a DML statement. If your application performs further processing on the buffer that contains the `RETURNING INTO` value, then consider changing the application because it may not expect to find a redacted value in the buffer.
- If a redacted column appears as the source in a DML or DDL operation, then Oracle Data Redaction considers this as an attempt to circumvent the policy and prevents it with an `ORA-28081: Insufficient privileges - the command references a redacted object` error unless you have the `EXEMPT REDACTION POLICY` system privilege. Internally, Oracle Data Pump issues these kinds of operations, so you may also need to grant the `EXEMPT REDACTION POLICY` system privilege to a user if they need to perform schema-level exports of tables that have redacted columns.

12.2 Oracle Data Redaction and Nested Functions, Inline Views, and the WHERE Clause

You can use Oracle Data Redaction with nested functions, inline views, and the `WHERE` clause in `SELECT` statements.

Oracle Data Redaction policies work as follows:

- **Nested functions are redacted innermost.** For example, in `SELECT SUM(AVG(TO_NUMBER((X))) FROM HR.EMPLOYEES WHERE ...`, the `TO_NUMBER` function is redacted first, followed by `AVG`, and then last the `SUM` function.
- **Inline views are redacted outermost.** For example, in `SELECT XYZ ... AS SELECT A... AS SELECT B... AS SELECT C...`, `SELECT XYZ` is redacted first, followed by `AS SELECT A`, then `AS SELECT B`, and so on. `AS SELECT C` is redacted last.

- **The WHERE clause is never redacted.** Data Redaction redacts only data in the column SELECT list.

12.3 Oracle Data Redaction and Database Links

Do not create Oracle Data Redaction policies on database views that reference database links.

You can find information about existing database links by querying the `DBA_DB_LINKS` data dictionary view.

See Also:

Oracle Database Administrator's Guide for detailed information about database links

12.4 Oracle Data Redaction and Aggregate Functions

Aggregate functions can affect performance overhead on Oracle Data Redaction policies.

Because Oracle Data Redaction dynamically modifies the value of each row in a column, certain SQL queries that use aggregate functions cannot take full advantage of database optimizations that presume the row values to be static.

In the case of SQL queries that call aggregate functions, it may be possible to notice performance overhead due to redaction.

12.5 Oracle Data Redaction and Object Types

You can use object types to model real-world entities such as customer accounts.

An object type is a user-defined type. You cannot redact object types. This is because Database Redaction cannot handle all of the possible ways that object types can be configured, because they are user defined. You can find the type that an object uses by querying the `OBJECT_NAME` and `OBJECT_TYPE` columns of the `ALL_OBJECTS` data dictionary view.

12.6 Oracle Data Redaction and XML Generation

You cannot use XML generation functions on columns that have Oracle Data Redaction policies defined on them.

Oracle XML DB Developer's Guide describes the kinds of SQL functions to which this restriction applies. This restriction applies irrespective of whether the Oracle Data Redaction policy has been enabled or disabled, or is active for the querying user.

12.7 Oracle Data Redaction and Editions

You cannot redact editioned views.

In addition to not being able to redact editioned views, you cannot use a redacted column in the definition of any editioned view. You can find information about editions by querying the `DBA_EDITIONS` data dictionary view.

12.8 Oracle Data Redaction in a Multitenant Environment

In a multitenant environment, Oracle Data Redaction policies apply only to the objects within the current pluggable database (PDB).

You cannot create a Data Redaction policy for a multitenant container database (CDB). This is because the objects for which you create Data Redaction policies typically reside in a PDB. You can find all the PDBs in a CDB by querying the `DBA_PDBS` data dictionary view.

As with the CDB root, you cannot create Data Redaction policies in an application root.

12.9 Oracle Data Redaction and Oracle Virtual Private Database

Oracle Data Redaction does not affect Oracle Virtual Private Database policies because the VPD inline view, which contains the VPD predicate, acts on actual values.

Oracle Data Redaction differs from Oracle Virtual Private Database in the following ways:

- Oracle Data Redaction provides more redacting features than Oracle Virtual Private Database, which only supports `NULL` redacting. Many applications cannot use `NULL` redacting, so Data Redaction is a good solution for these applications.
- Oracle Virtual Private Database policies can be static, dynamic, and context sensitive, whereas Data Redaction policies only allow static and context-sensitive policy expressions.
- Data Redaction permits only one policy to be defined on a table or view, whereas you can define multiple Virtual Private Database policies on an object.
- Data Redaction is when application users try to access an object that is protected by a Data Redaction policy using a synonym, but (unlike Oracle Virtual Private Database) Data Redaction does not support the creation of policies directly on the synonyms themselves.

12.10 Oracle Data Redaction and Oracle Database Real Application Security

Oracle Data Redaction differs from Oracle Database Real Application Security because of how security is implemented for applications.

Oracle Data Redaction differs from Oracle Database Real Application Security in that Real Application Security provides a comprehensive authorization framework for application security.

Column security within Real Application Security is based on application privileges that are defined by applications using the Real Application Security framework.

See Also:

Oracle Database Real Application Security Administrator's and Developer's Guide for information about how you can protect table columns with custom application privileges

12.11 Oracle Data Redaction and Oracle Database Vault

You can use Oracle Data Redaction in an Oracle Database Vault environment.

For example, if there is an Oracle Database Vault realm around an object, a user who does not belong to the authorized list of realm owners or participants cannot see the object data, regardless of whether the user was granted the `EXEMPT REDACTION`

POLICY privilege. If the user attempts a DML or DDL statement on the data, error messages result.

12.12 Oracle Data Redaction and Oracle Data Pump

Oracle Data Pump export operations can affect objects that have Oracle Data Redaction policies.

[Oracle Data Pump Security Model for Oracle Data Redaction](#) (page 12-5)

The DATAPUMP_EXP_FULL_DATABASE role includes the powerful EXEMPT REDACTION POLICY system privilege.

[Export of Objects That Have Oracle Data Redaction Policies Defined](#) (page 12-5)

You can export objects that have already had Oracle Data Redaction policies defined on them.

[Export of Data Using the EXPDP Utility access_method Parameter](#) (page 12-7)

Oracle Data Pump can export data from a schema that contains an object that has a Data Redaction policy.

[Import of Data into Objects Protected by Oracle Data Redaction](#) (page 12-7)

During an import operation, be careful that you do not inadvertently drop data redaction policies that protect imported data.

12.12.1 Oracle Data Pump Security Model for Oracle Data Redaction

The DATAPUMP_EXP_FULL_DATABASE role includes the powerful EXEMPT REDACTION POLICY system privilege.

Remember that by default the DBA role is granted the DATAPUMP_EXP_FULL_DATABASE role as well as DATAPUMP_IMP_FULL_DATABASE.

This enables users who were granted these roles to be exempt from Data Redaction policies. This means that, when you export objects with Data Redaction policies defined on them, the **actual data** in the protected tables is copied to the Data Pump target system without being redacted. Users with these roles, including users who were granted the DBA role, are able to see the actual data in the target system.

However, by default, all of the Data Redaction policies associated with any tables and views in the Data Pump source system are also included in the export and import operation (along with the objects themselves) and applied to the objects in the target system, so the data is still redacted when users query the objects in the target system.

Related Topics:

[Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)

You can exempt users from having Oracle Data Redaction policies applied to the data they access.

12.12.2 Export of Objects That Have Oracle Data Redaction Policies Defined

You can export objects that have already had Oracle Data Redaction policies defined on them.

[Finding Type Names Used by Oracle Data Pump](#) (page 12-6)

You must find the type names Oracle Data Pump uses before exporting objects that have Oracle Data Redaction policies defined on these objects.

[Exporting Only the Data Dictionary Metadata Related to Data Redaction Policies](#) (page 12-6)

You can export only the data dictionary metadata that is related to data redaction policies and full redaction settings.

[Importing Objects Using the INCLUDE Parameter in IMPDP](#) (page 12-7)

You can import objects using Oracle Database Pump.

12.12.2.1 Finding Type Names Used by Oracle Data Pump

You must find the type names Oracle Data Pump uses before exporting objects that have Oracle Data Redaction policies defined on these objects.

After you find these types, you should use these types as parameters for the `INCLUDE` directive to the `IMPDP` utility, to selectively export only metadata of these specific types to the dump file.

- To find type names, query the `DATABASE_EXPORT_OBJECTS` view.

For example:

```
SELECT OBJECT_PATH
FROM DATABASE_EXPORT_OBJECTS
WHERE OBJECT_PATH LIKE 'RADM_%';
```

Output similar to the following appears:

```
OBJECT_PATH
-----
RADM_FPTM
RADM_POLICY
```

12.12.2.2 Exporting Only the Data Dictionary Metadata Related to Data Redaction Policies

You can export only the data dictionary metadata that is related to data redaction policies and full redaction settings.

This kind of Data Pump export could, for example, be used if you must use the same set of Data Redaction policies and settings across development, test, and production databases. Because the flag `content=metadata_only` is specified, the dump file does not contain any actual data.

- To export only the data dictionary metadata related to data redaction policies and full redaction settings, enter an `EXPDP` utility command similar to the following:

```
expdp system/password \
full=y \
COMPRESSION=NONE \
content=metadata_only \
INCLUDE=RADM_FPTM,RADM_POLICY\
directory=my_directory \
job_name=my_job_name \
dumpfile=my_data_redaction_policy_metadata.dmp
```

See Also:

- *Oracle Database Utilities* for detailed information about the `INCLUDE` parameter of the `EXPDP` utility
 - *Oracle Database Utilities* for detailed information about metadata filters
-
-

12.12.2.3 Importing Objects Using the `INCLUDE` Parameter in `IMPDP`

You can import objects using Oracle Database Pump.

- To import the objects, include these names in the `INCLUDE` parameter in the `IMPDP` utility command, based on the output from querying the `OBJECT_PATH` column in the `DATABASE_EXPORT_OBJECTS` view.

12.12.3 Export of Data Using the `EXPDP` Utility `access_method` Parameter

Oracle Data Pump can export data from a schema that contains an object that has a Data Redaction policy.

If you are using Oracle Data Pump to perform full database export operations using the Data Pump default settings (`direct_path`), and if you receive error messages that you do not understand, then use this section to repeat the operation in such a way as to better understand the error.

If you try to use the Oracle Data Pump Export (`EXPDP`) utility with the `access_method` parameter set to `direct_path` to export data from a schema that contains an object that has a Data Redaction policy defined on it, then the following error message may appear and the export operation fails:

```
ORA-31696: unable to export/import TABLE_DATA:"schema.table" using client specified
DIRECT_PATH method
```

This problem only occurs when you perform a schema-level export as a user who was not granted the `EXP_FULL_DATABASE` role. It does not occur during a full database export, which requires the `EXP_FULL_DATABASE` role. The `EXP_FULL_DATABASE` role includes the `EXEMPT REDACTION POLICY` system privilege, which bypasses Data Redaction policies.

To find the underlying problem, try the `EXPDP` invocation again, but do not set the `access_method` parameter to `direct_path`. Instead, use either `automatic` or `external_table`. The underlying problem could be a permissions problem, for example:

```
ORA-28081: Insufficient privileges - the command references a redacted object.
```

See Also:

Oracle Database Utilities for more information about using Data Pump Export.

12.12.4 Import of Data into Objects Protected by Oracle Data Redaction

During an import operation, be careful that you do not inadvertently drop data redaction policies that protect imported data.

Consider a scenario in which the source tables that were exported using the Oracle Data Pump Export (`EXPDP`) utility do not have Oracle Data Redaction policies.

However, the destination tables to which the data is to be imported by using Oracle Data Pump Import (IMPDP) have Oracle Data Redaction policies.

During the Data Pump import operation, the status of the Data Redaction policies on the objects being imported depends on the `CONTENT` option of IMPDP command.

- If you use the `CONTENT=ALL` or `CONTENT=METADATA_ONLY` option in the IMPDP command, then the Data Redaction policies on the destination tables are dropped. You must recreate the Data Redaction policies.
- If you use `CONTENT=DATA_ONLY` in the IMPDP command, then the Data Redaction policies on the destination tables are not dropped.

See Also:

Oracle Database Utilities for more information about using Data Pump Export

12.13 Oracle Data Redaction and Data Masking and Subsetting Pack

Oracle Enterprise Manager Data Masking and Subsetting Pack can be used to create a development or test copy of a production database.

To accomplish this, you can mask this data in bulk, and then put the resulting masked data in the development or test copy.

You can still apply Data Redaction policies to the non-production database, in order to redact columns that contain data that was already masked by Oracle Enterprise Manager Data Masking and Subsetting Pack.

Remember that Oracle Enterprise Manager Data Masking and Subsetting Pack is used to mask data sets when you want to move the data to development and test environments. Data Redaction is mainly designed for redacting at runtime for production applications in a consistent fashion across multiple applications, without having to make application code changes.

See Also:

Oracle Data Masking and Subsetting Guide for more information about data masking and subsetting

12.14 Oracle Data Redaction and JSON

JavaScript Object Notation (JSON) can be used to create `is json` constraints on table columns.

However, you cannot create an Oracle Data Redaction policy on a table column that has the `is json` constraint. If you attempt to do so, an `ORA-28073 - The column column_name has an unsupported datatype error` is raised. As a workaround solution, Oracle recommends that you create a relational view that uses the `JSON_TABLE` row source operator on top of the JSON object, and then apply the Data Redaction policy to this view.

See *Oracle Database SQL Language Reference* for more information about `JASON_TABLE`.

Security Considerations for Oracle Data Redaction

Oracle provides guidelines for using Oracle Data Redaction.

[Oracle Data Redaction General Usage Guidelines](#) (page 13-1)

It is important to understand general guidelines for using Oracle Data Redaction.

[Restriction of Administrative Access to Oracle Data Redaction Policies](#)
(page 13-2)

You can restrict the list of users who can create, view and edit Data Redaction policies.

[How Oracle Data Redaction Affects the SYS, SYSTEM, and Default Schemas](#)
(page 13-2)

Both users SYS and SYSTEM automatically have the EXEMPT REDACTION POLICY system privilege.

[Policy Expressions That Use SYS_CONTEXT Attributes](#) (page 13-3)

Be careful when writing a policy expression that depends on a SYS_CONTEXT attribute that is populated by an application.

[Oracle Data Redaction Policies on Materialized Views](#) (page 13-3)

You can create Oracle Data Redaction policies on materialized views and on their base tables.

[Dropped Oracle Data Redaction Policies When the Recycle Bin Is Enabled](#)
(page 13-3)

You should check if the recycle bin is enabled before you drop Oracle Data Redaction policies.

13.1 Oracle Data Redaction General Usage Guidelines

It is important to understand general guidelines for using Oracle Data Redaction.

- Oracle Data Redaction is not intended to protect against attacks by regular and privileged database users who run ad hoc queries directly against the database. If the user can issue arbitrary SQL or PL/SQL statements, then he or she will be able to access the actual value.
- Oracle Data Redaction is not intended to protect against users who run ad hoc SQL queries that attempt to determine the actual values by [inference](#).
- Oracle Data Redaction relies on the database and application context values. For applications, it is the responsibility of the application to properly initialize the context value.

- Oracle Data Redaction is not enforced for users who are logged in using the SYSDBA administrative privilege.
- Certain DDL statements that attempt to copy the [actual data](#) out from under the control of a data redaction policy (that is, CREATE TABLE AS SELECT, INSERT AS SELECT) are blocked by default, but you can disable this behavior by granting the user the EXEMPT REDACTION POLICY system privilege.
- Oracle Data Redaction does not affect day-to-day database operations, such as backup and recovery, Oracle Data Pump exports and imports, Oracle Data Guard operations, and replication.
- Do not include any redacted columns in a SQL expression that is used in a GROUP BY clause in a SQL statement. Oracle does not support this behavior, and raises an ORA-00979: not a GROUP BY expression error. This happens because internally the expression in the SELECT list must be modified by Data Redaction, but this causes it to no longer be found when it comes time to process the GROUP BY clause (which is currently not updated by Data Redaction) leading to this unintended error message.

13.2 Restriction of Administrative Access to Oracle Data Redaction Policies

You can restrict the list of users who can create, view and edit Data Redaction policies.

To accomplish this, you can limit who has the EXECUTE privilege on the DBMS_REDACT package and by limiting who has the SELECT privilege on the REDACTION_POLICIES and REDACTION_COLUMNS views.

You also can restrict who is exempted from redaction by limiting the EXEMPT REDACTION POLICY privilege. If you use Oracle Database Vault to restrict privileged user access, then you can use realms to restrict granting of EXEMPT REDACTION POLICY.

See Also:

- [Exemption of Users from Oracle Data Redaction Policies](#) (page 10-50)
 - [Oracle Data Redaction and Oracle Database Vault](#) (page 12-4)
 - *Oracle Database Vault Administrator's Guide* for more information about Oracle Database Vault
-
-

13.3 How Oracle Data Redaction Affects the SYS, SYSTEM, and Default Schemas

Both users SYS and SYSTEM automatically have the EXEMPT REDACTION POLICY system privilege.

SYSTEM has the EXP_FULL_DATABASE role, which includes the EXEMPT REDACTION POLICY system privilege.

This means that the SYS and SYSTEM users can always bypass any existing Oracle Data Redaction policies, and will always be able to view data from tables (or views) that have Data Redaction policies defined on them.

Follow these guidelines:

- Do not create Data Redaction policies on the default Oracle Database schemas, including the SYS and SYSTEM schemas.
- Be aware that granting the EXEMPT DATA REDACTION system privilege to additional roles may enable users to bypass Oracle Data Redaction, because the grantee role may have been granted to additional roles.
- Do not revoke the EXEMPT DATA REDACTION system privilege from the roles that it was granted to by default.

13.4 Policy Expressions That Use SYS_CONTEXT Attributes

Be careful when writing a policy expression that depends on a SYS_CONTEXT attribute that is populated by an application.

The application might not always populate that attribute.

If the user somehow connects directly (rather than through the application), then the SYS_CONTEXT attribute would not have been populated. If you do not handle this NULL scenario in your policy expression, you could unintentionally reveal [actual data](#) to the querying user.

For example, suppose you wanted to create a policy expression that intends to redact the query results for everyone except users who have the client identifier value of SUPERVISOR. The following expression unintentionally enables querying users who have NULL as the value for their CLIENT_IDENTIFIER to see the real data:

```
SYS_CONTEXT('USERENV', 'CLIENT_IDENTIFIER') IS NOT 'SUPERVISOR'
```

A more rigorous policy expression redacts the result of the query if the client identifier is not set, that is, it has a NULL value.

```
SYS_CONTEXT('USERENV', 'CLIENT_IDENTIFIER') IS NOT 'SUPERVISOR' OR IS NULL
```

Remember that in SQL, comparisons with NULL are undefined, and are thus FALSE, but redaction only takes place when the policy expression evaluates to TRUE.

13.5 Oracle Data Redaction Policies on Materialized Views

You can create Oracle Data Redaction policies on materialized views and on their base tables.

However, ensure that the creator of the materialized view, or the user who performs the refresh of the materialized view, is not blocked by any Data Redaction policies. In other words, the user performing the materialized view creation or refresh operations should be exempt from the Data Redaction policy. As a best practice, when you create a new materialized view, treat it as a copy of the actual table, and then create a separate Data Redaction policy to protect it.

13.6 Dropped Oracle Data Redaction Policies When the Recycle Bin Is Enabled

You should check if the recycle bin is enabled before you drop Oracle Data Redaction policies.

If you drop a table or view that has an Oracle Data Redaction policy defined on it when the recycle bin feature is enabled, and if you query the REDACTION_COLUMNS or REDACTION_POLICIES data dictionary views before you purge the recycle bin, then you will see object names such as BIN\$. . . (for example, BIN\$1Xu5PSW5VaPgQxGS5AoAEA==\$0).

This is normal behavior. These policies are removed when you purge the recycle bin.

To find if the recycle bin is enabled, you can run the `SHOW PARAMETER RECYCLEBIN` command in SQL*Plus.

See Also:

Oracle Database Administrator's Guide for information about purging objects from the recycle bin

Glossary

actual data

In Oracle Data Redaction, the data in a protected table or view. An example of actual data could be the number 123456789, and the [redacted data](#) version of this number could be 999996789.

auto-login software keystore

A [software keystore](#) that is protected by a system-generated password and does not need to be explicitly opened by a security administrator. Auto-login software keystores are automatically opened when accessed and can be used on any computer that runs an Oracle database. For example, consider an Oracle RAC environment that has four nodes, and each node is on a different computer. This environment uses an auto-login keystore. When a REKEY operation is performed on node 1, the auto-login and password-based keystores must be copied to the computers that host nodes 2, 3, and 4. In this configuration, the auto-login keystores will be opened on all four nodes when required.

See also [local auto-login software keystore](#).

cipher suite

A set of authentication, encryption, and data integrity algorithms used to exchange messages between network nodes using Secure Sockets Layer (SSL). During an SSL handshake, for example, the two nodes negotiate to see which cipher suite they will use when transmitting messages back and forth.

ciphertext

Message text that has been encrypted.

See also [encrypted text](#).

data redaction

The ability to mask data with different values in real time, that is, at the moment a user tries to access the data. You can mask all of the data, a partial subset of the data, or display random values in place of the data. It does not change the actual data in the database.

decryption

The process of converting an encrypted message (the [ciphertext](#)), back to its original message ([plaintext](#)).

encrypted text

Text that has been encrypted, using an encryption algorithm and an encryption key; the output stream of an encryption process. The text is not readable or decipherable, without decrypting it first. Also called [ciphertext](#).

encryption

The process of converting an original message ([plaintext](#)) to an encrypted message ([ciphertext](#)).

hardware keystore

A container that stores a Transparent Data Encryption key for a hardware security module.

hardware security module

A physical device that provides secure storage for encryption keys.

inference

A query that is designed to find data by repeatedly trying queries. For example, to find the users who earn the highest salaries, an intruder could use the following query:

```
SELECT FIRST_NAME, LAST_NAME, SALARY FROM HR.EMPLOYEES WHERE SALARY > 16000 ORDER BY SALARY DESC;
```

FIRST_NAME	LAST_NAME	SALARY
-----	-----	-----
Steven	King	24000
Neena	Kochhar	17000
Lex	De Haan	17000

key pair

A [public key](#) and its associated [private key](#). See [public and private key pair](#).

keystore

A general term for any container that stores encryption keys, such as Transparent Data Encryption keys and other encrypted data. In previous releases, this container was referred to as a [wallet](#), which is specific to Oracle. Starting with Oracle Database 12c release 12.1, the term changed to keystore to encompass non-Oracle Database encryption key containers, such as hardware security modules.

See also [auto-login software keystore](#), [hardware keystore](#), and [local auto-login software keystore](#).

local auto-login software keystore

A [software keystore](#) that is local and restricted to the computer on which it was created.

See also [auto-login software keystore](#).

mask

The ability to redact data from a user or an application.

password-based software keystore

A [software keystore](#) that must be opened with a password before it can be accessed.

See also [keystore](#).

plaintext

Message text that has not been encrypted.

private key

In public-key cryptography, this key is the private key that is known only to its owner. It is primarily used for encrypting message digests used with digital signatures.

See [public and private key pair](#).

public key

One of two keys that are used in public key cryptography, the other key being the [private key](#). In typical public key cryptography usage, the public key is used to encrypt data or verify digital signatures. The the private key is used to decrypt data or generate digital signatures. The public key, unlike the private key, can be made available to anyone whereas the private key must remain secret.

See [public and private key pair](#).

public key encryption

The process where the sender of a message encrypts the encryption key of the recipient. Upon delivery, the message is decrypted by the recipient using its private key.

public and private key pair

A set of two related numbers used for [encryption](#) and [decryption](#), where one is called the [private key](#) and the other is called the [public key](#). Public keys are typically made widely available, while private keys are held by their respective owners. Data encrypted with either a public key or a private key from a [key pair](#) can be decrypted with its associated key from the key pair.

public key infrastructure (PKI)

Information security technology utilizing the principles of public key cryptography. Public key cryptography involves encrypting and decrypting information using a shared public and private key pair. Provides for secure, private communications within a public network.

redacted data

Masked data that is displayed to the querying user. For example, if the [actual data](#) is 3714-4963-5398-4321, then the redacted data could appear, depending on the Data Redaction policy, as [xxxx-xxxx-xxxx-4321](#).

salt

In cryptography, a way to strengthen the security of encrypted data. Salt is a random string that is added to the data before it is encrypted, making it more difficult for attackers to steal the data by matching patterns of ciphertext to known ciphertext samples. Salt is often also added to passwords, before the passwords are hashed, to avoid dictionary attacks, a method that attackers use to determine sensitive passwords. The addition of salt to a password before hashing makes it more difficult for intruders to match the hash values (sometimes called verifiers) with their dictionary list of common password hash values, because they do not know the salt beforehand.

software keystore

A container that stores a Transparent Data Encryption a TDE master encryption key for use as an [auto-login software keystore](#), a [local auto-login software keystore](#), or a [password-based software keystore](#).

tablespace encryption key

An encryption key for the encryption of a tablespace. The TDE tablespace encryption key encrypts the tablespace encryption key, which in turn encrypts and decrypts data in the tablespace.

TDE master encryption key

A key that is stored within a [software keystore](#) or a [hardware keystore](#). For table encryption, this key encrypts the TDE table key, and for tablespace encryption, it encrypts the tablespace encryption key.

See also [keystore](#).

TDE table key

An encryption key that is associated with a table whose columns are marked for encryption. The TDE master encryption key encrypts this table encryption key.

wallet

A data structure used to store and manage security credentials for an individual entity. Wallets are specific to Oracle Database only. A [Wallet Resource Locator \(WRL\)](#)

provides all of the necessary information to locate the wallet. For Transparent Data Encryption in Oracle Database Release 12c and later, the term for wallet is [keystore](#).

wallet obfuscation

The ability to store and access an Oracle [wallet](#) without querying the user for a password before access (supports single sign-on (SSO)).

Wallet Resource Locator (WRL)

A tool that provides all of the necessary information to locate a [wallet](#). It is a path to an operating system directory that contains a wallet.

A

ad hoc tools
 Oracle Data Redaction, [8-3](#)

administrative access to policies, restricting, [13-2](#)

aggregate functions
 affect on Data Redaction policy optimization, [12-3](#)

ALTER SYSTEM statement
 how compares with ADMINISTER KEY
 MANAGEMENT statement, [5-6](#)

APEX_UTIL.GET_NUMERIC_SESSION_STATE
 function
 Oracle Data Redaction policies (NV public
 function), [10-12](#)

APEX_UTIL.GET_SESSION_STATE function
 Oracle Data Redaction policies (V public
 function), [10-12](#)

applications
 database applications and Oracle Data Redaction,
 [8-3](#)
 modifying to use Transparent Data Encryption,
 [5-5](#)

auto login keystores
 and Transparent Data Encryption (TDE), [4-37](#)

Automatic Storage Management (ASM)
 moving software keystores from, [4-12](#)

C

CDBs
 cloning PDBs with encrypted data, [6-17](#)
 cloning PDBs with encrypted data, about, [6-17](#)
 Data Redaction masking policies, [12-3](#)
 moving PDB from one CDB to another, [6-11](#)
 PDBs with encrypted data, [6-14](#)
 preserving keystore passwords in PDB move
 operations, [6-11](#)
 TDE operations in root, [6-9](#)
 TDE operations in root and PDBs, [6-11](#)

change data capture, synchronous, [3-21](#)

closing hardware keystores, [4-21](#)

closing software keystores, [4-21](#)

column encryption
 about, [2-3](#)

column encryption (*continued*)
 changing algorithm, [3-28](#)
 changing encryption key, [3-28](#)
 creating encrypted table column with default
 algorithm, [3-23](#)
 creating encrypted table column with non-default
 algorithm, [3-24](#)
 creating index on encrypted column, [3-27](#)
 data loads from external file, [5-11](#)
 data types to encrypt, [3-20](#)
 existing tables
 about, [3-26](#)
 adding encrypted column to, [3-27](#)
 disabling encryption, [3-27](#)
 encrypting unencrypted column, [3-27](#)
 external tables, [3-25](#)
 incompatibilities, [7-1](#)
 limitations, [7-1](#)
 performance, optimum, [7-4](#)
 restrictions, [3-21](#)
 salt, [3-28](#)
 security considerations, [5-3](#)
 skipping integrity check, [3-24](#)

column sensitive type discovery
 enabling when creating a Data Redaction policy,
 [11-11](#)

compliance
 Transparent Data Encryption, [2-1](#)

compression of Transparent Data Encryption data, [5-1](#)

configuring software keystores
 creating local auto-login keystore, [3-8](#)

D

data at rest, [2-1](#)

data deduplication of Transparent Data Encryption
 data, [5-1](#)

data redaction
 See Oracle Data Redaction

Data Redaction supported functions, [10-7](#)

data storage
 Transparent Data Encryption, [5-4](#)

database close operations

- database close operations (*continued*)
 - keystores, [5-11](#)
- database links
 - with Oracle Data Redaction policies, [12-3](#)
- database roles
 - Data Redaction policies, [10-11](#)
- databases
 - about encrypting, [3-29](#)
 - encrypting existing, [3-48](#)
 - encrypting offline, [3-49](#)
 - encrypting online, [3-50](#)
- DDL statements
 - Oracle Data Redaction policies, [12-2](#)
- decryption
 - tablespaces, offline, [3-38](#), [3-40](#)
 - tablespaces, online, [3-43](#)
- DML statements
 - Oracle Data Redaction policies, [12-2](#)

E

- editing custom formats, [11-7](#)
- editing policies, [11-14](#)
- Editions
 - Transparent Data Encryption, [6-20](#)
- encrypted columns
 - data loads from external files, [5-11](#)
- encryption
 - cloning PDBs with encrypted data, [6-17](#)
 - databases offline, [3-49](#)
 - databases online, [3-50](#)
 - encrypting future tablespaces
 - about, [3-36](#)
 - existing databases, [3-48](#)
 - procedure, [3-37](#)
 - supported encryption algorithms, [3-43](#)
 - tablespaces, offline, [3-38](#)
 - tablespaces, online, [3-43](#)
 - See also* Transparent Data Encryption (TDE)
- encryption algorithms, supported, [3-43](#)
- EXEMPT REDACTION POLICY privilege
 - using with Database Vault, [13-2](#)
- expressions
 - LENGTH functions, character string, [10-9](#)
 - namespace functions, [10-8](#)
 - Oracle Application Express, [10-10](#)
 - Oracle Label Security functions, [10-10](#)
 - SUBSTR function, [10-9](#)
- external credential store, password-based software
 - keystores, [3-4](#)
- external files
 - loading data to tables with encrypted columns, [5-11](#)
- external keystores, [3-14](#)
- external store for passwords
 - open and close operations in CDB, [6-18](#)
 - operations in CDB environment, [6-8](#)

- external tables, encrypting columns in
 - ORACLE_DATPUMP, [5-11](#)
 - ORACLE_LOADER, [5-11](#)

G

- guidelines
 - ad hoc query attacks and Data Redaction, [13-1](#)
 - application context value handling by Data Redaction policies, [13-1](#)
 - day-to-day operations and Data Redaction, [13-1](#)
 - DDL statements and Data Redaction policies, [13-1](#)
 - exhaustive SQL queries and inference and Data Redaction, [13-1](#)
 - materialized views and Data Redaction, [13-3](#)
 - recycle bin and Data Redaction, [13-3](#)
 - SYS_CONTEXT values and Data Redaction, [13-3](#)

H

- hardware keystores
 - backing up, [4-7](#)
 - closing, [4-21](#)
- hardware security modules
 - backing up keystores, [4-7](#)
 - plugging PDBs, [6-16](#)
 - unplugging PDBs, [6-16](#)

I

- import/export utilities, original, [3-21](#)
- index range scans, [2-4](#)
- indexes
 - creating on encrypted column, [3-27](#)
- inline views
 - Data Redaction policies order of redaction, [12-2](#)
 - Data Redaction redaction, [12-2](#)
- intruders
 - ad hoc query attacks, [13-1](#)

J

- JSON
 - Oracle Data Redaction, [12-8](#)

K

- keystores
 - about, [2-6](#)
 - architecture, [2-3](#)
 - ASM-based, [4-23](#)
 - auto login, [4-37](#)
 - auto-login, open and close operations in CDBs, [6-18](#)
 - backing up password-based software keystores
 - about, [4-5](#)
 - backup identifier rules, [4-6](#)

keystores (*continued*)

- backing up password-based software keystores (*continued*)
 - procedure, [4-6](#)
- changing hardware keystore password, [4-4](#)
- changing passwords for password-based software
 - keystores, [4-3](#)
- closing hardware keystores, [4-21](#)
- closing in CDBs, [6-18](#)
- closing software keystores, [4-21](#)
- database close operations, [5-11](#)
- deleting, [4-24](#)
- external, [3-14](#)
- hardware keystore
 - configuration process, [3-13](#)
- master encryption key merge differing from
 - import or export, [4-44](#)
- merging
 - about, [4-8](#)
 - auto-login into password-based, [4-10](#)
 - one into another existing keystore, [4-10](#)
 - reversing merge operation, [4-11](#)
 - two into a third new keystore, [4-9](#)
- migrating
 - creating master encryption key for hardware
 - keystore-based encryption, [4-16](#)
 - hardware keystore to software keystore, [4-17](#)
 - keystore order after migration, [4-19](#)
 - password key into hardware keystore, [4-15](#)
- migration using Oracle Key Vault, [4-20](#)
- moving out of ASM, [4-12](#)
- moving software keystore to a new location, [4-11](#)
- opening hardware keystores, [3-15](#)
- opening in CDBs, [6-18](#)
- opening software keystores, [3-9](#)
- Oracle Database secrets
 - about, [4-46](#)
 - storing in hardware keystore, [4-49](#)
 - storing in software keystore, [4-47](#)
- password access, [4-2](#)
- password preservation in PDB move operations,
 - [6-11](#)
- possible states of, [3-9](#)
- using auto-login hardware keystore, [4-52](#)

keystores, software

- configuration process, [3-1](#)

L

LENGTH functions, character string

- expressions, [10-9](#)

M

masking

- See* Oracle Data Redaction

master encryption key

- See* TDE master encryption key

materialized views

- Data Redaction guideline, [13-3](#)
- Transparent Data Encryption tablespace
 - encryption, [6-4](#)

multitenant container databases

- See* CDBs

N

namespace functions

- expressions, [10-8](#)

nested functions

- Data Redaction policies order of redaction, [12-2](#)

NV public function

- (APEX_UTIL.GET_NUMERIC_SESSION_ST ATE function), Data Redaction policies, [10-12](#)

O

OLS_LABEL_DOMINATES public function

- Data Redaction policies, [10-12](#)

ONE_STEP_PLUGIN_FOR_PDB_WITH_TDE

- dynamic system parameter, [6-11](#)

opening hardware keystores, [3-15](#)

opening software keystores, [3-9](#)

ORA-00979

- not a GROUP BY expression error, [13-1](#)

ORA-28081

- Insufficient privileges - the command references a redacted object error, [12-2](#)

Oracle Application Express

- filtering using by session state in Data Redaction policies, [10-12](#)

Oracle Application Expression

- expressions, [10-10](#)

Oracle Call Interface

- Transparent Data Encryption, [6-20](#)

Oracle Data Guard

- Transparent Data Encryption, [6-4](#)

Oracle Data Pump

- encrypted columns, [6-2](#)
- encrypted data, [6-2](#)
- encrypted data with dump sets, [6-3](#)
- exported data from Data Redaction policies, [12-7](#)
- exporting Oracle Data Redaction objects, [12-6](#)
- imported data from Data Redaction policies, [12-7](#)
- Oracle Data Redaction security policy, [12-5](#)

Oracle Data Redaction

- about, [8-1](#)
- ad hoc tools, [8-3](#)
- aggregate functions, [12-3](#)
- benefits, [8-2](#)
- CDBs, [12-3](#)
- columns with XML-generated data, [12-3](#)
- creating custom format, [11-5](#)
- database applications, [8-3](#)
- DBMS_REDACT.ADD_POLICY procedure

Oracle Data Redaction (*continued*)

- DBMS_REDACT.ADD_POLICY procedure (*continued*)
 - using, [10-4](#)
- DBMS_REDACT.ALTER_POLICY procedure
 - about, [10-51](#)
 - example, [10-53](#)
 - parameters required for various actions, [10-52](#)
 - syntax, [10-51](#)
- DBMS_REDACT.DISABLE_POLICY
 - about, [10-57](#)
 - example, [10-57](#)
 - syntax, [10-57](#)
- DBMS_REDACT.DROP_POLICY
 - about, [10-59](#)
 - examples, [10-59](#)
 - syntax, [10-59](#)
- DBMS_REDACT.ENABLE_POLICY
 - about, [10-58](#)
 - example, [10-58](#)
 - syntax, [10-58](#)
- DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES
 - procedure
 - about, [10-27](#)
 - syntax, [10-27](#)
 - using, [10-28](#)
- deleting policies, [11-17](#)
- editing custom format, [11-7](#)
- editions, [12-3](#)
- Enterprise Manager Cloud Control, [11-5](#), [11-7](#), [11-10](#)
- Enterprise Manager Cloud Control workflow, [11-2](#)
- Enterprise Manager Cloud Control, about, [11-1](#)
- exporting data using Data Pump Export, [12-7](#)
- exporting objects using Data Pump, [12-6](#)
- full data redaction
 - about, [9-1](#)
 - creating policy for, [10-24](#)
 - examples, [10-25](#)
 - modifying default value, [10-27](#)
 - syntax, [10-25](#)
- functions used in expressions, [10-7](#)
- how differs from Oracle Database Real Application Security masking, [12-4](#)
- how differs from Oracle Virtual Private Database masking, [12-4](#)
- importing data using Data Pump Export, [12-7](#)
- inline views order of redaction, [12-2](#)
- JSON, [12-8](#)
- managing policies, [11-10](#)
- named policy expressions
 - about, [9-9](#)
- nested functions order of redaction, [12-2](#)
- no data redaction
 - about, [9-9](#), [10-49](#)
 - creating policies for, [10-49](#)
 - example, [10-49](#)
 - syntax, [10-49](#)
- Oracle Data Pump security policy, [12-5](#)

Oracle Data Redaction (*continued*)

- Oracle Enterprise Manager Data Masking and Subsetting Pack, [12-8](#)
 - partial data redaction
 - about, [9-2](#)
 - character types, policies for, [10-35](#)
 - data-time data types, [10-38](#)
 - example using character data type, [10-36](#)
 - example using data-time data type, [10-39](#)
 - example using fixed character format, [10-34](#)
 - example using number data type, [10-37](#)
 - formats, fixed character, [10-32](#)
 - number data types, [10-37](#)
 - syntax, [10-31](#)
 - policy expressions
 - about, [10-13](#)
 - creating, [10-15](#)
 - dropping, [10-16](#)
 - tutorial, [10-17](#)
 - updating, [10-16](#)
 - privileges for creating policies, [10-3](#)
 - random data redaction
 - about, [10-48](#)
 - creating policies for, [10-48](#)
 - example, [10-48](#)
 - randomized data redaction
 - about, [9-4](#)
 - regular expression data redaction
 - creating policies for, [10-39](#)
 - custom, creating policies for, [10-46](#)
 - example, [10-45](#)
 - example of custom, [10-47](#)
 - formats, [10-42](#)
 - formats, creating policies for, [10-42](#)
 - settings for, [10-46](#)
 - syntax, [10-40](#)
 - regular expression redaction
 - about, [9-3](#)
 - returning null values
 - about, [10-29](#)
 - example, [10-29](#)
 - syntax, [10-29](#)
 - SYS schema objects, [13-2](#)
 - SYSTEM schema objects, [13-2](#)
 - use cases, [8-2](#)
 - when to use, [8-2](#)
 - WHERE clause redaction, [12-2](#)
- ## Oracle Data Redaction formats
- creating in Cloud Control, [11-5](#)
 - deleting in Cloud Control, [11-9](#)
 - editing in Cloud Control, [11-7](#)
 - Enterprise Management Cloud Control, managing in, [11-5](#)
 - Enterprise Manager Cloud Control, sensitive column types, [11-2](#)
 - Enterprise Manager Cloud Control, viewing in, [11-8](#)

- Oracle Data Redaction partial redaction
 - creating policies for, [10-30](#)
- Oracle Data Redaction policies
 - about, [10-2](#)
 - altering, [10-51](#)
 - building reports, [10-60](#)
 - creating
 - examples, [10-26](#)
 - general syntax, [10-4](#)
 - procedure, [10-4](#)
 - creating in Cloud Control, [11-11](#)
 - deleting in Cloud Control, [11-17](#)
 - disabling, [10-57](#)
 - disabling in Cloud Control, [11-16](#)
 - dropping, [10-59](#)
 - editing in Cloud Control, [11-14](#)
 - enabling, [10-58](#)
 - Enterprise Manager Cloud Control, viewing in, [11-16](#)
 - exempting users from, [10-50](#)
 - expressions
 - by Application Express session state, [10-12](#)
 - by database role, [10-11](#)
 - by OLS label dominance, [10-12](#)
 - by user environment, [10-11](#)
 - filtering users
 - about, [10-7](#)
 - no filtering, [10-13](#)
 - finding information about, [10-61](#)
 - Oracle Enterprise Manager Cloud Control, [11-17](#)
 - redacting multiple columns in one policy, [10-56](#)
- Oracle Data Redaction policy expressions
 - Cloud Control, about, [11-18](#)
 - creating in Cloud Control, [11-18](#)
 - deleting in Cloud Control, [11-21](#)
 - editing in Cloud Control, [11-19](#)
 - viewing in Cloud Control, [11-20](#)
- Oracle Data Redaction, database links, [12-3](#)
- Oracle Data Redaction Enterprise Manager Cloud Control
 - deleting custom format, [11-9](#)
- Oracle Database Real Application Security
 - Data Redaction, [12-4](#)
- Oracle Database Vault
 - using with Data Redaction, [13-2](#)
- Oracle Enterprise Manager Cloud Control
 - creating custom formats, [11-5](#)
 - creating policy expressions, [11-18](#)
 - deleting policy expressions, [11-21](#)
 - disabling policies, [11-16](#)
 - editing policy expressions, [11-19](#)
 - Oracle Data Redaction, [11-5](#), [11-7](#), [11-16](#), [11-18-11-21](#)
 - Oracle Data Redaction formats, viewing in, [11-8](#)
 - Oracle Data Redaction, creating policies, [11-11](#)
 - Oracle Data Redaction, viewing details of a policy, [11-16](#)

- Oracle Enterprise Manager Cloud Control (*continued*)
 - policy expressions, about, [11-18](#)
 - viewing policy expressions, [11-20](#)
- Oracle Enterprise Manager Data Masking and Subsetting Pack
 - Oracle Data Redaction impact, [12-8](#)
- Oracle GoldenGate
 - storing secrets in Oracle keystores, [4-53](#)
- Oracle Key Vault
 - migration of keystores, [4-20](#)
- Oracle Label Security
 - functions using Data Redaction expressions, [10-10](#)
- Oracle Real Application Clusters
 - non-shared file systems to store TDE keystores, [6-5](#)
 - Transparent Data Encryption, [6-5](#)
- Oracle Recovery Manager
 - Transparent Data Encryption, [4-23](#)
- Oracle Securefiles
 - Transparent Data Encryption, [6-7](#)
- Oracle Virtual Private Database (VPD)
 - Data Redaction, [12-4](#)
- orapki utility
 - how compares with ADMINISTER KEY MANAGEMENT statement, [5-6](#)
- original import/export utilities, [3-21](#)

P

- passwords
 - access to for ADMINISTER KEY MANAGEMENT operations, [4-2](#)
 - preserving in PDB move operations, [6-11](#)
- PDBs
 - Data Redaction policies, [12-3](#)
 - finding TDE keystore status for all PDBs, [6-19](#)
 - master encryption keys
 - exporting, [6-12](#)
 - importing, [6-12](#)
 - Transparent Data Encryption, [6-8](#)
- performance
 - Transparent Data Encryption, [5-4](#)
- PKI encryption
 - backup and recovery operations, [5-10](#)
 - hardware keystores, [5-10](#)
 - master encryption key, [5-10](#)
 - tablespace encryption, [5-10](#)
- pluggable databases
 - See PDBs
- policy expressions, Oracle Data Redaction, [10-13](#)

R

- recycle bin
 - Data Redaction policies and, [13-3](#)
- reports
 - based Data Redaction policies, [10-60](#)

- returning null values
 - about, [9-4](#)
- rotating
 - master encryption key, [4-37](#)

S

- salt
 - removing, [3-28](#)
- salt (TDE)
 - adding, [3-28](#)
- secrets
 - storing Oracle Database secrets in keystore
 - about, [4-46](#)
 - storing in hardware keystore, [4-49](#)
 - storing in software keystore, [4-47](#)
- SecureFiles
 - Transparent Data Encryption, [6-7](#)
- software keystores
 - password-based using external keystore, [3-4](#)
- SUBSTR function
 - expressions, [10-9](#)
- synchronous change data capture, [3-21](#)
- SYS user
 - Data Redaction policies, [13-2](#)
- SYS_CONTEXT function
 - Data Redaction policies, [13-3](#)
 - SYS_SESSION_ROLES namespace used in Data Redaction, [10-11](#)
- SYS_SESSION_ROLES SYS_CONTEXT namespace
 - Data Redaction, [10-11](#)
- SYSTEM user
 - Data Redaction policies, [13-2](#)

T

- tablespace encryption
 - about, [2-4](#)
 - architecture, [2-4](#)
 - creating encrypted tablespaces, [3-35](#)
 - examples, [3-35](#)
 - incompatibilities, [7-1](#)
 - opening keystore, [3-33](#)
 - performance overhead, [5-4](#)
 - performance, optimum, [7-4](#)
 - procedure, [3-32](#)
 - restrictions, [3-31](#)
 - security considerations for plaintext fragments, [5-3](#)
 - setting tablespace key, [3-34](#)
 - storage overhead, [5-4](#)
- tablespace master encryption key
 - setting, [3-34](#)
- tablespaces
 - about encrypting, [3-29](#)
 - comparison between offline and online conversions, [3-29](#)

- tablespaces (*continued*)
 - rotating encryption algorithm, [4-39](#)
- tablespaces, offline decryption
 - procedure, [3-40](#)
- tablespaces, offline encryption
 - about, [3-38](#)
 - procedure, [3-39](#)
- tablespaces, online encryption
 - about, [3-43](#)
 - decrypting, [3-46](#)
 - finishing interrupted job, [3-47](#)
 - procedure, [3-41](#)
 - rekeying, [3-44](#)
- TDE
 - See* Transparent Data Encryption (TDE)
- TDE master encryption keys
 - activating
 - about, [4-29](#)
 - example, [4-30](#)
 - procedure, [4-29](#)
 - architecture, [2-3](#)
 - attributes, [4-31](#)
 - creating for later use
 - about, [4-26](#)
 - examples, [4-27](#)
 - procedure, [4-26](#)
 - custom attribute tags
 - about, [4-33](#)
 - creating, [4-33](#)
 - disabling not allowed, [4-35](#)
 - exporting, [4-40](#)
 - exporting in PDBs, [6-12](#)
 - finding currently used key, [4-32](#)
 - importing, [4-43](#)
 - importing in PDBs, [6-12](#)
 - keystore merge differing from import or export, [4-44](#)
 - resetting in keystore, [4-36](#)
 - rotating, [4-37](#)
 - setting in keystore, [4-35](#)
- Transparent Data Encryption (TDE)
 - about, [2-1](#)
 - benefits, [2-1](#)
 - CDBs
 - operations in root or PDBs, [6-11](#)
 - column encryption
 - about, [2-3](#), [3-20](#)
 - adding encrypting column to existing table, [3-27](#)
 - changing algorithm, [3-28](#)
 - changing encryption key, [3-28](#)
 - creating encrypted column in external table, [3-25](#)
 - creating index on encrypted column, [3-27](#)
 - creating tables with default encryption algorithm, [3-23](#)

- Transparent Data Encryption (TDE) *(continued)*
 - column encryption *(continued)*
 - creating tables with non-default encryption algorithm, [3-24](#)
 - data types supported, [3-20](#)
 - disabling encryption in existing column, [3-27](#)
 - encrypting columns in existing tables, [3-26](#)
 - encrypting existing column, [3-27](#)
 - encryption and integrity algorithms, [2-7](#)
 - restrictions, [3-21](#)
 - salt in encrypted columns, [3-28](#)
- columns with identity columns, [3-21](#)
- compatibility with application software, [7-1](#)
- compatibility with Oracle Database tools, [7-1](#)
- compression of encrypted data, [5-1](#)
- configuring hardware keystores
 - about, [3-14](#)
 - configuration step, [3-14](#)
 - opening, [3-15](#)
 - PKCS#11 library, [3-14](#)
 - reconfiguring software keystore, [3-19](#)
 - setting master encryption key, [3-17](#)
 - sqlnet.ora configuration, [3-14](#)
- configuring software keystores
 - about, [3-2](#)
 - creating auto-login keystore, [3-8](#)
 - creating password-based keystore, [3-7](#)
 - opening keystores, [3-9](#)
 - setting software master encryption key, [3-11](#)
 - sqlnet.ora file configuration, [3-2](#)
- data deduplication of encrypted data, [5-1](#)
- editions, [6-20](#)
- encryption and integrity algorithms, [2-7](#)
- finding information about, [3-51](#)
- frequently asked questions, [7-1](#)
- incompatibilities, [7-1](#)
- keystore management
 - ASM-based keystore, [4-23](#)
 - backing up password-based software keystores, [4-5](#)
 - changing hardware keystore password, [4-4](#)
 - changing password-based software keystore password, [4-3](#)
 - closing hardware keystores, [4-21](#)
 - closing software keystore, [4-21](#)
 - master encryption key attributes, [4-31](#)
 - merging keystores, about, [4-8](#)
 - merging keystores, auto-login into password-based, [4-10](#)
 - merging keystores, one into an existing, [4-10](#)
 - merging keystores, reversing merge operation, [4-11](#)
 - merging keystores, two into a third new keystore, [4-9](#)
 - migrating password key and hardware keystore, master encryption key creation, [4-16](#)

- Transparent Data Encryption (TDE) *(continued)*
 - keystore management *(continued)*
 - migrating password key and hardware keystore, reverse migration, [4-17](#)
 - migrating password key and hardware keystore, sqlnet.ora configuration, [4-15](#)
- keystores
 - about, [2-6](#)
 - benefits, [2-6](#)
 - types, [2-7](#)
- master encryption key
 - rotating, [4-37](#)
- master encryption key attributes
 - about, [4-33](#)
 - creating custom tags, [4-33](#)
- master encryption keys
 - exporting and importing, [4-40](#)
 - resetting in keystore, [4-36](#)
 - setting in keystore procedure, [4-35](#)
 - setting in keystore, about, [4-35](#)
- modifying applications for use with, [5-5](#)
- multidatabase environments, [6-21](#)
- Oracle Call Interface, [6-20](#)
- Oracle Data Guard, [6-4](#)
- Oracle Data Pump
 - export and import operations on dump sets, [6-3](#)
 - export and import operations on encrypted columns, [6-2](#)
- Oracle Data Pump export and import operations
 - about, [6-2](#)
- Oracle Real Application Clusters
 - about, [6-5](#)
 - non-shared file systems to store keystores, [6-5](#)
- Oracle Recovery Manager
 - keystores, [4-23](#)
- PDBs
 - about, [6-8](#)
 - finding keystore status for all PDBs, [6-19](#)
 - operations in root, [6-9](#)
- performance
 - database workloads, [7-4](#)
 - decrypting entire data set, [7-4](#)
 - optimum, [7-4](#)
 - worst case scenario, [7-4](#)
- performance overheads
 - about, [5-4](#)
 - typical, [7-4](#)
- PKI encryption, [5-9](#)
- privileges required, [2-2](#)
- SecureFiles, [6-7](#)
- security considerations
 - column encryption, [5-3](#)
 - general advice, [5-2](#)
 - plaintext fragments, [5-3](#)

- Transparent Data Encryption (TDE) *(continued)*
 - storage overhead, [5-4](#)
 - storing Oracle GoldenGate secrets, [4-53](#)
 - tablespace encryption
 - about, [2-4](#), [3-29](#)
 - creating, [3-34](#)
 - encryption and integrity algorithms, [2-7](#)
 - examples, [3-35](#)
 - opening keystore, [3-33](#)
 - restrictions, [3-31](#)
 - setting master encryption key, [3-34](#)
 - tablespace encryption, setting with COMPATIBLE parameter, [3-32](#)
 - views, [3-51](#)
- Transparent Data Encryption (TDE) keystores
 - deleting, [4-24](#)
 - moving software keystore to a new location, [4-11](#)
- Transparent Data Encryption (TDE) integrity
 - column encryption
 - creating tables without integrity checks (NOMAC), [3-24](#)
 - improving performance, [3-24](#)

- Transparent Data Encryption (TDE) integrity *(continued)*
 - column encryption *(continued)*
 - NOMAC parameter (TDE), [3-24](#)
 - transportable tablespaces, [3-21](#)
- tutorials
 - named Data Redaction policy expressions, [10-17](#)

U

- utilities, import/export, [3-21](#)

V

- V public function (APEX_UTIL.GET_SESSION_STATE function), Data Redaction policies, [10-12](#)
- V\$ENCRYPTION_WALLET view
 - keystore order after migration, [4-19](#)
- views
 - Data Redaction, [10-61](#)

X

- XML generation, [12-3](#)