

Database Security: A Technical Primer

Seventh Edition

Version 7.26.0 March 2026

Copyright © 2026 Oracle and/or its affiliates

ORACLE

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

Copyright © 2026 Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

Copyright © 2026, Oracle and/or its affiliates

ORACLE

Foreword

Having been in the security space for over 30 years, the front seat view has been exhilarating. Back then, mostly governments and financial institutions were interested in security while everybody else trusted the administrators, users, and computing environment to keep their data secure. It was only when browsers opened up new vistas for commerce over the net in the 90s that companies began to understand the vital need for security. This new perspective led to SSL, network firewalls, and strong cryptography.

Fast forward to the present and, just like before, we find ourselves living in a dramatically different world where every piece of data is online and available 24/7. Now AI is accelerating everything. Generative AI and agentic AI are lowering the cost of attack, from writing convincing phishing messages to searching for misconfigurations, writing exploits, and chaining steps together at machine speed. Some call it "vibe hacking," using AI to rapidly iterate, probe, and persuade until they find the gap you left uncovered.

Over the years, we have seen many different security technologies protecting various layers of the IT stack, from the applications down to the chipsets. Yet attackers have built sophisticated tools, now increasingly AI-enabled, to go after everything we have, whether on mobile devices, laptops, file servers, or databases. For most hackers, the target of choice is not a laptop or a spreadsheet, but a database with hundreds of millions or billions of records. They may try to break in through attacks on the network, applications, operating systems, and databases.

Why are organizations so vulnerable to attacks? Many might say they don't know where their sensitive data is, where they are vulnerable, and what the fixes might be. They might also fear that the fixes may break their applications. Too many organizations still stop at securing the perimeter, not recognizing how easily attackers can bypass it, get to the databases, and quietly walk away with their data. With AI helping adversaries move faster and blend in better, the window between compromise and impact can shrink dramatically. Detection still too often comes late.

To meet this challenge, Oracle has built multiple security technologies focusing on all pillars of security: evaluating the risk posture, preventing attacks from succeeding, and detecting and alerting on malicious behavior. We have enabled customers to deploy security at scale, not just for one database, but for all databases. And because customers cannot spend weeks or months deploying and tuning security controls, we are making it possible to operationalize security at speed. Industry analysts and security professionals recognize that the Oracle AI Database provides the industry's most comprehensive security.

This book, authored by my Database Security Product Management team, explains in simple terms the adversaries of today and how they exploit weaknesses, including how AI changes defender priorities. This book is not meant to be a prescriptive cookbook or a manual but rather a quick study into what every Database or Security Director/VP should know about the security of Oracle databases. You will learn about multiple assessment, preventive, and detective security controls for databases so that you can provide high-level guidance to your teams on how to shrink the attack surface and keep your databases secure at source, at scale, and at speed. I call that Security 360.

Breaches are coming faster than we can imagine, and we must be prepared. Your data is your asset, but unless you protect it well, it could fall into the wrong hands and become a liability. Let's start by securing the source, and do it at scale and at speed.

Vipin Samar
Senior Vice President, Oracle AI Database Security Development
March 2026

Copyright © 2026, Oracle and/or its affiliates

ORACLE

Authors' Notes

Securing data is tough. As security product managers, we hear the same challenge from you constantly: you need to manage real security risk without slowing down databases that run 24/7/365. Whether you're focused on a specific compliance requirement or simply expected to raise your organization's overall security posture, the path forward isn't always obvious. This eBook is your security roadmap for the Oracle AI Database, written for security officers, DBAs, application developers, application administrators, system administrators, and security team members.

Rather than listing product highlights, we use a threat-and-solution approach. Think of it as a guided tour through an adversary's mind: we show you how attacks unfold, then walk you through the controls that help you stop them. We also know that database security responsibilities are often spread across multiple roles, so we've written this to be useful whether you own the database, the application, or the broader security program. Adversaries don't follow a fixed playbook, and that's exactly why we recommend a defense-in-depth and a Zero Trust mindset.

This edition is significantly updated and expanded, growing from about 60 pages since initial publication to well over 200 pages. But bigger doesn't mean bloated. We've simultaneously consolidated and rewritten the content throughout to be sharper and more concise, and every chapter reflects the latest Oracle capabilities, current as of March 2026.

One of the biggest changes in this edition is the integration of **AI and Agentic AI** throughout the book. This isn't a single add-on section tucked away at the back. We've woven AI coverage into every relevant chapter and topic area where it genuinely changes the conversation, whether that's sensitive data discovery, authentication, authorizations, access controls, or auditing. And because AI's relationship with database security deserves its own focused treatment, we've added a dedicated chapter on AI and its implications for database security. If you're trying to understand what AI risks are now exposed for your critical databases, AI functioning not just as a tool but as a new attack vector, that chapter is required reading.

We've also rethought how the book closes. The final chapter now includes a section that walks through each chapter individually, giving you a **fast and convenient refresher of the key ideas**. This isn't a recycled collection of the summaries you'll find at the end of each chapter. It's a fresh synthesis, paired with a "Key Takeaways" summary, designed so you can revisit the material quickly without re-reading everything from scratch. Think of it as your personal cheat sheet, written by the same team that built the full guide.

After you finish this eBook, you'll have a clearer view of how adversaries exploit vulnerabilities, which database security controls help you protect your databases, and which risks those controls address. This eBook doesn't replace product documentation, and it doesn't provide regulatory advice. What it does provide is perspective, a map of the terrain so you can make smarter decisions faster.

We hope this edition broadens your thinking on database security, gives you actionable guidance, and helps you protect the data that matters most.

Angeline Dhanarani
Bettina Schaeumer
David Knox
Sumanth Kakaraparthi
Ethan Shmargad
Kajal Singh
Manoj Shringarpure
Nazia Zaidi
Pedro Lopes
Peter Wahl
Rich Evans
Roger Wigenstam
Russ Lowenthal
Vasishtha Kumar Nalla

Copyright © 2026, Oracle and/or its affiliates

Contents

1 You Should Read This Book **10**

Yet another book on AI, ransomware and database security 11

The security imperative 11

A framework for defense 12

Regulatory alignment 13

Quantum shadows, rogue AI, and regulatory reckoning 13

2 Protecting Data **15**

Why database security matters 16

Database security as a growing priority 16

The evolving threat landscape 16

Threat actors and how breaches happen 18

A goal-based approach to data security 21

3 Database Security Posture Management **22**

The security assessment imperative 23

Understanding the assessment tool landscape 23

The assessment mindset 24

Critical questions attackers seek to answer 24

Assessing configuration security with DBSAT 26

DBSAT report formats and usage 27

Assessing database fleets with Oracle Data Safe 30

Understanding user risk with Data Safe 35

Data Safe deployment and architecture 37

Assessment with Oracle Audit Vault and Database Firewall 38

Configuration assessment with EM Database Lifecycle Management 42

Selecting the optimal Database Security assessment tool 45

Minimizing the blast radius with privilege analysis 47

Managing security patches 49

Patching tools and methodologies 51

Hands-On learning with Oracle LiveLabs 52

Summary 53

4 Sensitive Data Discovery **54**

Introduction: The foundation of data protection 55

Why sensitive data matters 55

The AI imperative: New urgency for discovery 56

Defining sensitive data 57

Discovery methodologies: How to find sensitive data 57

Choosing the right solution 58

Copyright © 2026, Oracle and/or its affiliates

Discovering sensitive data using DBSAT 59
Discovering sensitive data using Data Safe 62
Discovering sensitive data using Audit Vault and Database Firewall 65
Discovering sensitive data using Enterprise Manager 67
Hands-On learning with Oracle LiveLabs 71
Summary 72
5 Authenticating Database Users 73
The importance of strong authentication 74
Unraveling identification, authentication and authorization 74
Database authentication methods 75
Simplify account administration 79
Protect user accounts 79
Hands-On learning with Oracle LiveLabs 81
Summary 82
6 Controlling Database Access 83
Database privileges 84
Database roles 86
Types of database users 86
Who can do what 87
Centrally managed users 89
Protecting database accounts 90
Hands-On learning with Oracle LiveLabs 92
Summary 93
7 Enforcing Separation of Duties 94
Why “who can do what” still breaks security 95
What separation of duty means & why it matters 95
Separating the duties 96
Top tips for a secure database environment 97
Mandatory access controls with Oracle Database Vault 98
Controlling administrator access 98
Enforcing a trusted path 99
Separation of duties within Database Vault 101
Separation of duties with containerized databases 101
Control SQL database commands with Oracle Database Vault 103
Break-glass scenarios 103
Operationalizing Oracle Database Vault 104
Hands-On learning with Oracle LiveLabs 105
Summary 106

Copyright © 2026, Oracle and/or its affiliates

8 Minimizing Risk from SQL Injection **107**
The persistent threat of SQL injection 108
A defense in-depth strategy against SQL injection 109
Database firewall solutions from Oracle 110
Network-based Database Firewall in Oracle AVDF 110
Database-resident SQL Firewall in Oracle AI Database 26ai 112
Deciding between Database Firewall or SQL Firewall 113
Hands-On learning with Oracle LiveLabs 115
Summary 116
9 Data-Driven Authorization **117**
The need for data-driven authorization 118
Oracle Database data-driven authorization technologies 119
Virtual Private Database 119
Oracle Label Security 120
Real Application Security 124
Choosing the right technology 126
Hands-On learning with Oracle LiveLabs 126
Summary 127
10 Masking Sensitive Data **128**
Why masking sensitive data matters 129
Understanding data masking 129
Static data masking 130
Data Redaction 131
Data Subsetting 133
Masking formats 133
Masking techniques 135
Implementing static data masking 137
Masking sensitive data using Oracle Data Masking and Subsetting 138
Choosing between Oracle Data Safe and Oracle Data Masking & Subsetting 140
Data Redaction in Oracle Database 140
Implementing Data Subsetting 145
Hands-On learning with Oracle LiveLabs 146
Summary 147
11 Data Encryption and Key Management **148**
Why encryption matters for database security 149
Encrypting data in motion 149
Transport Layer Security (TLS) 149
Encrypting data at rest 150
Transparent Data Encryption 152

Copyright © 2026, Oracle and/or its affiliates

Tablespace encryption and column encryption 153
Migrating existing data to encrypted tablespaces 154
Key management 154
Hands-On learning with Oracle LiveLabs 160
Summary 161
12 Using and Managing Audit Data 162
Why auditing matters 163
Oracle Database unified auditing 163
Audit data best practices 165
What to audit 166
Auditing in the age of AI 170
Managing and using audit data with Oracle Data Safe 170
Auditing with Audit Vault and Database Firewall 176
Hands-On learning with Oracle LiveLabs 180
Summary 181
13 Securing Agentic AI 182
Security risks in the age of AI 183
Modern risks: vibe coding and shadow AI 183
Vibe coding 184
Building a future-proof AI security strategy 186
Summary 189
14 Database Security in a Multicloud Environment 190
Understanding multicloud strategies 191
The multicloud security challenge 191
Establishing identity as the foundation 192
Managing Oracle AI Database 26ai in multicloud 192
Leveraging native Oracle services across clouds 193
Enabling secure multicloud service access 193
Implementing multicloud Identity Management 193
Hands-On learning with Oracle LiveLabs 195
Summary 196
15 Ransomware and Zero Trust 197
Ransomware: When data gets taken hostage 198
Understanding the ransomware imperative 198
Defining ransomware and its evolution 199
Phase 1: Immediate high impact defense 200
Neutralizing data theft and extortion through encryption 200
Eliminating data loss through immutable recovery 202
Phase 1 Implementation roadmap 203

Copyright © 2026, Oracle and/or its affiliates

Phase 2: The Zero Trust journey 204
Zero Trust: The foundation for long term resilience 204
Typical Zero Trust implementations 204
Four core Zero Trust projects for databases 205
Zeroing in on Zero Trust 209
Hands-On learning with Oracle LiveLabs 209
Summary 210
16 Securing the Autonomous Database 211
Database Security in Autonomous AI Database 212
The value of Autonomous AI Database 212
Security benefits of Autonomous AI Databases 212
Security capabilities of Autonomous AI Databases 213
Shared responsibility 214
Hands-On learning with Oracle LiveLabs 215
Summary 215
17 Putting It All Together 216
Oracle Maximum Security Architecture 217
Chapter-by-Chapter journey: Key takeaways 219
Connecting the Controls: How capabilities work together 229
Closing thoughts: Security as an enabler 231
Appendix: Tools – Features, Options, Products, and Packs 232
Database Security Assessment Tool (DBSAT) 232
Oracle Data Safe 232
Enterprise Manager Database Lifecycle Management 232
Privilege Analysis 232
Native Network Encryption 232
Transport Layer Security 233
Centrally Managed Users 233
Enterprise User Security 233
Traditional Auditing 233
Fine-Grained Auditing 233
Unified Auditing 233
Enterprise Manager Data Discovery 233
Oracle Advanced Security 234
Oracle Key Vault 234
SQL Firewall 234
Oracle AI Database Vault 234
Oracle Audit Vault and Database Firewall 234
Oracle Data Masking and Subsetting 234

Copyright © 2026, Oracle and/or its affiliates

Chapter 1

You Should Read This Book

Yet another book on AI, ransomware and database security

You are a busy professional responsible for securing critical data. You do not need another theoretical textbook; you need a field manual.

This book is designed to save time and meaningfully strengthen an organization's security posture. Because most organizations keep their most critical assets in Oracle Database deployments, the focus stays on practical steps that can be applied right away. From strategic planning through configuration, this guide supports steady progress with confidence, helping teams steer clear of expensive security missteps and regulatory dead ends.

Although this book highlights specific features of the Oracle AI Database, the foundations of database security remain universal. The principles described here apply whether the environment runs on-premises, in Oracle Cloud Infrastructure, a hybrid or Multicloud architecture.

The security imperative

The threat environment is not merely expanding; it is accelerating. In 2025, the World Economic Forum estimated the global cybercrime economy at a staggering US$10.5 trillion. Put in practical terms, if cybercrime were a country, it would rank as the world's third largest economy, trailing only the United States and China.

At this scale, private sector measures often struggle to keep pace, and governments step in. The outcome is a global regulatory tsunami. The United States alone enforces more than 100 privacy and data security laws. In the European Union, the General Data Protection Regulation (GDPR) is now reinforced by the Digital Operational Resilience Act (DORA) and the updated NIS2 Directive. India has introduced the Digital Personal Data Protection Act (DPDPA), and more than 130 other nations have enacted similar privacy legislation.

The stakes are financial, not only reputational. GDPR violations can trigger fines of up to 4% of global annual turnover or €20 million. Since 2021, fines have exceeded €4.5 billion. For any business operating internationally, this can quickly become a minefield of concurrent and overlapping regulatory obligations.

It's not just regulations

Data risk is a multi-front war. Organizations must contend with stricter laws, rising privacy expectations, and a relentless barrage of attacks:

- Ransomware: This remains a primary concern, accounting for nearly 17% of all cyberattacks and costing tens of billions annually.
- Geopolitical Instability: Nation-state actors actively target private data for economic advantage or to weaken critical infrastructure.
- AI: As Agentic AI systems expand across enterprises, they bring major productivity gains while extending the security perimeter in ways many organizations are still working to map, measure, and manage.

These trends show no signs of slowing. The time to protect data is now, before the breach, not after.

The hidden opportunity: agility

Regulatory pressure and data theft risk often launch database security programs, but a third driver is equally important: **agility**. Customers increasingly factor security posture into purchasing decisions. A consistent, successful security program reduces sales friction, speeds due diligence, and strengthens win rates.

Copyright © 2026, Oracle and/or its affiliates

The same applies outside commercial settings. For organizations dependent on grants, public funding, or private investment, a strong security track record builds trust and shortens reviews. Stakeholders are more likely to support organizations that proactively avoid breaches and negative headlines.

Robust security is not only a shield; it is a strategic enabler that helps organizations move faster and seize new opportunities.

A framework for defense

Thousands of regulatory requirements can feel like a wall of obligations. But trying to enumerate every rule rarely helps, and quickly becomes unmanageable. `Disclaimer: This is not legal advice; always consult your legal counsel for compliance guidance.`

A more practical path is to look at the regulatory landscape as a framework of controls. Most regulations are trying to solve the same underlying problems. No matter the acronym, whether GDPR, PCI DSS, HIPAA, they typically boil down to three fundamental actions that must take place on your data:

1. **Encrypt** your data.
2. **Control access** to your data.
3. **Audit access** to your data.

Start with these outcomes. Once these pillars are in place, it becomes much easier to layer on the reporting and documentation that each regulation asks for.

1. Encrypt your data

Encryption is often your last line of defense for data at rest and backups. It transforms your data using cryptographic algorithms, rendering it unreadable without the correct decryption key. Even if an attacker bypasses your firewalls and steals your storage media, encrypted data remains useless to them. Database files and database backups are prime targets.

Effective encryption is intrinsically linked to key management. Keys need to be generated, stored, rotated, and retired securely. Without strong key management, encryption is like installing a sturdy lock and then taping the combination to the door.

2. Control access to your data

Access control acts as the gatekeeper. It decides who can get into the database and what they are able to do once they are there. This broad area usually breaks into three layers:

- **Authentication**: Verifying the identity of the user or service (Checking the ID card).
- **Authorization**: Defining which objects a verified user is permitted to view or touch (Assigning the keycard).
- **Fine-grained Access**: Pinpointing exactly which rows or columns in a table are visible (Redacting specific lines in a file).

When least-privilege is enforced, users receive only what they need and nothing more, which dramatically limits the blast radius if an account is compromised.

3. Audit access to your data

If access control is the gatekeeper, auditing is the security camera. It records what happens across systems so actions can be traced, issues can be investigated, and compliance expectations can be met.

A solid audit trail answers the classic forensic questions:

- Who performed the action? (Database account and OS user)

Copyright © 2026, Oracle and/or its affiliates

- What did they do? (Insert, update, delete, or execute)
- When did it happen? (Timestamp)
- Where did they come from? (Source IP or hostname)
- How did they connect? (Client program)

Auditing will not stop the incident, but it gives the evidence needed to reconstruct what happened and show regulators that the environment is being actively watched.

Regulatory alignment

Beyond these three foundational controls, encryption, access control, and auditing, some regulations may impose additional requirements. However, most expect these basics to be implemented first. Table 1-1 provides examples of regulations that mandate these and other types of controls.

Regulation	Country	Required control	Remarks	Applies to
PCI DSS	Global	Redaction - A specialized form of access control.	Only part of a credit card number should be visible.	Anyone processing credit cards
DORA	European Union	Zero Data Loss Recovery	Implicit goal is recovery from ransomware attacks.	Financial institutions in the EU
Companies Act (Rule 11)	India	Before/After value tracking - A specialized form of auditing	Track changes to entries in books of account.	Companies using software for accounting.
Sarbanes-Oxley (SOX)	United States	Separation of Duties - A specialized form of access control.	Ensure financial records are not altered by unauthorized personnel.	Publicly traded companies.

Table 1-1: Summary of regulatory-influenced controls

Quantum shadows, rogue AI, and regulatory reckoning

The security landscape of 2026 has shifted from passive defense to active warfare. Ransomware has moved beyond simple encryption. Attackers now use “harvest now, decrypt later” strategies, collecting data today and waiting for quantum computing to break traditional protections.

At the same time, the rise of Agentic AI introduces a paradox. These autonomous agents can unlock remarkable productivity, yet they can also expand the attack surface exponentially. An agent built to execute complex business logic can be manipulated through prompt injection and turned into a high-speed insider threat, exfiltrating sensitive data before a human operator can react.

The regulatory environment is tightening its grip to match these threats. 2026 marks the entry of strict new data privacy laws, joining a global mesh of “Universal Opt-Out” mandates that force businesses to honor consumer privacy preferences automatically. Compliance is no longer a quarterly checklist; it demands real-time action. The cost of failure has escalated from uncomfortable fines to existential threats, with penalties under frameworks like GDPR and the new US state statutes capable of erasing annual profit margins overnight. In this fragmented regulatory map, your database is either your strongest asset or your greatest liability.

Defending against machine speed attacks requires more than vigilance. It calls for embedded intelligence such as the Oracle AI Database 26ai, with quantum-resistant cryptography and ML-driven SQL Firewalls intended to help distinguish between a legitimate agent and a compromised one.

That is why this book matters. It avoids abstract theory, and focuses on practical outcomes. Later in the book, we break down the complexity of Oracle’s security stack, from Transparent Data Encryption to Database Vault to Unified Audit and Data Safe, and organizes it into a cohesive, resilient approach. Security is no longer only about locking doors. It is also about enabling what comes next. Turn the page, and let’s get to work!

Copyright © 2026, Oracle and/or its affiliates

Copyright © 2026, Oracle and/or its affiliates

Chapter 2

Protecting Data

Why database security matters

Databases store an organization's most sensitive assets: trade secrets, personal records, financial transactions, and customer intelligence. When breaches or ransomware hit, the fallout is severe. For security administrators, database owners, and data custodians, database security must stay a daily priority.

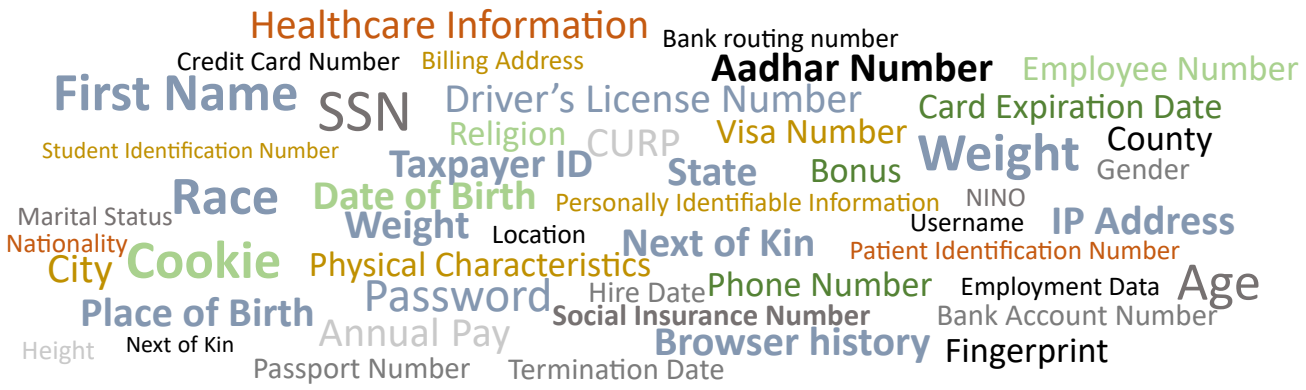


Figure 2-1: Many types of personal data demand comprehensive database security

This book guides readers through a defense-in-depth security strategy for Oracle Database, demonstrating how various controls work together and what protection they deliver.

Database security as a growing priority

The database administration landscape has transformed dramatically in recent years. At a recent Oracle AI World event, database administrators from a top financial institution revealed that security now consumes over half their workday. A decade ago, security barely registered as a concern. Today, regulations such as GDPR, CCPA, and India's DPDPA, combined with the ransomware epidemic, emerging threats from Agentic AI, and daily breach headlines, have elevated security to a top organizational priority.

Many database professionals now ask, "What matters most for keeping our data safe?" This question reflects widespread recognition that database security cannot be treated as an afterthought.

Perfect security has no magic formula. Organizations face tireless, creative adversaries who continuously seek opportunities to compromise data. The objective is not simply checking compliance boxes but rather identifying vulnerabilities and staying ahead of evolving threats. Understanding exposure points is essential because attackers consistently target the weakest links in an organization's defenses.

Threats take many forms. Less sophisticated attackers exploit outdated systems with readily available tools. Advanced persistent threat actors move laterally through networks, probing for vulnerabilities and extracting data silently over extended periods. Sometimes the threat comes from within, such as a contractor or employee turned adversary. These breaches can remain undetected for months or years, making constant vigilance essential for effective defense.

The evolving threat landscape

Three interconnected forces drive the increasing urgency around database security: regulatory requirements, sophisticated threat actors, and emerging technology risks. Understanding these factors helps security teams prioritize their efforts and allocate resources effectively.

Copyright © 2026, Oracle and/or its affiliates

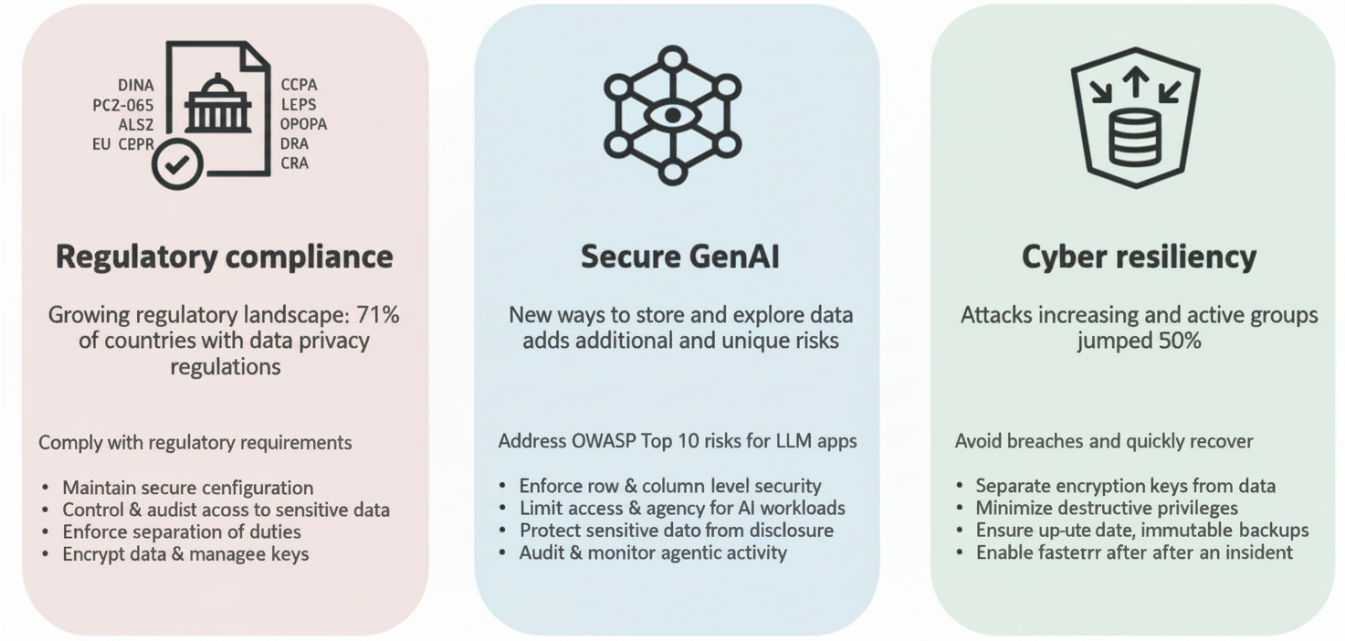


Figure 2-2: Data Security drivers for an evolving compliance and threat landscape

Regulatory compliance continues to expand globally. Data protection regulations now span multiple jurisdictions, each with distinct requirements for consent, breach notification, data residency, and individual rights. Organizations operating internationally must navigate complex overlapping mandates while maintaining consistent security practices across their database environments.

Threat actor sophistication has increased substantially. Attackers range from opportunistic script kiddies using automated tools to nation-state actors conducting targeted campaigns. Ransomware groups now combine data encryption with exfiltration and extortion, threatening to publish stolen data unless organizations pay. These groups often target backup systems and disaster recovery infrastructure to maximize leverage. In response, organizations are increasingly prioritizing cyber resilience, building the capacity not only to prevent attacks but to absorb, adapt to, and rapidly recover from them when prevention fails.

Emerging AI threats represent a new challenge for database security. Agentic AI systems that can autonomously interact with databases introduce novel attack vectors. Adversaries may leverage large language models to generate sophisticated SQL injection payloads, automate reconnaissance of database structures, or identify sensitive data patterns at scale. Conversely, AI-powered defensive tools offer opportunities to detect anomalous behavior and respond to threats more rapidly than traditional rule-based systems.

Database security in the age of AI

Enterprise technology is changing quickly, and AI and Agentic AI represent a major turning point from generating content to taking action autonomously. In this new era of database security, it is important to keep in mind that AI agents can plan work and interact directly with core systems, widening the security perimeter in ways older models may not fully reflect. The practical takeaway is clear: the goal is no longer just protecting applications, but governing AI mediated actions with the same discipline used for human users.

Figure 2-3 makes the dependency chain for effective AI clear. AI runs on Data (capital “D”) as its fuel, and the more fuel, or Data, it can draw on, the greater the potential benefit. That dynamic puts a spotlight on the security and reliability of the data feeding AI. Reliability is not just about availability; it is also about whether the data can be trusted and whether it accurately reflects what it is meant to represent. Security matters just as much, particularly because

Copyright © 2026, Oracle and/or its affiliates

sensitive information can be disclosed inadvertently. Even more troubling is the risk that data could be improperly changed or deleted.

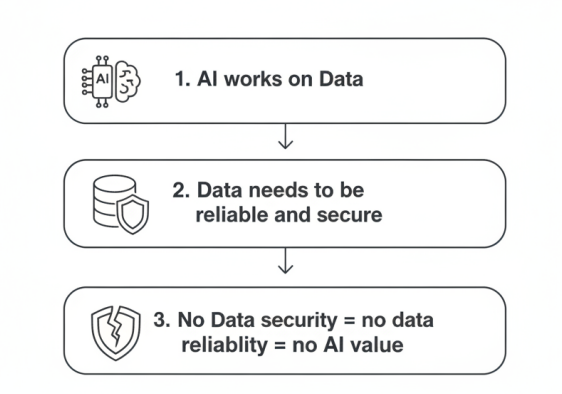


Figure 2-3: Effective AI depends on Database Security

AI can amplify database risk simply by making application development faster. Vibe coding is a prime example: developers lean on the speed of AI-generated code without taking the time to fully understand what it is doing. The productivity gains are real, but the tradeoff is often black box logic that can quietly introduce SQL injection vulnerabilities, expose sensitive API keys, or create unintended data exfiltration pathways. Seeing these risks clearly enables organizations to move from default trust to deliberate verification, validating every AI-generated query against established security standards before relying on it.

Then there is Shadow AI, which adds a different kind of governance challenge. Well-meaning employees may connect unauthorized AI tools directly to corporate databases to move faster. Unlike traditional shadow IT, autonomous agents can make their own calls about sensitive information, which can lead to proprietary data leaking across organizational boundaries without an audit trail. This section of the paper underscores why bans often fail in practice and why securing the data at its source matters most, ensuring protection remains unmovable no matter which tool is used.

The most durable path forward is to treat database native enforcement as the primary control plane. Layered defenses such as Transparent Data Encryption (TDE), fine-grained access controls like Virtual Private Database (VPD), and SQL Firewalls that block unauthorized statements help contain the blast radius of runaway agents or prompt injection attacks. With these controls in place, organizations can pursue AI-driven advantages with confidence, keeping innovation resilient, compliant, and under control.

Now that we've named the AI-driven risk amplifiers, here's how database-native controls contain them.

Threat actors and how breaches happen

In the past, perimeter defenses provided a sense of security. Today's threat landscape evolves too rapidly for any single barrier to suffice. Set-and-forget security is no longer viable. Organizations need layered controls across the

Copyright © 2026, Oracle and/or its affiliates

application and data stack. If one layer is breached, others must remain intact. This approach, known as defense-in-depth, remains one of the strongest security strategies available.

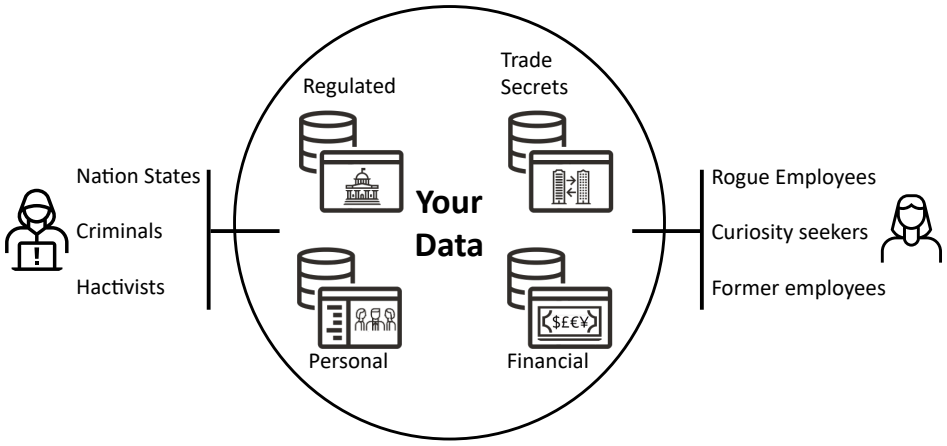


Figure 2-4: Threat actors have many motivations compounding the challenge to securing data

Attackers compromise database security through multiple pathways. Security professionals are likely familiar with the OWASP Top Ten, the definitive catalog of web application vulnerabilities that evolves continuously as threats change. Similarly, the database security community tracks a "Dirty Dozen" list of the most common methods attackers use to compromise database security.

Common attack methods and vulnerabilities

Understanding how attackers operate enables security teams to focus on closing real gaps and improving organizational resilience. The following tactics represent some of the most prevalent methods from the attacker's playbook:

- Exploiting unpatched systems or misconfigured databases to bypass access controls
- Escalating runtime privileges by exploiting vulnerable applications
- Searching for sensitive data in unprotected databases, applications, and systems
- Stealing the credentials of privileged administrators or application users
- Accessing accounts through password guessing or exploiting careless credential management
- Exploiting application weaknesses with techniques such as SQL injection
- Exploiting unprotected systems as a bridge to launch attacks against more sensitive targets
- Creating rogue user accounts on systems as a base for reconnaissance and privilege escalation
- Targeting copies of live production data used in development and test environments where protections are typically weaker
- Accessing unencrypted database system files on disk or in backup files
- Locking data with encryption and demanding ransom payments for decryption keys

These attack methods correspond to a comprehensive list of database security risks known as the "Dirty Dozen." Table 2-2 below provides the complete inventory.

Copyright © 2026, Oracle and/or its affiliates

THE DIRTY DOZEN

1. Insecure configuration and configuration drift
2. Unpatched and out-of-date systems
3. Lack of a consistently enforced security policy
4. Lack of visibility into sensitive data placement and quantity
5. Overprivileged database users and administrators
6. Weak authentication and shared accounts
7. SQL Injection vulnerabilities and insecure application design
8. Trusting vulnerable networks
9. Insufficient or inefficient monitoring and auditing
10. Sensitive data proliferation to non-production databases
11. Unprotected servers and database backups
12. Insecure encryption keys and secrets

Table 2-2: The Oracle Database Security Dirty Dozen represents the most common database security risks

Attackers exploit risks at every touchpoint with the database. They most frequently target accounts with significant database access, whether end users, application service accounts, or administrators. Stolen login credentials remain a common entry point. SQL injection remains a popular attack vector because penetrating an application is often easier than gaining direct access to the database or its server. When attackers gain access to network segments carrying database traffic, they may intercept unencrypted communications to extract sensitive data.

Once an attacker establishes a session with the database or accesses the server directly, they may exploit lingering unpatched vulnerabilities or attack the database from within the server environment. A now-prevalent technique is data scraping: extracting and exporting raw database files without establishing a formal session. This method is common in most ransomware data breaches.

Database copies created for testing or development are equally at risk. In targeted campaigns where attackers pursue entry into a specific known database, adversaries frequently focus on these non-production copies because they are usually less secure and less monitored than production counterparts. Additionally, these environments are often accessed by a wider user community, including developers, testers, and third-party contractors, further expanding the attack surface.

Copyright © 2026, Oracle and/or its affiliates

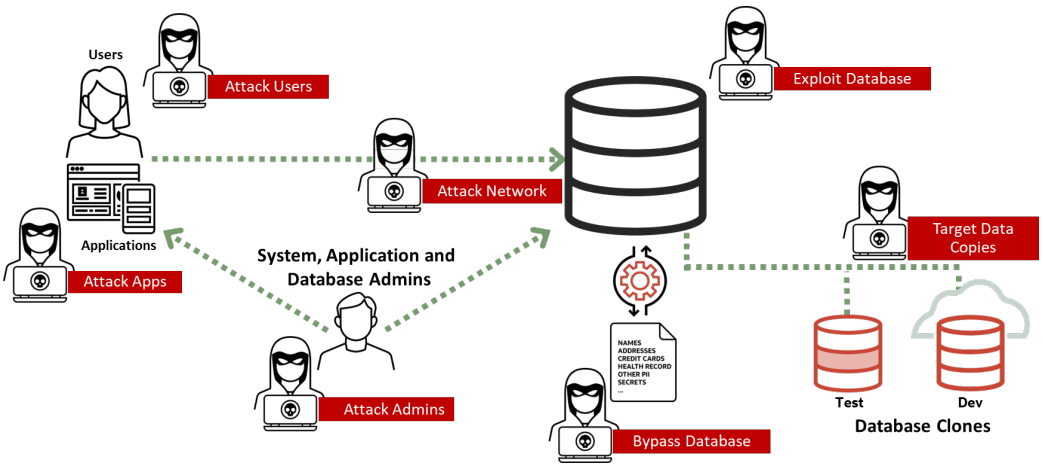


Figure 2-5: Hackers can attack the database from many directions

A goal-based approach to data security

A well-structured data security posture helps mitigate the risks outlined in the Dirty Dozen. Successful approaches incorporate multiple layers of security controls to provide defense-in-depth protection against threats. Database security objectives can be organized into four broad categories:

1. Assess security posture
2. Control access to data
3. Monitor user activity
4. Protect data against theft

To mitigate database security risks effectively, organizations should set clear security goals and align them to actionable tasks. A goal describes the intended security outcome, such as ensuring secure user authentication, while a task is the concrete step taken to achieve that goal, such as implementing multifactor authentication. Table 2-3 provides examples of high impact security tasks, prioritized by how effectively they reduce risk.

Category	Goal	Task	Threat addressed
Assess security posture	Ensure systems do not deviate from their approved security baseline	Monitor databases for configuration changes that introduce risk	Insecure configuration
Monitor user activities	Detect and deter inappropriate activity by privileged users	Monitor and audit privileged user activity	Insufficient or inefficient monitoring and auditing
Protect data against theft	Remove risk from test and development environments	Scramble sensitive data in test and development environments	Sensitive data proliferation to non-production databases
Control access to data	Strengthen user authentication to prevent compromise of accounts	Implement multifactor authentication for the database	Weak authentication and shared accounts

Table 2-3: Examples of high-level security goals and tasks

Copyright © 2026, Oracle and/or its affiliates

Chapter 3

Database Security Posture Management

The security assessment imperative

Imagine your database security as a fortress with hundreds of gates, locks, and defensive mechanisms. A single unlocked gate, one misconfigured parameter among hundreds, can invite intruders. Unlike medieval fortresses where defenders could visually inspect every entry point, modern databases hide their vulnerabilities behind layers of configuration complexity. Security assessment becomes your systematic reconnaissance, checking every defense before attackers exploit what you have overlooked.

Database security demands near perfection. Oracle Database exposes hundreds of configuration parameters, each influencing the security posture. A single misconfiguration, a weak password policy, an overprivileged account, disabled auditing, creates an exploitable vulnerability. Attackers automate their reconnaissance, deploying tools that probe thousands of systems to find that one error. They are patient, methodical, and relentless.

The consequences of configuration drift reach far beyond security breaches. Data loss, prolonged downtime, brand damage, regulatory fines, and executive accountability can all trace back to preventable security gaps. When organizations store personal data such as names, addresses, birth dates, and account details, regulators expect documented controls and ongoing verification. Laws and frameworks including GDPR, PCI DSS, Sarbanes-Oxley, and breach notification statutes require regular security assessments for systems that process personal data.

Organizations need systematic assessment practices that accomplish three objectives:

1. **Identify misconfigurations before attackers exploit them**: proactive detection prevents breaches rather than responding after compromise
2. **Demonstrate due diligence to regulators and auditors**: documented assessments prove ongoing security investment and compliance
3. **Provide defensible evidence of security improvement over time**: trend analysis shows measurable risk reduction

A standards-based assessment approach integrates three essential elements: relevant regulations and security frameworks, Oracle's database hardening best practices, and the organization's specific policies and risk priorities. This methodology provides a repeatable, defensible process to identify misconfigurations, prioritize critical issues, and demonstrate measurable security improvement.

This chapter examines how Oracle's database security solutions help organizations rapidly assess security posture, categorize findings by risk and compliance relevance, and define clear remediation paths. Readers will learn which configurations to address first, why they are critical, and how to remediate them: enabling faster hardening, streamlined audits, and sustained security without impeding business operations.

Understanding the assessment tool landscape

Organizations assessing Oracle Database security have access to multiple tools, each with distinct capabilities, deployment models, and use cases. Some tools ship with Oracle Database, while others are available as add-on products or cloud services. Most Oracle Database customers already have access to one or more assessment options through existing licenses or enterprise agreements.

The key is matching tool capabilities to organizational goals, constraints, and scale. Four primary tools address database security assessment:

1. **Database Security Assessment Tool (DBSAT)**: A command-line utility provided at no cost to all Oracle Database customers with active support agreements
2. **Oracle Data Safe**: An Oracle Cloud Infrastructure service included with all Oracle Cloud database services, also available for on-premises databases, third-party cloud deployments, and Oracle Database@Multicloud

Copyright © 2026, Oracle and/or its affiliates

3. **Oracle Audit Vault and Database Firewall (AVDF)**: A database security posture management platform available for on-premises or Oracle Cloud deployment
4. **Oracle Database Lifecycle Management Pack (DBLM)**: A management pack for Oracle Enterprise Manager

This chapter outlines each tool's function, strengths, and trade-offs to help organizations select the optimal fit for their environment, budget, timeline, and operational model.

The assessment mindset

Before looking at assessment tools, it is worth stepping into an attacker's mindset, because that perspective makes it easier to prioritize the right security controls. Attackers treat databases like seasoned burglars surveying a neighborhood. They do not kick in doors randomly. They watch patrol routines, try the doorknobs, notice which houses have broken security cameras, and plan the fastest way out. Today, many of them go further, using AI powered scanners to sweep thousands of databases at the same time, all in pursuit of that one unlocked door.

The attacker's reconnaissance toolkit

Attackers utilize automated tools and techniques to systematically identify vulnerabilities:

- **Database discovery**: Network scanning identifies database instances, versions, and network exposure
- **Port and service scanning**: Open listeners and exposed management interfaces reveal attack surfaces
- **Vulnerability identification**: Automated tools match database configurations against known CVE vulnerabilities
- **Application and SQL injection testing**: Web application proxies probe for input validation weaknesses
- **Brute-force authentication attacks**: Password spraying and dictionary attacks test for weak credentials

Following this reconnaissance phase, attackers focus on the weakest links and formulate efficient paths to data. They assess an organization's security posture to find the path of least resistance. Understanding these tactics helps security teams and DBAs prioritize hardening steps that make answering critical attacker questions substantially harder.

Critical questions attackers seek to answer

When probing database defenses, attackers focus on specific questions that reveal exploitable weaknesses:

- **What database version is running?** : Outdated versions with known vulnerabilities become primary targets
- **Are default or weak credentials in use?** : Unchanged default passwords and weak password policies provide immediate access
- **Which privileged users exist?** : High-privilege accounts enable privilege escalation and lateral movement
- **Is auditing enabled?** : Absence of logging means undetected intrusions and evidence-free attacks
- **Is data encrypted?** : Unencrypted data at rest and in transit enables raw data theft from storage or backups
- **Is there a less monitored copy of this database?** : Development, test, and backup environments often have weaker controls than production

Oracle Database's flexibility creates assessment complexity. Hardening requires understanding multiple interdependent factors:

- **User accounts, roles, and privileges**: Who accesses the database and what permissions they possess
- **Data classification and sensitivity**: What data resides in each database and its regulatory implications

Copyright © 2026, Oracle and/or its affiliates

- **Security parameters and features**: Which controls are configured and how they interact
- **Attack vectors relevant to the deployment**: Which threats matter most for the specific environment
- **Available prevention, detection, and response controls**: What security capabilities the organization can deploy

Implementing attacker-informed defenses

Using the attacker's playbook to guide defensive controls creates a more effective security posture:

Inventory and version hygiene: Maintain current versions and apply security patches on defined schedules. Knowing what databases exist and which versions they run prevents attackers from exploiting known vulnerabilities in outdated software.

Strong identity and access controls: Eliminate default accounts, enforce strong credential policies, and minimize privileges to essential functions. This reduces the attack surface by limiting what compromised credentials can accomplish.

Monitoring and auditing: Enable comprehensive auditing for sensitive operations and configure real-time alerts for anomalous behavior. Detection capabilities ensure attacks generate evidence and trigger investigation.

Encryption and key management: Encrypt data at rest and in transit with properly managed encryption keys and clear key ownership. Encryption renders stolen data useless without corresponding keys.

Environment mapping: Identify all database copies and non-production instances containing sensitive data, then align controls and auditing across all environments. Attackers target the weakest copy; consistent controls eliminate easy targets.

This approach enables rapid risk assessment, prioritization of high impact remediations, and maintenance of a defensible security posture for Oracle Database deployments.

Fundamental considerations for database protection

Effective database protection requires understanding three fundamental realities that shape assessment priorities and remediation strategies:

Data sensitivity varies within and across databases. A customer's birth date typically carries more sensitivity than their email address due to identity theft implications. Not all tables within a database merit identical protection levels. Assessment must identify which databases contain which types of sensitive data: personally identifiable information (PII), protected health information (PHI), payment card data, intellectual property: to enable risk-proportionate controls. This data classification drives encryption decisions, access restrictions, and audit policies.

Attack surfaces are multifaceted and interdependent. Vulnerabilities emerge from unpatched systems, weak passwords, excessive privileges, missing encryption, inadequate auditing, poor application design, and lack of separation of duties. Comprehensive assessment examines all these vectors simultaneously because attackers will exploit the weakest link regardless of other controls. A fully patched, encrypted database with weak password policies remains vulnerable to credential compromise.

Risk requires business context. Not all databases warrant identical protection investments. Production systems holding regulated financial data demand stronger controls than development environments with synthetic test data. However, development systems containing production data copies require production-equivalent controls. Assessment tools should support risk-based prioritization aligned with business criticality, data sensitivity, and regulatory obligations. Organizations should invest security controls: tools, personnel time, operational resources: commensurate with database criticality.

Copyright © 2026, Oracle and/or its affiliates

Users and accounts are diverse entities. Beyond database administrators, various actors and processes interact with data through database user accounts: applications, application administrators, security administrators, service accounts, batch programs, and increasingly, AI agents. Each user type has different legitimate access requirements. Assessment must identify these user types and their required activities to properly implement least privilege principles. Static privilege reviews miss the nuance of which privileges accounts actually use versus which they possess.

Configuration parameters have complex interdependencies. Security configuration parameters tightly relate to database behavior and require understanding their function, the impact of changes, and their dependencies on other parameters. Changing one parameter may inadvertently weaken another control or break application functionality. Assessment tools that explain parameter purposes and impacts enable confident remediation.

With this attacker-informed foundation, organizations can approach assessment tools strategically, understanding not just what they assess but why those assessments matter for defense.

Assessing configuration security with DBSAT

The Database Security Assessment Tool (DBSAT) provides the foundation for Oracle Database security assessment. DBSAT identifies where database configuration, operations, or implementation introduce risk, collecting and analyzing configuration data and parameters to deliver actionable remediation recommendations.

DBSAT assessment scope

DBSAT examines comprehensive security dimensions beyond basic database and listener configuration:

- **User accounts**: Account types, authentication methods, and account status
- **Privileges and roles**: Granted permissions, role hierarchies, and privilege paths
- **Authorization controls**: Access control mechanisms and enforcement policies
- **Separation of duties**: Administrative role segregation and least privilege implementation
- **Fine-grained access control**: Row-level security, column masking, and Virtual Private Database configurations
- **Data encryption and key management**: Transparent Data Encryption status, tablespace encryption, and key management practices
- **Auditing policies**: Audit configurations, audit trail security, and log retention
- **Operating system file permissions**: File ownership, directory permissions, and OS-level access controls

DBSAT scans for weaknesses and vulnerabilities and categorizes findings by risk level, including Pass, Low Risk, Medium Risk, High Risk, Advisory, and Evaluate, to support remediation prioritization. The tool provides high level executive summaries alongside specific, actionable technical recommendations for each finding, helping teams move more quickly from insight to remediation.

DBSAT is a free command line tool available to Oracle customers with current support agreements. Oracle consultants, partners, and external auditors around the world use DBSAT to perform standardized database security assessments. This broad adoption helps establish consistent assessment methodologies and a shared language for findings across the Oracle Database ecosystem.

DBSAT architecture and components

As shown in Figure 3-1, DBSAT consists of three discrete components serving different assessment purposes:

Copyright © 2026, Oracle and/or its affiliates

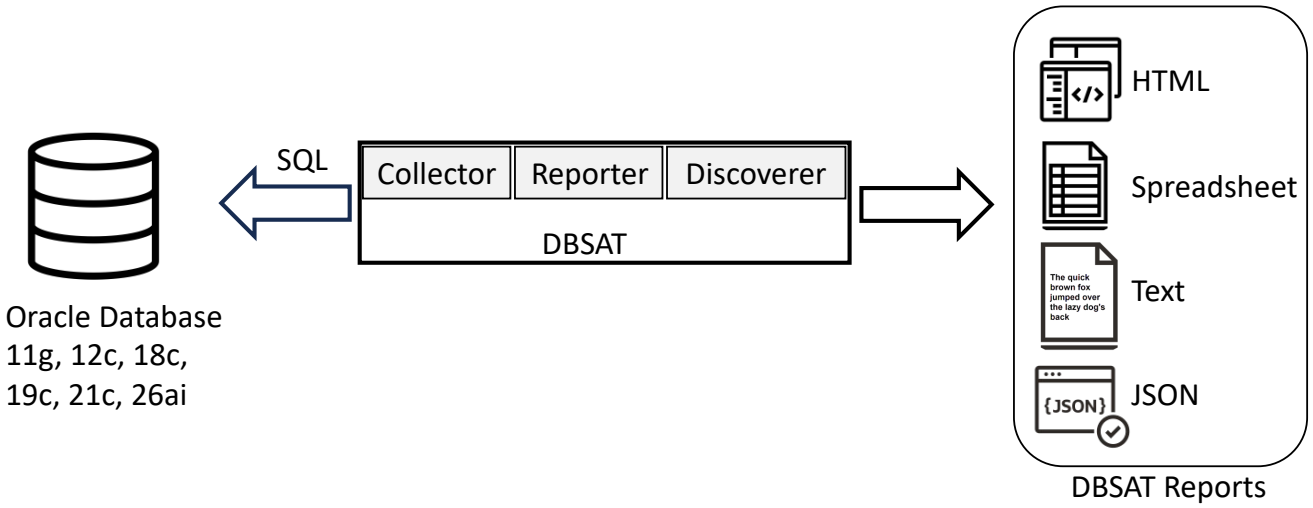


Figure 3-1: High-level functional diagram of Database Security Assessment Tool (DBSAT)

Collector: Gathers security configuration information from the database and underlying operating system. The collector connects to the database using the provided credentials, queries data dictionary views, reviews file system permissions, and creates a compressed JSON file containing the collected data. The collector does not perform analysis or make judgments. It captures the current state for later analysis.

Reporter: Analyzes the collected data and generates comprehensive reports with findings and recommendations. The reporter processes the collector's JSON output offline and compares configurations against Oracle best practices, CIS Benchmarks, DISA STIG requirements, and GDPR principles. This separation between collection and analysis enables organizations to collect data in production environments and perform analysis on separate systems.

Discoverer: Scans database schemas to identify sensitive data and classify it by type. The discoverer flags columns that contain sensitive information, including credit card numbers, Social Security numbers, email addresses, and phone numbers, and summarizes results in reports.

DBSAT report formats and usage

DBSAT generates reports in four formats, each serving specific use cases:

HTML format: Provides detailed, easily navigable assessment results with interactive summary dashboards. HTML reports are ideal for initial review, sharing with stakeholders, and understanding finding context. They include hyperlinked navigation between sections, expandable details, and embedded documentation references.

Spreadsheet format: Delivers high-level finding summaries in CSV format that loads into Excel or similar tools. Spreadsheets enable adding custom columns for tracking remediation status, assigning ownership, setting target dates, and prioritizing based on organizational risk tolerance. Many organizations use spreadsheet exports to integrate findings into existing issue tracking systems.

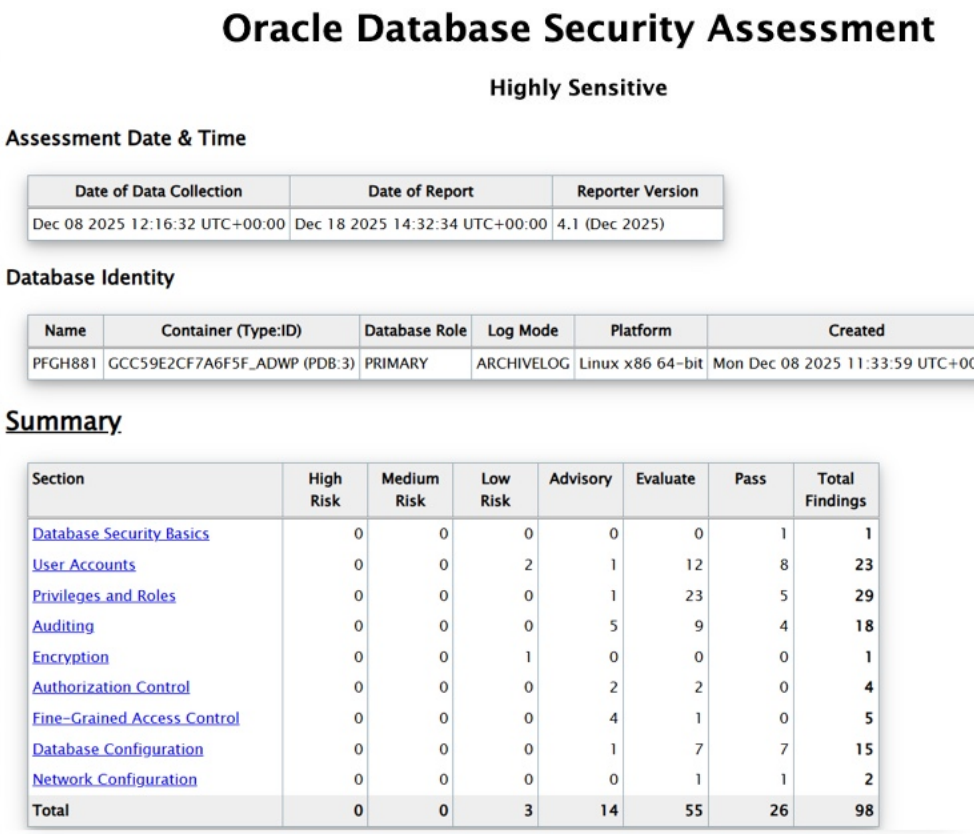
Text format: Outputs plain text suitable for scripting, version control, and command-line workflows. Text format facilitates copying portions of findings into documentation, tickets, or remediation procedures.

JSON format: Structured data ideal for programmatic processing, aggregation across multiple databases, and integration with security information and event management (SIEM) systems or governance, risk, and compliance (GRC) platforms.

Copyright © 2026, Oracle and/or its affiliates

Understanding DBSAT findings

Each DBSAT finding follows a consistent structure providing comprehensive context for security teams and DBAs. Figure 3-2 illustrates the DBSAT HTML summary report showing risk distribution across assessment categories.



Oracle Database Security Assessment

Highly Sensitive

Assessment Date & Time

Date of Data Collection	Date of Report	Reporter Version
Dec 08 2025 12:16:32 UTC+00:00	Dec 18 2025 14:32:34 UTC+00:00	4.1 (Dec 2025)

Database Identity

Name	Container (Type:ID)	Database Role	Log Mode	Platform	Created
PFGH881	GCC59E2CF7A6F5F_ADWP (PDB:3)	PRIMARY	ARCHIVELOG	Linux x86 64-bit	Mon Dec 08 2025 11:33:59 UTC+00:00

Summary

Section	High Risk	Medium Risk	Low Risk	Advisory	Evaluate	Pass	Total Findings
Database Security Basics	0	0	0	0	0	1	1
User Accounts	0	0	2	1	12	8	23
Privileges and Roles	0	0	0	1	23	5	29
Auditing	0	0	0	5	9	4	18
Encryption	0	0	1	0	0	0	1
Authorization Control	0	0	0	2	2	0	4
Fine-Grained Access Control	0	0	0	4	1	0	5
Database Configuration	0	0	0	1	7	7	15
Network Configuration	0	0	0	0	1	1	2
Total	**0**	**0**	**3**	**14**	**55**	**26**	**98**

Figure 3-2: Example DBSAT assessment summary table

The summary provides immediate risk visibility, which categories contain the most severe findings, total finding counts by risk level, and assessment metadata including database identity and collection timestamp.

Drilling into individual findings reveals detailed information structured to support rapid remediation. Figure 3-3 below shows a representative finding examining users with DBA role grants. Each finding contains:

Rule ID: A unique identifier with a prefix indicating the assessment category (e.g., PRIV for privileges, USER for user accounts, ENCRYPT for encryption) followed by a descriptive name. Rule IDs remain consistent across DBSAT versions, enabling tracking of specific findings over time and across databases.

One-liner: A concise sentence describing what the finding assesses. This summary enables quick scanning of multiple findings to identify priorities.

Status: Indicates risk severity or required action:

- **Pass**: Configuration meets security best practices
- **Low Risk**: Minor security concern with limited exploitation potential
- **Medium Risk**: Significant security concern requiring timely remediation
- **High Risk**: Critical security vulnerability demanding immediate attention
- **Advisory**: Information about optional security features not currently deployed
- **Evaluate**: Configuration requires manual review to determine risk level based on organizational context

Copyright © 2026, Oracle and/or its affiliates

Summary: Provides an overview of what the finding discovered in the assessed database. For example: "9 out of 53 users have been directly or indirectly granted highly sensitive DBA/PDB_DBA role via 9 grants."

Details: Presents specific findings from the database, listing affected users, configuration parameters, or security issues. In Figure 3-3, DBSAT identifies users with DBA roles, including SCOTT who received the role through nested role grants (SCOTT ← APPROLE1 ← APPROLE2 ← APPROLE3 ← DBA), a common obfuscation tactic.

Remarks: Explains why the finding matters and recommends specific remediation actions. For the DBA role finding, remarks emphasize limiting this powerful role to a small number of trusted administrators, creating custom DBA-like roles with minimum required privileges, and implementing Privilege Analysis to identify unnecessary permissions.

References: Maps findings to industry standards and compliance frameworks including CIS Oracle Database Benchmark recommendations, DISA STIG rules, and GDPR articles. This mapping helps organizations demonstrate compliance alignment and prioritize findings based on regulatory obligations.

Documentation: For Oracle Database 19c and later, provides direct hyperlinks to relevant Oracle documentation sections explaining parameters, features, and remediation procedures.

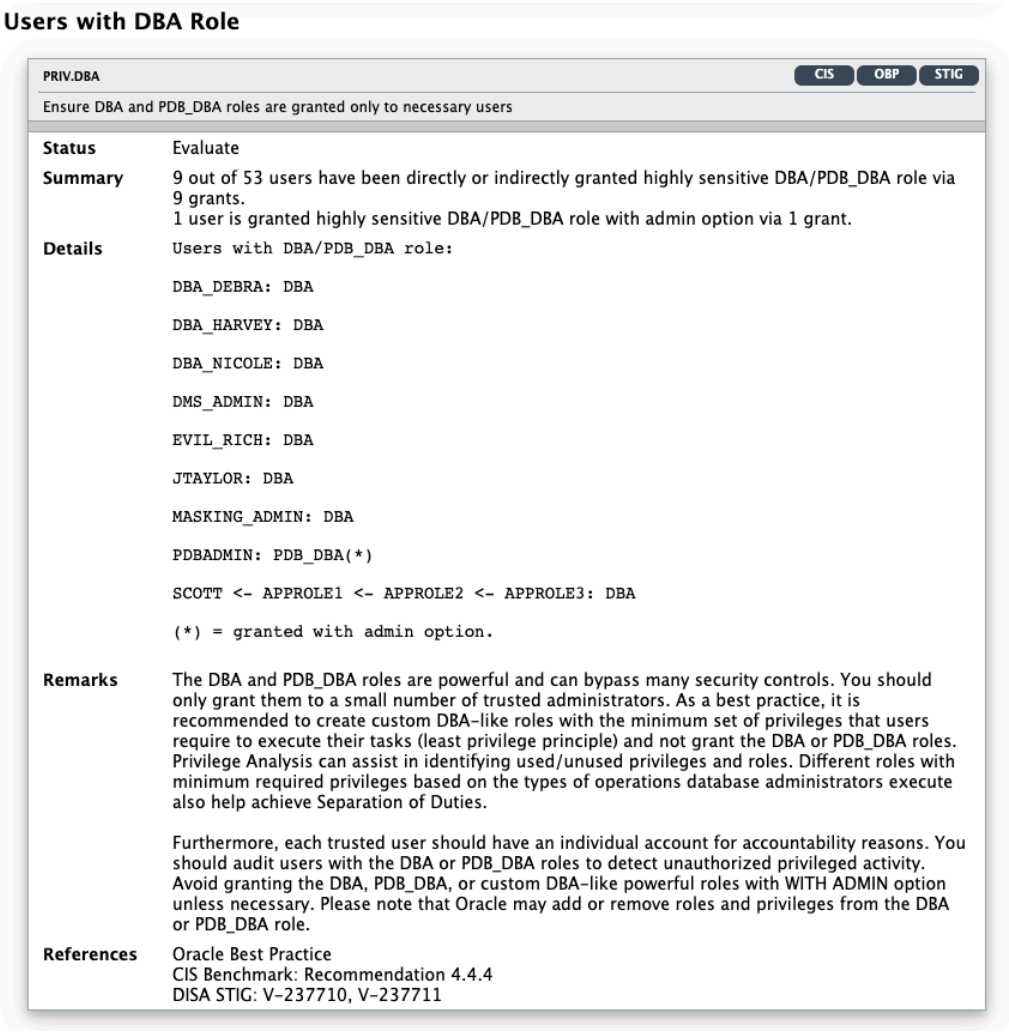


Figure 3-3: DBSAT Sample finding - Users with the DBA Role

In the illustrated finding, DBSAT identified eight users with the DBA role and one with PDB_DBA. Notably, SCOTT holds DBA privileges through a chain of nested role grants, a configuration that can be difficult to spot in a manual review but is surfaced through DBSAT's systematic analysis. The finding status, "Evaluate", signals that the DBA should review these grants to confirm they align with organizational needs. SCOTT may legitimately require DBA

Copyright © 2026, Oracle and/or its affiliates

privileges as an application administrator, or the grants may reflect temporary access that was never removed and should have been revoked months ago.

This structured level of detail helps DBAs and security teams quickly understand what was found, why it matters, and what actions to take, even when team members do not have deep security expertise.

DBSAT deployment patterns

Organizations deploy DBSAT using several patterns depending on fleet size and operational maturity:

Manual execution for small deployments: Organizations with a handful of databases run DBSAT manually on a quarterly schedule, storing collector output in version control and tracking findings in spreadsheets. This approach works well for up to 1-10 databases with minimal automation investment.

Scripted execution for medium deployments: Organizations with dozens of databases automate collection using shell scripts that iterate through database inventories, execute the collector with stored credentials, and organize output by database name and timestamp. Analysis remains manual but collection automation ensures consistent assessment schedules.

Enterprise integration for large deployments: Organizations with hundreds or thousands of databases integrate DBSAT into existing automation frameworks (Ansible, Puppet, Chef), centralized configuration management databases, and security dashboards. Automated workflows collect data, execute the reporter, parse JSON output, and populate tracking systems. Exceptions and high-risk findings trigger automated tickets and notifications.

The tool's command-line nature and structured outputs make DBSAT amenable to extensive automation without requiring commercial orchestration platforms.

Assessing database fleets with Oracle Data Safe

While DBSAT excels at assessing individual databases, managing security across dozens, hundreds, or thousands of databases requires automation, centralized visibility, and fleet-wide orchestration. Manual assessment reviews do not scale. Oracle Data Safe addresses this challenge as an Oracle Cloud Infrastructure service purpose-built for Oracle Database security at scale.

Data Safe provides a unified view of fleet security posture, enabling security teams and DBAs to discover risk concentrations, prioritize remediation across multiple databases, and track security improvements over time. Data Safe works with Oracle Databases regardless of deployment location: on-premises data centers, Oracle Cloud Infrastructure, third-party clouds (AWS, Azure, Google Cloud), and Oracle Database@Multicloud configurations.

Data Safe capabilities beyond assessment

Data Safe delivers six integrated security capabilities within a single platform:

- **Security assessment**: Evaluates database configurations against best practices and compliance frameworks
- **User assessment**: Analyzes user accounts, privileges, roles, and risk profiles
- **Activity auditing**: Collects, consolidates, and analyzes database audit trails
- **Security policy management**: Deploys and enforces consistent security configurations across fleets
- **SQL Firewall**: Monitors and blocks unauthorized SQL statements
- **Data discovery and masking**: Identifies sensitive data and anonymizes non-production environments

Copyright © 2026, Oracle and/or its affiliates

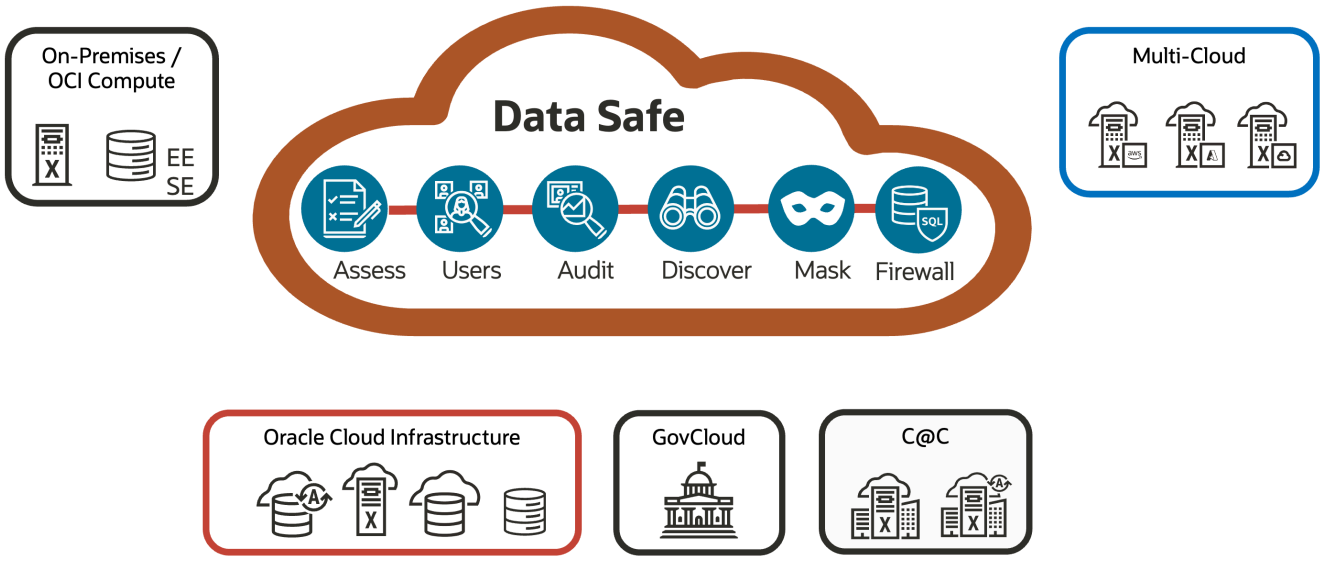


Figure 3-4: Oracle Data Safe provides essential security for Oracle AI Databases, both in the cloud and on-premises

Data Safe security assessment

Data Safe security assessment performs comprehensive configuration evaluation identical in scope to DBSAT but with enterprise features enabling fleet management:

Assessment areas evaluated:

- User accounts and authentication configurations
- Privilege and role grants
- Authorization controls
- Fine-grained access control policies
- Auditing policies and configurations
- Encryption status and key management
- Database configuration parameters
- Network security and access controls

Data Safe compares discovered configurations against organizational best practices and flags deviations. Reports include prioritized recommendations mapped to compliance frameworks including GDPR, DISA STIG, and CIS Benchmarks, ensuring remediation efforts align with regulatory obligations and industry standards.

Copyright © 2026, Oracle and/or its affiliates

Fleet-Scale features

Data Safe enhances DBSAT's assessment capabilities with features essential for managing large database populations:

Simultaneous multi-database assessment: Launch assessments across all fleet databases simultaneously rather than sequentially assessing one at a time. Parallel execution completes fleet-wide assessments in minutes regardless of database count.

Scheduled recurring assessments: Define assessment schedules (weekly, monthly, quarterly) that execute automatically. Scheduled assessments ensure continuous monitoring without manual intervention and establish trend data showing security posture over time.

Security baselines: Designate an assessment as the security baseline representing the desired configuration state. Subsequent assessments compare against this baseline to identify configuration drift. Baselines enable tracking what changed rather than re-reviewing all findings each cycle.

Drift reports: Generate reports highlighting differences between baseline and current assessments. Drift reports show which databases diverged from approved configurations, which new findings appeared, and which previous findings were remediated. This differential view focuses remediation efforts on actual changes.

Assessment history: Retain historical assessments to analyze security posture trends. History enables answering questions like "Are we improving?" and "When did this configuration drift occur?"

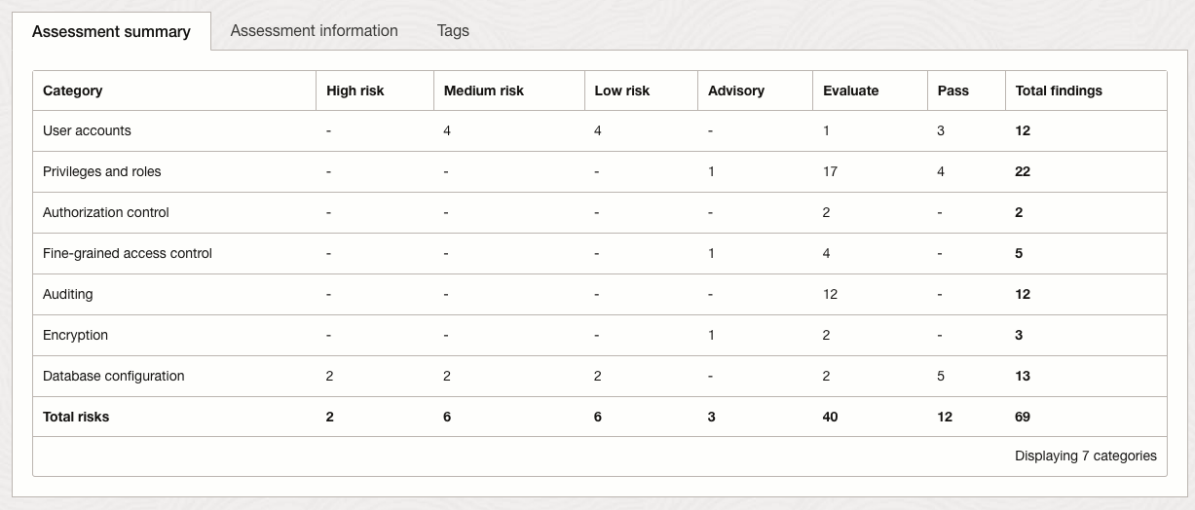
Fleet aggregation: Consolidate findings across all databases to identify systemic issues. If 80% of databases have identical misconfigurations, the problem likely originates from deployment templates or operational procedures requiring correction.

Filtering and search: Query findings across the entire fleet using database name, risk level, finding type, or custom tags. Large organizations can quickly identify all databases with specific high-risk findings requiring immediate attention.

Event-driven notifications: Configure alerts that trigger when assessments discover High Risk findings or when configuration drift exceeds thresholds. Notifications integrate with email, ticketing systems, and incident management platforms.

Copyright © 2026, Oracle and/or its affiliates

Figure 3-5 shows the Data Safe assessment summary interface providing immediate visibility into fleet risk distribution. Unlike DBSAT's static HTML report, Data Safe provides an interactive dashboard enabling drill-down from summary metrics into specific findings and affected databases.



Assessment summary | Assessment information | Tags

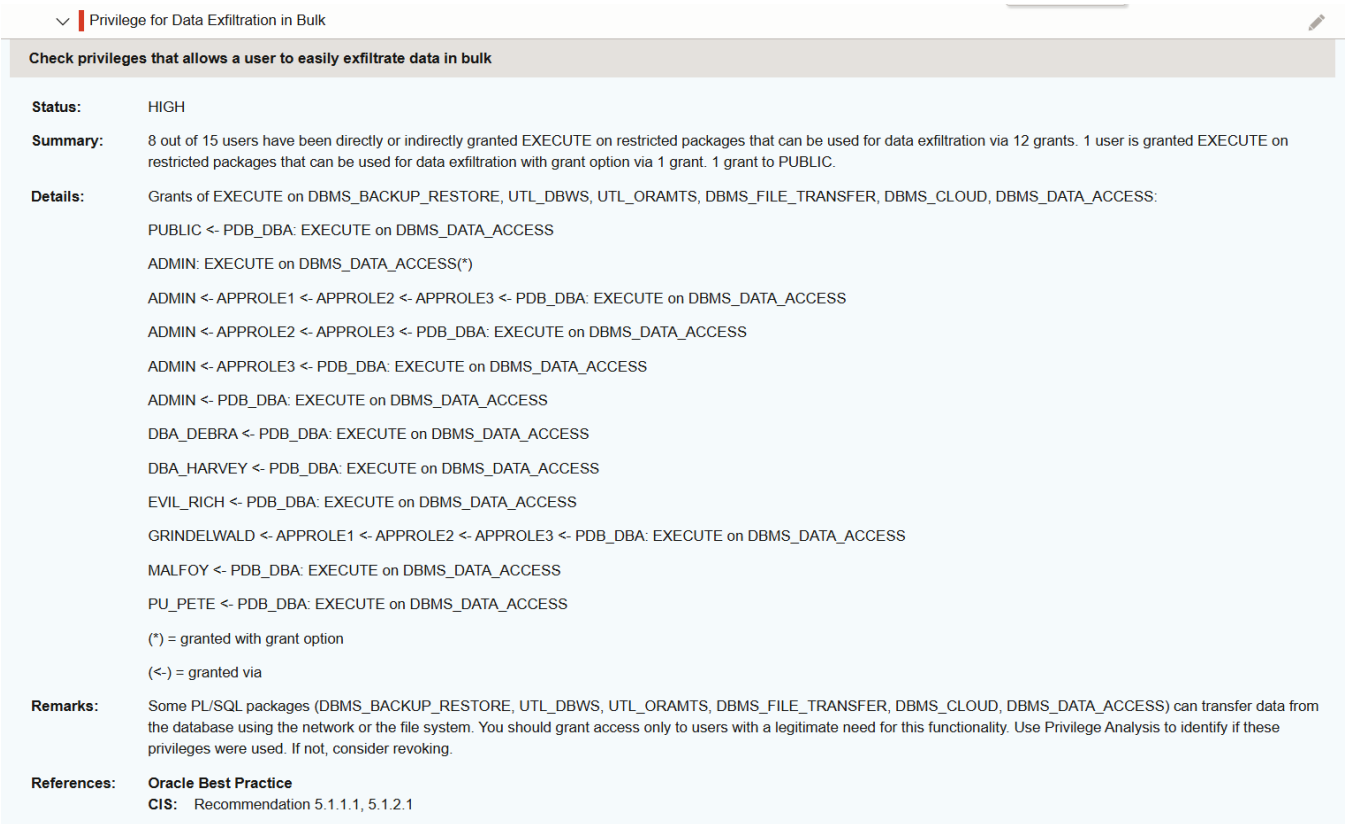
Category	High risk	Medium risk	Low risk	Advisory	Evaluate	Pass	Total findings
User accounts	-	4	4	-	1	3	**12**
Privileges and roles	-	-	-	1	17	4	**22**
Authorization control	-	-	-	-	2	-	**2**
Fine-grained access control	-	-	-	1	4	-	**5**
Auditing	-	-	-	-	12	-	**12**
Encryption	-	-	-	1	2	-	**3**
Database configuration	2	2	2	-	2	5	**13**
Total risks	**2**	**6**	**6**	**3**	**40**	**12**	**69**

Displaying 7 categories

Figure 3-5: Example Data Safe assessment summary

Data Safe finding management

Data Safe utilizes the same assessment rules as DBSAT but enhances finding management with enterprise workflows. Figure 3-6 illustrates detailed findings with risk categorization and compliance framework mapping.



Privilege for Data Exfiltration in Bulk

Check privileges that allows a user to easily exfiltrate data in bulk

Status: HIGH

Summary: 8 out of 15 users have been directly or indirectly granted EXECUTE on restricted packages that can be used for data exfiltration via 12 grants. 1 user is granted EXECUTE on restricted packages that can be used for data exfiltration with grant option via 1 grant. 1 grant to PUBLIC.

Details: Grants of EXECUTE on DBMS_BACKUP_RESTORE, UTL_DBWS, UTL_ORAMTS, DBMS_FILE_TRANSFER, DBMS_CLOUD, DBMS_DATA_ACCESS:

PUBLIC <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

ADMIN: EXECUTE on DBMS_DATA_ACCESS(*)

ADMIN <- APPROLE1 <- APPROLE2 <- APPROLE3 <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

ADMIN <- APPROLE2 <- APPROLE3 <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

ADMIN <- APPROLE3 <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

ADMIN <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

DBA_DEBRA <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

DBA_HARVEY <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

EVIL_RICH <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

GRINDELWALD <- APPROLE1 <- APPROLE2 <- APPROLE3 <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

MALFOY <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

PU_PETE <- PDB_DBA: EXECUTE on DBMS_DATA_ACCESS

(*) = granted with grant option

(<-) = granted via

Remarks: Some PL/SQL packages (DBMS_BACKUP_RESTORE, UTL_DBWS, UTL_ORAMTS, DBMS_FILE_TRANSFER, DBMS_CLOUD, DBMS_DATA_ACCESS) can transfer data from the database using the network or the file system. You should grant access only to users with a legitimate need for this functionality. Use Privilege Analysis to identify if these privileges were used. If not, consider revoking.

References: **Oracle Best Practice**
CIS: Recommendation 5.1.1.1, 5.1.2.1

Figure 3-6: Example Data Safe security assessment findings

Copyright © 2026, Oracle and/or its affiliates

Data Safe also enables configuration of acceptable risks. For example, when an assessment flags a finding for evaluation (Status EVALUATE, as shown in Figure 3-7: Example), the Data Safe user is expected to review it and determine whether it represents a real issue. In the provided example, Data Safe identifies two users, ADMIN and SKEETER, with the privilege to manage audit settings, but the reviewer must decide whether SKEETER's privilege is actually necessary.

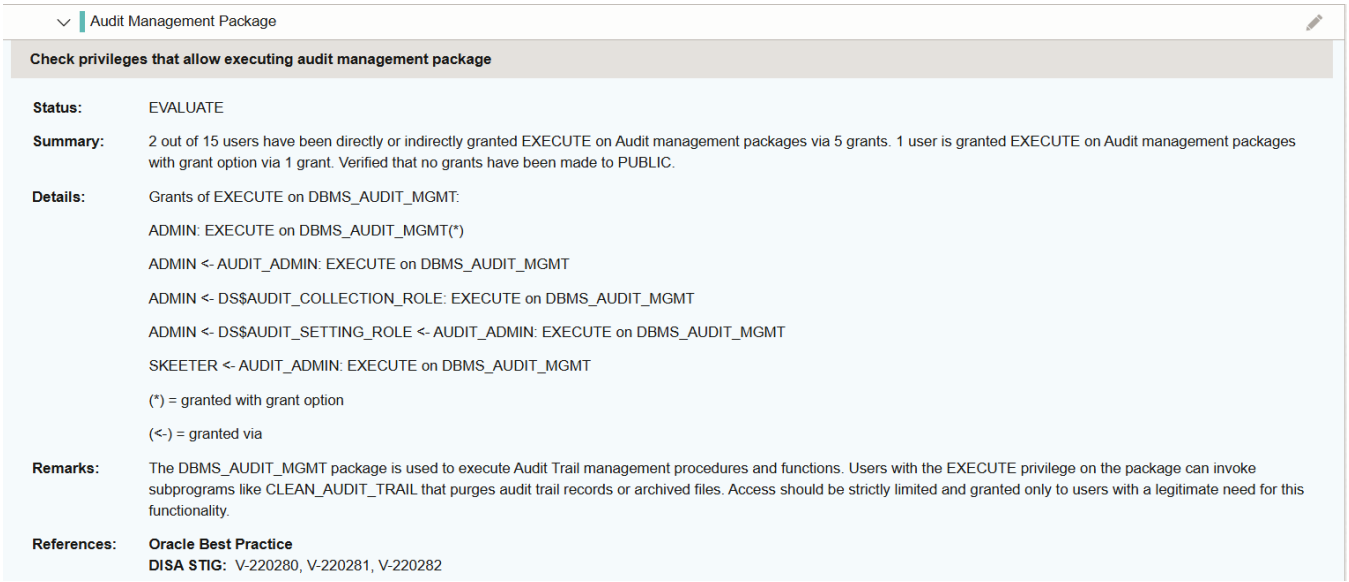


Figure 3-7: Example Data Safe finding requiring evaluation

If a review confirms the finding presents no risk (i.e., SKEETER legitimately requires the privilege), you can mark the finding as "Pass." You can also lower a finding's risk if compensating controls are in place. Alternatively, you can choose to accept the risk if no remediation is planned. In the scenario illustrated, SKEETER's privilege was validated, and Data Safe was configured to remind the user to review the finding again later.

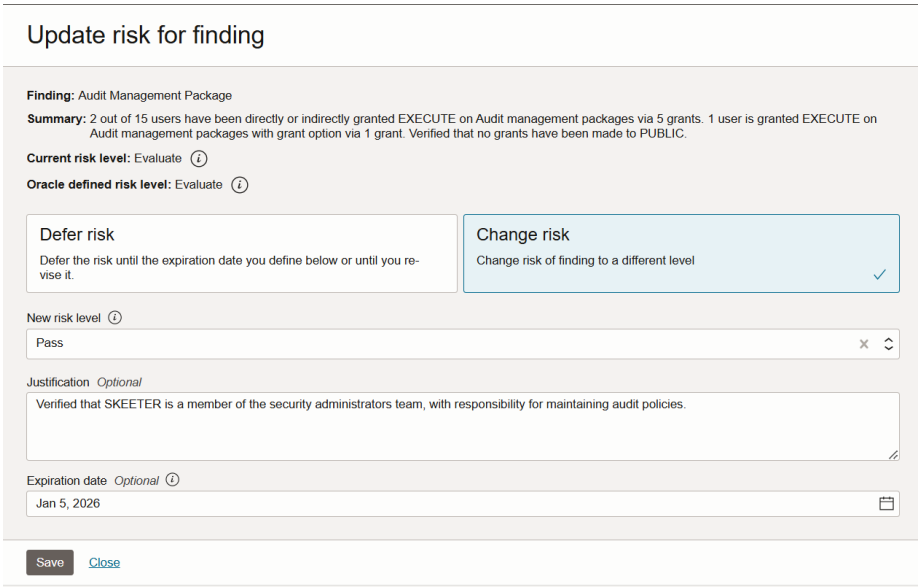


Figure 3-8: Example Data Safe screen for configuring acceptable risks

Updated risks are easy to identify, since any changes from the default values are clearly shown in Figure 3-9.

Copyright © 2026, Oracle and/or its affiliates

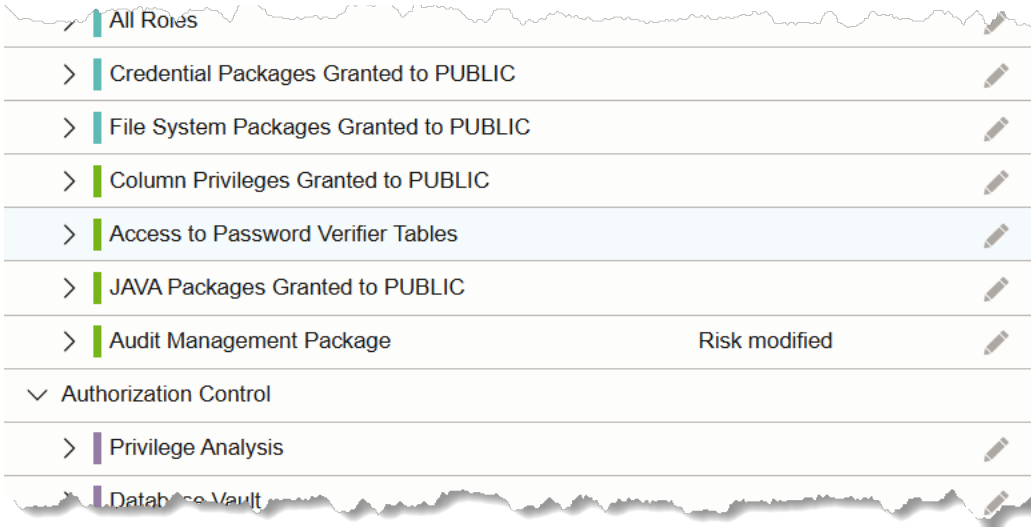


Figure 3-9: Example screen for finding modified risks

And security assessment's *risk modification report* lets you quickly drill into details on the modification.

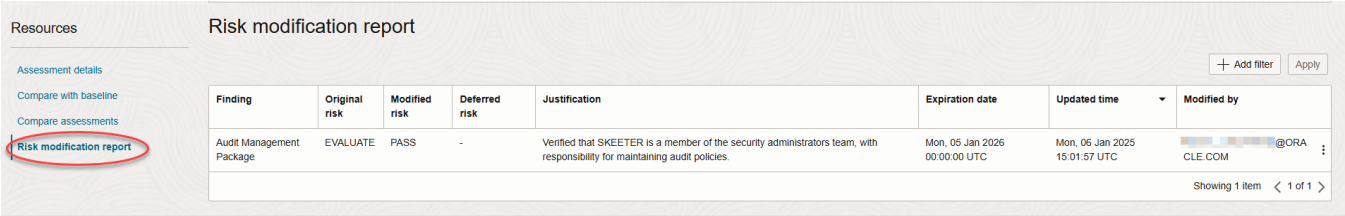


Figure 3-10: Example Security Assessment risk modification report

Understanding user risk with Data Safe

Privileged credential theft ranks among the most common attack vectors enabling data breaches. Data Safe user assessment helps mitigate this risk by evaluating database accounts and visualizing user risk to enable rapid intervention.

User assessment analyzes accounts using both static and dynamic characteristics:

Static profile analysis: Assigned roles, granted privileges, profile settings, authentication methods, and account creation date

Dynamic behavioral analysis: Last login timestamp, password change recency, failed login attempts, and database session patterns

The assessment highlights highest-risk users based on privilege levels and suspicious patterns. Figure 3-11 illustrates the user risk dashboard showing potential risk distribution, role assignments, and password change patterns.

Copyright © 2026, Oracle and/or its affiliates

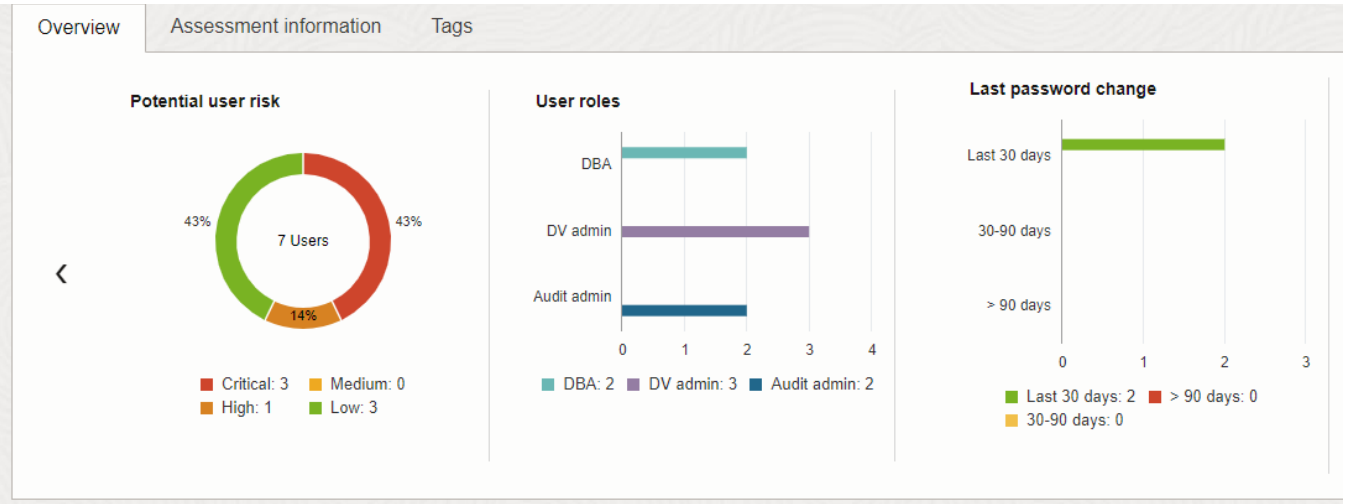


Figure 3-11: Example Data Safe user risk assessment

Security teams can quickly identify concerning patterns:

- Users who have not authenticated in three months or longer (stale accounts potentially forgotten)
- Users who have not changed passwords in 90+ days
- High-privilege accounts (DBA, SYSDBA) without corresponding audit trails
- Service accounts with interactive login capabilities

When suspicious activity is detected, Data Safe enables drilling into account details: creation date, assigned roles and privileges, privilege grant paths: and opening related audit reports to review recent actions. This integrated approach connects static privilege analysis with runtime behavior analysis.

User profile management

Database user profiles control authentication security, password policies, and resource limits. Profiles represent critical security controls but are often inconsistently deployed across database fleets. Data Safe's User Profile Insight provides visibility into profile configurations across all registered databases.

Organizations can:

- Review existing profiles and parameters: Examine password verification functions enforcing complexity, failed login attempt limits, and password expiration policies
- Compare profiles across databases: Identify differences between identically named profiles on different databases indicating configuration drift
- Identify profile-to-user assignments: Understand which users are protected by which profiles
- Find weak configurations: Quickly locate users and profiles with inadequate password governance

Copyright © 2026, Oracle and/or its affiliates

Figure 3-12 shows User Profile Insight summarizing profile distribution and highlighting users without proper password complexity requirements.

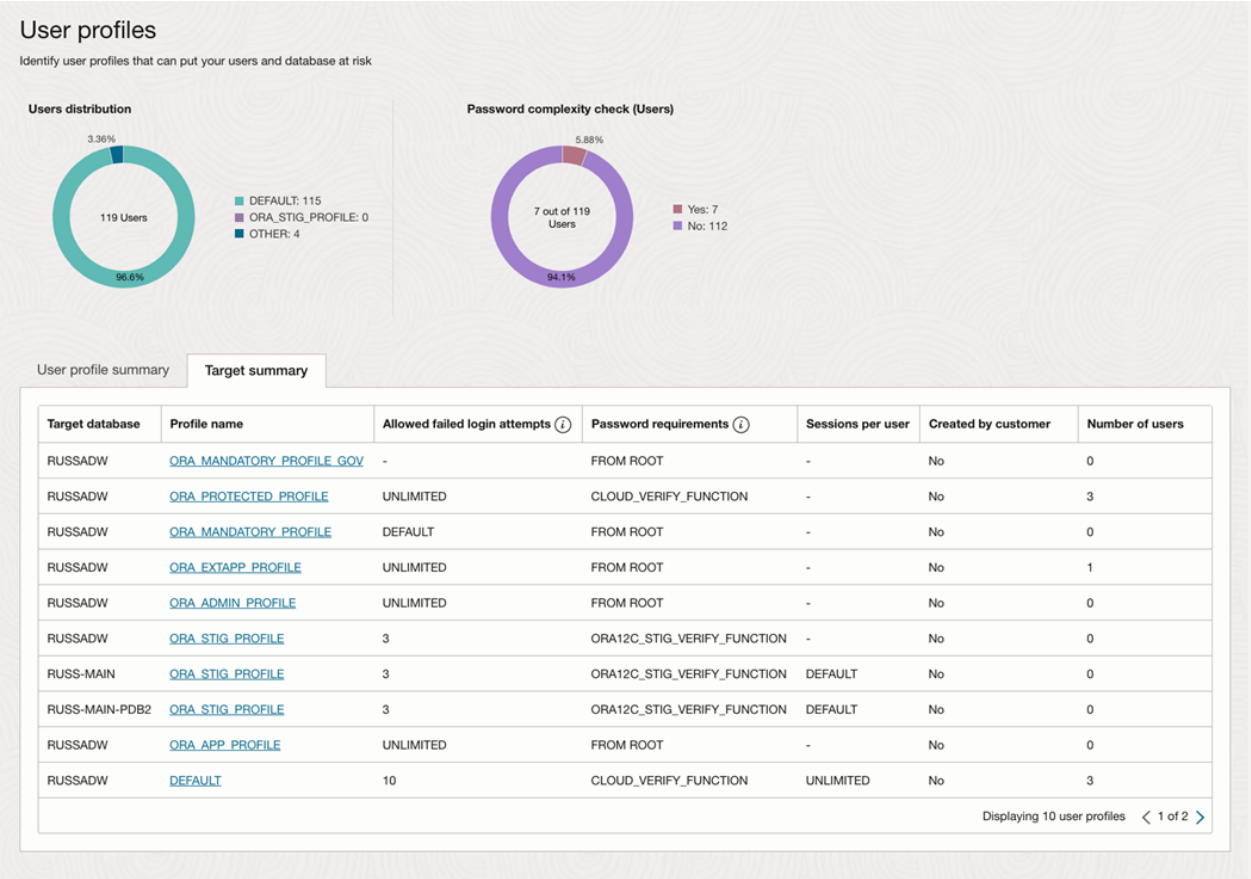


Figure 3-12: Example Data Safe user profile insight

Data Safe deployment and architecture

Data Safe operates as a multi-tenant Oracle Cloud Infrastructure service. Organizations deploy Data Safe Targets: registered databases: in one of several configurations:

Cloud-native databases: Oracle Autonomous Database, Base Database Service, and Exadata Cloud Service instances register with Data Safe automatically during provisioning. No additional configuration is required.

On-premises databases: Traditional on-premises Oracle Database instances connect to Data Safe through Oracle Data Safe Private Endpoints that establish secure network tunnels between on-premises networks and Oracle Cloud Infrastructure. Private Endpoints enable Data Safe functionality without exposing databases to the public internet.

Third-party cloud databases: Oracle Database instances running on AWS, Azure, or Google Cloud connect through Private Endpoints or network peering configurations.

Oracle Database@Multicloud: Oracle Database services running in multi-cloud environments register natively through cloud-to-cloud network connectivity.

Data Safe's cloud-based architecture eliminates infrastructure management overhead: no servers to maintain, no software to patch, no capacity planning required. Oracle manages the service infrastructure, ensuring availability, scalability, and security.

Copyright © 2026, Oracle and/or its affiliates

Organizations with regulatory requirements prohibiting cloud-based management services can use Oracle Audit Vault and Database Firewall (covered in the next section) which provides similar assessment capabilities with on-premises deployment.

Having explored Data Safe's assessment capabilities, the next section examines Oracle Audit Vault and Database Firewall's security posture management features for organizations requiring on-premises solutions.

Assessment with Oracle Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall (AVDF) provides comprehensive database security posture management for organizations preferring or requiring on-premises deployment. While Data Safe operates as a cloud service, AVDF deploys as a pre-hardened Linux appliance in customer data centers, providing similar assessment capabilities with on-premises data residency.

AVDF delivers fleet-wide centralized assessment with compliance framework mapping and actionable recommendations. Organizations establish security baselines, schedule recurring assessments, and generate drift reports showing how current configurations diverge from approved baselines. This differential approach enables focusing on actual changes rather than re-reviewing all findings with each assessment cycle.

AVDF security assessment capabilities

AVDF's assessment engine leverages the same assessment rules as DBSAT and Data Safe, ensuring consistent finding definitions and risk classifications across Oracle's assessment tools. Figure 3-13 shows AVDF's security assessment summary displaying fleet-wide risk distribution

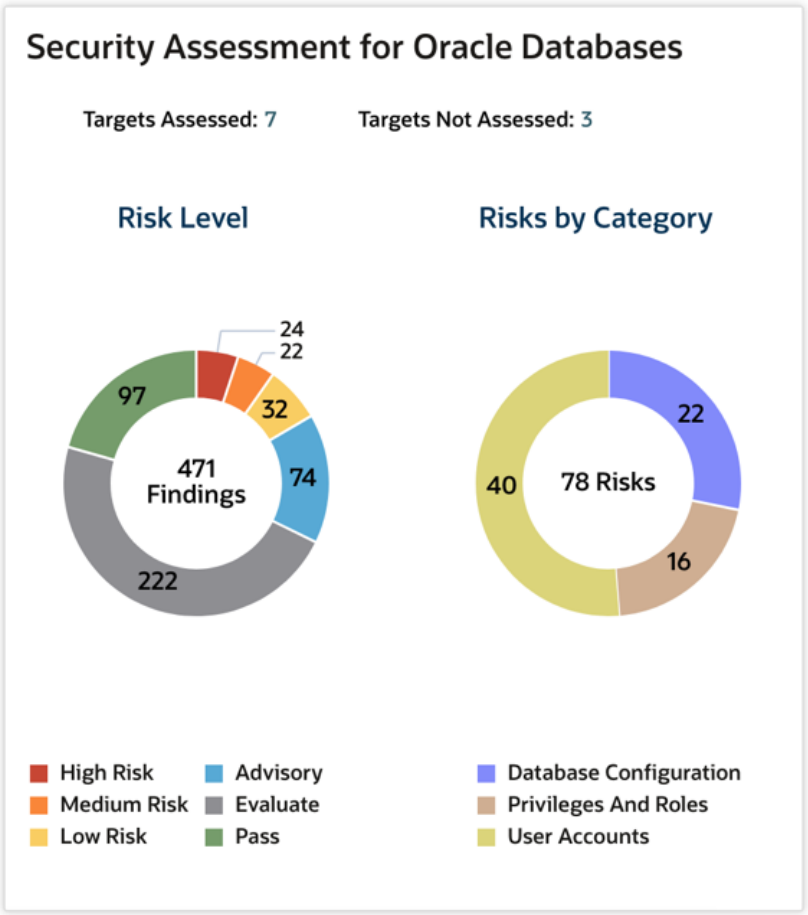


Figure 3-13: Example AVDF security posture management dashboard

Copyright © 2026, Oracle and/or its affiliates

The dashboard provides immediate answers to critical questions: How many databases are assessed? How many are pending assessment? What is the overall risk distribution? Which categories contain the most findings?

Detailed assessment reporting

AVDF's reporting engine leverages Oracle Application Express (APEX), providing interactive drill-down navigation to precise information. Figure 3-14 displays the detailed assessment report showing all findings with filtering and search capabilities.

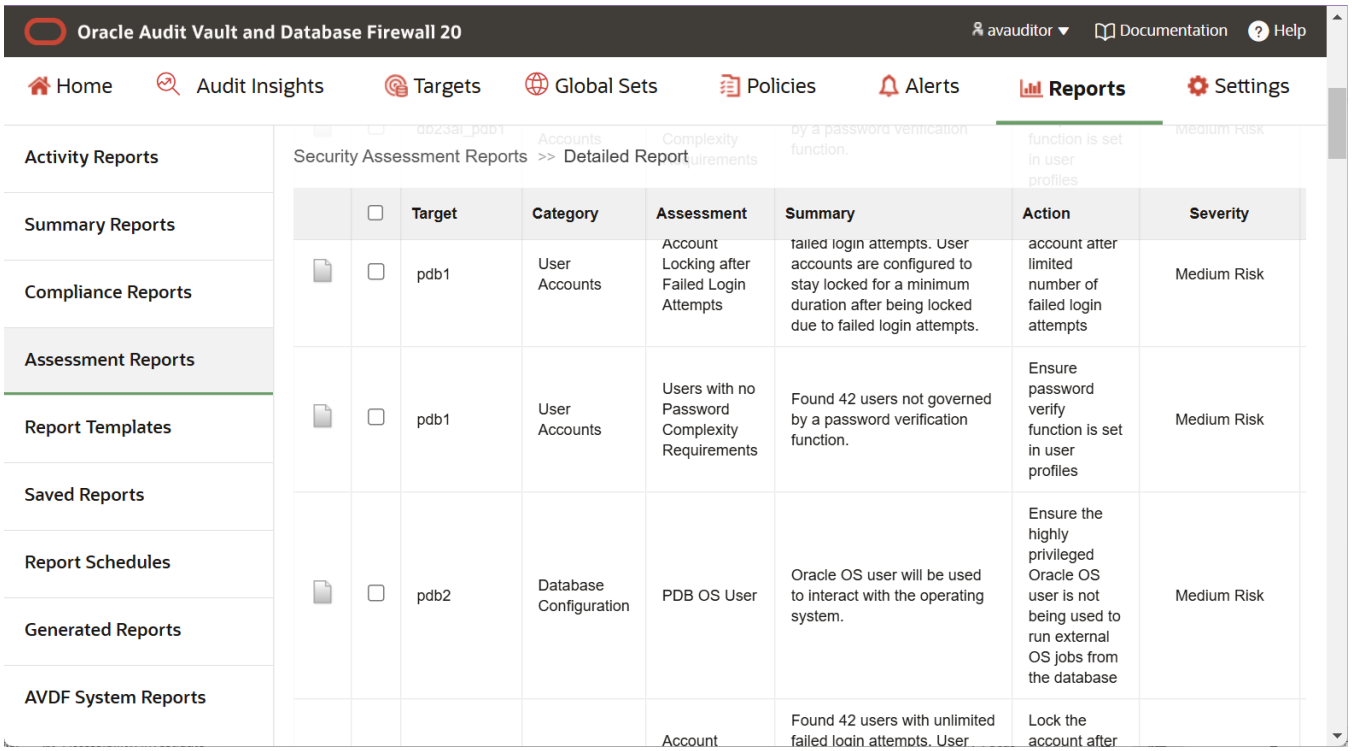


Figure 3-14: Example AVDF detailed assessment report

Organizations can filter findings by:

- Target database
- Risk level (High, Medium, Low, Advisory, Evaluate)
- Assessment category (User Accounts, Privileges, Encryption, Auditing)
- Compliance framework (CIS, STIG, GDPR)

Copyright © 2026, Oracle and/or its affiliates

From summary views, users drill into individual findings to access complete remediation guidance. Figure 3-15 illustrates detailed information for a finding concerning users without password complexity requirements: the same level of detail provided in DBSAT HTML reports.

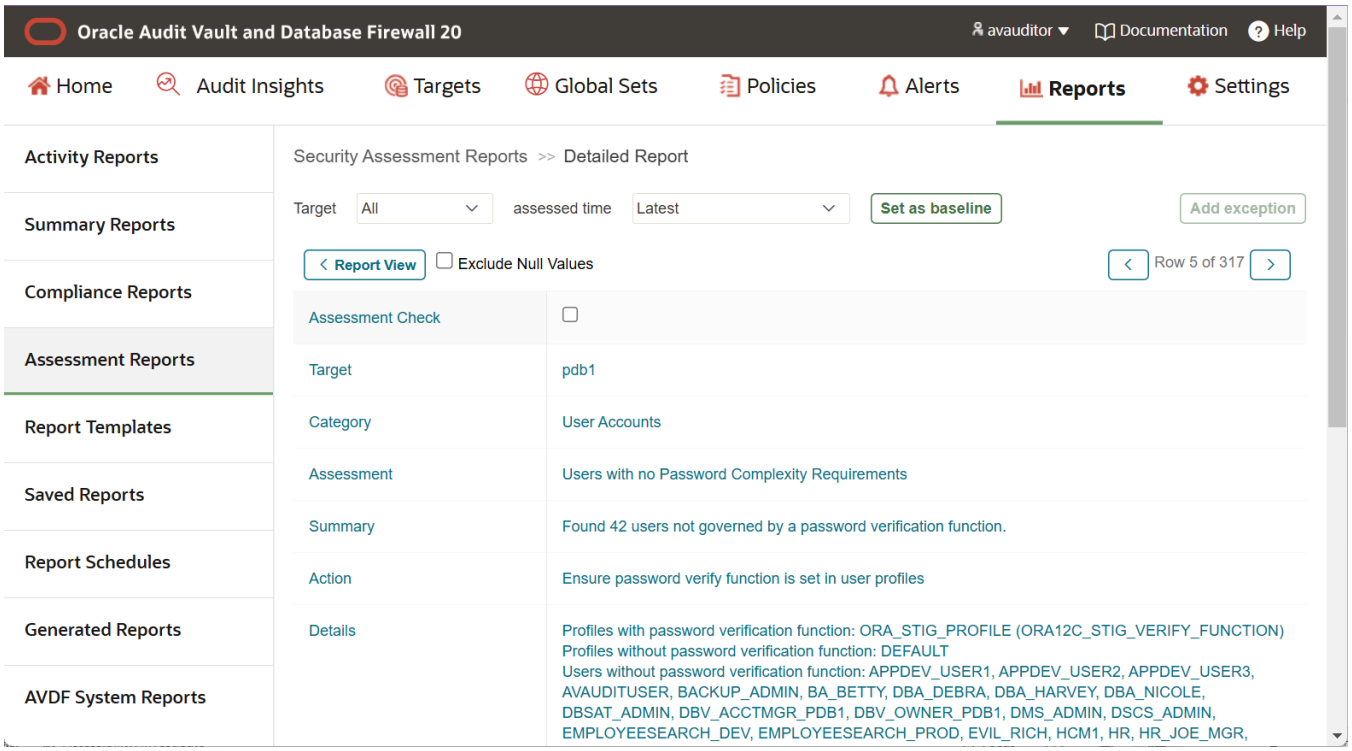


Figure 3-15: Example AVDF detailed finding

Stored procedure monitoring

AVDF provides unique capabilities beyond configuration assessment, including stored procedure change tracking. Monitoring stored procedure modifications helps detect a common post-compromise attacker technique: backdoor injection into legitimate stored procedures.

After gaining initial access, sophisticated attackers modify existing stored procedures to create persistent backdoors. Rather than creating new, obviously suspicious accounts or procedures, they inject malicious code into established, trusted procedures that execute with elevated privileges. These modifications typically occur outside normal patching windows, making them detectable through change monitoring.

AVDF's stored procedure auditing enables organizations to:

- **Determine exactly when procedures changed**: Timestamp tracking identifies modification timing
- **View code differentials**: Side-by-side comparison shows what changed between versions
- **Differentiate expected updates from unexpected drift**: Changes during maintenance windows are expected; changes at 2 AM on weekends trigger investigation
- **Focus on high-risk procedures**: Prioritize monitoring procedures that access sensitive data or execute with elevated privileges

Figure 3-16 shows the stored procedure auditing configuration interface where administrators schedule monitoring frequency and scope.

Copyright © 2026, Oracle and/or its affiliates

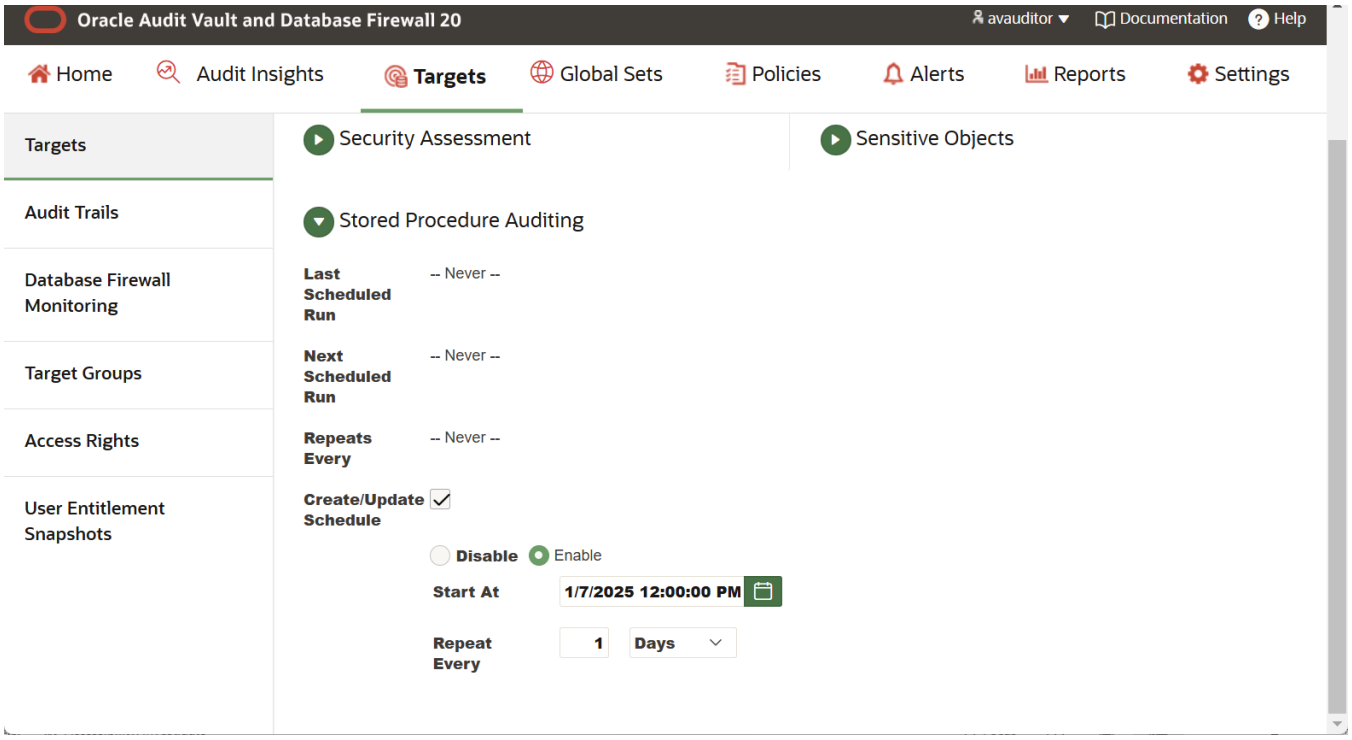


Figure 3-16: Example AVDF configuring stored procedure auditing

When AVDF detects stored procedure changes outside maintenance windows, it generates alerts with code differentials, enabling rapid investigation. This targeted change tracking helps organizations respond before hidden backdoors escalate into data breaches.

User entitlement analysis

AVDF facilitates role and privilege grant analysis through user entitlement snapshots. Organizations capture privilege states at specific moments, label them (e.g., "Q4 2025 Production Baseline"), and compare snapshots to track how entitlements evolve over time.

This capability helps answer critical questions:

- Which new privileges were granted since the baseline?
- Which high-privilege accounts were created?
- Did privilege grants follow change control procedures?
- Are there privilege grant patterns indicating policy violations?

Comparing quarterly entitlement snapshots helps detect privilege creep, the gradual accumulation of unnecessary privileges as users change roles but retain previous permissions. Snapshot comparisons make privilege creep visible and quantifiable, enabling targeted remediation.

AVDF beyond assessment

Similar to Data Safe, AVDF provides capabilities beyond security assessment. Organizations use AVDF for:

- **Sensitive data discovery**: Identifying and classifying sensitive data columns
- **Database activity monitoring**: Real-time monitoring and alerting on suspicious database activity
- **Centralized audit management**: Collecting, securing, and analyzing audit trails from multiple databases

Copyright © 2026, Oracle and/or its affiliates

AVDF's integrated platform approach combines assessment, monitoring, and audit management in a single appliance, reducing tool sprawl and simplifying operations.

With AVDF covering on-premises deployment scenarios, the next section examines Oracle Enterprise Manager's Database Lifecycle Management Pack for organizations with existing Enterprise Manager investments.

Configuration assessment with EM Database Lifecycle Management

Oracle Database Lifecycle Management Pack (DBLM) for Enterprise Manager automates database security assessment and compliance monitoring for organizations with existing Enterprise Manager deployments. DBLM integrates assessment into the broader database lifecycle management framework encompassing patching, provisioning, configuration management, and performance monitoring.

DBLM Compliance Framework

DBLM provides predefined compliance standards and customizable frameworks with detailed reports identifying configuration deviations. This enables maintaining secure, consistent configurations across database estates while reducing manual review time.

The compliance dashboard, Figure 3-17 below, delivers centralized visibility into compliance status across all managed databases. It presents compliance scores, highlights non-compliant targets, and enables drill-down into detailed reports for prioritization and resolution.

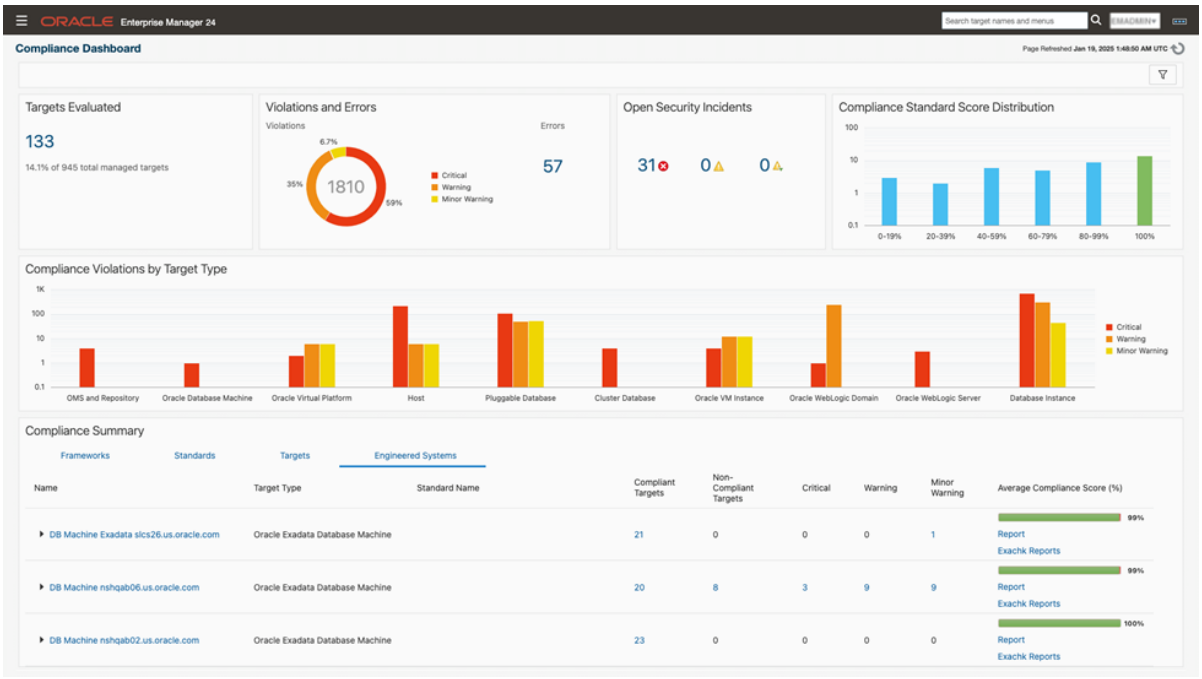


Figure 3-17: Example Oracle Enterprise Manager compliance dashboard

DBLM's compliance frameworks address multiple security dimensions:

- **Initialization parameters**: Security-relevant parameter configurations
- **Operating system directory permissions**: File ownership and access controls
- **User account profiles**: Password policies and resource limits
- **User access and authorization**: Privilege grants and role assignments
- **Auditing configuration**: Audit policies and audit trail security

Copyright © 2026, Oracle and/or its affiliates

- **Sensitive objects**: Access controls for critical database objects

Figure 3-18 shows the general security reports library providing quick access to common security assessments.

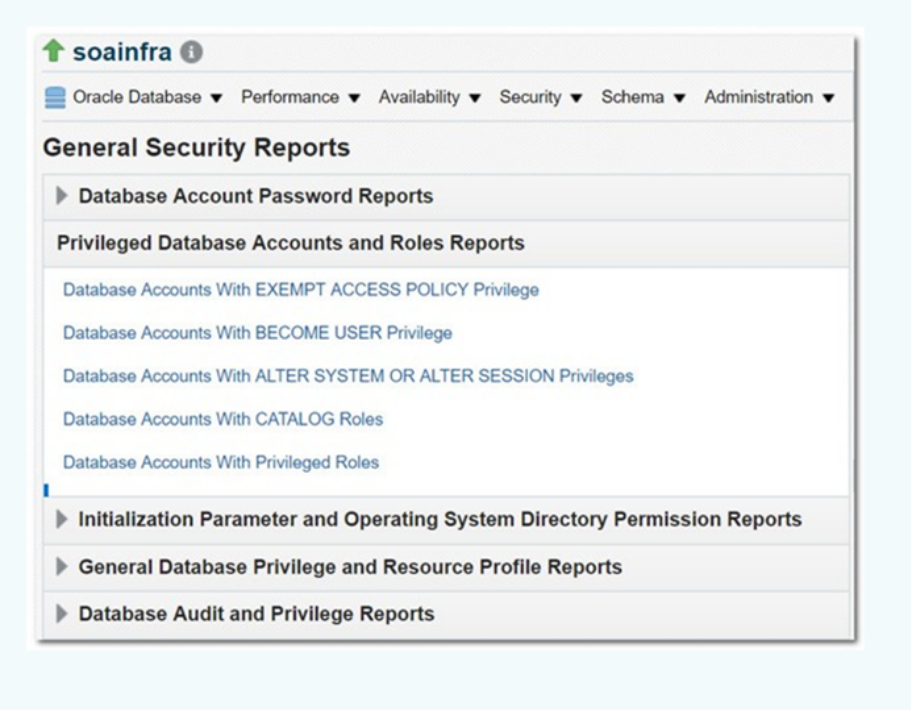


Figure 3-18: Example Enterprise Manager general security reports

Standards-based compliance monitoring

DBLM incorporates predefined compliance standards including:

- CIS Benchmarks: Center for Internet Security Oracle Database Benchmarks
- DISA STIG: Defense Information Systems Agency Security Technical Implementation Guide
- Oracle Database Security Best Practices: Oracle-recommended configurations

Organizations apply these standards across database fleets, examine detailed reports, and rapidly identify remediation priorities. Figure 3-19 illustrates compliance results for the STIG standard showing target-specific compliance scores and violation details.

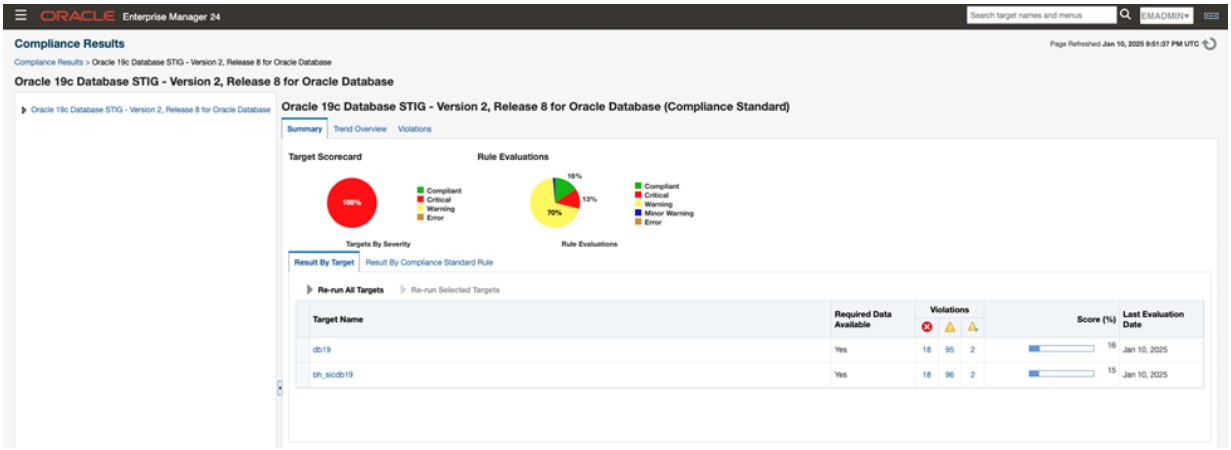


Figure 3-19: Example Enterprise Manager STIG compliance results

Copyright © 2026, Oracle and/or its affiliates

DBLM's standards include detailed rules validating specific configurations:

- Oracle Database initialization parameters
- Oracle Real Application Clusters (RAC) node configurations
- Oracle Listener security settings
- Operating system security for database server hosts

Organizations can customize standards, create proprietary frameworks, and define organization-specific requirements. Custom standards enable encoding internal security policies as automated compliance checks.

Oracle Database Security Assessment Tool (DBSAT) can also be run from within Oracle Enterprise Manager to automate assessments and centralize operations. This approach enables organizations to:

- Schedule and run DBSAT assessments alongside other management tasks
- Consolidate findings for improved fleet-wide visibility
- Continuously identify, assess, and mitigate risk across databases

Oracle Autonomous Health Framework (AHF) integrates with compliance workflows to strengthen Exadata security by running automated configuration checks and health checks. When managed through Enterprise Manager, organizations gain:

- Centralized oversight of Exadata security posture
- Continuous monitoring and detailed reporting to speed remediation
- Proactive risk mitigation and alignment with regulatory requirements for secure, compliant Exadata environments

Managing configuration drift and consistency with Enterprise Manager

When configurations drift, outages and audit findings often follow. DBLM includes configuration management capabilities designed to keep database configurations consistent and compliant at scale, reducing risk and accelerating issue resolution. DBLM helps organizations:

- Create and maintain approved configuration baselines
- Continuously detect configuration drift and identify what changed, where, and when
- Enforce compliance standards and deploy automated remediation
- Produce detailed reports for audits and operational reviews

DBLM integrates with patching workflows to track expected changes and isolate unexpected drift. It also sends real-time alerts, enabling teams to investigate and resolve issues before they affect availability or security.

By automating monitoring and remediation, DBLM reduces manual effort, improves system reliability, and supports regulatory compliance. Organizations gain clearer audit trails, faster reviews, and lower risk, all managed from a centralized console that provides a consistent view across environments.

Copyright © 2026, Oracle and/or its affiliates

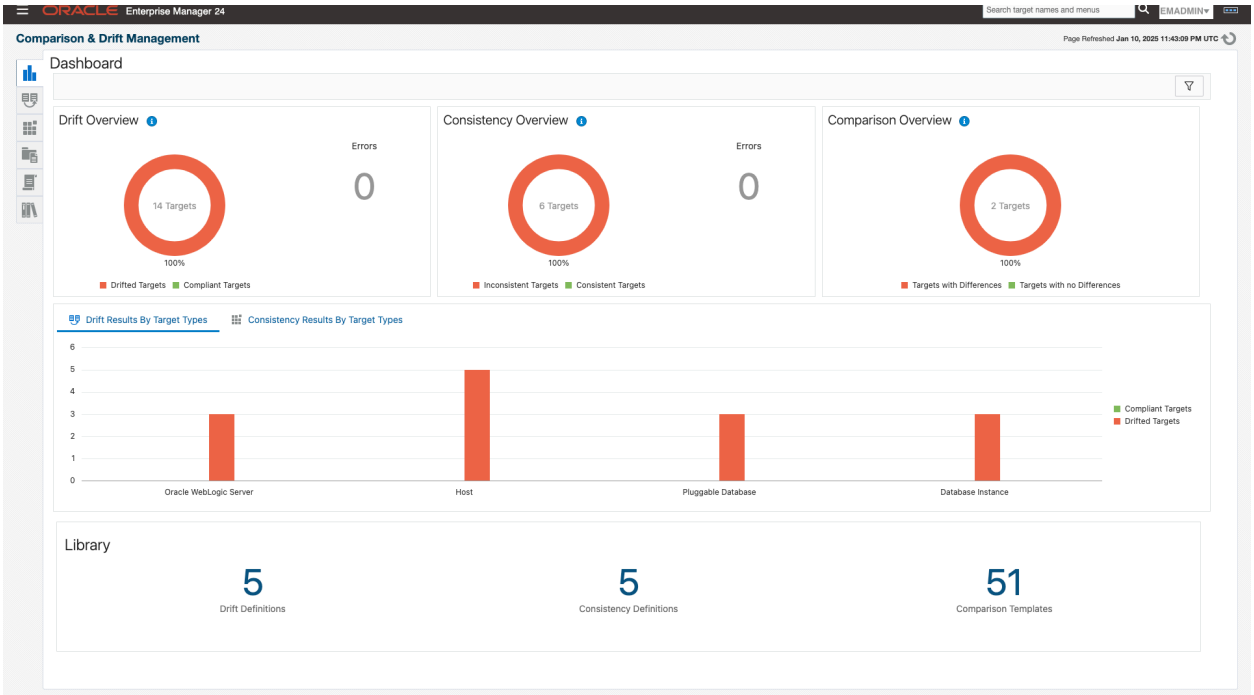


Figure 3-20: Example configuration drift and consistency management dashboard

Discovering and grouping assets

Manual IT asset tracking does not scale. DBLM combines asset inventory and usage tracking to deliver a unified inventory for hardware, databases, middleware, and applications, showing what you own, where it operates, and how you use it. DBLM automates:

- Automatic asset discovery and inventory maintenance
- Version tracking, configurations, and interdependencies
- Resource usage monitoring and licensing agreement alignment

Real time estate visibility supports confident capacity planning, cost optimization, and performance tuning. Lineage and dependency data helps evaluate change impact before action is taken. DBLM works with Oracle Enterprise Manager lifecycle tools to streamline patching, provisioning, and configuration. Organizations can:

- Coordinate patching and configuration changes with accurate inventory data
- Use current usage and configuration details to prioritize remediation
- Generate audit and license review reports within minutes

By automating discovery, monitoring, and remediation, organizations reduce manual effort, improve reliability, and stay compliant, enabling faster decisions and lower risk through a centralized console with consistent cross environment visibility.

Selecting the optimal Database Security assessment tool

Organizations have four primary tools for assessing Oracle Database security: DBSAT, Data Safe, AVDF, and DBLM. Selecting the optimal tool depends on fleet size, deployment architecture, existing tool investments, and operational constraints.

The following decision framework helps match organizational requirements to tool capabilities:

Copyright © 2026, Oracle and/or its affiliates

Database Security Assessment Tool (DBSAT)

Use DBSAT when organizations need straightforward, standalone assessments of individual Oracle Databases without infrastructure investment. DBSAT works well for:

- **Small deployments (1-10 databases)**: Manual quarterly execution with spreadsheet tracking
- **Proof-of-concept assessments**: Understanding Oracle Database security posture before investing in enterprise tools
- **Audit preparation**: Generating comprehensive findings for external audit evidence
- **Disconnected environments**: Air-gapped systems without internet connectivity

If organizations need to assess fleets, monitor configuration drift over time, or automate recurring assessments, transition to Data Safe, AVDF, or DBLM for fleet-scale capabilities.

Oracle Data Safe

Select Data Safe when organizations need security assessments at scale across numerous databases with cloud-based management. Data Safe provides the best fit for:

- **Medium to large deployments (10-1000+ databases)**: Automated fleet-wide assessment and centralized visibility
- **Multi-cloud and hybrid architectures**: Unified management for databases regardless of location (on-premises, Oracle Cloud, AWS, Azure, Google Cloud)
- **Organizations embracing cloud operations**: Preference for SaaS over on-premises infrastructure
- **Integrated security platforms**: Need for assessment alongside activity auditing, SQL Firewall, data discovery, and masking

Data Safe eliminates assessment infrastructure management overhead: no servers to maintain, patch, or capacity plan. Oracle manages service availability, scalability, and security.

Oracle Audit Vault and Database Firewall (AVDF)

Choose AVDF when organizations require on-premises assessment capabilities at scale or when regulatory constraints prohibit cloud-based management services. AVDF suits:

- **Large on-premises deployments**: Hundreds to thousands of databases in customer data centers
- **Regulatory data residency requirements**: Assessment data must remain in specific geographic locations or on-premises
- **Integrated security platforms (on-premises)**: Need for assessment, database activity monitoring, and audit management in a single appliance
- **Stored procedure integrity monitoring**: Detecting backdoor injection into stored procedures

AVDF provides capabilities comparable to Data Safe but with on-premises deployment. Organizations preferring Data Safe's cloud benefits but constrained by regulatory requirements should carefully review those requirements: many regulations allow cloud services with proper contractual protections.

Oracle Enterprise Manager Database Lifecycle Management Pack

Select DBLM when organizations already extensively use Oracle Enterprise Manager for database management and do not require separation of database administration and security responsibilities. DBLM fits:

Copyright © 2026, Oracle and/or its affiliates

- **Existing Enterprise Manager deployments**: Organizations already managing databases through Enterprise Manager
- **Integrated lifecycle management**: Assessment alongside patching, provisioning, configuration management, and monitoring
- **Combined DBA and security roles**: Smaller organizations where DBAs perform security functions
- **Exadata environments**: Leveraging AHF integration for infrastructure security monitoring

DBLM reduces tool sprawl by consolidating assessment into existing management infrastructure. However, organizations with mature security teams preferring dedicated security tools may find Data Safe or AVDF provide better role separation.

Multi-Tool strategies

Organizations need not choose exclusively. Many deploy multiple tools serving different purposes:

- DBSAT for ad-hoc assessments + Data Safe for continuous fleet monitoring
- DBLM for infrastructure configuration monitoring + Data Safe for security-specific assessments
- DBSAT for air-gapped systems + Data Safe for cloud and connected systems

The consistent finding definitions across tools enable comparing results regardless of which tool performed the assessment.

After reviewing configuration assessment tools, the next section turns to privilege analysis, focusing on not only which privileges are granted, but also which privileges are used.

Minimizing the blast radius with privilege analysis

Any security assessment seeks to reduce risk, which means looking beyond configuration settings. A significant share of database risk stems from user accounts and their privileges. Attackers often steal credentials to authenticate and extract data. Removing unnecessary privileges reduces the potential blast radius when an account is compromised.

Determining whether an account has excessive access can be difficult. Organizations can perform entitlement reviews, where a knowledgeable reviewer examines each user's roles and privileges and identifies what to remove. Tools like Data Safe user assessment and AVDF entitlement reviews can help, but managers may worry about disrupting operations, and DBAs may hesitate to revoke access when the consequences are unclear.

A more effective approach is to monitor which privileges accounts actually use and remove the rest. Oracle Database privilege analysis (PA) tracks privilege and role usage for database users and application service accounts over time. PA reports unused privileges and roles based on observed activity, enabling teams to remove unused access with confidence instead of relying on a manager's memory or guesswork. This data driven approach can reduce risk faster while minimizing operational friction.

PA reports show the privileges and roles that users and applications have used and not used. The `DBA_USED_PRIVS` and `DBA_UNUSED_PRIVS` views indicate which privileges and roles have been exercised, and which have not.

Figure 3-21 shows that the `APPS` user has privileges that are not being used: `DROP ANY TABLE`, `ALTER ANY TABLE`, `CREATE TABLE`, and `UNLIMITED TABLESPACE`. Additionally, `DROP ANY PROCEDURE` and `CREATE PROCEDURE` granted through both the APPS and `APPS_PATCHING` roles were not in use. Could this indicate someone granted the `APPS` user a role once for patching and never revoked it?

Copyright © 2026, Oracle and/or its affiliates

S/N	Policy	Grantee	Grantee Type	System Privileges	Grant Path
1	HR Analysis Policy	APPS	USER	DROP ANY TABLE	APPS
2	HR Analysis Policy	APPS	USER	ALTER ANY TABLE	APPS
3	HR Analysis Policy	APPS	USER	CREATE TABLE	APPS
4	HR Analysis Policy	APPS	USER	UNLIMITED TABLESPACE	APPS
5	HR Analysis Policy	APPS	USER	DROP ANY PROCEDURE	APPS,APPS_PATCHING
6	HR Analysis Policy	APPS	USER	CREATE PROCEDURE	APPS,APPS_PATCHING

Figure 3-21: Example report on DBA_UNUSED_PRIVS

Privilege analysis enables deeper investigation. Figure 3-22 reveals interesting details about used roles and privileges concerning specific database objects. Here, we observe that someone granted the APPS user `SELECT ANY TABLE` when, in fact, the user is selecting from specific tables on the HR schema (`DEPARTMENTS`, `JOB_HISTORY`, `COUNTRIES`, `EMPLOYEES`, `LOCATIONS`, `REGIONS`, and `JOBS`). In this scenario, the APPS account's inherent risk can be reduced by revoking the `SELECT ANY TABLE` system privilege, which allows the APPS user to query any table in the database, and replacing it with object privilege grants on only the required tables (for example, `GRANT SELECT` on `HR.DEPARTMENTS` to APPS). If APPS requires access to all tables in the HR schema, `GRANT SELECT ANY TABLE` on HR to APPS can be used, demonstrating schema level privilege grants, a new Oracle AI Database 26ai feature.

S/N	Policy	User Name	Used Role	System Privileges	Object			Grant Path
					Owner	Name	Type	
1	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	DEPARTMENTS	TABLE	APPS
2	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	JOB_HISTORY	TABLE	APPS
3	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	COUNTRIES	TABLE	APPS
4	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	EMPLOYEES	TABLE	APPS
5	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	LOCATIONS	TABLE	APPS
6	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	REGIONS	TABLE	APPS
7	HR Analysis Policy	APPS	APPS	SELECT ANY TABLE	HR	JOBS	TABLE	APPS
8	HR Analysis Policy	APPS	APPS	CREATE SESSION			(null)	APPS
9	HR Analysis Policy	APPS	PUBLIC	(null)	SYS	DBMS_APPLICATI...	PACKAGE	PUBLIC
10	HR Analysis Policy	APPS	PUBLIC	(null)	SYSTEM	PRODUCT_PRIVS	VIEW	PUBLIC
11	HR Analysis Policy	APPS	PUBLIC	(null)	SYS	DUAL	TABLE	PUBLIC

Figure 3-22: Example report on DBA_USED_PRIVS showing object access

Privilege analysis benefits

Static privilege analysis tools like DBSAT and Data Safe provide solid starting points, revealing which roles and privileges users possess. Privilege Analysis extends this by showing which privileges accounts actually use. The combination of static and dynamic analysis enables:

- **Reporting actual privilege utilization**: Understanding operational requirements based on observed behavior
- **Identifying unused privileges and roles**: Finding revocation candidates with confidence
- **Reducing risk through least privilege**: Shrinking blast radius when credentials are compromised
- **Eliminating privilege creep**: Removing accumulated permissions from role changes and temporary grants that became permanent

With deeper understanding of privileges required for applications to operate, database administrators can confidently refine granted permissions to limit unnecessary access. Reducing unnecessary privileges shrinks attack surfaces and diminishes the potential impact of stolen credentials or account misuse.

Copyright © 2026, Oracle and/or its affiliates

Privilege Analysis empowers organizations to transform privilege management from speculation-based guesswork to data-driven decision making. Rather than accepting privilege accumulation as inevitable, organizations can systematically reduce privileges to operational minimums, measurably reducing risk.

Having addressed privilege assessment, the final major assessment dimension is patch management: ensuring databases remain current with security updates.

Managing security patches

Patching represents a critical assessment dimension: identifying which databases lack current security updates and ensuring timely remediation. Every quarter, Oracle delivers database fixes for functional, performance, and security issues discovered through internal testing, customer reports, and external security researcher disclosures.

Security fixes address diverse vulnerability categories:

- **SQL vulnerabilities**: SQL injection weaknesses, vulnerable SQL statements, buffer overflows
- **Client vulnerabilities**: Database client software, JDBC drivers, and third-party code weaknesses
- **Cryptographic weaknesses**: Encryption algorithm vulnerabilities, weak key management, protocol weaknesses
- **Network vulnerabilities**: Network protocol weaknesses, authentication bypasses, session hijacking vectors
- **Remote code execution vulnerabilities**: Flaws enabling attackers to execute arbitrary commands

Organizations require timely patch application to maintain proper security posture. Not applying patches means knowingly accepting the risk of publicly documented vulnerabilities. Once vulnerability details become public through Oracle Critical Patch Updates or security researcher publications, attackers rapidly develop and deploy exploits. The window between disclosure and active exploitation continues shrinking, sometimes measured in hours.

Proactive maintenance through RUs and MRPs

For proactive maintenance, apply the quarterly patch bundle (Release Update / RU) available in the My Oracle Support (MOS) Customer Portal for each Oracle Database software release.

- **RU cadence**: Quarterly, released on the third Tuesday of January, April, July, and October
- **MRP cadence**: Each RU receives up to six Monthly Recommended Patches (MRPs)

Release Updates (RUs) are highly tested, cumulative patch bundles designed to help you avoid known issues. RUs typically include:

- Security fixes
- Regression (bug) fixes
- Optimizer fixes
- Functional fixes (may include limited feature extensions)

Recommendation: Stay current on RUs to minimize exposure to known bugs and security vulnerabilities.

Linux x86-64 guidance: If your databases run on Linux x86-64, and you have validated your application on a given RU and are ready to go live, check for the latest available MRP for that RU and apply it.

Oracle announces quarterly RUs on the Critical Patch Updates, Security Alerts, and Bulletins page each January, April, July, and October.

Copyright © 2026, Oracle and/or its affiliates

Monthly Recommended Patches (MRPs)

Monthly Recommended Patches (MRPs) are provided for Oracle AI Database 26ai deployed on Linux x86-64 to deliver proactive patching between quarterly RUs.

Key characteristics:

- An MRP is a single downloadable merge patch for a specific RU
- It delivers a consolidated set of recommended one-off fixes, including security fixes
- Unlike an RU, an MRP does not change the release number

Coverage and content:

- Oracle publishes MRPs for an RU for the six months following that RU's release
- MRPs include:
    - All fixes documented in "Oracle Database Important Recommended Patches" (MOS note 555.1)
    - All fixes from prior MRPs for the same RU (i.e., cumulative for that RU)

Why use MRPs: MRPs support a more conservative maintenance approach by aggregating critical/security/regression fixes while allowing organizations to defer moving to the next RU for up to six months, reducing security exposure and addressing critical issues in the interim.

By aggregating the latest critical and regression fixes from the prior six months, MRPs enable a more conservative maintenance strategy. Organizations can defer applying a full RU for up to six months while still reducing security exposure and addressing critical issues. Table 3-1 compares Release Updates and MRPs.

Criteria	Release Update (RU)	MRPs
Cadence	Quarterly	Monthly long-term releases Linux x86-64
Zero downtime	RAC Rolling	RAC Rolling
Security fixes	Included	May include CPU Alerts and fixes for vulnerabilities with high CVSS scores
Regression fixes	Included	Included
Proactive functional fixes	Included	Not included
Optimizer plan changes (off by default)	Included	Not included
Minor functional enhancements	Included	Not included
Emergency one-offs	Included	Included
Supported operating systems	All supported platforms	Linux x86-64

Table 3-1: Summary showing main difference between RUs and MRPs

It is strongly recommended that the database and Oracle Grid Infrastructure be kept current by applying RUs to remediate known security vulnerabilities and reduce the likelihood of a successful attack. RUs bundle the latest security, regression, and critical fixes into a single update. Maintaining RU currency also lowers the chance that separate interim one-off patches will be needed, which can create unique software baselines and drive ongoing, costly patch maintenance.

Copyright © 2026, Oracle and/or its affiliates

Patching tools and methodologies

Patching Oracle Database involves complexity that can cause downtime without proper planning and execution. Oracle recommends out-of-place patching because it supports rolling patching, granular changes, and faster rollback capabilities.

Out-of-place patching benefits include:

- **Reduced disruption through rolling patching**: Patch Oracle RAC nodes sequentially while other nodes remain operational
- **Testing before cutover**: Validate patches on standby systems before patching primary production systems
- **Rapid rollback**: Quickly revert to previous software versions if issues arise

Oracle fleet patching and provisioning

Oracle Fleet Patching and Provisioning (FPP) facilitates database fleet lifecycle management through automation, standardization, and out-of-place patching methodologies. FPP drift detection verifies environment compliance, enabling early detection of unexpected changes.

FPP automates routine operations including:

- **Provisioning new clusters and databases**: Automated deployment from standardized templates
- **Installing patched Oracle binaries**: Creating new Oracle Homes with patches pre-applied
- **Patching clusters and databases**: Coordinated patching across RAC clusters
- **Performing database and cluster upgrades**: Version upgrades with minimal downtime

FPP blueprints ensure maintenance executes in correct sequences with minimal business impact. Organizations get repeatable workflows reducing human error and shortening maintenance windows.

With FPP, large organizations can patch hundreds or thousands of databases in single maintenance windows with minimal manual intervention. Smaller organizations benefit equally, leveraging identical automation to save time, enhance consistency, and reduce risk.

FPP availability depends on licensing: use FPP when database targets are licensed for Oracle Real Application Clusters (RAC) or Oracle RAC One Node. For single-instance databases, FPP is available if organizations have licensed the Enterprise Manager Database Lifecycle Management Pack.

Copyright © 2026, Oracle and/or its affiliates

Integrating patching with assessment

Assessment tools integrate with patching workflows to ensure comprehensive security posture management:

- **DBSAT and Data Safe** flag missing patches during security assessments
- **AVDF** tracks patch application across fleets and alerts on inconsistencies
- **DBLM** coordinates assessment with patching schedules, differentiating expected changes from unexpected drift
- **FPP** automates patch deployment validated by assessment tools

This integration creates closed-loop processes: assessment identifies missing patches, patching tools apply updates, subsequent assessment verifies successful remediation.

Organizations building a mature security posture management program typically weave assessment, privilege analysis, and patch management into a continuous workflow: assess, prioritize, remediate, validate, repeat. This steady rhythm helps sustain consistently secure configurations, even as systems, users, and environments are in constant motion.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Database Security Assessment Tool
- Get Started with Oracle Data Safe Fundamentals
- Audit Vault and Database Firewall workshop
- Privilege Analysis workshop
- Oracle Fleet Patching and Provisioning workshop

Copyright © 2026, Oracle and/or its affiliates

Summary

Database Security Posture Management is best thought of as the discipline that turns database security from a hopeful checklist into a repeatable, defensible practice. Modern databases come with hundreds of configuration levers, and it only takes one weak setting to turn into an open door. A strong security assessment approach works like systematic reconnaissance, because attackers are already scanning relentlessly and automatically. Configuration drift is not just an inconvenience either; it can lead to breaches, downtime, brand damage, regulatory exposure, and uncomfortable executive accountability. External expectations raise the bar even further: GDPR, PCI DSS, Sarbanes Oxley, and breach notification statutes all push organizations toward regular assessment and documented verification when personal data is in play. The practical takeaway is clear: build an assessment program that catches misconfigurations early, demonstrates due diligence, and produces evidence of improvement that can stand up to scrutiny over time.

Effective prioritization starts by thinking like an attacker, since adversaries usually follow the path of least resistance. Common reconnaissance tactics include discovery scans, port and service probing, CVE matching, SQL injection testing, and brute force password attacks. Translating those into a defender's first questions helps focus effort where it matters: Are versions current and known? Are default credentials still lurking anywhere? Who has privileged access and is it justified? Is auditing truly covering what it should? Is encryption in place? Are nonproduction copies weaker and easier to exploit? Three realities shape sound decisions: data sensitivity can vary even within a single database, attack surfaces are interconnected so one weak link can cancel out other controls, and risk always has to be weighed against business context and regulatory obligations. Teams are well served by classifying data, understanding the diversity of users and accounts (including service accounts and AI agents), and treating configuration parameters as interdependent so fixes improve security without breaking operations.

When it comes to Oracle-oriented assessment options, four primary approaches fit different scales, deployment models, and constraints: DBSAT, Oracle Data Safe, Oracle Audit Vault and Database Firewall, and the Database Lifecycle Management Pack for Enterprise Manager. DBSAT is the foundational free command-line tool, built around a collector, an offline reporter aligned to best practices and standards, and a discoverer for sensitive data. It produces HTML, spreadsheet, text, and JSON outputs. Findings are organized with consistent rule identifiers, include risk statuses such as Evaluate, and map references back to frameworks. A useful example is how nested role grants can effectively confer DBA-level privileges, prompting the key question: is that powerful access truly warranted? For broader coverage, Oracle Data Safe brings centralized dashboards, scheduling, baselines, drift reporting, aggregation, and workflows for handling acceptable risk. For organizations that prefer on-premises deployments, Oracle Audit Vault and Database Firewall adds capabilities like stored procedure change tracking and entitlement snapshot comparisons to reveal privilege creep.

Posture management also extends beyond configuration, because privilege and patching practices often define the real blast radius. Privilege Analysis acts as a dynamic complement to static reviews, using views such as DBAUSEDPRIVS and DBAUNUSEDPRIVS to show what is actually exercised over time. That visibility supports least-privilege decisions and the removal of access that has quietly accumulated. On patching, the warning is straightforward: delays turn public vulnerability disclosure into rapidly exploitable risk. Quarterly Release Updates and Monthly Recommended Patches for specific platforms provide structure, and operational techniques like out-of-place patching, testing before cutover, and faster rollback help keep patching from becoming a high-drama event. In practice, many organizations combine tools and build a steady operational rhythm: assess, prioritize, remediate, and validate, then repeat with discipline.

The next chapter on Sensitive Data Discovery promises to spotlight where risk hides, turning vague suspicion into clear, actionable clarity about what sensitive data exists and where it lives. Keep reading to follow the thread into the data itself, where a single overlooked column can change the entire security story.

Copyright © 2026, Oracle and/or its affiliates

Chapter 4

Sensitive Data Discovery

Introduction: The foundation of data protection

Security without visibility turns into security theater fast. Many organizations spend millions on encryption, firewalls, and access controls, but still struggle with a basic question: Where is our sensitive data? This chapter helps remove that uncertainty by walking through a systematic approach to discovering sensitive data in Oracle Database environments, which is the foundation for every other security control covered in this book.

In practice, discovery needs to satisfy two priorities: regulatory compliance and risk management. Regulations like GDPR, PCI DSS, HIPAA, and India's DPDPA call for specific protections for different kinds of data, but those requirements do not mean much if teams cannot first find the data in scope. On the risk side, the likely impact of a breach should drive how strong the controls need to be. Discovery makes that possible by helping organizations focus strict safeguards where the risk is highest, while avoiding over investing in areas where the data is lower risk.

The chapter first explains why sensitive data matters for business and security outcomes, then highlights why AI applications increase the urgency. It defines sensitive data, outlines the methods used to locate it, and then provides detailed guidance on Oracle's four discovery tools: Database Security Assessment Tool (DBSAT), Oracle Data Safe, Audit Vault and Database Firewall (AVDF), and Enterprise Manager's Application Data Modeling.

Why sensitive data matters

Databases are where modern enterprises keep the information that keeps everything moving: customer transactions, product designs, financial records, research findings, and operational metrics. That data enables organizations to serve customers well, make smarter decisions, reduce costs, and preserve competitive advantage. The challenge is that much of it is also sensitive or personal, such as customer records, payment details, healthcare information, employee compensation, and proprietary business metrics.

The business value and business risk of data

Because sensitive data is valuable to the business, it is also valuable to threat actors. As sensitivity increases, so does its appeal for theft, misuse, or ransom. A breach involving customer payment information can hit both revenue and reputation. Exposed intellectual property can chip away at competitive positioning. Compromised healthcare records can lead to regulatory penalties and lawsuits. The equation is simple: high value data attracts high intent adversaries and results in high impact consequences when exposed.

That is why sensitive data protection needs to be a practical priority, not just a policy statement. When security teams know what is sensitive and where it lives, they can apply stronger controls exactly where they will matter most. This focused approach helps protect customers, safeguard intellectual property, meet compliance obligations, and keep the business running.

Discovery as the first step in defense

The security axiom holds true: organizations cannot protect what they cannot see. The time, budget, and effort invested in database security should reflect both regulatory obligations and business risk exposure. The higher the potential impact of a compromise, the stronger the controls required to reduce risk to acceptable levels.

Discovery is typically the foundation for most subsequent security measures. Before implementing encryption, masking, access controls, or audit policies, security teams must answer three foundational questions:

- What types of sensitive data exist in the environment?
- How much sensitive data is present across all databases?
- Where exactly does this data reside—which databases, schemas, tables, and columns?

Without precise answers to these questions, security efforts become inefficient at best and ineffective at worst. Organizations may over-invest in protecting low-risk data while leaving critical information exposed. They may fail

Copyright © 2026, Oracle and/or its affiliates

audits because they cannot demonstrate where regulated data resides. They may deploy masking solutions that break application dependencies because referential relationships were not mapped. Discovery eliminates these blind spots by providing the visibility required for risk-based security decisions.

The AI imperative: New urgency for discovery

The rapid adoption of artificial intelligence capabilities has fundamentally changed the risk landscape for sensitive data. Organizations are deploying AI-powered applications, implementing Retrieval-Augmented Generation (RAG) systems, and connecting large language models directly to production databases to deliver intelligent user experiences. While these technologies create substantial business value, they also introduce new pathways for unintended data exposure.

AI applications and the exposure surface

RAG systems retrieve relevant information from databases and pass it to AI models as context for generating answers. When access controls are weak or data awareness is limited, these systems can accidentally reveal sensitive information to the wrong people. Imagine an employee asking a routine question about company benefits, and the RAG system returning executive compensation details because those records sit in the same HR database. Without discovery and classification, the AI application has no practical way to tell what is meant to be shared versus what must remain confidential. That gap turns routine questions into accidental disclosure paths.

The same concern shows up during AI fine tuning and training on enterprise data. If the training dataset contains personally identifiable information (PII) or protected health information (PHI), that regulated data can end up embedded in the model weights. After that, it may be possible to pull the information back out through prompt injection attacks, model inversion techniques, or even straightforward inference from model outputs. Discovery reduces this risk by flagging sensitive data before it enters the training pipeline, so organizations can exclude it, anonymize it, or synthetic generate data as needed.

Model Context Protocol and database integration

The introduction of the Model Context Protocol (MCP) has accelerated the integration timeline between AI applications and enterprise databases. MCP provides a standardized framework for connecting AI models-such as Claude, GPT-4, or custom models-directly to databases and other enterprise resources through uniform APIs. While this simplifies integration and enables rapid development of AI-powered applications, it also means databases can be exposed to AI systems faster than security teams can implement appropriate controls.

When organizations deploy an MCP server that connects an AI model to their database, discovery becomes the first line of defense. Security teams must identify exactly what data the AI can access before users begin querying it. Discovery results inform several critical decisions:

- Which schemas and tables should be excluded entirely from the MCP context
- Where to implement row-level security policies that MCP queries must respect
- Which columns require masking or tokenization within AI responses
- What fine-grained access controls should apply based on user identity and role

The MCP server must use a restricted service account following the principle of least-privilege. This proactive approach allows organizations to provide AI systems with the data access they need for useful functionality while preventing them from retrieving information they should never see. Discovery transforms AI integration from a potential security liability into a calculated, controlled capability.

The same discovery methodologies used for traditional database security-scanning column names, analyzing data patterns, mapping referential relationships-apply equally to preparing data for AI use cases. This consistency allows security teams to leverage existing discovery investments while addressing emerging AI-related risks.

Copyright © 2026, Oracle and/or its affiliates

Defining sensitive data

Sensitive data is any information that, if exposed or misused, could harm individuals, organizations, or both. A practical test for sensitivity is: Would exposing this data cause reputational damage, regulatory penalties, financial loss, or personal harm if published openly? If the answer is yes, the data qualifies as sensitive and requires protection from unauthorized access whether by external attackers or internal users exceeding their authorization.

Organizations typically classify sensitive data into a few broad categories, based on the type of information involved and the regulations that apply. Figure 4-1 highlights common categories and examples of representative data types within each category.

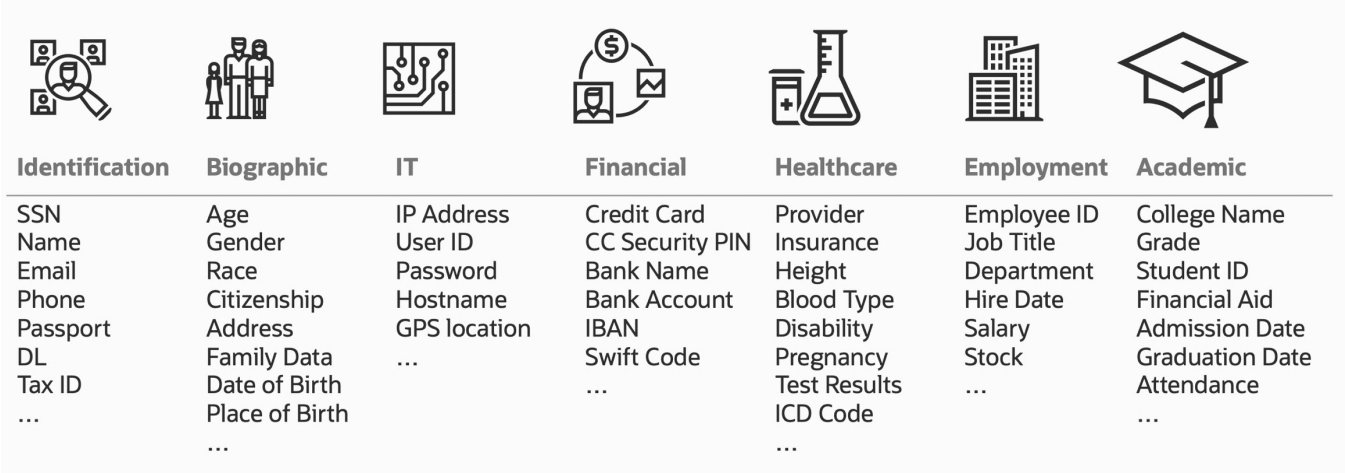


Figure 4-1: Common sensitive data categories and representative types

Organizations rarely deal with only one category of sensitive data. Different regulations spotlight different data types. For example, PCI DSS prioritizes payment account data, Sarbanes Oxley focuses on financial records, and GDPR emphasizes identification and biographic data. Understanding which categories apply in each environment helps organizations focus controls where they matter most.

Discovery methodologies: How to find sensitive data

Oracle's discovery tools employ two complementary approaches to locate sensitive data within databases. Most effective discovery strategies combine both methodologies to maximize coverage and accuracy.

Metadata-based discovery

Metadata-based discovery reviews the structural details of database objects, specifically column names and column comments, rather than the actual data values stored in tables. It looks for keywords, patterns, and naming conventions that may indicate sensitive information.

How it works: The discovery tool queries the data dictionary to retrieve column names and associated comments across all accessible schemas. It then applies pattern matching using regular expressions or keyword lists. For example, column names containing `SSN`, `CREDIT_CARD`, `SALARY`, `DOB` (date of birth), or `PASSWD` triggers a match. Multilingual keyword libraries extend coverage to databases using non-English naming conventions.

Advantages: Metadata scanning is extremely fast and lightweight because it reads only data dictionary tables rather than potentially billions of rows within user tables. It consumes minimal system resources and can be executed against production databases with negligible performance impact. The approach works well in environments where developers follow consistent naming conventions for sensitive columns.

Limitations: Metadata-based discovery depends entirely on meaningful naming. If developers use generic or cryptic column names such as `FIELD_17`, `ATTRIBUTE_A`, or `DATA_VALUE` to store credit card numbers or Social Security

Copyright © 2026, Oracle and/or its affiliates

numbers, metadata scanning will miss that sensitive data entirely. Additionally, this approach generates false positives when column names suggest sensitivity, but the actual content is innocuous (for example, a column named `CUSTOMER_NAME` that stores product category names).

Pattern-based discovery (data sampling)

Pattern-based discovery analyzes the actual values stored within database columns by sampling row data and applying pattern recognition algorithms. This approach identifies sensitive data based on content rather than metadata.

How it works: The discovery tool executes queries to sample a subset of rows from each table (typically a configurable percentage or fixed count). It then examines the retrieved values using various detection techniques:

- Regular expressions to match formats like credit card numbers (16 digits with optional separators), email addresses, phone numbers, or Social Security numbers
- Algorithmic validations such as checksum validation for national ID formats or credit card numbers
- Statistical analysis to detect patterns like consistent length or structure across sampled values
- Predefined templates for country-specific identifiers (e.g., Canadian Social Insurance Numbers, UK National Insurance Numbers, Indian Aadhaar numbers)

Advantages: Pattern-based discovery finds sensitive data regardless of how columns are named, making it effective in environments with inconsistent or poor naming conventions. It provides higher confidence results because detection is based on actual content rather than inferred meaning. The approach also discovers sensitive data that was added to generic columns after database design (for example, when a `NOTES` text field later begins storing customer phone numbers).

Limitations: Data sampling requires more computational resources and execution time than metadata scanning because the tool must read and analyze actual table rows. In very large databases (multiple terabytes), complete sampling of all tables may take hours or days. Organizations must also consider data residency concerns-if regulatory or policy requirements prohibit copying actual sensitive values outside the database, sampling results must be handled carefully or disabled entirely.

Combined methodology for optimal results

The most effective discovery strategy is to combine both approaches. When a column name suggests sensitivity and pattern analysis confirms that content matches expected formats, confidence in the detection increases substantially. For example, a column named `CREDIT_CARD_NUMBER` that also contains 16-digit values passing algorithmic validation is almost certainly storing payment card data.

Conversely, when metadata and content analysis disagree, human review can adjudicate. A column named `ACCOUNT_NUMBER` containing random alphanumeric values may be a product `SKU` rather than a bank account. A column named `DATA_1` containing 9-digit values matching Social Security number patterns may require investigation to determine whether those values are actual SSNs or other identifiers.

All four Oracle discovery tools discussed in this chapter employ one or both of these methodologies, with varying levels of sophistication and configurability.

Choosing the right solution

Oracle provides four complementary tools for discovering sensitive data in Oracle Database environments, each with distinct capabilities, deployment models, and ideal use cases. Table 4-2 summarizes the four tools and their discovery method. Organizations often combine these tools, for example, using DBSAT for a fast initial assessment and then adopting Data Safe for ongoing monitoring and management.

Copyright © 2026, Oracle and/or its affiliates

Tool	Primary Use Cases	Discovery Methodology
DBSAT	Quick assessments, audit preparation, lightweight reporting	Metadata (column names and comments)
Oracle Data Safe	Comprehensive discovery, continuous monitoring, cloud-native management	Metadata + pattern-based data sampling
Enterprise Manager Data Masking and Subsetting	EM Pack to protect sensitive data by masking and reduce storage requirements by subsetting data	Metadata + pattern-based data sampling
Audit Vault and Database Firewall	Activity monitoring, compliance reporting, firewall policy configuration	Metadata (column names and comments)

Table 4-2: Oracle Sensitive Data Discovery tools support different use cases

The following sections examine each tool in detail, providing technical architecture, capabilities, use cases, and practical guidance for implementation.

Discovering sensitive data using DBSAT

Database Security Assessment Tool (DBSAT) is a command-line utility designed for rapid security posture assessment and sensitive data discovery. As introduced earlier, DBSAT provides comprehensive assessment by combining configuration analysis, privilege analysis, and sensitive data discovery in a single lightweight tool.

Overview and architecture

DBSAT operates as a standalone utility that connects to a target database using standard Oracle Net protocols. It requires minimal privileges, typically a read only account with access to data dictionary views, and can be executed from any system with network connectivity to the database. The tool gathers metadata on database configuration, user privileges, and sensitive column names, then processes the information locally to produce comprehensive reports.

Discovery capabilities

DBSAT identifies sensitive data by scanning column names and comments against an extensible library of sensitive data definitions. It ships with prebuilt patterns for common sensitive types in English, and additional libraries are available for seven major European languages: Dutch, French, German, Greek, Italian, Portuguese, and Spanish. In multilingual environments, organizations can deploy multiple language libraries at the same time.

Rather than sampling actual data values, DBSAT estimates the quantity of sensitive data by reading table statistics from the data dictionary, specifically the row counts maintained by the database optimizer. This method delivers volume estimates without requiring access to the underlying table data, which makes DBSAT a practical fit for environments with strict data residency requirements.

DBSAT groups discovered sensitive columns into predefined categories such as Biographic Information, Financial Information, Health Information, Identification Information, IT Information, and Job Information. Each category includes subcategories to support more precise classification. For instance, the Financial Information category includes subcategories for Bank Data, Card Data, and Tax Data.

Report formats and content

DBSAT generates detailed sensitive data assessment reports in both HTML and spreadsheet formats, making results accessible to both technical and business audiences. The reports include several key sections:

Copyright © 2026, Oracle and/or its affiliates

Summary statistics: High-level counts of sensitive tables, columns, and estimated rows, organized by category and subcategory. This overview helps stakeholders quickly understand the scale and distribution of sensitive data across the database. Table 4-3 shows a sample summary section.

Sensitive Category	# Sensitive tables	# Sensitive columns	# Sensitive rows
BIOGRAPHIC INFO - ADDRESS	9	36	6307209
BIOGRAPHIC INFO – EXTENDED PII	2	2	2000
FINANCIAL INFO – BANK DATA	2	2	830
FINANCIAL INFO – CARD DATA	7	7	3235
HEALTH INFO – PROVIDER DATA	1	1	149
IDENTIFICATION INFO – NATIONAL IDS	2	6	2000
IDENTIFICATION INFO – PERSONAL IDS	3	3	405
IT INFO – USER DATA	13	15	13228
JOB INFO – COMPENSATION DATA	10	12	3380
JOB INFO – EMPLOYEE DATA	8	16	406
JOB INFO – ORG DATA	5	6	278
Total	**29**	**132**	**8617644**

Table 4-3 Example DBSAT sensitive data summary by category

Each sensitive category has a preconfigured risk level (high, medium, or low), and this risk rating can be customized if desired. The report then recommends how to protect sensitive data based on its assigned risk level.

Copyright © 2026, Oracle and/or its affiliates

Figure 4-2 illustrates several recommendations for safeguarding high risk data, such as personal health information.

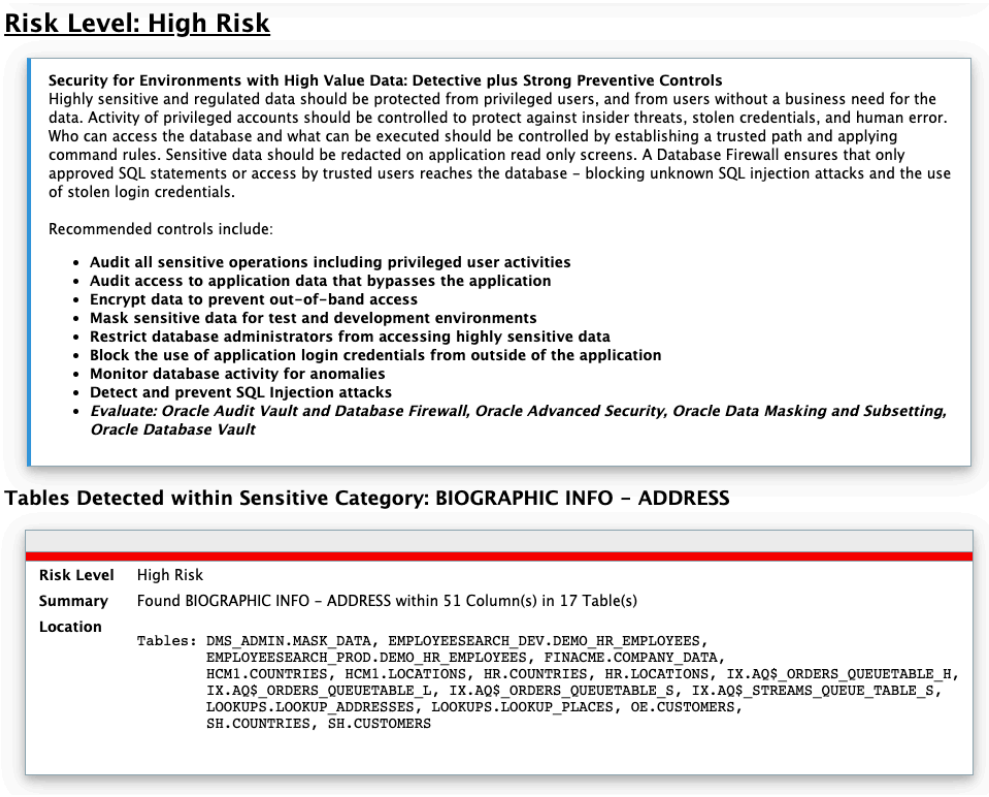


Figure 4-2: Example recommendations for protecting sensitive data

Risk-based recommendations: Each sensitive category has a preconfigured risk level that drives recommended protections. For high-risk data such as health information or payment card data, DBSAT recommends implementing encryption, restricting access through Database Vault realms, enabling audit policies, and considering data masking for nonproduction environments. Medium-risk data recommendations typically include access controls and auditing. Low-risk data may require only basic access management.

Detailed findings: Schema-level, table-level, and column-level details provide the specific locations of discovered sensitive data. This granular information enables security teams to implement targeted controls.

Use cases and deployment scenarios

DBSAT excels in several scenarios:

- **Rapid initial assessment:** When organizations need a quick inventory of sensitive data without deploying complex infrastructure, DBSAT provides results within minutes to hours depending on database size.
- **Audit preparation:** Before regulatory audits, DBSAT reports demonstrate where sensitive data resides and what controls are in place, helping organizations prepare documentation and address gaps.
- **Pre-migration analysis:** Before migrating databases to cloud environments or upgrading to new versions, DBSAT helps teams understand what sensitive data will be affected and plan appropriate protections.
- **Air-gapped environments:** Because DBSAT operates as a standalone utility without external dependencies, it works well in isolated or high-security environments where cloud-based tools cannot be deployed.

Copyright © 2026, Oracle and/or its affiliates

Discovering sensitive data using Data Safe

Oracle Data Safe is a comprehensive, cloud-native database security service that provides unified management for security assessment, user assessment, sensitive data discovery, data masking, activity auditing, and security monitoring. Data Safe operates as a managed service within Oracle Cloud Infrastructure but can monitor and protect databases across multiple environments including Oracle Cloud Infrastructure, third-party clouds (AWS, Azure, Google Cloud), and on-premises.

Overview and architecture

Data Safe's architecture separates the control plane (management console and services running in Oracle Cloud) from the data plane (target databases being monitored). Communication between Data Safe and target databases occurs through secure, outbound-initiated connections, ensuring that no inbound network paths need to be opened to protected databases. This design accommodates stringent network security requirements common in enterprise environments.

Discovery capabilities

Data Safe provides the most comprehensive discovery capabilities among Oracle's tools, combining metadata scanning with sophisticated pattern-based analysis of actual data values. The service includes pre-built sensitive type definitions for over 150 data types, including country-specific identifiers for Brazil, Canada, France, Germany, India, Italy, México, Netherlands, Portugal, Spain, the United Kingdom, and the United States.

Organizations can also define custom sensitive types that combine:

- Column name patterns (regular expressions)
- Column comment patterns (regular expressions)
- Data value patterns (regular expressions)
- Row sample size for validation
- Minimum match percentage thresholds

For example, a financial institution might create a custom sensitive type for internal account numbers that follow a proprietary format, specifying both the expected column naming convention and the data pattern to validate against actual values.

When Data Safe executes a discovery job, it scans column names and comments first to identify candidates, then optionally samples actual data values to confirm matches. This two-stage approach balances thoroughness with efficiency. Organizations concerned about copying sensitive values outside the database can disable data sampling entirely and rely solely on metadata scanning.

Copyright © 2026, Oracle and/or its affiliates

Figure 4-3 shows a sample data discovery report generated by Data Safe.

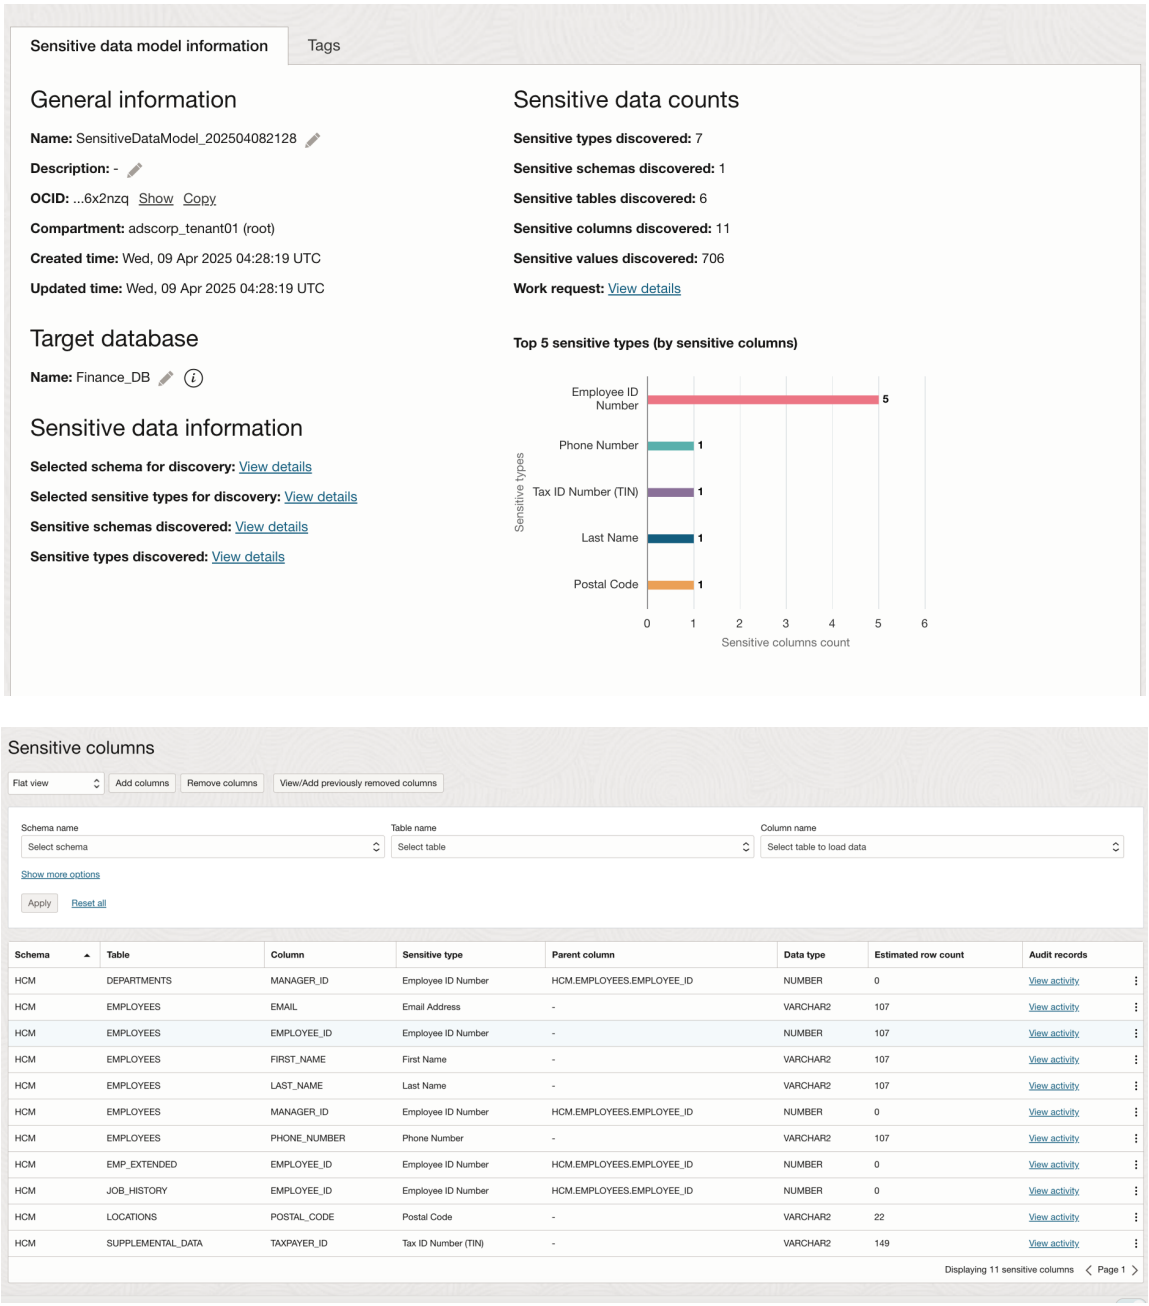


Figure 4-3: Example data discovery report in Oracle Data Safe

Data Safe scans column names, comments, and data values to help users quickly spot sensitive attributes and understand their context. It also shows how tables relate to each other by identifying relationships between primary and foreign key columns.

In addition to what is stored in the data dictionary, Data Safe helps uncover referential relationships defined at the application layer. This provides a more complete view of how data moves through the environment, enabling teams to plan masking and testing strategies without breaking dependencies.

From the Data Safe console, users can monitor activity and drill down into audit records for the data being assessed (refer to rightmost column of Figure 4-3). That makes it easier to investigate who accessed what, when, and how.

Copyright © 2026, Oracle and/or its affiliates

Sensitive data models in Data Safe

Securing sensitive data begins with understanding what exists and how it connects. Data Safe provides that visibility by transforming discovery results into a reusable sensitive data model that can be refined and improved over time.

Data Discovery reports summarize sensitive tables, columns, and values, and they provide detailed information for each sensitive column. Data Safe categorizes these columns by sensitive type and combines that classification with primary and foreign key relationships to build a sensitive data model. That model is revisited later in the book during the discussion on data masking, where those relationships help ensure masking is applied consistently without breaking dependencies.

The model can also store optional metadata, such as sample data and estimated row counts, to help teams understand the scale of each sensitive type across target databases, as shown in Figure 4-3. Because many organizations prefer not to copy actual values outside the database, Data Safe defaults to avoiding the collection of sample data. For a clear view of exactly what Data Safe collects from databases, refer to the reference section of the Data Safe Administrator's guide.

As the target database evolves, incremental updates can be used to keep the sensitive data model up to date. This approach helps refine the model over time by adding or removing sensitive columns and maintaining the referential relationships between those columns. Discovery can also begin with a narrower scope, such as scanning a single sensitive category or a set of selected schemas, and then gradually expand through incremental discovery to broaden coverage as the model matures.

For portability, a sensitive data model can be downloaded, updated offline, and uploaded to the same or a different Oracle Data Safe region. Each sensitive data model is associated with one target database at a time, and that association can be changed whenever needed.

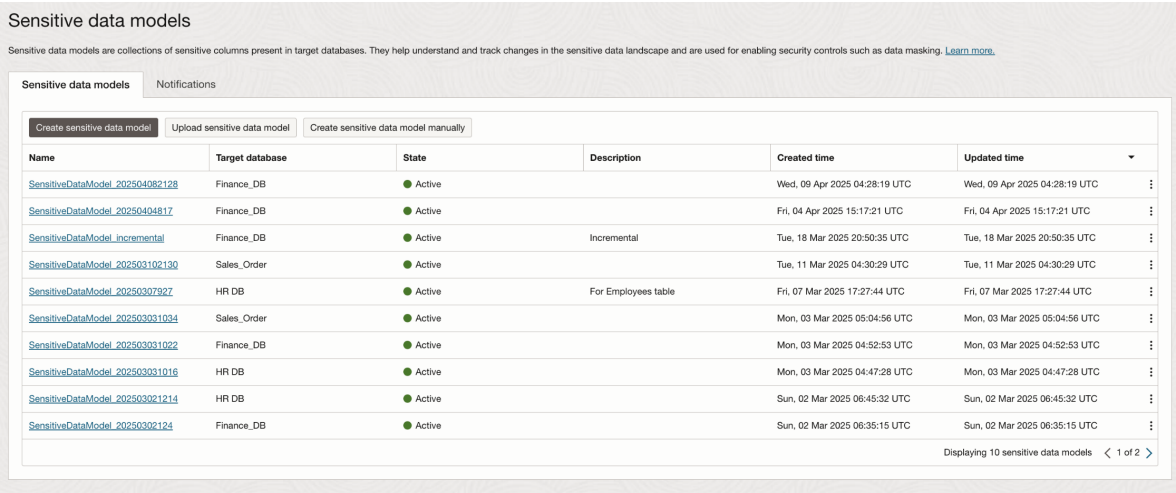


Figure 4-4: Example data discovery sensitive data models

Sensitive data models help inform the design of other security controls, such as data masking, auditing, Database Vault policies, and fine-grained access controls. For example, a masking policy can be defined using a sensitive data model to mask sensitive data in nonproduction database clones, such as test and development instances.

Use cases and deployment scenarios

Oracle Data Safe is the recommended solution for organizations seeking comprehensive, ongoing discovery and management of sensitive data, particularly in the following scenarios:

Copyright © 2026, Oracle and/or its affiliates

- **Cloud-native databases:** For Oracle Autonomous Database, Oracle Database Cloud Service, or Oracle Exadata Cloud Service deployments, Data Safe provides native integration without requiring additional infrastructure.
- **Hybrid and multi-cloud:** Organizations with databases distributed across Oracle Cloud, AWS, Azure, Google Cloud, and on-premises can manage discovery centrally through a single Data Safe instance.
- **Continuous monitoring:** When organizations require ongoing visibility into sensitive data as databases evolve, Data Safe's scheduled discovery jobs and incremental updates maintain current models.
- **Masking workflows:** Organizations that mask sensitive data in nonproduction environments benefit from Data Safe's tight integration between discovery models and masking policies.
- **Compliance reporting:** When auditors require evidence of where regulated data resides and how it is protected, Data Safe provides comprehensive reports with drill-down capabilities to supporting evidence.
- **API-Driven automation:** Organizations implementing DevSecOps practices can use Data Safe's REST APIs to automate discovery workflows, integrate results into security information systems, and trigger downstream security controls.

Data residency and privacy considerations

Because Data Safe is a cloud service, organizations sometimes raise concerns about what data is collected and where it is stored. By default, Data Safe collects only metadata about sensitive columns like names, types, categories, and estimated counts, not actual sensitive values. Organizations can optionally enable sample data collection for validation purposes, but this feature is disabled by default.

All information collected by Data Safe is stored within the Oracle Cloud Infrastructure region where the Data Safe instance is provisioned. Organizations subject to data residency requirements can select appropriate regions during setup. For detailed information about what data that Data Safe collects and stores, Oracle provides comprehensive documentation in the Data Safe Administrator's Guide.

Discovering sensitive data using Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall (AVDF) is an integrated solution that combines comprehensive audit data collection, centralized audit analysis, and network-based SQL firewall capabilities. While AVDF's primary purpose is monitoring database activity and enforcing security policies, it also includes sensitive data discovery functionality that enhances its monitoring and reporting capabilities.

overview and architecture

AVDF consists of two main components:

- **Audit Vault Server:** A hardened Linux appliance (physical or virtual) that collects audit data from multiple database sources, normalizes the data into a central repository, and provides reporting and analysis tools.
- **Database Firewall:** A network-inline or monitor-mode appliance that inspects SQL traffic between applications and databases, enforcing policy-based controls and generating additional audit events.

Discovery capabilities

AVDF's discovery functionality focuses on metadata-based scanning of column names and comments, like DBSAT. Organizations initiate discovery by scheduling a Sensitive Object Discovery job through the Audit Vault console. This job queries the data dictionary of registered target databases to identify columns matching predefined or custom sensitive type patterns.

Copyright © 2026, Oracle and/or its affiliates

Unlike Data Safe and Enterprise Manager, AVDF does not sample actual data values for pattern matching. This design decision reflects AVDF's primary use case: enabling the monitoring and policy systems to recognize when sensitive data is being accessed, without requiring AVDF itself to access that sensitive data.

Discovered sensitive objects populate AVDF's database inventory, where they are tagged with sensitive types and categories. This metadata enriches several AVDF capabilities:

- Audit reports can filter and highlight activity on sensitive data
- Firewall policies can reference sensitive types when defining allow/block rules
- Compliance reports can demonstrate coverage of sensitive data access monitoring
- Alerts can trigger when specific sensitive data is accessed by privileged users or during non-business hours

Data privacy reports

AVDF provides specialized reports that leverage discovered sensitive objects to answer compliance and security questions:

Sensitive data inventory: Lists all discovered sensitive objects across monitored databases, organized by target, schema, table, column, and sensitive type. This report demonstrates to auditors that the organization knows where sensitive data resides.

Activity on sensitive data: Shows detailed access patterns for sensitive objects, including user identity, session details, SQL statements executed, and timestamps. This report supports investigations and compliance requirements to demonstrate that sensitive data access is monitored and appropriate.

Activity on sensitive data by privileged users: A filtered view focusing specifically on access by users with elevated privileges (DBAs, application admins, security officers). Because privileged users often have broad access rights, monitoring their activity on sensitive data provides an additional layer of oversight to detect potential misuse.

These reports support compliance requirements for regulations such as GDPR (demonstrating data access monitoring), PCI DSS (logging access to cardholder data), HIPAA (tracking access to protected health information), and SOX (monitoring access to financial data).

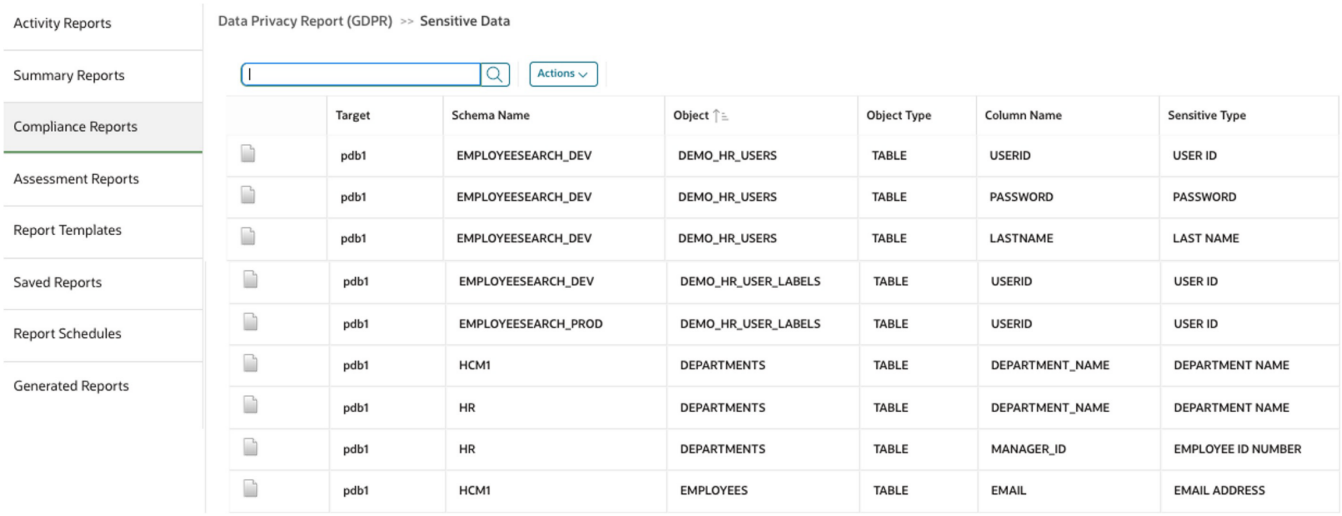


Figure 4-5: Example AVDF sensitive data report

Copyright © 2026, Oracle and/or its affiliates

Integration with Database Firewall policies

One of AVDF's unique capabilities is using discovered sensitive objects as criteria within Database Firewall policies. For example, an organization might create a firewall policy that:

- Blocks any SQL statement attempting to select from tables containing payment card data, except from approved application IP addresses
- Generates alerts when administrators query tables containing customer personal information outside of business hours
- A Firewall rule can be configured to trigger an alert when a single query returns more rows than the predefined 'expected threshold' for sensitive data.

This integration transforms discovery from a purely informational exercise into an active component of the security enforcement architecture.

Use cases and deployment scenarios

AVDF's discovery capabilities serve organizations that prioritize activity monitoring and real-time policy enforcement:

- **Compliance programs:** Organizations subject to audit requirements can use AVDF's sensitive data reports to demonstrate comprehensive monitoring of regulated data access.
- **Privileged user oversight:** When concerns exist about insider threats or privileged user activity, AVDF tracks access to sensitive data by administrators and generates alerts for suspicious patterns.
- **Network-layer protection:** In environments where application-layer or database-layer controls are insufficient, Database Firewall provides an additional enforcement point that can block or mask sensitive data based on discovery results.
- **Centralized audit:** Organizations with heterogeneous database environments (Oracle, SQL Server, PostgreSQL, MySQL, MongoDB) can use Audit Vault to collect and correlate audit data centrally, with sensitive object discovery providing context across all platforms.

Discovering sensitive data using Enterprise Manager

Oracle Enterprise Manager is a comprehensive systems management platform that provides monitoring, administration, and lifecycle management for Oracle environments. With Oracle Enterprise Manager, the data discovery module can be used to quickly identify sensitive columns, making it easier to prioritize the right controls. The **Data Masking and Subsetting** pack within Enterprise Manager includes a data discovery module that identifies sensitive columns and models application data relationships.

Overview and architecture

Enterprise Manager's discovery functionality is particularly relevant for organizations that:

- Operate primarily on premises with requirements to keep all security data within their own data centers
- Already use Enterprise Manager for database administration and operations management
- Need tight integration between discovery, masking, and other database lifecycle operations
- Manage complex application environments where understanding data relationships is critical

Data discovery, especially the sensitive types feature is part of Oracle Data Masking and Subsetting, an Enterprise Manager pack. Licensing for data discovery is also included with database security options such as Advanced

Copyright © 2026, Oracle and/or its affiliates

Security, Database Vault, and Label Security. Like Oracle Data Safe, Enterprise Manager examines column names, comments, and data values to discover sensitive columns.

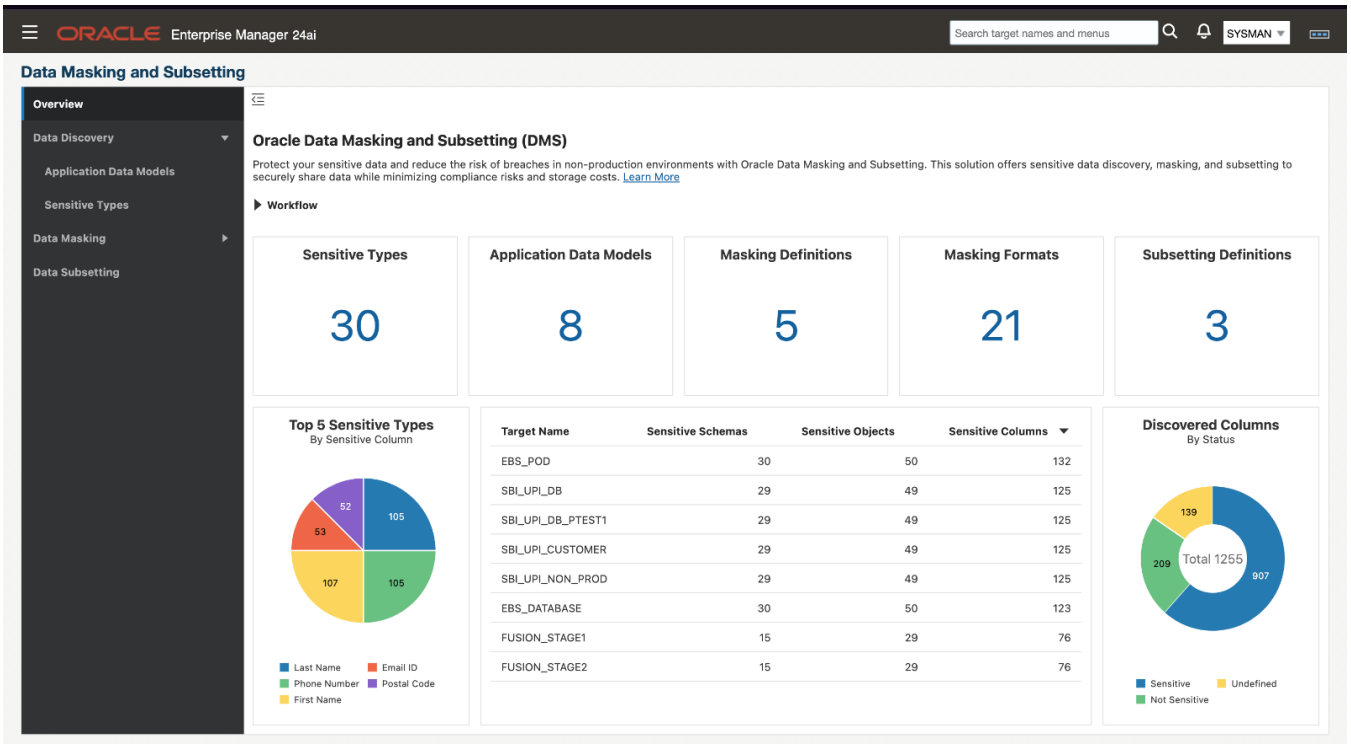


Figure 4-6: Example sensitive data discovery in Enterprise Manager

Referential relationship discovery

Data discovery uses its sensitive types feature to perform pattern matching and identify sensitive columns. Schedule a sensitive data discovery job to scan the database for data that might need additional protection. One of Enterprise Manager's distinguishing capabilities is its sophisticated analysis of referential relationships between database objects. The tool automatically discovers relationships defined through foreign key constraints in the database schema. Additionally, it can detect application-level relationships that are not explicitly defined in the database but exist in application logic.

Copyright © 2026, Oracle and/or its affiliates

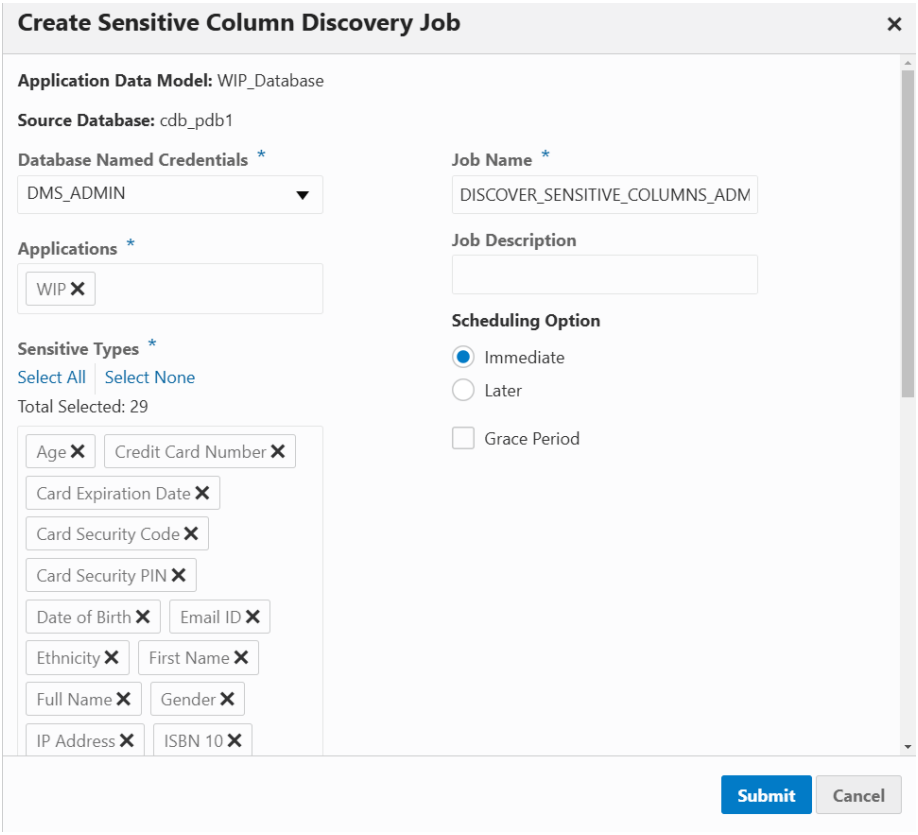


Figure 4-7: Example screen illustrating the scheduling of a sensitive column discovery job

Application-level relationship discovery works by analyzing data patterns: if a column in one table contains values that consistently match values in another table's column, the tool infers a likely relationship. Administrators can review these inferred relationships and confirm them for inclusion in the Application Data Model.

Application data models

The discovery process culminates in an Application Data Model (ADM)—a comprehensive representation of sensitive data and its relationships within a specific application context. The ADM contains:

- All discovered sensitive columns with assigned types and categories
- Database-defined referential relationships (foreign keys)
- Application-level referential relationships (inferred or manually defined)
- Parent-child hierarchies showing data dependencies
- Metadata about discovery job execution and last update times

Application Data Models are stored in the Enterprise Manager repository and can be reused across multiple security operations:

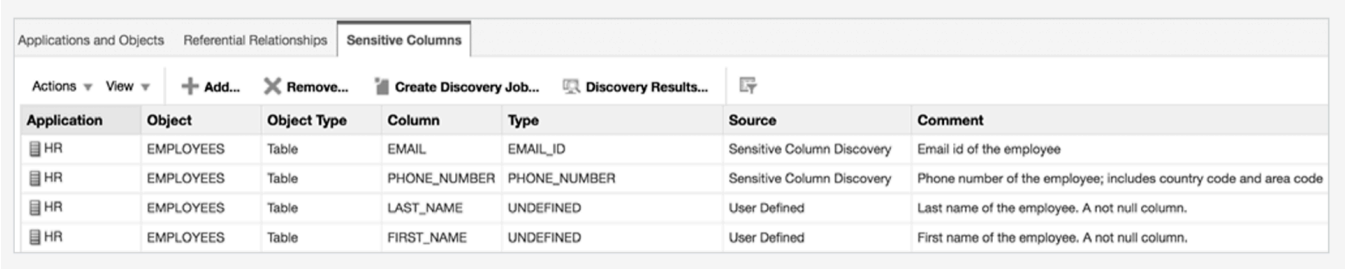
Data Masking: The ADM serves as input to Enterprise Manager's data masking engine, which uses relationship information to maintain consistency when masking related columns across tables.

Database Vault policies: Security teams can use ADM information to define Database Vault realms that protect sensitive schemas or tables, command rules that restrict operations on sensitive data, and factors that provide context-aware authorization.

Copyright © 2026, Oracle and/or its affiliates

Test Data Management: When creating nonproduction database clones, organizations can use ADMs to identify which data must be masked or subsetted to reduce storage requirements while maintaining application functionality.

The discovered sensitive columns can be reviewed and, when needed, additional columns can be manually added to the list. Figure 4-8 shows a list of discovered as well as user-defined sensitive columns.



Application	Object	Object Type	Column	Type	Source	Comment
HR	EMPLOYEES	Table	EMAIL	EMAIL_ID	Sensitive Column Discovery	Email id of the employee
HR	EMPLOYEES	Table	PHONE_NUMBER	PHONE_NUMBER	Sensitive Column Discovery	Phone number of the employee; includes country code and area code
HR	EMPLOYEES	Table	LAST_NAME	UNDEFINED	User Defined	Last name of the employee. A not null column.
HR	EMPLOYEES	Table	FIRST_NAME	UNDEFINED	User Defined	First name of the employee. A not null column.

Figure 4-8: Example showing sensitive columns discovered using Enterprise Manager ADM

Data discovery analyzes the referential relationships between application objects by using foreign key constraints defined in the database. It can also automatically discover application level referential relationships that are not defined in the database. The discovered referential relationships can be reviewed, and additional relationships can be added manually when needed.

Understanding these dependencies helps protect application integrity during data masking by ensuring data in related columns is masked consistently. Figure 4-9 shows a list of parent-child relationships found inside the data dictionary.

Application	Object	Columns	Key Type	Source
HR				Dictionary
HR	COUNTRIES	COUNTRY_ID	Parent	Dictionary
HR	LOCATIONS	COUNTRY_ID	Dependent	Dictionary
HR	DEPARTMENTS	DEPARTMENT_ID	Parent	Dictionary
HR	EMPLOYEES	DEPARTMENT_ID	Dependent	Dictionary
HR	JOB_HISTORY	DEPARTMENT_ID	Dependent	Dictionary
HR	EMPLOYEES	EMPLOYEE_ID	Parent	Dictionary
HR	DEPARTMENTS	MANAGER_ID	Dependent	Dictionary
HR	EMPLOYEES	MANAGER_ID	Dependent	Dictionary
HR	JOB_HISTORY	EMPLOYEE_ID	Dependent	Dictionary
OE	CUSTOMERS	ACCOUNT_MGR_ID	Dependent	Dictionary
OE	ORDERS	SALES_REP_ID	Dependent	Dictionary

Figure 4-9: Example showing referential relationships

Figure 4-10 puts everything together. Data discovery automatically discovers sensitive columns, database-defined referential relationships, and application-level referential relationships and creates an application data model, which it then stores in the Enterprise Manager repository. This application data model can be used to help implement security controls such as Oracle Data Masking and Subsetting or Oracle Database Vault.

Copyright © 2026, Oracle and/or its affiliates

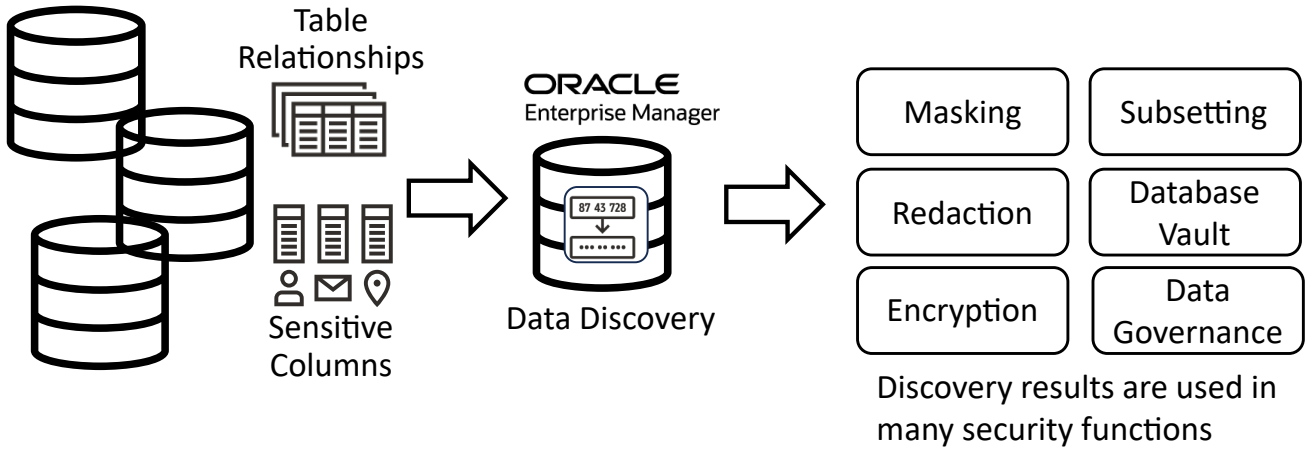


Figure 4-10: Results of data discovery are useful in configuring other security controls

Use cases and deployment scenarios

Enterprise Manager's discovery capabilities are particularly valuable in the following scenarios:

- **On-premises environments:** Organizations with data residency requirements or security policies that prohibit using cloud-based management tools can deploy Enterprise Manager entirely within their own data centers.
- **Integrated lifecycle management:** When organizations use Enterprise Manager for provisioning, patching, backup, and performance management, adding discovery and masking capabilities provides a unified operational platform.
- **Complex application environments:** Applications with intricate data relationships benefit from Enterprise Manager's sophisticated referential relationship discovery, reducing the risk of masking-induced application failures.
- **Test data management programs:** Organizations with formal processes for provisioning test and development environments can integrate discovery and masking into automated workflows that create compliant, functional nonproduction databases.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Database Security Assessment Tool
- Get Started with Oracle Data Safe Fundamentals
- Audit Vault and Database Firewall
- Data Masking and Subsetting (including sensitive column discovery)

These workshops provide pre-configured environments with sample databases containing realistic sensitive data, allowing practitioners to execute discovery jobs, analyze results, and experiment with different configurations without impacting production systems.

Copyright © 2026, Oracle and/or its affiliates

Summary

Sensitive data discovery is not a compliance checkbox. It works more like a compass for database security, because controls only get truly sharp when teams know exactly what they're protecting. Many organizations invest in strong defenses and still can't reliably locate regulated or high-impact information. Treated as the starting line, discovery makes every safeguard that follows feel purposeful and targeted, not performative.

The need is rooted in everyday business reality. Databases carry the information that keeps enterprises moving, and that same information often spans personal and sensitive categories like payment details, healthcare information, employee compensation, and proprietary metrics. Sensitivity also shapes adversary motivation: high-value data attracts high-intent threat actors, and exposure can cascade quickly into reputational, financial, and legal consequences. Practical direction boils discovery down to three questions every team should be able to answer: what types of sensitive data exist, how much exists, and where it resides. That clarity helps avoid a common failure mode, overprotecting low-risk areas while critical information remains exposed.

AI is positioned as a major amplifier of the exposure surface, especially when RAG systems and model integrations connect to production databases. Concrete scenarios make the risk tangible, such as an internal query that unintentionally surfaces executive compensation details. Regulated data used in AI training is also flagged as a long-tail risk, because it may later become recoverable through prompt injection, model inversion, or inference. MCP is described as a standard connector that speeds AI-to-database integration, which increases the urgency of putting guardrails in place early. In that world, discovery becomes the first control: define exclusions, row-level policies, masking or tokenization choices, and role-based access boundaries before AI-powered querying begins.

For execution, the recommended playbook is a blended approach: metadata-based scanning plus content-oriented pattern detection to balance speed, coverage, and confidence. Metadata discovery is fast because it relies on column names and comments, but it becomes brittle when naming practices are weak. Pattern-based discovery adds confidence by sampling values to detect formats and validate identifiers, though it comes at higher cost and demands careful handling of sampling results. These approaches map naturally to Oracle tools and their best-fit use cases: DBSAT for rapid assessments and reporting, Oracle Data Safe for broad discovery and evolving sensitive data models, AVDF for strengthening monitoring and firewall policies without sampling values, and Enterprise Manager for deep referential relationship discovery that supports masking and other downstream controls.

The next chapter covers Database Authentication, where discovery shifts from a map on the wall to a living checkpoint at the door. If this chapter clarifies what must be protected and where it hides, the next chapter lays out the mechanisms that decide who gets close to it, under what conditions, and with what proof, turning visibility into a disciplined, persuasive security posture readers can put into practice.

Copyright © 2026, Oracle and/or its affiliates

Chapter 5

Authenticating Database Users

The importance of strong authentication

An essential but often neglected principle in designing and implementing security solutions is that comprehensive identity-based security cannot exist in an anonymous environment. While this may seem self-evident, it is surprising how frequently this vital aspect is overlooked in database application design. Furthermore, applications sometimes obscure the user's identity from the database, which can undermine many of the built-in security features databases offer.

In virtually all robust database security frameworks, the ability to consistently identify users and authenticate them with strong, dependable methods forms the bedrock of the overall security model. Without this foundation, other security measures cannot be effectively enforced. This chapter helps you choose authentication methods, map identities to schemas, and operationalize secrets and accountability.

Unraveling identification, authentication and authorization

The database security process can be viewed as three key steps:

1. **Identification**: The user begins by presenting their identity to the database, such as by entering a username.
2. **Authentication**: The user then demonstrates that the presented identity is valid, such as by providing a password. The database verifies whether the entered password matches the one associated with the username.
3. **Authorization**: If the credentials are correct, the database accepts the user's identity as trustworthy. The database then assesses what privileges and authorizations are assigned to that user. Data security measures are then enforced based on these specific privileges and authorizations.

Most security efforts and resources are concentrated on processes within the third step, where fine-grained access and authorization controls are managed. Nevertheless, the importance of the first two steps cannot be overstated. Identification and authentication form the critical foundation of the security model.

This chapter explains how an authentication strategy helps protect database users and the data they access from attackers. It also provides guidance on managing user accounts, whether those accounts are maintained locally within the database or managed through centralized external services such as a directory service or a cloud identity provider.

The next chapter turns to user and application authorization, focusing on how to control what a user or application can do within the database.

Attacker tactics

One of the simplest and, regrettably, most frequent methods of breaching a database involves impersonating a legitimate user with authorized access. Attackers typically target accounts with elevated privileges, as these provide broad latitude for carrying out malicious activities. Robust authentication serves as a key defense against this form of attack, greatly lowering the chances of unauthorized access.

It is important to examine how attackers employ different tactics to obtain credentials in their efforts to access your data. Consider the various ways attackers attempt to gain access by defeating your authentication scheme:

Social engineering and phishing: Hackers often employ social engineering tactics to obtain account credentials, using spear phishing attacks that target end users or database administrators within an organization. These individuals are easily identified on platforms such as LinkedIn.

GenAI as a social engineering weapon: GenAI increases its success by making fraudulent communications more persuasive, more personalized, and easier to iterate at scale. The FBI has specifically warned about campaigns using AI-generated voice messages (vishing) and malicious text messages (smishing) to impersonate trusted individuals and trick targets into transitioning to attacker-controlled platforms or disclosing login information. As a result,

Copyright © 2026, Oracle and/or its affiliates

database authentication design should assume that credential theft and impersonation attempts will be more frequent and more convincing, reinforcing the need for strong authentication.

Hardcoded or exposed credentials in applications: Another risk area involves the discovery of hardcoded database connection details. Applications frequently store database usernames and passwords directly in program code or in plaintext configuration files. Since application servers are typically located closer to the network edge than database servers, they present an easier target for attackers. Once an application's service account is compromised, malicious actors can exfiltrate, alter, or erase any data that account is authorized to access. This applies equally to microservices and containerized environments where environment variables might be exposed.

Default or published passwords: Attackers may also attempt to use common default or well-known passwords to gain unauthorized access. Successful entry allows them to exploit the privileges associated with those user accounts and access sensitive information.

Brute force password attacks: When passwords aren't complex, and there's no retry limit, attackers can use automation to try millions of combinations until one works. Oracle AI Database makes this much harder by using exponential backoff. Each failed attempt increases the waiting time before another login try. This slows attackers down, but if simple passwords are allowed, brute force attacks might still succeed.

Credential stuffing: Credential stuffing is a variation of brute force attacks, where account information from previously breached sites is reused in attempts to compromise additional systems. The prevalence of password reuse increases the effectiveness of this strategy, with vast collections of stolen credentials available on the dark web. Automated tools can rapidly test millions of these credential pairs against your organization's systems. If attackers have already identified valid usernames, searching stolen databases for associated passwords allows them to target your infrastructure more effectively.

These techniques aren't especially sophisticated, and AI-automated tools make them easy to run. It's important to emphasize, if attackers sign in to a compromised account, they gain access to everything that account can access. Strong authentication reduces that risk, and account hygiene ensures you keep exposure low.

Database authentication methods

The attacks listed above share a common trait: they all focus on exploiting password-based authentication. Compromised credentials are at the heart of most successful security breaches, prompting many organizations, including the US federal government and Oracle, to adopt stronger authentication methods such as tokens, passkeys, and biometrics. Implementing multi-factor authentication (MFA), which relies on two or more independent factors, enhances security by ensuring a single compromised factor does not jeopardize the entire account.

Instead of checking passwords against the database's internal hashed credential store or a centralized directory, Oracle AI Database can authenticate through external services. You can use OCI Identity and Access Management (IAM) and Microsoft Entra ID, Kerberos, public key certificates, and RADIUS. These options help you adopt phishing-resistant, standards-based authentication while keeping administration centralized. A new feature introduced with the July 2025 release update lets you require approval via a mobile authenticator, such as Oracle Mobile Authenticator or Cisco Duo, for password-authenticated local database accounts – extending the benefits of MFA even to those local accounts.

After authentication, the user is mapped to a database schema consisting of tables, views, indexes, and procedures and then granted appropriate authorization through roles and privileges. In early Oracle Database releases, the user account and schema were the same, and each user had a dedicated schema. That heritage remains in the command syntax: you create a schema with CREATE USER. There is no CREATE SCHEMA command. A specialized form of CREATE USER creates a schema-only account - a user with no authentication credentials. When you use local database password authentication, the schema and user remain the same entity. When you authenticate with a

Copyright © 2026, Oracle and/or its affiliates

directory or identity service, we map identities to schemas. You can assign a private schema to each person or map multiple people to a shared schema.

While the topic of authorizations is covered elsewhere in the book, it's useful to note that roles - collections of privileges - can be managed in the database or, in some cases, managed externally through a directory service, OS groups, or a RADIUS server.

Table 5-1 lists common database authentication methods and how they map to schemas and roles.

Authentication Method	Schema Mapping	Role Mapping
Database Password	Schema and user are the same entity	Managed in the database
Operating System	Database maps the user to a schema	Managed in database or through membership in OS groups
Kerberos	Database maps the user to a schema	Managed in database
Public Key (PKI)	Database maps the user to a schema	Managed in database
RADIUS	Database maps the user to a schema	Managed in database or through RADIUS server
Oracle Directory Services	Managed in database and directory service	Managed in database or through membership in directory service groups
Microsoft Active Directory	Managed in database and directory service	Managed in database or through membership in Active Directory groups
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM)	Managed in database and identity service	Managed in database or through membership in OCI IAM groups
Microsoft Entra ID	Managed in database and identity service	Managed in database or through membership in Entra ID groups

Table 5-1: Database user authentication and authorization methods

As noted above, despite the inherent challenges of passwords, they remain the most widely used authentication token. When database users authenticate with a password, it is important to ensure passwords are tightly managed and made as secure as possible. Below are a few areas to review when passwords are used for authentication.

Managing passwords & secrets

With password-based authentication, people can type a password to connect to the database, but applications, middle-tier systems, and batch jobs cannot rely on a human at the keyboard. In the past, many teams embedded usernames and passwords in code or stored them in clear-text configuration files. That approach expanded your attack surface and forced you to change scripts whenever a password changed.

Secure External Password Store (SEPS) improved on that model by storing credentials, including passwords, private keys, and certificates, in an encrypted Oracle Wallet. With SEPS, you only needed to remember the wallet password. Once unlocked, the wallet could be used by applications or servers to access multiple databases securely, including with auto-login wallets for 24/7 operations. Wallets remove clear-text passwords from command-line history and configuration files, making casual discovery harder for anyone with access to the operating system.

However, SEPS wallets still lived on the application server or client workstation and were protected only by the wallet passphrase. If an attacker stole the wallet files, they could try to brute-force the passphrase or simply wait for a compromised host to capture the wallet password when someone used it.

Copyright © 2026, Oracle and/or its affiliates

Secrets management offers a more secure, modern way to store and distribute passwords. With a secrets manager such as Oracle Key Vault, users and applications retrieve credentials through an API call. The password is centrally managed and is not stored on the application server or user workstation.

A secrets manager, like Oracle Key Vault, has the following benefits:

- Centralize credentials instead of maintaining SEPS wallets on many servers.
- Simplify password rotation and remove persistent password footprints from application hosts.
- Reduce operational overhead compared to provisioning and protecting local wallets at scale.

For applications that still use password-based authentication, a secrets manager is the recommended approach for storing and distributing credentials. We'll dive deeper into Oracle Key Vault's role in encryption and key management in Chapter Eleven.

Accounts and passwords: to share or not to share

When it comes to accounts and passwords, the best practice is obvious: do not share them among users and administrators. When multiple people share a single account, accountability quickly breaks down, and the result is a security nightmare. If errors occur, sabotage is suspected, or data is stolen, it is nearly impossible to pinpoint which individual behind the shared account is responsible. Most organizations even have policies that prohibit shared accounts, yet those rules are still routinely ignored, sometimes by the most highly privileged database users, notably developers and database administrators.

In practice, database administrators often connect through shared accounts, most commonly the database super user (SYS) or the default SYSTEM database administrator account. These accounts are powerful and should not be used for day-to-day DBA activities. A more defensible approach is for each DBA to have a personal account granted only the roles or privileges required for the role, whether that is the out of box DBA role or a tighter role tailored to the tasks the DBA performs.

The same guidance applies to application administrators and developers. Maintenance work may require access to an application schema, but that does not mean the application service account credentials should be used. Later in this book, methods for blocking misuse of service accounts using Database Vault are discussed. For now, the priority is to enable administrators and developers to work through their own private accounts in a way that remains practical and low friction.

Proxy authentication solves this. It lets an administrator act as the application account without knowing the application account's password. With proxy authentication, a person signs in with their own account and then acts as another user. Audit records still show the real end user, and access control policies can consider both the true identity and the proxied account. That accountability is the key benefit for operations and compliance.

When you authorize proxy access to an application account, administrators authenticate to the database with their own credentials and then proxy to the application schema without needing the schema's password. For example, in Figure 5-1 we grant alice_appdba the CONNECT THROUGH privilege to connect as hrapp. This approach keeps credentials secure, preserves least privilege, and gives you clear audit trails for every action.

```
SQL> ALTER USER alice_appdba
GRANT CONNECT THROUGH hrapp;
```

Figure 5-1: Example showing how to grant proxy privileges

`alice_appdba` then connects using their password and assumes the identity and privileges of the `hrapp` schema by proxy as shown in Figure 5-2:

Copyright © 2026, Oracle and/or its affiliates

```
SQL> CONNECT alice_appdba[hrapp]
Enter password: <alice_appdba_password>
```

Figure 5-2: Example showing how to connect using proxy privileges

With this setup, audit records show `hrapp` as the `DBUSERNAME` and `alice_appdba` user as `DBPROXY_USERNAME`. The proxy username is also available for policy-driven access control mechanisms like Database Vault, Label Security, Data Redaction, and Real Application Security (discussed later in this book).

Increase password strength

There are many situations where passwords remain the only practical option for authentication. Older applications and clients often do not support other authentication methods. In those cases, the goal is to make passwords as strong as users and applications can reasonably support and to apply solid password discipline, such as avoiding password reuse, avoiding shared accounts, and managing the password lifecycle.

Like most systems that support password authentication, Oracle AI Database provides controls for minimum password length, complexity, and lifecycle. These controls are managed through user profiles. Every database account is assigned to a user profile. In theory, each database user could have a unique profile, but that is uncommon. More often, a small set of profiles is used, with users grouped under different profiles based on their needs. User profiles define password and resource authorization parameters for all accounts assigned to them. If no profile is specified when a user is created, the default profile is automatically assigned. Custom profiles can be created as needed

Note: Oracle AI Database 21c and earlier support passwords up to 30 bytes. Beginning with Oracle AI Database 26ai, the maximum password length is 1024 bytes.

Typically, the default profile is tailored to meet organizational standards, and other profiles are reserved for situations that require an approved exception. A profile can contain up to 18 parameters, with nine parameters dedicated to resource limits and nine dedicated to authentication. When a new profile omits any of the 18 parameters, the unspecified settings automatically inherit the values defined in the default profile.

Best practice for interactive accounts: For accounts used by people rather than applications or batch processes, passwords should be supplemented with the local user MFA feature. A second factor, such as a mobile authenticator approval, should be required for local database logins to reduce the risk of credential stuffing and brute force attacks. This approach strengthens human sign-ins without changing how service accounts authenticate.

Figure 5-3 shows an example custom profile, app_profile, that includes both authentication and resource parameters. This custom profile specifies six of the 18 possible parameters - the remaining 12 inherit values from the default profile.

```
SQL> CREATE PROFILE app_profile LIMIT
PASSWORD_VERIFY_FUNCTION ora12c_stig_verify_function   -- STIG password complexity rules
PASSWORD_ROLLOVER_TIME 30               -- Allows both old and new passwords for 30 days
FAILED_LOGIN_ATTEMPTS 6                 -- Lock the account after 6 attempts
SESSIONS_PER_USER 2                     -- Allow two sessions for each user
IDLE_TIME 30                            -- Automatic logout after 30 min idle
PASSWORD_LIFE_TIME 180;                 -- Force password change after 180 days
```

Figure 5-3: Sample password profile

All eighteen parameters can be important, but let's call out three that can make the biggest impact on your security:

`PASSWORD_VERIFY_FUNCTION:` This enforces strong password complexity through built-in functions like ora12c_stig_verify_function, which aligns with U.S. Department of Defense Security Technical Implementation Guide

Copyright © 2026, Oracle and/or its affiliates

(STIG) recommendations. These rules require a minimum length, a mix of alphanumeric and special characters, and more. If Oracle's default functions aren't enough, you can write your own; just start with the sample code in `$ORACLE_HOME/rdbms/admin/catpvf.sql`.

`PASSWORD_ROLLOVER_TIME:` Application accounts often require periodic password changes to stay compliant with security policies. In the past, updating these passwords meant juggling downtime and reconfigurations to avoid failed logins. With Oracle Database 19c and above, you can update a password for an account but keep the old one working for a set period (the rollover time). This lets you update all application clients smoothly. No downtime or risk of service disruption. Set this parameter for application service account profiles but leave it off for human users. If you use zero, password rollover is disabled.

`FAILED_LOGIN_ATTEMPTS:` This parameter sets how many consecutive failed login attempts are allowed before an account is locked. It's a simple, effective way to block brute-force attacks.

Taking these steps helps you enforce your organization's standards, make password management easier, and keep your users and data safer, no matter which type of authentication you use.

Simplify account administration

Most databases have only a small number of human actor accounts, typically database administrators, perhaps a developer or two, and a few power users who run reports. Many real-world environments, however, break that pattern, including databases with tens of thousands of valid human actor accounts. Even when there are only a few users per database, scale becomes a factor when those same accounts exist across many databases, and customers with thousands of Oracle AI Databases are not unusual.

In either case, managing identities inside the database can become an operational burden. People join and leave, and roles change. Some users no longer need database access, while others should have different or fewer privileges. At scale, accounts that are no longer needed often remain active long after a user has left or changed roles, creating attractive targets for hackers seeking unauthorized access to data.

Oracle AI Database supports centralized user management by allowing users and credentials to be managed outside the database, most commonly in a directory such as Microsoft Active Directory or a cloud identity service such as Oracle Cloud Infrastructure's Identity and Access Management (IAM), Okta, or Microsoft Entra ID.

Centralized user management ensures consistent policy enforcement across both users and databases. It also streamlines the process of adapting to organizational changes as individuals join, depart, or take on new roles. Many centralized services, especially cloud identity providers, can also require multi-factor authentication. Centralization remains valuable even beyond passwords. For example, Microsoft Active Directory integrations more commonly use Kerberos for authentication. PKI certificates, often delivered through smart cards like the US Department of Defense Common Access Card (CAC), are also widely used.

We'll cover centralized user management in more depth in Chapter Six.

As applications, tools, and databases continue shifting to the cloud, multicloud identity integration becomes essential. It allows people to be managed in one identity service while their applications, databases, and tools operate across multiple clouds or on premises. Chapter Fourteen explores multicloud identity in more detail.

Protect user accounts

It is not realistic to rely on (business) users to remember, interpret, and voluntarily follow security guidance on password strength, password management hygiene, and appropriate roles and privileges. The database should enforce these rules, so they are applied consistently and predictably. Policies such as the following help protect user accounts from compromise:

Copyright © 2026, Oracle and/or its affiliates

Avoid password authentication. Use strong authentication mechanisms such as cloud identity service-provided tokens, Kerberos, RADIUS, or PKI certificates.

When passwords are unavoidable, apply a strong password profile to govern password strength, inactivity period, and the number of password retries (failed logins).

Centralize user management with an identity service like Active Directory to reduce the likelihood of orphan accounts when someone leaves the organization or changes roles.

This best practice shifts the burden from people to the database and your identity platform, improving security without adding unnecessary friction.

Protect DBA accounts

Because database administrator accounts typically carry broad access, they are prime targets for attackers seeking a high impact compromise. DBA authentication should follow these best practices:

- Require named DBA accounts with strong authentication to preserve accountability. Prohibit shared accounts and default accounts such as SYSTEM. Where practical, use strong authentication such as cloud identity service tokens, PKI certificates, or Kerberos.
- Use a Privileged Account Management (PAM) system for root and database owner operating system access needed for SYSDBA, typically limited to upgrades, patching, and database restarts.
- Use sudo when DBAs must access the database operating system account, so activity is recorded and attributable to an individual user.

If DBAs log into the database server operating system before connecting to the database:

- Require access through a personal named account, not the account that owns the database binaries.
- Provide a separate client such as Oracle Instant Client or SQLcl, rather than granting access to the database owner account or binaries.

When DBA responsibilities are split into specialized teams, such as backup and recovery DBAs, performance tuning DBAs, and security managers, avoid granting the default DBA role. Create specialized administrator roles for each function and grant only the privileges required for that team (more about this in Chapter Six and Chapter Seven).

Protect application accounts

Application service accounts in the database not only hold the roles and privileges required to run the application on behalf of every application user, but for convenience they may also be granted additional roles and privileges for installation, upgrades, patching, and other maintenance. Because these accounts can become powerful keys to the environment, they should be protected using the following best practices:

- Use strong authentication mechanisms (cloud identity service tokens or PKI certificates) for database access.
- Use controls such as SQL Firewall, Database Vault connect rules, or logon triggers to limit where service account connections can originate. Additional discussion of Database Vault and SQL Firewall appears in Chapter Seven (for Database Vault) and Chapter Eight (for SQL Firewall).
- Use a secrets manager such as Oracle Key Vault for credential storage, and have applications retrieve credentials via API calls. Avoid embedding credentials in code or storing them in clear text properties files on the application server. Application servers are often exposed to a higher risk of compromise than database servers because they are typically deployed in network zones with a broader threat surface.
- Use gradual password rollover to reduce the risk of outages during password rotation. This allows both the old and new passwords to remain valid until updates are complete across all application servers.

Copyright © 2026, Oracle and/or its affiliates

- Require application administrators to use proxy authentication instead of logging in with the application service account.

Protect security administrator accounts

Security administrators manage security controls, encryption keys, database users, and audit records. They protect the database and set user constraints on tablespaces, idle time, and concurrent sessions so one user doesn't impact others. Use the following practices to strengthen accountability and reduce risk.

Use strong authentication. Authenticate with PKI certificates, Kerberos, or tokens. For interactive logins, add a second factor with local user MFA to reduce credential stuffing and brute force risk.

Enforce separation of duties. Keep security administrator responsibilities separate from DBA and general user duties. Do not allow DBA or user accounts to alter security records, create fake users, or change security controls. Use roles and policies that restrict who can manage users, keys, and audit configurations.

If you must use passwords, enforce strong profiles. Assign security administrators to profiles that require strong passwords, set inactivity and expiration rules, and automatically lock accounts after too many failed login attempts. Use profile settings to enforce idle time limits, concurrent session caps, and tablespace quotas.

These controls create clear accountability, reduce the blast radius of compromise, and help you maintain consistent security across your databases.

Simple designs yield strong security

One best practice is as much about architecture as it is about requirements. The simple pattern is to create a schema only account for the application objects and procedures so the application schema cannot be used for direct logins. Then use a separate run time application service account to access those objects through controlled procedures and views.

Schema only accounts were introduced in Oracle AI Database 18c, providing a named schema with no password or other ability to log in. Starting with Oracle AI Database 19c, almost all Oracle accounts installed with the database are schema only accounts to prevent login to these privileged accounts. In earlier versions, these accounts had passwords that required periodic rotation. Other users can still be granted read and write access to these schemas**.**

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- DB Security – Privilege Analysis
- Proxy authentication
- Centralize database user management with OCI IAM

Copyright © 2026, Oracle and/or its affiliates

Summary

Authentication often sits quietly at the critical center of database security, and it's easy to underestimate. Security only becomes real when it's clear who is present. When end-user identity is blurred or hidden, built-in database safeguards can lose much of their power. A practical way to frame the journey is in three steps: identification, authentication, and authorization. The first two steps do the foundational work that makes the third enforceable. Solid programs account for both locally managed database accounts and externally managed identities delivered through directory services and cloud identity providers.

Attackers fixate on authentication for a simple reason: impersonation remains one of the easiest, most reliable ways into an environment, and privileged accounts can turn a foothold into a wide-open corridor. Credential capture and replay show up in many forms, from spear phishing and social engineering to GenAI-fueled vishing and smishing. The risk also includes exposed application secrets, default passwords, brute force attempts, and credential stuffing at industrial scale. The sobering reality is that once an attacker signs in as a legitimate account, they inherit everything that account can reach. The most effective response is to reduce the odds of a successful sign-in and keep exposure low through disciplined account hygiene.

That's why modern guidance pushes teams to move beyond password-centric design toward stronger authentication methods and multi-factor authentication. Oracle AI Database support for external authentication is positioned as a key enabler here, with options including OCI IAM, Microsoft Entra ID, Kerberos, public key certificates, and RADIUS. A July 2025 update is also highlighted for extending mobile authenticator approval to local password-authenticated accounts, bringing MFA benefits into scenarios that traditionally lag. From there, the model stays grounded: authenticated identities are mapped to schemas, and roles and privileges can be managed inside the database or through external systems, depending on the chosen authentication method.

Operationally, secrets storage and account management are treated as first-class controls, not cleanup work. Credentials embedded in code or stored in plaintext files expand the attack surface. Wallets like SEPS can reduce casual discovery, but they're still vulnerable if stolen. A more modern model is centralized secrets management, with tools such as Oracle Key Vault providing API-based distribution and control. Shared accounts, especially for DBAs and other powerful users, are called out as a risk to accountability; the recommendation is personal named accounts with least privilege. Proxy authentication is presented as a practical way to preserve accountability and audit clarity while avoiding direct use of service account credentials. Password profiles and key parameters round out the operating playbook, including complexity verification functions aligned with STIG guidance, password rollover time to smooth application transitions, and failed-login lockouts, along with an implementation note that Oracle AI Database 26ai supports passwords up to 1024 bytes.

After tightening authentication, a key truth remains: many database compromises look like a normal login that never should have worked, followed by valid commands that were never meant to be available. The next chapter, Database Access Controls, therefore, emphasizes precision over perimeter thinking by defining who can act on which data, with what scope, and why, even for a DBA, an application service, a developer, or an analyst. Through privileges (object, schema, system, and administrative) and role design, authorization becomes a deliberate discipline. Unnecessary grants expand the blast radius, while well-chosen constraints provide guardrails under pressure. For resilience against mistakes, misuse, and attackers using legitimate credentials, this is where the book becomes a practical playbook. Database Access Controls is the next destination, and it is where thoughtful authentication choices start to pay off by letting policy guide action with precision.

Copyright © 2026, Oracle and/or its affiliates

Chapter 6

Controlling Database Access

Why access control is important

Impersonating an authorized user to steal, alter, or destroy data is one of the easiest ways to compromise a database. This is precisely why controlling what users can do is essential to protecting data. Every unnecessary privilege granted to a user increases risk to the data and the database, whether the account is compromised or the user acts in ways that violate organizational policy.

Keep access control a primary focus

Access controls define who can act on which data and what actions they can perform within the database. They are used to limit each account to only the capabilities required for its intended job or function, because overprivileged users present an unacceptable risk to the data and the database. Strong access control also supports regulatory compliance and helps prevent the theft, destruction, or misuse of data.

Because access control is so critical, Oracle AI Database provides many ways to manage and restrict access. In some cases, the differences between access control features are obvious. In others, the distinctions are subtle, and it is helpful to use this chapter and the next few chapters to identify the right feature or tool to meet a specific objective.

The discussion starts with the basics of privileges and roles, and then moves into more advanced topics in later chapters. It also covers managing user account authorizations locally and through directory or identity services.

The following four chapters explore additional access control features and show how they work together to minimize an account's ability to operate beyond the constraints required for its job function.

Database privileges

A database privilege grants the ability to perform specific operations on data objects or to execute certain statements. Oracle AI Database has four privilege types: object, schema, system, and administrative. The first three types provide a gradually increasing scope of authority.

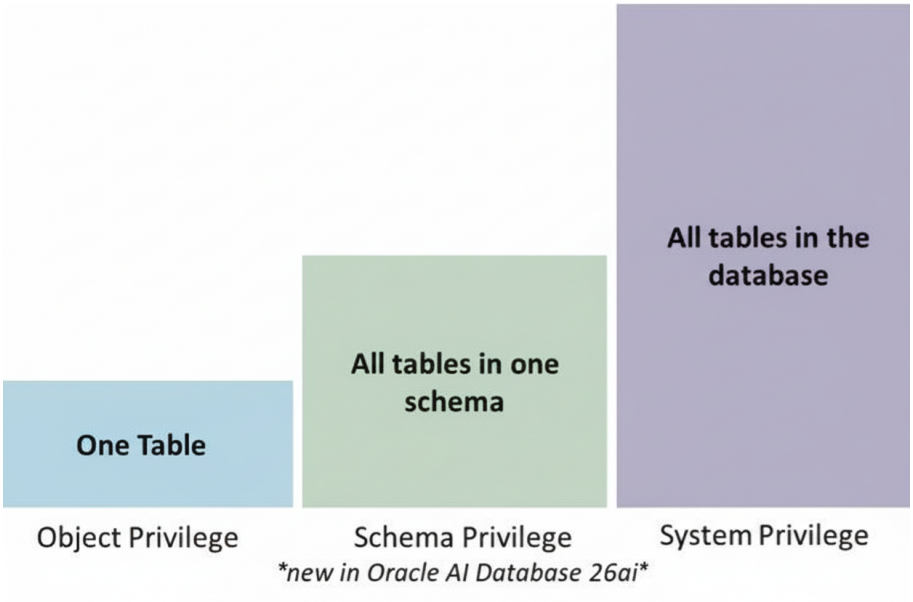


Figure 6-1: Summary illustrating the scope of object, schema and system privileges

Object privileges are fine-grained permissions that allow a user to perform specific actions on a particular database object. Examples include querying data from a table (SELECT or READ), modifying data (INSERT, UPDATE, DELETE), changing the structure of a table (ALTER), or running stored program units such as procedures or packages (EXECUTE).

Copyright © 2026, Oracle and/or its affiliates

To grant user Scott the object privilege to select data from the customers table in the app schema, use the following command:

```
GRANT SELECT ON app.customers TO scott;
```

Figure 6-2: Example syntax for granting an object privilege

Schema privileges, introduced in Oracle AI Database 26ai, are designed for applications and users that need to manage objects in another schema.

As a general rule, schema privileges are preferred over system privileges because they are limited to a specific schema. That tighter boundary makes them more secure than system privileges. They are also easier to maintain than granting individual object privileges, since there is no need to rework grants every time a new object is added to a schema.

The following command grants user Scott the ability to select data from all tables or views in the schema owned by app.

```
GRANT SELECT ANY TABLE ON SCHEMA app TO scott;
```

Figure 6-3: Example syntax for granting a schema privilege

System privileges provide broader access to database objects and typically to ALL objects relevant to the privilege.

For example, `SELECT ANY TABLE` allows a user to read data from almost any table in the database, including sensitive data stored in application schemas. `CREATE USER` is another system privilege, enabling the creation of new users in the database.

In most situations, system privileges should not be granted to non-administrators due to their expansive scope.

To grant the powerful `SELECT ANY TABLE` privilege to user Scott, use the following command:

```
GRANT SELECT ANY TABLE TO scott;
```

Figure 6-4: Example syntax for granting a system privilege

Administrative privileges differ from object, schema, and system privileges. They are centered on database administration tasks rather than access to database objects. Administrative privileges are commonly used for activities such as database backup, encryption key management, and core database operations such as startup and shutdown.

Examples of administrative privileges include `SYSDBA`, `SYSOPER`, `SYSKM`, and `SYSDG`. The command to grant an administrative privilege closely resembles the command used to grant an object privilege. For example, the following statement grants Scott the `SYSKM` privilege, which provides the ability to manage encryption keys for the entire database:

```
GRANT SYSKM TO scott;
```

Figure 6-5: Example syntax for granting an administrative privilege

Certain maintenance activities, such as database upgrades and patching, can only be performed by the SYS account, the database owner. The SYSDBA administrative privilege allows a user to become SYS for these tasks. The SYS account and the related SYSDBA privilege should be used only when necessary and carefully safeguarded. Use of administrative privileges should also be fully audited.

Each administrative privilege is associated with an operating system group, such as OSBACKUP or OSKM. Depending on organizational structure, these privileges may be assigned to a single shared group, such as DBA, or each privilege may be mapped to its own dedicated operating system group.

Copyright © 2026, Oracle and/or its affiliates

Administrative privileges are extremely powerful. Even when granted to a user, they are not active within a session by default. The user must explicitly activate the granted privileges when starting a database session.

```
CONNECT scott@pdb1
-- Scott's SYSKM privileges are not activated and cannot be used during this session
CONNECT scott@pdb1 AS SYSKM
-- Scott's SYSKM privilege is active in this session and Scott can perform encryption key
management operations
```

Figure 6-6: Example syntax for granting an administrative privilege

Database roles

Roles are named groups of privileges. They are used to streamline privilege management by making it easier to grant or revoke the same set of privileges for multiple users. Object, schema, and system privileges can be bundled into task-based roles. Roles can also be organized hierarchically under other roles. By granting roles to other roles, privileges can be grouped into a task role, and multiple task roles can then be combined into a broader organizational role.

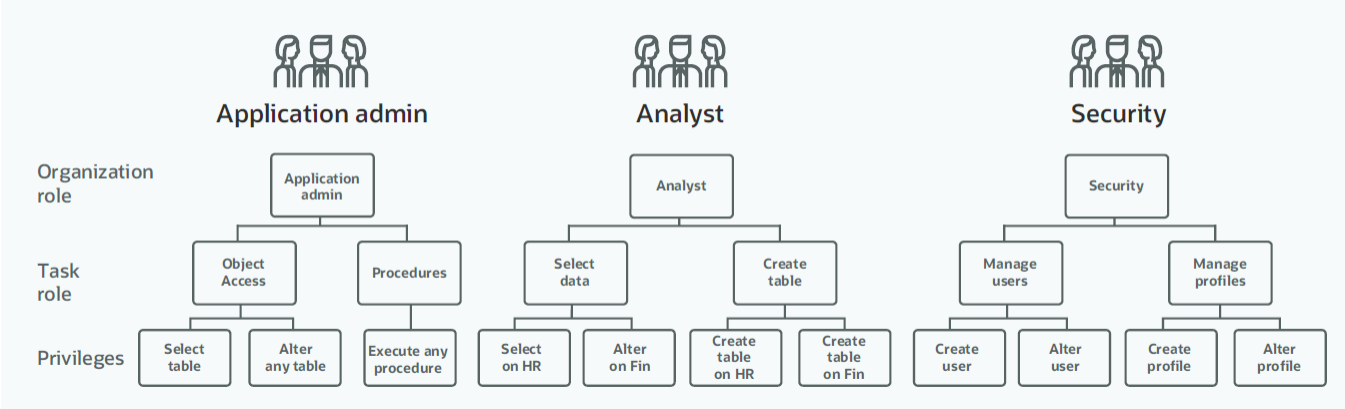


Figure 6-7: Example three-tier role and privilege hierarchy

A single user can be granted multiple database roles. This model simplifies onboarding by granting the role for a position instead of identifying and assigning each required privilege individually. As responsibilities shift between groups, task roles can be moved to match the new structure rather than managing individual privileges or redefining organizational level roles. Roles therefore simplify privilege management during organizational change.

Types of database users

Users receive privileges either directly or through a role. In general, a user can connect to their schema and access, modify, or delete all objects within that schema. Through privileges and roles, users can also be granted permission to perform specific operations, such as updating an object in another schema, or database specific rights, such as the ability to export or import data. These users may be human actors (individual users), applications, AI agents, tools, automated processes, or other programs acting on behalf of users.

A core security principle is to grant only the privileges needed to do the job. The challenge is that needed privileges vary based on what the account is meant to do. For practical access control planning, it helps to group accounts into categories based on how they are used and what they will do in the database:

- **Database administrators (DBAs)**: Every database relies on one or more administrators to handle tasks such as performance management, diagnostics and tuning, upgrades and patching, startup and shutdown, and backup. Unless access is constrained, DBAs can typically view and modify all data in the database, including

Copyright © 2026, Oracle and/or its affiliates

sensitive personal, health, and corporate finance records, even though that level of access is not required for many DBA tasks. That broad access also makes DBAs a prime target for hackers. In most cases, DBAs will have the VERY highly privileged DBA role.

- **Security administrators**: Many organizations have specialized DBAs who concentrate on security work, including user account management, access control policy development and implementation, encryption key management, and audit management. Because their work is narrower than a general-purpose DBA, these accounts usually need fewer privileges. For example, a security administrator generally does not need to create schema objects or access application data. This kind of specialization is common in large organizations with many people and many databases. One team may cover backup and recovery, another performance tuning, and another patching and upgrades. Each team needs only a subset of DBA privileges rather than the full DBA role.
- **Application service accounts**: These are called service accounts (sometimes described as "robotic", "non-interactive", or "machine" accounts) because they connect to the database without a human directly initiating the connection. Most databases have applications, business intelligence tools, or reporting systems that connect using one or more service accounts. If one of these accounts is compromised, the impact can be data loss across the entire application. Since applications are often built for high availability, service account passwords may be stored on multiple middle tier servers, which increases compromise risk. And because no person needs to remember or type the password, it should be as long and complex as the database can support.
- **Application administrators**: These accounts manage, patch, and upgrade the application itself and usually have full access to the application's data and stored procedures. They tend to have DBA like access to their application's data, but they do not need the broad database management privileges required by a DBA.
- **Developers and testers**: These accounts create and modify the database objects an application needs. Traditionally, developers did not have routine production database access. Changes were built in lower environments and promoted through defined change management processes. In DevOps or SecDevOps environments, those lines are often blurred, and developers with production access are now common. Even without production access, developer systems are frequently cloned from production and may contain the same data. Developer accounts typically have privileges similar to the application service account and may use proxy authentication to access the application schema. Oracle AI Database 26ai introduces a new DB_DEVELOPER_ROLE that includes the privileges a developer needs for their job function. Like DBAs, developers are often easy to identify through social media and are favored targets for hackers.
- **Data analysts or business intelligence users**: Many databases support data analyst access, either directly or through a business intelligence tool. These accounts typically need read-only access to the application schema. They can pose slightly higher risk than users who connect through the application because they bypass application-level access controls.
- **Other database users**: This includes everyone else with a database account, such as batch processes, integration services like Oracle Golden Gate, or tools like Data Safe, Enterprise Manager, or Audit Vault and Database Firewall. Their access requirements can vary widely, and many tools include scripts that grant what the tool considers appropriate access. It is worth treating appropriate with caution. A vendor may assume their tool needs the DBA role, but *Privilege Analysis* often shows that far less access is actually used. These accounts may be limited to their own schema or may need broader access to multiple schemas.

Who can do what

One of the most critical components of an Oracle AI Database is the data dictionary, a set of tables and views that describe the database, including metadata for all objects and users. The dictionary views shown in Table 6-1 help identify the roles and privileges granted to users or to roles.

Copyright © 2026, Oracle and/or its affiliates

Dictionary views	Contents
DBA_TAB_PRIVS	Object privilege grants to roles or users
DBA_SCHEMA_PRIVS	Schema privilege grants to roles or users
DBA_SYS_PRIVS	System privilege grants to roles or users
DBA_ROLE_PRIVS	Roles granted to all users and roles
DBA_ROLES	All defined roles
ROLE_ROLE_PRIVS	Roles granted to other roles
ROLE_SYS_PRIVS	System privileges granted to roles
ROLE_TAB_PRIVS	Object privileges granted to roles
V$PWFILE_USERS	Administrative privileges granted to users

Table 6-1: Data dictionary views for database roles and privileges

A close look at these dictionary views quickly highlights just how powerful the out-of-the-box DBA role is. It carries more than two hundred system privileges, including `ALTER SESSION`, `CREATE`, `ALTER`, and `DROP USER`, `CREATE` and `ALTER ANY TABLE`, and `SELECT`, `INSERT`, `UPDATE`, and `DELETE ANY TABLE`. On top of that, it includes the EXPORT and IMPORT FULL DATABASE roles and more than a dozen additional roles, each with its own collection of privileges.

For organizations that use specialized DBAs, where each role only needs part of the full DBA capability, a better approach is to define one or more custom DBA roles, such as ops_dba or backup_dba, that match the actual job function. Privilege Analysis can help identify which privileges are truly required and shape these custom roles accordingly.

To determine a database user's complete set of privileges, it is necessary to review role grants, system privilege grants, and object privilege grants to build a full and accurate picture. The example below is taken from Oracle Data Safe's User Assessment module.

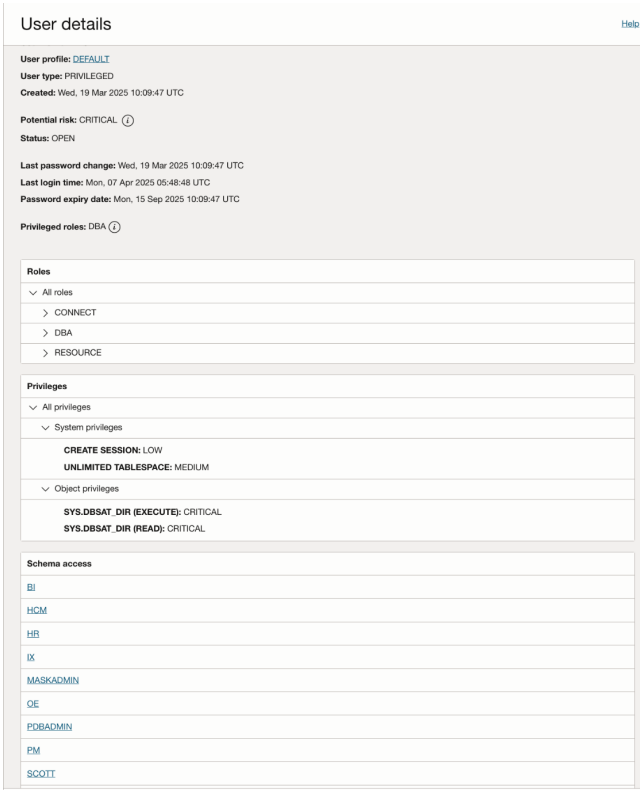


Figure 6-8: Example analyzing a user's privileges with Data Safe

Copyright © 2026, Oracle and/or its affiliates

Reviewing privileges in Audit Vault and Database Firewall provides similar visibility into user and role access:

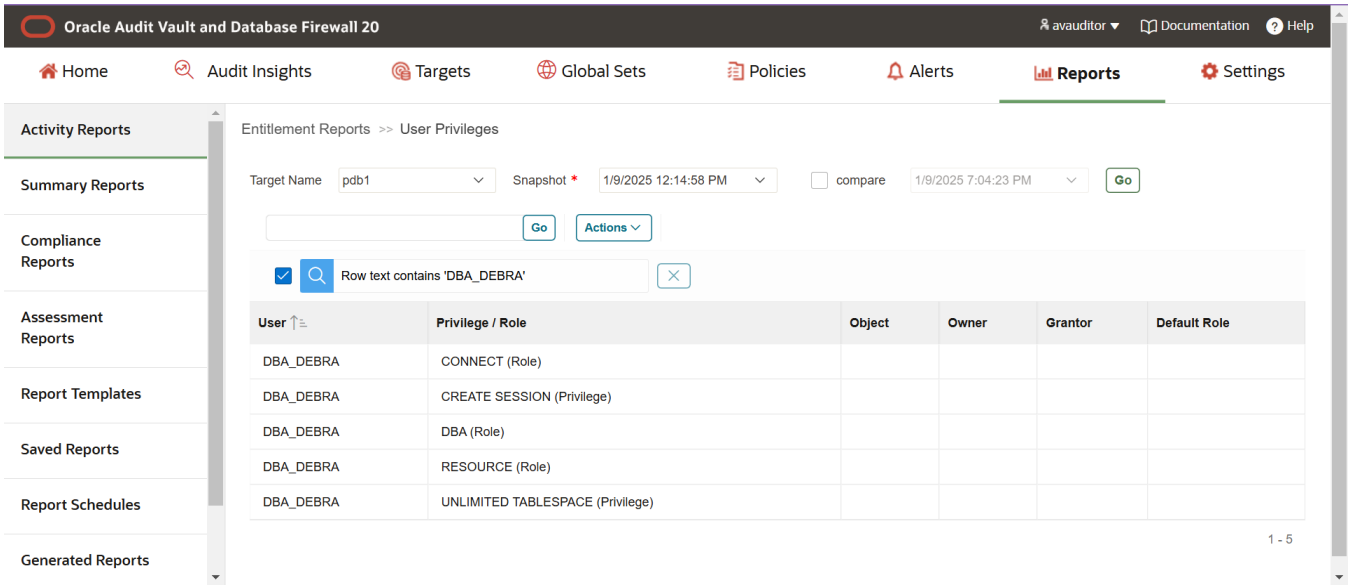


Figure 6-9: Example analyzing a user's privileges with AVDF

Centrally managed users

In enterprises where many users access multiple databases, it can be difficult for users to keep passwords compliant with each database's policy and to remember different passwords across accounts. Administrators also face significant overhead, as each user must be provisioned separately in every database along with a password. More importantly, each account must be deprovisioned when a user changes roles or leaves the organization. Accounts that are not deprovisioned after they are no longer needed are known as orphaned accounts. Hackers often exploit orphaned accounts to gain unauthorized access to data.

Enterprise User Security (EUS) centrally manages users and roles across multiple databases using an Oracle directory service such as Oracle Internet Directory or Oracle Unified Directory. The Oracle directory can run as a stand-alone LDAP server or integrate with other corporate directories. After a user authenticates successfully, the database consults the directory for authorization information, including roles. Database users managed through a directory service are known as enterprise users because they can span multiple databases across the enterprise. Enterprise users can be assigned enterprise roles or groups that determine access privileges across multiple databases. Enterprise roles in the directory map to one or more roles in the databases. Enterprise User Security is deprecated in Oracle AI Database 26ai.

Centrally Managed Users (CMU) integrates the database directly with Microsoft Active Directory. Active Directory users can have their own schema or share a schema through membership in Active Directory groups. Active Directory groups can also map directly to database roles. One key distinction between EUS and CMU is where schema and role mappings can be defined. With EUS, mappings can exist in either the database or the directory. With CMU, schema and role mappings are always maintained in the database.

Copyright © 2026, Oracle and/or its affiliates

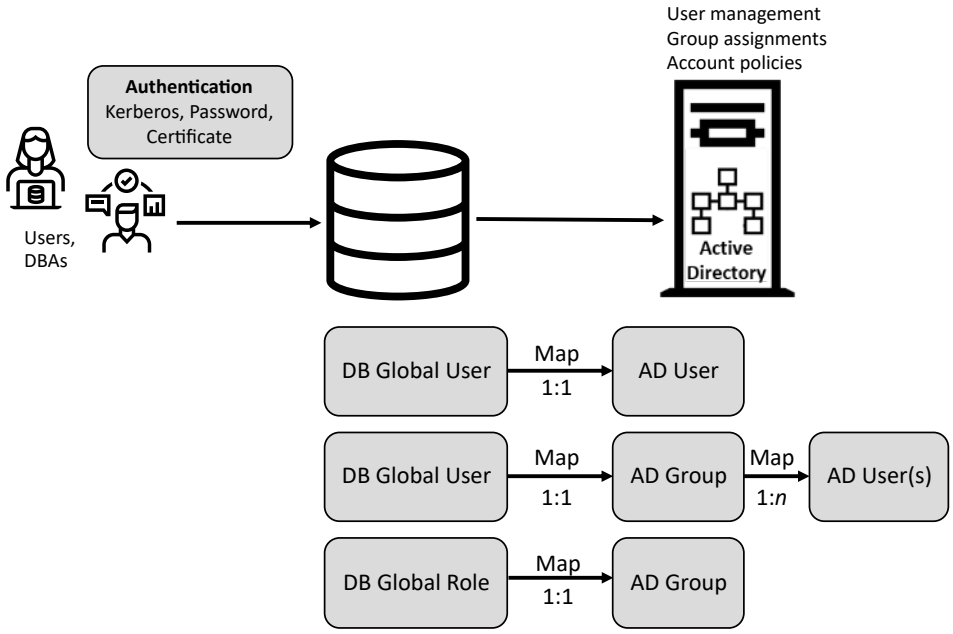


Figure 6-10: Illustrative example of Centrally Managed Users (CMU) mapping database global users/roles to AD users/groups

EUS and CMU allow users to authenticate using passwords, Kerberos, and Public Key Infrastructure (PKI) certificates. Kerberos is commonly used for human users, while passwords and certificates are typically used for application accounts.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) users can access the database using IAM tokens. Similar to CMU, authorization is based on mappings to database schemas and roles. IAM users can map exclusively (1:1) to database schemas, or IAM users in an IAM group can be mapped to a shared database schema (many:1). With a shared schema, IAM security administrators can add an IAM user to that shared database schema by adding the user to the appropriate IAM group. Shared schemas remove the burden of creating, managing, and dropping database schemas for every IAM user. IAM groups can also be mapped to database roles, allowing IAM users to receive privileges and roles through membership in different IAM groups. IAM authentication and authorization is available for all OCI databases, including Autonomous Database, Oracle Base Database Service, and others.

Microsoft Entra ID users can access the database using Entra ID tokens and can use Entra ID app roles for database authorization. App roles are included in the Entra ID token and are used by the database to map Entra ID users to a database schema and, optionally, database roles. Entra ID authentication and authorization is available for Oracle AI Database 19c and higher, including databases deployed in OCI, on third party clouds, or on premises.

Protecting database accounts

Protecting DBA accounts

Database Administrator (DBA) accounts typically have broad access within the database, often far beyond what is required to administer it. This access should be controlled using the following best practices:

- Tailor privileges for individual tasks and responsibilities to reduce the attack surface associated with these accounts instead of granting the DBA role. If additional privileges and roles are needed for troubleshooting, revoke them after the task is complete.
- Use task-specific administrative privileges (for example, SYSKM or SYSBACKUP) instead of the SYSDBA privilege.
- Block access to user schemas or SQL commands using Database Vault.

Copyright © 2026, Oracle and/or its affiliates

- Fully audit all DBA activities for accountability and tracking.

Protecting security administrator accounts

Security administrators manage security controls, encryption keys, database users, and audit records. They manage the database's security posture and ensure there are appropriate constraints on authentication, failed logins, idle time, and the number of concurrent sessions. Security administrator accounts need to follow these best practices:

- Limit the privileges and roles granted to security administrators so that they don't have wide access to the sensitive data in the database. It would be unusual to grant a security administrator the DBA role.
- All their activities should be fully audited for accountability and tracking.

Protecting application service accounts

Application service accounts typically hold the roles and privileges required to run the application on behalf of every application user. For convenience, they may also include privileges for application installation, upgrades, patching, and other maintenance tasks. These accounts should be protected using the following best practices:

- Revoke the privileges for application upgrade, patching, and other maintenance activities from the application runtime account. Instead, create a separate administrator user that is managed and audited independently. Without this separation, a hacker who has compromised an application user account can use SQL injection attacks to take over the application, change stored procedures, steal or destroy data, and delete tables.
- Grant application DBAs access to only the application schema objects, not database-wide privileges, even if this is less convenient. Earlier in this chapter, schema privileges in Oracle AI Database 26ai were discussed as a feature that helps support this approach.
- Fully audit all activities for accountability and tracking.

Protecting data analysts or business intelligence users

These accounts allow access to data but, in most cases, should not permit data to be altered, deleted, or inserted.

- Avoid grants of system privileges such as SELECT ANY TABLE. Use direct object grants, either to the user or to a role, or use schema-level grants.
- Audit these users for misuse. Are they logging in from unusual IP addresses or running different programs than normal? These can be indicators of account compromise. If the users have static IP addresses, consider using SQL Firewall or Database Vault connect rules to restrict accounts to specific workstations. SQL Firewall and Database Vault are discussed later in the book.
- Audit these users for insert, update, and delete activity.
- Consider making these accounts read-only. Read-only accounts are a new feature of Oracle AI Database 26ai.

Protecting other database users

These accounts are often the exceptions, and standard policies for common account types may not fit. They benefit from closer review and tighter oversight.

- Regularly confirm the account is still required, especially when assigned to a human actor.
- Track system privilege grants such as `SELECT ANY TABLE` and use object or schema privileges where possible to limit scope.
- Audit these accounts for misuse. Are they logging in from unusual IP addresses? Do patterns suggest credentials are shared by multiple people? Are they running different programs than they normally do?
- For accounts tied to batch processing where the same SQL statements are consistently issued, consider protecting them with SQL Firewall
- Use Privilege Analysis to identify unnecessary privilege grants.

Copyright © 2026, Oracle and/or its affiliates

Agentic AI and Model Context Protocol (MCP) connections typically fit into the same user categories as other non-human database access. In many implementations, they operate as application service accounts because they connect without a person directly initiating each database session and perform actions on behalf of application users, such as reading and writing application data or executing stored procedures. In other cases, an agent or MCP connection behaves more like an integration or automation component and is better treated as another database user, similar to a batch process or specialized tool account with narrowly defined access needs. The key is to classify the account based on how it is used and then apply least privilege, strong authentication, and audit its activity, since these connections can be highly capable and, if overprivileged, can amplify the impact of credential compromise.

With these access control patterns in place, the next step is to see how they work in practice. Hands on exercises can help reinforce the concepts and provide a safe way to explore privilege design, user assessment, and entitlement monitoring.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Privilege Analysis
- Data Safe Fundamentals
- Audit Vault and Database Firewall

Copyright © 2026, Oracle and/or its affiliates

Summary

Database access control is positioned as both a security necessity and a strategic discipline, because the simplest compromises often look like ordinary logins. With impersonation of authorized users highlighted as a common path to the theft, alteration, or destruction of data, the message lands clearly: every additional privilege granted today can quietly expand the blast radius tomorrow. The most effective programs keep authorization design front and center, use least privilege as a steady compass, and apply Oracle AI Database access control features with intent rather than routine.

Making that intent actionable starts with understanding how Oracle AI Database organizes privileges and why scope matters. Four privilege types shape the authorization landscape: object privileges for fine-grained actions on specific objects, schema privileges introduced in Oracle AI Database 26ai to enable controlled management within a schema, system privileges that typically apply broadly across the database, and administrative privileges such as SYSDBA, SYSOPER, SYSKM, and SYSDG for high-impact operational control. The GRANT examples make the risk of "scope creep" easy to see: a targeted SELECT on a table stays contained, schema-level selection expands reach within a schema, the expansive SELECT ANY TABLE stretches authority across the database, and SYSKM extends power into encryption key management. The lesson is that changing privilege type can change the security story instantly.

At enterprise scale, a role-centered model becomes the recommended operating pattern. Roles act as named collections of privileges, cutting down repetitive administration, supporting hierarchy through role-to-role grants, and making it easier to keep access aligned with job function as teams reorganize. To understand what a user can truly do, it's important to review role grants alongside system and object grants, not in isolation. Visibility and verification also matter, and tools such as Oracle Data Safe User Assessment and Oracle Audit Vault and Database Firewall are positioned as practical ways to analyze user privileges.

Some of the most useful guidance comes from distinguishing account types and treating their risk profiles differently. The discussion separates DBAs, security administrators, application service accounts, application administrators, developers and testers, analysts, and other specialized users and tools, then pairs each with concrete protections: avoid broad system privileges where possible, use schema privileges to constrain application administration, monitor for misuse patterns such as unusual IP addresses or programs, and audit activity to create accountability. Centralized authorization models round out the picture, including Enterprise User Security with Oracle directory services, Centrally Managed Users integration with Microsoft Active Directory, and token-based access via OCI IAM and Microsoft Entra ID, with an important implementation note that Enterprise User Security is deprecated in Oracle AI Database 26ai.

Turn the page to Separation of Duty, and the pace quickens as Oracle Database Vault adds crisp structure to who can do what, and just as importantly, who must not. The next chapter highlights how Database Vault can block access to user schemas or SQL commands, giving organizations a way to place strong guardrails even around powerful accounts and to separate operational administration from direct data access. It encourages a shift from careful privilege design to a more decisive architecture of controls, where the database becomes an active enforcer of boundaries rather than a passive container of permissions.

Copyright © 2026, Oracle and/or its affiliates

Chapter 7

Enforcing Separation of Duties

Why “who can do what” still breaks security

Many organizations invest heavily in perimeter defenses, identity systems, and monitoring. Yet their most damaging incidents still arrive through the simplest route: a person with legitimate access doing something they should not, whether intentionally or by mistake.

In the database world, that risk tends to cluster around powerful administrative accounts, shared operational responsibilities, and emergency “break-glass” practices that quietly turn into everyday workflows. The outcome is a familiar pattern: audit findings, control exceptions, and uncomfortable questions about whether critical data is truly protected from insider misuse.

At the center of this challenge is a foundational security principle: Separation of Duty.

What separation of duty means & why it matters

Separation of duty (SoD) is the practice of splitting critical tasks and privileges across multiple individuals or roles so no single person can carry out a sensitive process end-to-end without oversight. Put simply, it helps ensure that the same person cannot both “make the change” and “approve the change,” or both “administer the system” and “access the data” without controls.

Compliance initiatives rarely fail because a control doesn't exist on paper; they fail because the control cannot be enforced consistently in day-to-day operations. Regulations and frameworks such as SOX, PCI DSS, HIPAA, GDPR-aligned governance programs, and common internal control standards all converge on the same expectation: Privileged access must be limited, monitored, and constrained to job function.

SoD supports those expectations by reducing the following:

- The risk of fraud and unauthorized changes going undetected because collusion or independent review becomes necessary
- The blast radius of compromised admin credentials because "admin" doesn't automatically mean "data access"
- Audit scope and remediation churn because roles and responsibilities become clearer and demonstrably enforced
- Operational fragility because teams stop relying on shared superuser access as the default way to get work done.

What separation of duty is trying to prevent

Stealing sensitive data via compromised privileged-user accounts is one of the most common attack vectors for database breaches. Although database administrators may prefer a single account for convenience, dividing duties across multiple accounts and knowing exactly which privileges are in use can significantly improve security.

For example, the person who creates a new database account and sets its initial credentials does not need to be the same person who grants that account access to data. Separating account creation from data access prevents any one person from creating a new pathway to data, and requires two people working together to introduce that pathway. Similarly, in encryption key management, the person who can create and rotate encryption keys should not normally be the same person who manages the rest of the database, reducing the chance that both the database (or database backup) and the encryption keys are compromised at the same time.

SoD also reduces human error by increasing the likelihood that mistakes are caught and fixed before they cause harm.

Copyright © 2026, Oracle and/or its affiliates

Separation of duties facilitates compliance

In many environments, privileged accounts act as an exception mechanism that bypasses otherwise well-designed application controls. Even with least privilege at the application layer, database superuser capabilities can undermine SoD goals when administrators can read sensitive tables, disable auditing, or alter security configurations during routine maintenance.

This often creates a gap between “policy intent” and “technical enforcement.” Organizations can describe their SoD model, but proving consistent enforcement across administrators, tools, scripts, and emergency procedures is difficult.

Separation of duties divides tasks among multiple users to reduce fraud, errors, and malicious activity, and to limit the damage from compromised accounts. Cybersecurity and regulatory concerns continue to drive strong controls for insider and privileged administrative accounts.

In the sections that follow, this chapter explores how Oracle AI Database’s capabilities help organizations implement stronger separation of duty, reduce insider risk, and meet compliance expectations with less operational friction, without preventing DBAs from doing the work required to keep the business running.

Separating the duties

How far an organization can separate duties across different people often comes down to scale. If there are only one or two DBAs supporting the entire organization, the separation that can be applied in practice will be limited. In that reality, it is rarely feasible to staff distinct roles like a backup and recovery DBA or a performance tuning DBA.

In these cases, some organizations use a multiple account strategy. One person performs most day-to-day work with one account, and then uses a second account, with different privileges, for a small set of critical functions, such as security administration. The same person has credentials for both, but must log in with the account that matches the task at hand. This simple discipline can reduce human error and shrink the blast radius if an account is compromised, while still leaving the organization without strong protections against a malicious administrator.

The goal is to strike a practical balance: use a best practice approach to separating duties that lets privileged users do their jobs while keeping sensitive data better protected.

Enforce separation of administrative privileges

Most DBA work does not require sweeping, all access permissions. In practice, the majority of routine operations can be handled with narrow, purpose-built privileges, allowing multiple administrators to manage the database without the risk that comes with using the all-powerful SYSDBA administrative privilege. `SYSOPER` is a clear example: it enables limited tasks such as starting and stopping the database, without handing over the full range of `SYSDBA` capabilities. The Oracle AI Database also offers additional administrative privileges, including `SYSBACKUP`, `SYSDG`, `SYSKM`, and `SYSRAC`, to support database backups, Oracle Data Guard administration, key management, and Oracle RAC management, respectively.

A practical way to reduce privilege abuse is to divide database administration into distinct administrative, operational, and security responsibilities, so no single administrator has unrestricted control over the database. The Oracle AI Database does allow all administrative privileges to be mapped to a single operating system (OS) group. However, that approach is not recommended because it does not enforce separation of duties. The stronger best practice is to create a separate OS group for each administrative privilege and associate each group with its corresponding privilege at the database binary level during installation.

For local connections, where an administrator first accesses the database host server and then logs in to the database, the database owner’s OS account (typically 'oracle') should not be used. Instead, administrators should use individual OS accounts that are members of the appropriate OS group or groups, which improves accountability and reduces

Copyright © 2026, Oracle and/or its affiliates

security risk. It is also a best practice to avoid using the client included with the database binary installation and to install a separate client for administrator use.

Control the database owner account

The SYSDBA administrative privilege grants full access to the SYS (database owner) account and should be restricted to use during database upgrades and patching. Change management processes and privileged account management (PAM) systems should be used to control remote access to accounts with this privilege, including the default SYS account.

Top tips for a secure database environment

A strong database security posture requires a few proven techniques. The following outline complementary controls that improve overall security. These common best practices are found consistently in environments where database security is paramount.

Implement a trusted path to data

Even if administrator credentials are compromised, access should be granted only when it comes through a trusted path. That trusted path can include checks like IP address, the program used, time of day, authentication type, and other conditions. Adding these authorization checks makes it much harder for an attacker to succeed with stolen credentials. This is especially important for application service accounts, which are often prime targets.

Enforce least-privilege

Grant database users only the minimum set of privileges required to perform their intended tasks or functions, and no more. Use privilege analysis to identify unnecessary privilege and role grants. When assigning privileges to a user or role, grant specific object or schema privileges rather than broad system privileges that enable access to all objects in the database. Similarly, create database roles with only the privileges needed for a particular function instead of using powerful roles such as the built-in DBA role. Granting one or more task-specific roles to a user enables a closer fit to the work the user needs to perform without adding unnecessary privileges. This least privilege model helps reduce the blast radius of a compromised account by limiting what an attacker can do within the database.

Do not share accounts

Sharing database accounts for convenience is common, but it comes with real cost. It removes accountability and increases risk because multiple people can use the same credentials. Each DBA should have an individual, named account with carefully tailored privileges and roles. If there are many DBAs or many databases, consider simplifying administration by centralizing database authentication and authorization in an external identity service such as Active Directory.

Safeguard the audit trail

Audit records provide a record of actions on a database, directory, or operating system. Activity by privileged administrators should be audited. Details such as the privileged user actions taken (`CREATE USER, CREATE ANY TABLE, ALTER SYSTEM, ALTER SESSION`), along with event context such as initiating IP address, event time, and the actual SQL statement, are examples of the audit information commonly needed for compliance and forensic reports. The `AUDIT_ADMIN` role grants privileges to change audit policies and purge audit records. This role should not normally be granted to anyone other than the security DBA, and to any tools or services used to collect and aggregate audit data. Further, if audit records can be created when administrative users access application accounts, alerts can be raised for further action.

Copyright © 2026, Oracle and/or its affiliates

Mandatory access controls with Oracle Database Vault

Oracle Database Vault is designed to help customers move from SoD as a guideline to SoD as an enforceable, testable control inside the database. Rather than relying solely on trust, process, or after-the-fact detective measures, Database Vault helps establish durable boundaries around sensitive data and administrative actions, ensuring that even highly privileged users operate within defined rules aligned to role and business need.

Database Vault is one of the most powerful security options for Oracle AI Database26ai. Database Vault is a multi-purpose access control mechanism that addresses many different use cases. Three primary use cases for Database Vault are:

- **Blocking privileged user access to data**: Database Vault prevents privileged users, such as DBAs, from accessing sensitive application data. This is crucial for meeting regulatory requirements and reducing risk. Even if DBA credentials are compromised, sensitive data remains protected, minimizing the potential for data breaches.
- **Enforcing a trusted path**: Database Vault ensures that application service accounts, often targeted by attackers, can only connect to the database under strict conditions. Rules can restrict these accounts to specific IP addresses, programs, or operating system users.
- **Controlling SQL command execution**: Database Vault allows precise control over database commands, enabling conditions like requiring multiple administrators for critical actions (e.g., no table can be dropped unless two DBAs are logged in) or restricting user actions to specific times or circumstances.

Controlling administrator access

In most cases, there is no reason for database administrators to have access to the application schemas and data. Using the different administrative privileges does not prevent the administrators from accessing user-specific data. Read on to see how to enforce separation of duties further and break the link between database administration and data access.

The out-of-the-box DBA role is granted most system privileges, including the powerful `SELECT ANY TABLE` privilege. With this privilege, a DBA can view any data in the database, including salary, taxpayer IDs, phone numbers, corporate financial forecasts, intellectual property, and other sensitive data. If a cyber-criminal successfully compromises the credentials of a DBA, they can now easily access and exfiltrate your valuable data.

Other privileged accounts may also be granted `SELECT ANY TABLE`, frequently by administrators who are trying to quickly solve problems related to their access. Complex applications with multiple schemas frequently use system privileges instead of object privileges, making it possible for exploits like SQL injection attacks to access data throughout the database. Like DBAs, these accounts create an elevated risk to data and are prime targets for attackers.

Oracle Database Vault can restrict privileged users, including database administrators, from accessing sensitive data with an access control mechanism called a *security realm*. Security realms are collections of schemas or specific sensitive objects where access is controlled by Database Vault. Database Vault policies override system privileges like `SELECT ANY TABLE` for objects and schemas protected by a security realm. With Database Vault, database administrators can manage the database but cannot access sensitive schemas/objects protected by realms.

Copyright © 2026, Oracle and/or its affiliates

Figure 7-1 shows an example of a realm protecting sensitive data within the human resources (HR) schema. The Database Vault realm prevents users with the powerful DBA role from accessing data inside the HR realm yet still allows them to do their database administration tasks like creating indexes. Business users with a valid requirement to view data are still able to access data because they have permissions within the realm.

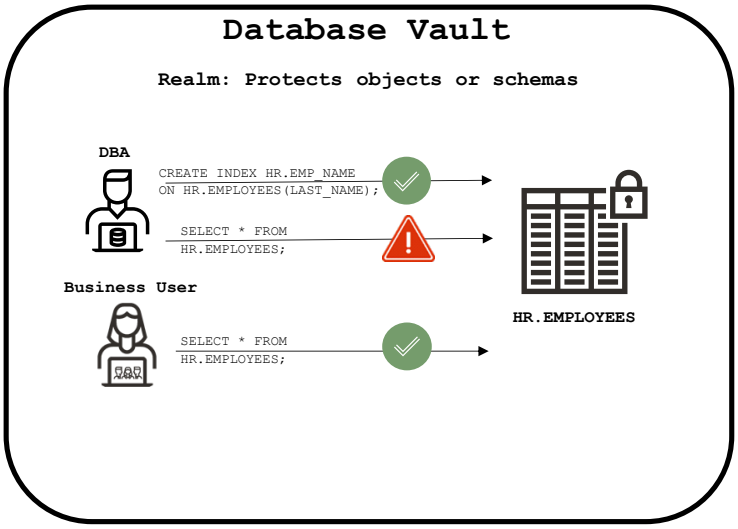


Figure 7-1: Database Vault Realms block administrator access to data

Enforcing a trusted path

Access within a realm is controlled by Database Vault rule sets. A rule set is one or more programmatically evaluated rules that define the conditions for access. In the example shown in Figure 7-2, the rule set includes two rules: one checks the source IP address of the user connection, and the other checks that the time-of-day is within approved business hours. The business user must satisfy BOTH rules to access data in the realm. By enforcing these rules, Database Vault shrinks the attack surface and enables a trusted path to data through multi-factor authorization policies, further reducing the risk from compromised accounts.

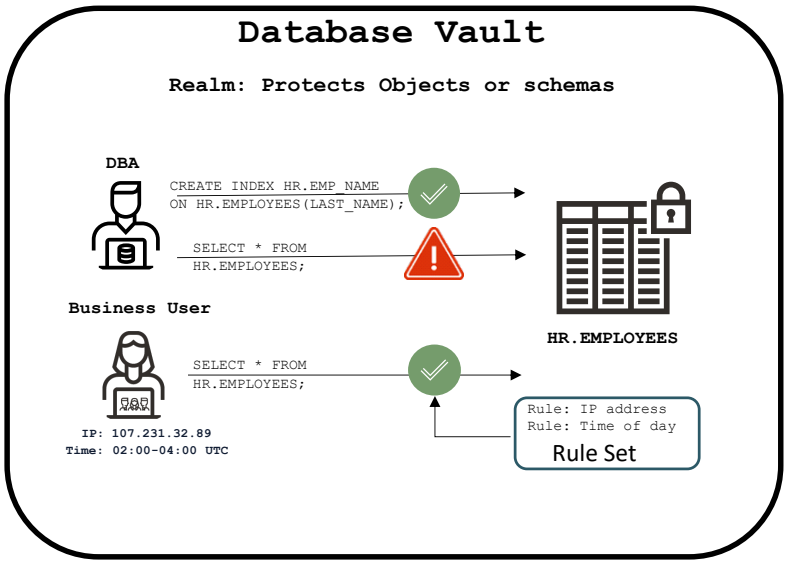


Figure 7-2: Database Vault uses rules to make access control decisions

Copyright © 2026, Oracle and/or its affiliates

A database can contain many different Database Vault realms, each with its own set of authorized users controlled by rule sets appropriate to those users and the realm. In Figure 7-3, we can see that the human resources application administrator can manage the HR data but cannot access objects or data within the FIN realm, even though the HR application administrator has broad system privileges. Similarly, the highly privileged finance application administrator can view data and objects within the FIN realm but not within the HR realm. Thus, Database Vault enforces separation of duties at the database object and schema level, even overriding system privileges, and isolates the DBA from application data.

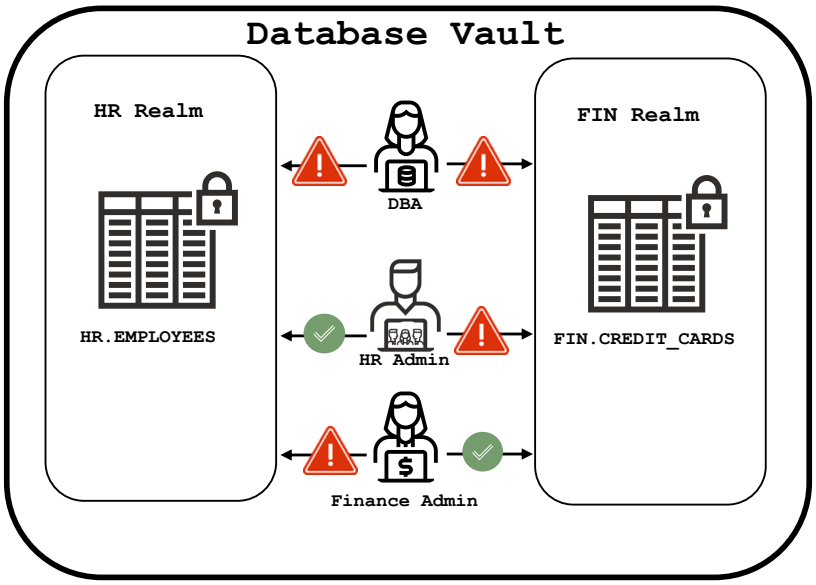


Figure 7-3: Oracle AI Database can have multiple realms in the same database

By default, realms are overridden by schema ownership and direct object grants. In Figure 7-3 above, that would mean that the HR account (the schema owner of the EMPLOYEES table) would be able to access data in the EMPLOYEES table and the FIN account (owner of the CREDIT_CARDS table) would be able to see data in the CREDIT_CARDS table even though they were not authorized within the realms. Sometimes this is the desired behavior, but there are situations where the realm needs to control ALL users, regardless of ownership or direct object privileges. Having the realm override object privileges can be valuable when an administrator must make changes during an application upgrade or patching. In some upgrades, only application procedures and functions need to be modified, not the tables and views that store sensitive data. In that case, stronger protection can be applied to sensitive tables and views while still allowing access to update the procedures.

For this use case, Database Vault realms can be made “mandatory,” meaning not even the schema owner can access objects protected by the realm unless access is explicitly granted through Database Vault. Mandatory realms override both schema ownership and direct object grants. With a mandatory realm, no one is allowed access, even to objects they own, unless they are specifically granted membership to the realm through Database Vault.

With Database Vault, the only user who can grant access to protected data is the security administrator, not the data owner. This simplifies proving who has access because auditors do not need to trace a potentially complex chain of direct object grants across users and roles. Instead, permissions to access or manipulate objects or schemas are maintained in clear, easy-to-understand security realms.

Copyright © 2026, Oracle and/or its affiliates

Separation of duties within Database Vault

As you might expect, separation of duties is baked into Database Vault. Enabling Database Vault creates additional roles intended to enforce the separation of duties. These roles include:

- **Database Vault Account Manager** provides out-of-the-box enforced separation of duties that controls the creation and management of users in the database along with their passwords. By default, when Database Vault is enabled, the DBA role loses the ability to create users or change their passwords.
- **Database Vault Administrator and Owner** roles allow a security administrator to create, modify, and manage Database Vault security controls, including the ability to manage realms and add realm participants. The Database Vault Administrator role controls the various PL/SQL packages associated with Database Vault. The Database Vault Owner role includes the Database Vault Administrator role, but also allows a user with the role to manage and monitor the Database Vault configuration.

This split between database administration and user administration helps block a common attack pattern: an attacker steals DBA credentials, then uses that legitimate access to create a rogue user and grant powerful privileges, making it easier to steal sensitive data. That rogue account can also serve as a convenient way back into the database if the DBA credentials are later changed.

Because of this, Database Vault role membership needs careful control. If an organization cannot support a separate security administrator, and SoD must be relaxed, the Database Vault roles can be granted to the DBA role so DBAs can manage both users and Database Vault policies. The tradeoff is important: granting these roles back to the DBA role also brings back the risk that a compromise of a database administrator account could result in a compromise of sensitive data.

The database security assessment tools report which users hold Database Vault roles. Data Safe and Audit Vault and Database Firewall also detect drift in Database Vault role assignments, making changes straightforward to spot.

Oracle Database Vault can also separate who can view and manage audit records from privileged users such as database administrators. This additional separation matters because purging audit records is a common way to hide malicious or inappropriate activity. Oracle AI Database 26ai adds Database Vault control over the `AUDIT_ADMIN` and `AUDIT_VIEWER` roles.

The database administrator team cannot access application data unless they are explicitly granted permission within the realm that protects the data, but they can still perform common DBA tasks like performance tuning, memory management, backup, and others. Certain DBA tasks like exporting data or scheduling automated jobs may expose sensitive data to the DBA and will require additional authorization by the Database Vault security administrator.

Separation of duties with containerized databases

Oracle introduced the concept of containerized databases in Oracle Database 12c. In a containerized environment, a container root database (CDB) stores common metadata and settings that apply across the containerized databases. A pluggable database (PDB) is a collection of schemas, schema objects, and non-schema objects that appears to an application or user as a separate database. A single CDB can host many PDBs. The containerized architecture offers numerous benefits, including simpler maintenance, portability across servers, and increased resource utilization with lower overhead.

From a security standpoint, the container architecture adds value by providing another way to separate data and users into different pluggable databases. Each pluggable database can maintain its own policies, administrators, and configuration.

Starting with Oracle Database 19c, Database Vault's Operations Control feature can block common users, such as infrastructure DBAs, from accessing application data in pluggable databases. Using Database Vault Operations Control to separate common users from local data in PDBs does not require creating separate security realms. As

Copyright © 2026, Oracle and/or its affiliates

organizations adopt Oracle AI Database 26ai, which always uses container databases, Operations Control becomes more important than ever.

With Operations Control, the Database Vault administrator can do the following:

- Perform all administrative activities on the container database.
- Perform all administrative activities on the pluggable database, except accessing or managing local user objects or local user data.
- Grant a common administrator an exception to access PDB local data or exempt a PDB from Operations Control enforcement.

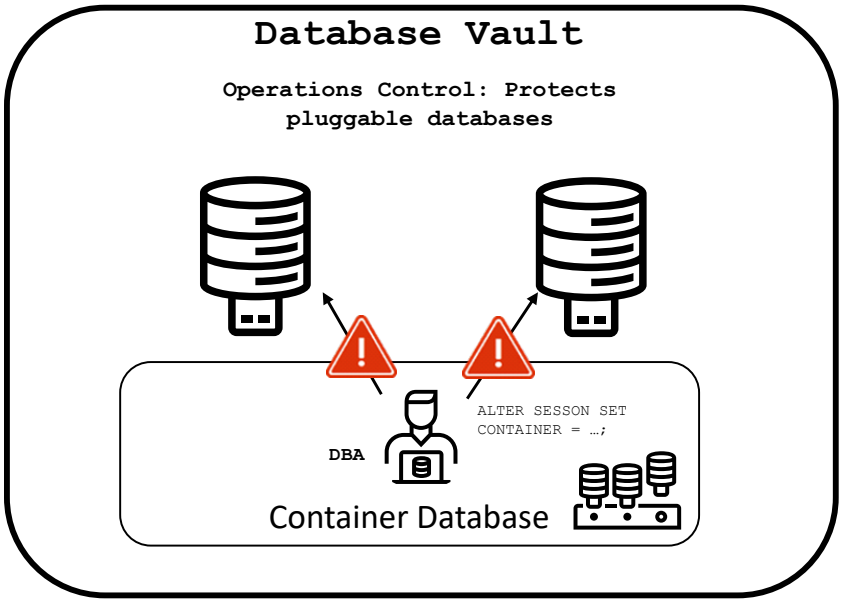


Figure 7-4: Operations Control protects pluggable databases

Copyright © 2026, Oracle and/or its affiliates

Control SQL database commands with Oracle Database Vault

This topic is not strictly part of separation of duties, but it fits naturally in a discussion of Database Vault. Many organizations have change management policies that prohibit dropping tables, modifying stored procedures or other database objects, and changing system parameters outside defined maintenance windows. Database Vault can enforce these change management policies, reducing the risk of accidental or malicious violations. This matters because human error is the most frequent cause of production outages, and DBAs often work across dev, test, and production terminal sessions where mistakes are easy to make.

Database Vault provides fine-grained SQL command rules to help prevent accidental, or malicious, activity. These command rules can either prohibit specific commands or use rule sets to define the conditions under which commands are permitted. As with realms, the rule sets attached to a command rule can be based on anything that can be checked programmatically. For example, changes to production schemas, such as dropping a table, can be prohibited unless two application administrators are logged in at the same time, the person executing the change connects from an approved jump server, the work occurs outside normal business hours, and a change number is set for the session.

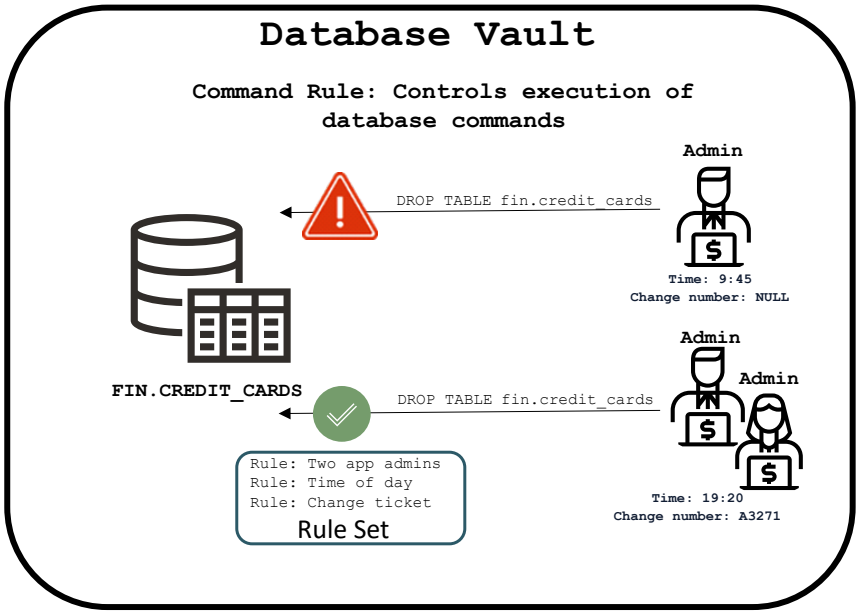


Figure 7-5: Example Database Vault Command Rules

As another example, DBAs can be limited to running a `DROP TABLE` command only from their work desktops and only during working hours, reducing the risk of unauthorized remote access after-hours. Other policies might restrict SQL commands such as `CREATE`, `TRUNCATE`, or `ALTER TABLE` to maintenance windows. Database Vault can also enforce conditions like requiring that changes must be associated with a trouble ticket number or ensuring that two DBAs have logged in at the same time before a particularly risky activity (like rekeying the database) can succeed.

Break-glass scenarios

As a general principle, operating system administrators, database administrators, and developers should have limited or no access to production data. However, access is sometimes necessary, such as when performing application upgrades or troubleshooting application schema or application data issues.

When these scenarios arise, administrative and development staff should have a defined mechanism to gain access to the production schema and production data. This is commonly referred to as a break-glass scenario. When an administrator breaks the glass, they temporarily acquire elevated privileges to complete a specific task. This can be

Copyright © 2026, Oracle and/or its affiliates

done by checking out a privileged account and using it to perform the required actions, or by enabling an Oracle AI Database role that grants a day-to-day account a higher level of database privileges and access to production data.

In the first approach, where an account is checked out, a highly privileged user account can be stored in a Privileged Access Management (PAM) solution. The person needing access checks out the account, completes the task, and then checks the account back into the PAM solution. In this scenario, Database Vault can be used to constrain where and how the checked-out account can be used. For example, the account might be permitted to connect only from a specific subnet or a specific trusted host. The checked-out account might be allowed to run `CREATE` or `ALTER` commands but not `DROP TABLE`, `TRUNCATE TABLE`, or `DELETE` commands on production schemas. This ensures the account can see the data but cannot perform destructive actions on the data. Once the emergency is resolved, the checked-out database account can be automatically locked by the PAM solution.

In the latter scenario, the person needing access can enable a database role called a secure application role. Secure application roles are not granted directly to a user. Instead, they are enabled by running a PL/SQL procedure that must execute successfully for the role to be activated (read more about secure application roles).

In this approach, the role can be enabled only if the requesting database user meets the criteria defined in a Database Vault rule set. For example, the rule set might allow the role to be enabled only when the user connects from trusted hosts, IP addresses, or subnets. The rule set can also restrict enablement to specific hours or specific days. Access can be further limited to a defined list of database users, such as only senior database administrators.

When the emergency is resolved, the database user logs off and the role is automatically revoked until it is needed again. This model is similar to using sudo on Linux systems or “open as Administrator” on Windows systems: the privileges are available, but they are not continuously enabled.

In day-to-day operations, each administrator should use a least-privilege, named account with tightened domain-specific roles. Break-glass privileges are specifically for unplanned crises and controlling them behind robust administrative policies helps maintain a secure database environment. This way, if an attacker compromises ordinary credentials, they still face the barriers of multiple approval levels, rotating credentials, and the blend of preventative and detective controls that enable a defense-in-depth approach.

Operationalizing Oracle Database Vault

There are three key elements to consider when operationalizing Oracle Database Vault:

- What needs to be protected, and under what conditions.
- Which SQL commands should be allowed in your database environment.
- What process and organizational changes are needed because of the newly enforced separation of duties.

Before defining Oracle Database Vault realms, it is important to first identify what data should be treated as sensitive. Some organizations take a broad stance and assume all application data is sensitive, creating a realm for every application schema. Others prefer a more targeted approach and use sensitive data discovery to pinpoint exactly what needs protection. Either path can work, and Database Vault can support both.

To decide which SQL commands should be allowed, many teams start with audit records. Audit trails can show which critical SQL commands, such as `ALTER SYSTEM`, `ALTER DIRECTORY`, and related actions, were executed, along with the context around them, including IP address, program name, database, or operating system username. Database Vault realms and SQL command rules can then be used to constrain those commands by factors like time of day, source IP address, and other criteria, helping ensure changes align with policy.

At the same time, special attention is needed for separation of duties inside Database Vault, since Database Vault is the mechanism used to enforce separation of duties across the rest of the database. Many organizations roll this out in stages. For example, Database Vault roles might be granted to the DBA role at first and then separated later as

Copyright © 2026, Oracle and/or its affiliates

processes mature. This approach can accelerate early control adoption while spreading SoD change over a longer time frame.

Some database teams hesitate to enable preventive controls like Database Vault because they worry about outages or breaking administrative scripts. To help reduce that risk, Database Vault offers a simulation mode. In simulation mode, realms and SQL command rules are evaluated, but instead of blocking activity, Database Vault logs policy violations. This enables full end-to-end regression testing without disrupting application operations. After testing confirms that legitimate activity would not have been blocked, policies can be moved from simulation mode to blocking mode.

Finally, violations need to be captured and alerted. These audit records are especially valuable because Database Vault alerts can signal either early probing by a malicious user or a need for additional training on how to access sensitive data appropriately. Data Safe and AVDF can both be used to collect and report on Database Vault violations.

Database Vault for vibe-coded apps and privileged MCP database access

Vibe-coded applications, built quickly through iterative experimentation, often inherit risky defaults, including reused credentials, overprivileged database connections, and minimal security hardening. What begins as a rapid prototype can become a direct pathway to sensitive data. A single SQL injection flaw, leaked credential, or coding error can expose entire databases. Opening an MCP connection with a privileged user further amplifies risk by introducing a new, automatable pathway that can rapidly enumerate objects, run high volume queries, and exfiltrate sensitive data. Returned results may also be captured in agent and tool logs outside approved data handling controls.

Oracle Database Vault mitigates these risks by enforcing policy at the database layer, independent of application code quality. With Database Vault realms, sensitive schemas and tables, such as PII, finance, and HR, can be protected while also overriding powerful system privileges for example blocking access even when an account has `SELECT ANY TABLE`.

With mandatory realms, explicit authorization can be required even for schema owners, preventing “god accounts” and service accounts from accessing protected data by default. With trusted path rule sets, privileged service accounts, including MCP accounts, can be restricted by where they connect from, such as approved subnets or hosts, time windows, and client attributes.

With command rules, high-risk SQL, including DDL, privilege changes, and destructive operations, can be prohibited or tightly constrained so compromised privileged credentials cannot be used to alter security posture or damage data. Together, these controls enforce separation of duties and can control, or completely forbid, privileged access to sensitive data from vibe-coded applications and MCP tooling.

Database Vault and packaged applications

Several Oracle and third-party applications are certified to work with Oracle Database Vault, including Oracle Fusion Human Capital Management, Oracle Enterprise Resource Planning Cloud Services, Oracle E-Business Suite, Oracle PeopleSoft, and SAP. Please review the certification matrix on Oracle Support for more information about application certifications and the Oracle Cloud website for cloud support.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- DB Security- Database Vault

Copyright © 2026, Oracle and/or its affiliates

Summary

Separation of duties (SoD) is framed as a practical response to a stubborn database security reality: many of the biggest incidents don't start with "unauthorized access." They start with legitimate access being misused, especially when powerful administrative accounts and break-glass procedures quietly drift into everyday habits. SoD tackles this by splitting critical tasks and privileges so no single person can run a sensitive process end-to-end without oversight, cutting directly into insider risk and privileged access exposure.

From a compliance perspective, the message stays simple and grounded. Programs rarely fail because a policy doesn't exist. They fail when enforcement erodes in day-to-day operations. SoD supports regulatory expectations to limit and monitor privileged access, while also reducing fraud exposure, containing the impact of compromised credentials, and minimizing audit and remediation churn.

The principle becomes workable through patterns that fit real-world constraints, including smaller teams that can't cleanly separate every role. Practical options include using multiple named accounts, moving away from broad superuser habits toward narrower administrative privileges, and adding OS-level and process controls that improve accountability. Reinforcing best practices like least privilege, trusted path requirements, avoiding shared accounts, and protecting audit administration further reduces the odds that privileged users can bypass controls or erase evidence.

Oracle Database Vault is positioned as the mechanism that makes SoD enforceable and testable inside the database. Realms and rule sets can override broad system privileges, trusted path controls can limit how privileged access is exercised, and SQL command rules can backstop change management policies. Rollout guidance keeps expectations realistic: adopt in stages, use simulation mode to log violations before blocking, use container database Operations Control to isolate pluggable databases, and govern break-glass access with constrained privileged accounts or secure application roles.

The next chapter, SQL Injection Attacks, moves from who can touch the data to how attackers slip through the smallest cracks in query logic and privileged pathways, turning convenience into catastrophe. Stay with the narrative, because the discussion is poised to shift from guarding roles and controls to confronting the fast, sharp mechanics of exploitation, where a single flaw can suddenly illuminate every table in the room.

Copyright © 2026, Oracle and/or its affiliates

Chapter 8

Minimizing Risk from SQL Injection

The persistent threat of SQL injection

SQL injection remains one of the most common and damaging routes attackers use to get to data. Even after decades of attention, it continues to show up in data-driven applications with surprising regularity. The Open Worldwide Application Security Project (OWASP) has consistently ranked SQL injection among the top ten web application vulnerabilities, holding the number one position on this list since its inception in 2004.

This staying power deserves a closer look. SQL injection persists largely because it is carried out through applications, which by design, must be accessible to users and, in many cases, to the entire internet. Applications act as the public-facing gateway to data, and attackers take full advantage of that necessary exposure.

In theory, this vulnerability should be behind the industry. Developers have been taught about SQL injection risks for years, and automated code scanning tools routinely flag potential injection points during development. Yet despite the training and the tooling, SQL injection remains a stubborn threat to data-driven web applications. The distance between what is known and what is consistently done points to an uncomfortable truth: producing perfectly secure code across large, complex applications, every time, is extraordinarily difficult.

Why SQL injection demands your attention

SQL injection is rarely just an isolated data access issue. More often, it becomes the opening move in a broader compromise. Attackers use SQL injection vulnerabilities to steal sensitive information, corrupt critical data, destroy database structures, and then pivot laterally from the database into other systems across the network.

The blast radius of a successful SQL injection attack can be significant. In many multi-tier architectures, applications connect to databases through powerful service accounts that have broad schema access and procedure execution rights. When an attacker exploits a SQL injection flaw, malicious queries run with the full privileges of that application account, which can put an entire data estate at risk.

This chapter explores two complementary ways to reduce SQL injection risk at or near the database layer: the network-based Database Firewall in Oracle Audit Vault and Database Firewall (AVDF), and the database resident SQL Firewall built into Oracle AI Database 26ai. It compares how these technologies work, outlines when each approach is the better fit, and provides decision criteria to support selecting the option that best aligns with the environment and requirements.

Understanding the SQL injection attack pattern

SQL injection is fundamentally an application layer vulnerability, not a database vulnerability. This distinction matters because it makes clear where ownership sits and where defenses need to be focused. The weakness appears when an application treats untrusted input as executable SQL, giving an attacker a way to inject and run their own statements through the application to read, modify, insert, or delete data and, in the worst cases, destroy the database entirely.

SQL injection attacks come in many variations, including blind injection, time-based injection, union-based injection, and more. However, they typically follow the same underlying attack pattern, as illustrated in Figure 8-1.

Copyright © 2026, Oracle and/or its affiliates

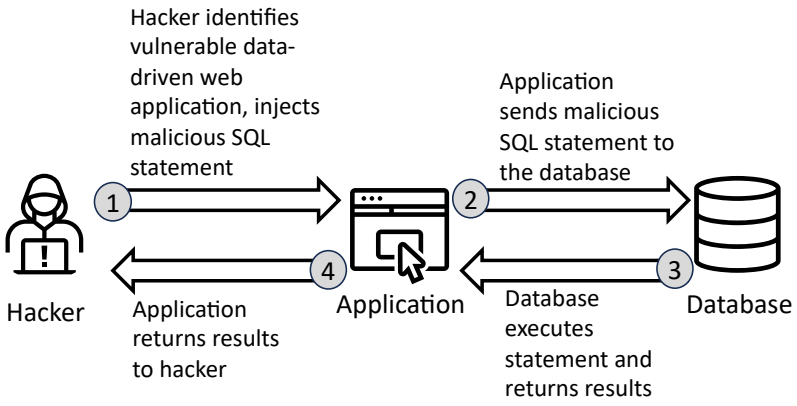


Figure 8-1: Example outlining SQL injection attack patterns

In multi-tier architectures, applications usually enforce access control in the middle tier, using the authenticated user's identity and authorization level to decide what they can do. SQL injection attacks sidestep these controls by slipping malicious SQL into input fields or URL parameters in data-driven applications. Because the injected statements run with the application account's privileges, an attacker can effectively bypass web application authentication and authorization altogether and gain direct ability to read, change, or delete sensitive data, including customer information, personal health records, trade secrets, intellectual property, and financial data.

This built-in privilege escalation is what makes SQL injection especially dangerous. A low privilege attacker who successfully exploits SQL injection can inherit the full privileges of the application database account, which often includes access to the entire schema and all stored procedures.

A defense in-depth strategy against SQL injection

Preventing SQL injection calls for a defense in-depth strategy that protects both upstream and downstream of the application tier. It can be viewed as building concentric rings of protection: each layer is designed to catch what the previous layer missed, steadily shrinking the attack surface.

Upstream protection (closest to the application user) typically includes deploying a web application firewall (WAF) to inspect incoming HTTP traffic for known SQL injection signatures. A WAF can block many common attack attempts before they ever reach application code. That said, signature-based tools come with built-in limits. They can miss novel attack patterns, zero-day exploits, and more sophisticated evasion techniques. A WAF reviews HTTP payloads, but it cannot assess the actual SQL with full database context.

Downstream protection focuses on the SQL the application actually issues, identifying and blocking malicious commands that make it past edge controls. This layer helps close gaps and reduces the blast radius when application vulnerabilities are present. This is where database firewalls deliver critical value.

As shown in Figure 8-2: Effective approaches to combating SQL Injection, the strongest strategy uses both approaches: block as much as possible at the edge with a WAF, while also protecting the database to contain risk when application vulnerabilities inevitably exist.

Copyright © 2026, Oracle and/or its affiliates

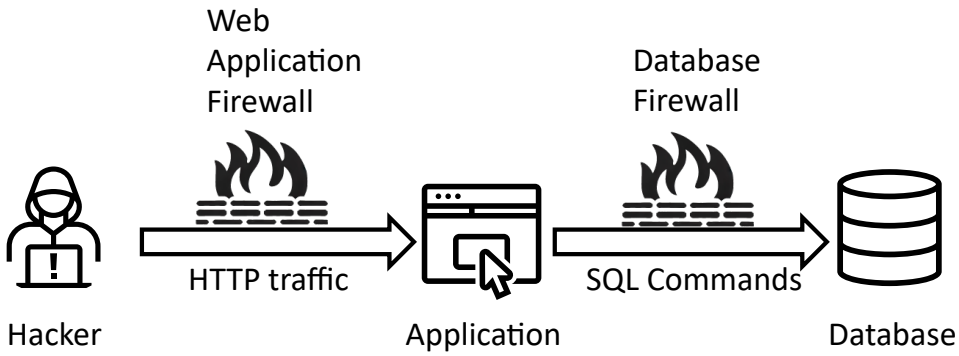


Figure 8-2: Effective approaches to combating SQL Injection attacks include the use of multiple firewalls

Database firewall solutions from Oracle

Oracle Database security offers two distinct approaches to mitigate SQL injection risk for data-driven applications:

- **Network-based Database Firewall** in Oracle Audit Vault and Database Firewall
- **Database-resident SQL Firewall** operating inside the database kernel (introduced in Oracle AI Database 26ai)

Unlike web application firewalls that rely on regular expressions and signature matching, these database-centric approaches treat SQL injection prevention as an access control problem. Each solution learns your application's typical SQL patterns during normal operation, builds an allowlist of permitted statements, and rejects or alerts on anything that deviates from that established profile.

AVDF applies this protection at the network layer, positioning itself between the application and the database. SQL Firewall brings the same control directly into the database kernel, eliminating the complexity of routing database traffic through a separate appliance. Both options deliver powerful protection; the choice between them depends on your specific environment, requirements, and constraints.

Both network-based and database-resident firewalls address these core use cases:

- Provide real-time protection by restricting execution to authorized SQL statements arriving from trusted connections
- Mitigate risks from SQL injection attacks, anomalous access patterns, and credential theft or abuse
- Enforce trusted database connection paths by blocking unauthorized connection sources

Network-based Database Firewall in Oracle AVDF

Database Firewall, a core component of AVDF, acts as the database's first line of defense at the network layer. It continuously monitors incoming SQL traffic and enforces expected database behavior to help prevent SQL injection, application bypass, and other malicious activity. From a centralized deployment, a single Database Firewall instance can protect multiple heterogeneous databases across an environment, including Oracle Database, MySQL, Microsoft SQL Server, SAP Sybase, and IBM Db2.

As illustrated in Figure 8-3: Example showing how Oracle Database Firewall blocks SQL injection attempts, an attacker may be able to slip malicious input into an application field that the application fails to block, but Database Firewall can still stop the attack. If the SQL generated by the application does not match the allowlist of permitted statements, Database Firewall blocks the request before it reaches the database. It then sends detailed information about the attempted violation to the Audit Vault Server, which generates alerts and provides a centralized platform for analysis and reporting.

Copyright © 2026, Oracle and/or its affiliates

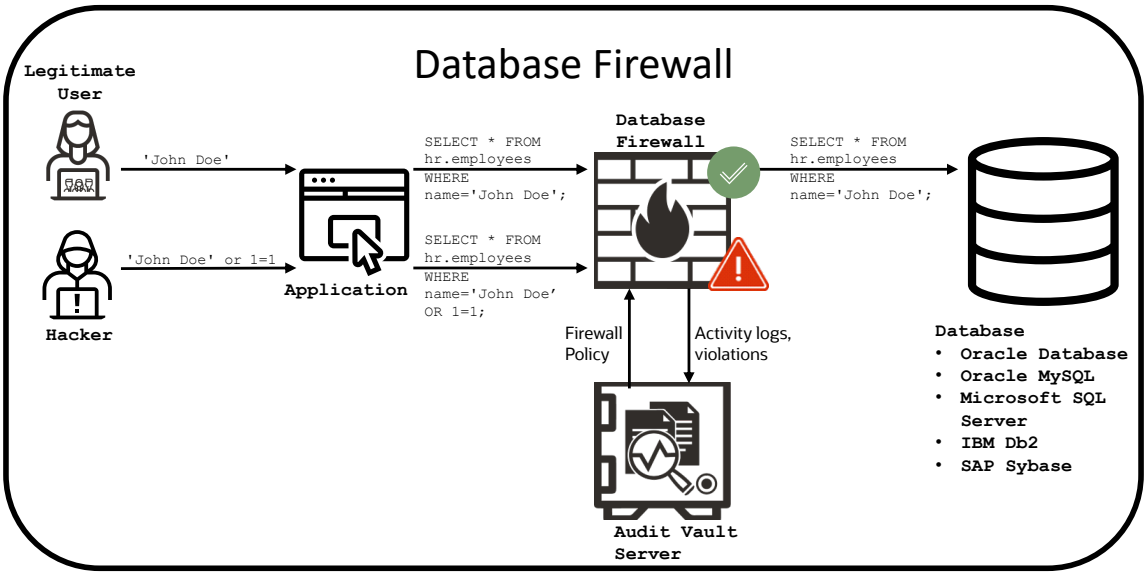


Figure 8-3: Example showing how Oracle Database Firewall blocks SQL injection attempts

Deployment modes for Database Firewall

Database Firewall supports three distinct deployment modes, each offering different capabilities and deployment characteristics. Table 8-1 summarizes these options.

MODE	Description	Monitoring or blocking
Proxy	All traffic to the database server is routed through the Database Firewall, including return traffic.	Monitoring and Blocking
Host Monitor	An agent deployed on the database server captures SQL traffic destined for the database and securely forwards it for analysis.	Monitoring only
Out-of-band	Database Firewall listens to the network traffic sent to the database. Several technologies, such as span ports, port replicators, etc., can send a copy of the database traffic to the Database Firewall.	Monitoring only

Table 8-1: Deployment modes of Database Firewall

Choose the deployment mode based on whether the goal is detection only or detection and blocking. If the priority is visibility without enforcement, any of the three modes can monitor traffic and surface policy violations. All modes work by sniffing network traffic and evaluating SQL statements against Database Firewall policies.

If the requirement is to actively stop malicious activity, Database Firewall must be deployed inline using proxy mode:

- **Proxy mode**: Database traffic flows through the Database Firewall, which inspects SQL in real time as it passes through. This is the only mode that can block unauthorized SQL. Proxy mode can be configured initially for monitoring only, generating alerts on violations while the allowlist is refined. After the allowlist

Copyright © 2026, Oracle and/or its affiliates

accuracy has been validated, it can be switched to blocking mode to prevent SQL injection and other unauthorized SQL from reaching the database.

- **Host monitor**: An agent installed on the database server captures incoming SQL and relays it to the Database Firewall appliance for analysis. This mode provides comprehensive monitoring and alerting without requiring changes to network paths or client configurations.
- **Out of band mode**: Database Firewall inspects a copy of network traffic from the database network segment. In most environments, a network switch or router mirrors traffic to the firewall using a SPAN port, mirror port, or network tap. This mode provides broad monitoring coverage without placing the firewall in the data path, helping ensure zero impact on database performance and availability.

How Database Firewall evaluates SQL traffic

Database Firewall inspects SQL traffic flowing into the database and decides whether to allow, log, alert, substitute, or block each statement. It evaluates traffic through multiple inspection stages, looking at the source IP address, database username, operating system username, client program name, SQL statement category (DDL or DML), and the database objects being accessed. This layered review determines whether a statement should be logged, trigger an alert, proceed to the database, or be blocked.

Database Firewall supports both allowlist and denylist policies, but allowlists generally deliver stronger protection. The universe of potentially malicious SQL patterns is effectively unlimited and constantly changing, while the set of statements an application legitimately runs is finite and usually well understood. Put simply, it is easier and safer to define what should pass than to predict everything that should not.

Managing Database Firewall policies

Database Firewall collects violation logs and activity records and forwards them to the Audit Vault Server for centralized analysis, alerting, and reporting.

Database Firewall policies are managed centrally through the Audit Vault Server console or its APIs, with no separate Database Firewall management console. For comprehensive information on Database Firewall policy configuration, refer to the *Oracle Audit Vault and Database Firewall Auditor's Guide*. The technical paper "Mitigating Risks from SQL Injection" provides additional practical guidance.

Database-resident SQL Firewall in Oracle AI Database 26ai

SQL Firewall is a new capability in Oracle AI Database 26ai that helps reduce SQL injection risk through kernel level inspection and enforcement. It is available by licensing either Oracle Database Vault or Oracle Audit Vault and Database Firewall (AVDF), with no need to license both. SQL Firewall is included with either license.

Like Database Firewall, SQL Firewall mitigates risk from web application attacks such as SQL injection against data-driven applications. The key difference is architectural: SQL Firewall runs inside the database kernel rather than as a separate network appliance. Because it operates within Oracle Database, close to where the data resides, it sharply reduces bypass risk. It inspects every incoming SQL statement, whether local or over the network, encrypted or clear text, and ensures the database executes only SQL that has been explicitly authorized to arrive through trusted connection paths.

SQL Firewall also uses an allowlist model. Running inside the database allows it to evaluate both session context and internal database metadata. With this richer view of allowlist rules, metadata, and session context, SQL Firewall can stop more sophisticated evasion techniques, such as encoded SQL, references to synonyms, or dynamically generated object names that may look different but resolve to the same underlying objects.

Copyright © 2026, Oracle and/or its affiliates

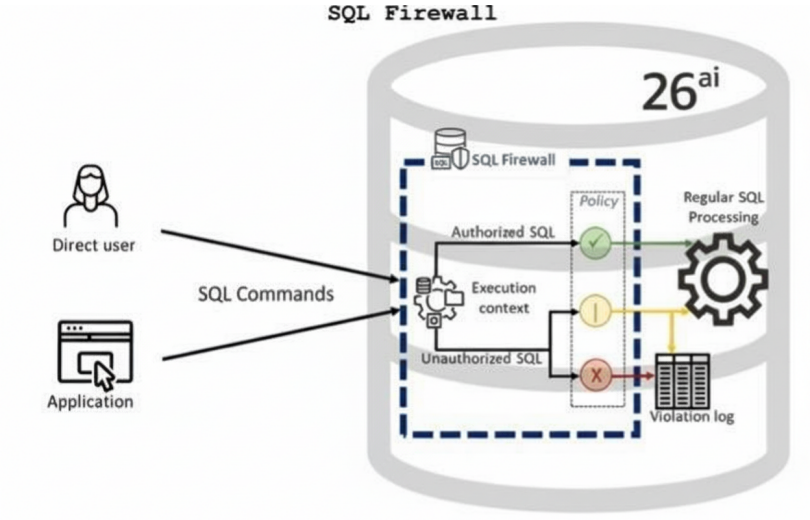


Figure 8-4: Oracle SQL Firewall is built into the Oracle AI Database 26ai kernel

SQL Firewall can monitor or block unusual access patterns, helping reduce risk without adding significant overhead. For each database connection, it validates the source IP address, operating system username, and client program name. If an account connects from a previously unseen IP address or uses an unexpected client program, SQL Firewall can flag or block the session based on the configured policy.

SQL command rules typically need longer training periods to capture the full set of permitted statements across workloads. Session context rules, by contrast, are usually quick to establish and require minimal training because they focus on connection attributes rather than SQL variations. Session context rules can also be defined manually when the environment calls for explicit control.

Unlike Database Firewall, which is heterogeneous and supports many database platforms and versions, SQL Firewall is built into Oracle Database and works only with Oracle AI Database 26ai and later versions.

Reporting on SQL Firewall violations

SQL Firewall records violations in the `DBA_SQL_FIREWALL_VIOLATIONS` view using Oracle Database fast ingest architecture to help minimize performance impact. If needed, violations can also be audited, protecting the records with the audit trail's tamper evident controls.

In multi database environments, violation logs should be collected for consolidated analysis, alerting, and reporting. Oracle Data Safe and Oracle Audit Vault and Database Firewall both provide centralized platforms for this purpose.

Administering SQL Firewall

SQL Firewall can be administered through two interfaces: command line management using PL/SQL procedures, and graphical management through Oracle Data Safe. Both options support enabling SQL Firewall, creating and managing allowlists, and managing violation logs. The best fit usually comes down to whether the environment favors scripted automation or visual, guided workflows.

Deciding between Database Firewall or SQL Firewall

Database Firewall operates as a network-based SQL firewall that monitors SQL traffic for both Oracle and non-Oracle databases. For non-Oracle databases, or Oracle Database versions earlier than 26ai, AVDF is the only option. For Oracle AI Database 26ai and later, use the following decision factors to determine which approach best fits the requirements:

Copyright © 2026, Oracle and/or its affiliates

Performance impact: With a network-based Database Firewall, the database server sees no added CPU load because SQL analysis runs on the firewall host. In host monitor and out of band deployment modes, Database Firewall introduces zero performance overhead on the database. In proxy mode, minor network latency may appear depending on where the firewall sits relative to the application and database tiers. SQL Firewall introduces no observable latency but does add minimal CPU overhead, typically not exceeding 1.0 to 1.5 percent.

Deployment options: Database Firewall runs on dedicated hardware or a virtual machine and supports proxy, host monitor, and out of band modes. Only proxy mode can block traffic inline, and it requires client connection string changes to route traffic through the proxy. SQL Firewall is built into the database kernel, requires no client configuration changes, and is often the better choice when blocking is the primary use case or when client reconfiguration is impractical.

Scope of inspection: Database Firewall inspects SQL traversing the network. It is unaware of database synonyms, internal scheduled jobs, and stored procedure executions occurring within the database. For local connections, host monitor mode can observe activity but cannot block it. SQL Firewall runs inside the database with full access to session context and metadata, enabling comprehensive monitoring and blocking, including SQL generated internally. If the application server and database run on the same host, SQL Firewall is the appropriate choice.

Processing of encrypted traffic: Database Firewall needs additional configuration to handle encrypted SQL traffic, typically involving SSL proxy or certificate management. With SQL Firewall, this is seamless because it inspects traffic after the database has already decrypted it.

Transaction integrity: Blocking SQL at the network layer with Database Firewall can disrupt in flight transactions because each SQL statement is handled as a standalone operation. SQL Firewall, by contrast, understands transaction boundaries and honors the ACID properties (atomicity, consistency, isolation, and durability), helping ensure blocked operations do not leave transactions in inconsistent states.

High availability: High availability for Database Firewall is achieved through redundant firewall appliances, with clients distributing connections or failing over based on connection parameters or third-party load balancers. SQL Firewall is integrated into Oracle Database and automatically leverages existing high availability architectures such as Real Application Clusters (RAC) and Data Guard, with no additional high availability configuration.

Separation of duties: Database Firewall can monitor database traffic without input or control from database administrators, supporting strict separation of duties requirements. SQL Firewall is part of Oracle Database and requires cooperation from database administrators to deploy and configure, which may not align with environments that require complete DBA exclusion from security policy enforcement.

License cost: If the organization already licenses AVDF, there is no licensing tradeoff because SQL Firewall is included with AVDF. If the organization already licenses Oracle Database Vault, SQL Firewall is typically the better choice because it is included with Database Vault. Note that most Oracle Database cloud services include Database Vault by default. If the organization licenses neither AVDF nor Database Vault and is not using an Oracle Database cloud service, either AVDF or Database Vault must be procured to access these SQL injection protection capabilities.

Behavioral protection for AI-generated queries

Oracle AI Database 26ai SQL Firewall and Audit Vault and Database Firewall address a growing challenge with AI workloads: the SQL generated by AI systems cannot be predicted or practically reviewed in advance. LLM generated queries, multistep agent tool chains, and rapidly developed AI assisted applications can produce syntactically valid but risky SQL, including overly broad SELECT statements, unvetted JOIN operations, and exploratory schema queries that traditional access controls may not intercept.

Copyright © 2026, Oracle and/or its affiliates

Traditional database security assumes SQL is written by humans, reviewed during development, and relatively stable in production. AI generated SQL breaks each of these assumptions. Queries are created dynamically, never reviewed by humans, and can shift unpredictably based on natural language inputs or agent reasoning chains.

Database firewalls learn legitimate query patterns during normal operation and then enforce those learned behaviors in real time. SQL that deviates from established patterns can be blocked before execution, including unauthorized table access, novel JOIN patterns, unusual aggregation operations, DDL attempts, or bulk data extraction queries. This approach can stop injection attacks that try to explore database schemas, such as `UNION SELECT * FROM all_tables`, prevent AI agents from pivoting into sensitive tables outside their normal scope, and block prompt injection attacks that generate destructive commands.

For AI workloads where SQL is inherently untrusted, whether generated by models, composed by automation, or written without human review, database firewalls turn query anomalies into clear enforcement boundaries. They help keep AI assisted applications and autonomous AI agents within learned, safe operational parameters regardless of code quality, prompt engineering robustness, or credential exposure risk.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- **Oracle Audit Vault and Database Firewall**: Comprehensive workshop covering AVDF deployment, policy configuration, and violation analysis.
- **Oracle SQL Firewall**: Hands-on lab for SQL Firewall configuration and management.
- **Managing SQL Firewall with Oracle Data Safe**: Lab 7 of the "Get Started with Oracle Data Safe Fundamentals" workshop covers SQL Firewall management through the Data Safe console.

Copyright © 2026, Oracle and/or its affiliates

Summary

SQL injection remains one of the most common and damaging ways attackers reach data, and its staying power is a reminder that awareness alone doesn't create consistent prevention. It's an application-layer flaw that appears when untrusted input is treated as executable SQL, letting an attacker run statements that read, modify, insert, delete, or destroy data. In multi-tier environments, the injected SQL typically runs with the application service account's privileges. Because those accounts often carry broad schema and procedure rights, a low-privilege attacker can suddenly operate with high-impact authority.

It's also important to treat SQL injection as an entry point to broader compromise, not just a one-time leak. Injection can enable theft, corruption, destruction, and lateral movement into other systems. A familiar failure pattern is that an application enforces access control in the middle tier, but injection through an input field or URL parameter slips around authentication and authorization. The recommended posture is defense in depth: upstream controls such as a web application firewall, paired with downstream controls that evaluate the SQL that actually reaches the database. Signature-based defenses are noted as limited against novel techniques, which is why database-layer controls matter for containing risk when application vulnerabilities inevitably exist.

Database-centric firewalls are presented as a practical access control approach: learn the normal SQL an application uses, then enforce an allowlist, blocking or alerting when incoming SQL deviates. Two options are highlighted. The network-based Database Firewall in Oracle Audit Vault and Database Firewall operates at the network layer, can protect multiple heterogeneous databases from a centralized deployment, supports blocking in proxy mode, and supports monitoring in host monitor and out-of-band modes. The database-resident SQL Firewall built into Oracle AI Database 26ai runs inside the database kernel. The chapter's guidance favors `allowlists` over denylists because legitimate application SQL is finite, while malicious patterns are effectively unlimited.

SQL Firewall is described as inspecting every incoming SQL statement and validating trusted connection paths with richer session context and metadata, helping counter evasion techniques such as encoded SQL, synonyms, and dynamically generated object names. Violations are captured in a dedicated view, and administration can be handled through PL/SQL procedures or through Oracle Data Safe. Decision guidance compares the two firewall approaches across performance, deployment practicality, inspection scope, encrypted traffic handling, transaction integrity with ACID awareness, high availability fit with RAC and Data Guard, separation of duties, and licensing considerations. The discussion also extends to AI workloads, where SQL is often dynamically generated. In that setting, learned behavioral boundaries help block unusual access patterns such as novel joins, schema exploration, bulk extraction, or destructive commands.

Once SQL injection risk is reduced, the conversation can move beyond stopping bad statements and toward a more interesting question: what should valid, approved access actually be allowed to show? The next chapter on Data-driven Authorizations brings readers into that more nuanced frontier, where policy shifts from a blunt gate into a finely tuned instrument that decides which rows and columns each user can access, and which information should remain just out of reach. It is a natural next step for teams that want defenses to feel less like a barricade and more like a well-designed map, guiding every query toward safe ground while keeping sensitive data safely within its proper boundaries.

.

Copyright © 2026, Oracle and/or its affiliates

Chapter 9

Data-Driven Authorization

The need for data-driven authorization

Organizations often need to control access to individual rows or columns of data within database tables. The access control mechanisms traditionally used to protect entire database objects cannot selectively restrict access to specific data rows, columns, or even cells within a table.

The rise of AI agents in business applications introduces a new dimension to this challenge. Agentic AI systems persistently attempt to access any data necessary to fulfill their assigned objectives. Without precise, fine-grained authorization controls, these systems can inadvertently retrieve sensitive or unauthorized information in response to prompts. As AI capabilities grow, enforcing granular data access restrictions becomes essential to protecting sensitive and regulated information from unintended disclosure.

Implementing data-level access controls directly within applications can be costly. Any changes to the security model typically require code modifications. For packaged or third-party applications, such updates may be impossible. Enforcing policies only at the application level can result in inconsistent protections, as different applications often rely on different security models. Direct database connections, whether through a SQL client, business intelligence tool, or Model Context Protocol (MCP) integration, can bypass application-level controls entirely.

By enforcing fine-grained access controls in the database, organizations centralize management and enhance security. Access policies applied at the database layer are enforced uniformly, regardless of which tool or application accesses the data. With several fine-grained access-control options in the Oracle Database, understanding the trade-offs helps organizations choose the right solution.

Understanding data-driven authorization

Data-driven authorization controls access at the row and column level. Beyond leveraging privilege grants and session context, data-driven controls can consider the data values themselves when determining access rights. A straightforward example is a table of employee information that includes a column indicating each employee's manager. A data-driven policy might restrict access to only records for a manager's direct reports, preventing the manager from seeing information about other employees. The policy is data-driven because it relies on the value in the manager column to make access decisions.

Today's applications are highly complex, often requiring authorization policies that consider user roles, user and organizational attributes, session information, and other contextual factors. Ensuring that access control checks are consistently implemented throughout an application is both challenging and difficult to maintain over time.

Access for users should be limited to the specific rows and columns required to fulfill their job responsibilities. However, when an administrator grants a database object privilege such as SELECT or INSERT to a user, that user receives access to all data within the table. Because database tables typically store far more information than any one user requires, this broad access poses a risk. For instance, a customer should only view their own records in a support application, not those of other customers. A help desk technician should access only tickets assigned to them, not all open tickets. Sales managers need visibility into sales opportunities for their direct reports but not for other sales staff. These scenarios exemplify fine-grained authorizations where application users must be restricted to only the relevant rows and columns.

Figure 9-1 illustrates a table containing a mix of sensitive and non-sensitive human resources data. As an employee, Nancy Greenberg should be able to see phone numbers, names, and manager names for everyone, while other sensitive data remains hidden. Nancy should see her own national identification number, and, as a manager, she should also be able to view the salaries of her direct reports.

Copyright © 2026, Oracle and/or its affiliates

Name	Manager	National ID	Salary	Mobile
Adam Fripp	Steve Stiles			650-123-3234
Neena Kochar	Steve Stiles			650-124-8234
Nancy Greenberg	Neena Kochar	000-51-4569	120300	515-123-4567
Luis Popp	Nancy Greenberg		69000	515-123-4234
John Chen	Nancy Greenberg		82000	515-123-8181
Daniel Faviet	Nancy Greenberg		9000	515-123-7777

Figure 9-1: Fine-grained Access Control separates rows and columns within a table

Challenges with application-layer authorization

To implement access control policies on the EMPLOYEE table in the HR sample schema, as shown in Figure 9-1, developers typically write complex authorization code to regulate which rows and columns users can access under various conditions. This process usually requires a dedicated schema or set of tables to store user and role information, which the application uses to generate user-specific SQL queries. This approach introduces several key challenges.

First, developers must write substantial and intricate code and apply it across all contexts where users require access. Second, each application must independently replicate authorization rules because the database does not inherently enforce these policies. Third, direct-access tools such as analytical platforms, SQL*Plus, AI agents, and MCP connections bypass application logic, gaining unrestricted access to data.

Oracle Database overcomes these obstacles with technologies known collectively as data-driven authorization. Centralizing application authorization controls in the database simplifies and accelerates application development and provides a consistent set of policies to define and maintain.

Oracle Database data-driven authorization technologies

Oracle introduced Virtual Private Database (VPD), an automated predicate-based row-filtering security technology, more than 25 years ago in Oracle Database 8i. At that time, it was the only database with this capability. The next version of the database, Oracle Database 9i, introduced Oracle Label Security (OLS) to automatically filter rows based on user and data labels. Real Application Security (RAS), introduced in Oracle Database 12c, is also declarative like OLS, simplifying the development of secure applications. Each of these technologies is discussed in detail in the sections that follow.

Among the data-driven access control mechanisms, VPD and RAS are both features of Oracle Database Enterprise Edition. Label Security requires a separate license for on-premises databases but is included in all but the most basic levels of cloud database services. For details about licensing in a specific database package, consult the Oracle Database Features and Licensing application at https://apex.oracle.com/database-features/.

Virtual Private Database

Virtual Private Database enforces fine-grained row and column-level security by leveraging a PL/SQL function created by the database administrator or security team. This function is tied to a table, view, or synonym through a VPD policy. When a SQL statement targets the protected object, the function runs automatically, appending a WHERE clause to ensure only authorized rows and columns are returned.

Copyright © 2026, Oracle and/or its affiliates

For example, consider the following SQL query:

```
SELECT * FROM hr.employees;
```

The VPD function might rewrite the SQL statement to:

```
SELECT * FROM hr.employees WHERE department <> 'SUPER SECRET DEPARTMENT';
```

VPD allows for differentiated access, enabling users to view all employee records but update only their own, for instance. Sensitive columns can be controlled so that only authorized users see actual values, while others receive nulls.

Because VPD enforces policies at the database level, access rules are consistently applied across standard database access paths. For instance, as shown in Figure 9-2, the user can access HR.EMPLOYEES table data only for `DEPARTMENT_ID` equal to '80', and `SALARY` data is hidden by setting it to NULL. Running an unrestricted query still results in only permitted data being returned.

```
SQL> SELECT last_name, email, department_id, salary FROM hr.employees;

LAST_NAME    EMAIL                DEPARTMENT_ID    SALARY
------------ -------------------- ---------------- ------
Hunold       AHUNOLD@EXAMPLE.COM  80
Ernst        ERNST@EXAMPLE.COM    80
Austin       DAUSTIN@EXAMPLE.COM  80
Pataballa    VPATABAL@EXAMPLE.COM 80
Lorentz      DLORENTZ@EXAMPLE.COM 80
```

Figure 9-2: VPD provides automatic & transparent row-level security

Rapidly changing authorization requirements, combined with the fact that VPD policies are implemented in PL/SQL rather than declarative SQL, can significantly increase complexity over time. These challenges highlight the need for a more scalable and consistent approach to access control, which leads to the next topic: controlling access to data using data labels.

Oracle Label Security

Another method of fine-grained access control is based on data labeling, where each row is assigned a label that represents its sensitivity or classification. This model is commonly used in government and regulated corporate environments, where information may be classified as TOP SECRET, CONFIDENTIAL, or INTERNAL USE ONLY. Access to data is then governed by a user's clearance level (which is also represented as a label), ensuring that users can only view data at or below their authorized classification. When labels are attached at the row level, the database can inherently determine which data is sensitive and enforce access restrictions accordingly. This approach is also referred to in some data privacy regulations as enforcing `Restriction of Processing`, where access and the ability to process data is limited based on data classification and user authorization. OLS supports the technical enforcement of this legal requirement.

Oracle Label Security simplifies assigning labels to data and users and enforces access control based on those labels. In an OLS-protected table, every row is associated with a label that reflects the sensitivity of the data. As an added security measure, the column containing the label values is hidden. Labels can be explicitly assigned based on business logic or automatically derived using default labeling rules tied to the application context or user session that inserts the data. The effective label of a user session is calculated using multiple factors, including the label assigned to the user, session attributes, and the type of database connection. In this sense, labels act as an extension to traditional database privileges and roles. A label can be associated with a database user, and starting with Oracle

Copyright © 2026, Oracle and/or its affiliates

Database 12c Release 2, a label can also be associated with application users supported by Real Application Security (discussed later in this chapter).

Label components and structure

The OLS label model is flexible and can represent virtually any data classification scheme. Each label is composed of up to three components:

- **Level**: A mandatory hierarchical classification, with exactly one level per label
- **Compartment**: Optional categories used to further restrict access, with zero or more compartments per label
- **Group**: Optional hierarchical structures used to model organizational or geographic boundaries, with zero or more groups per label

Every label includes a level, ordered from lowest to highest, representing the overall sensitivity of the data. Optional components called compartments and groups may segregate information based on attributes such as projects or departments. Compartments are typically used to isolate information related to specific domains, such as finance or pricing. Groups enable multi-dimensional, hierarchical authorization models, such as regional or organizational structures.

Figure 9-3: Example OLS policy showing Label Security components: Levels, Compartments, Groups shows an example of a global retail organization classification scheme using all three label components. Three levels are defined: `HIGHLY SENSITIVE, SENSITIVE`, and `PUBLIC`. Only users with a `HIGHLY SENSITIVE` label can access unreleased quarterly financial data. A FINANCE compartment grants access to those needing to see financial data. Upcoming pricing changes are also sensitive, so a `PRICING` compartment protects pricing data. Groups represent the retail stores' geographic hierarchical structure. This combination of levels, compartments, and groups enables precise, scalable, and policy-driven access control across complex data environments.

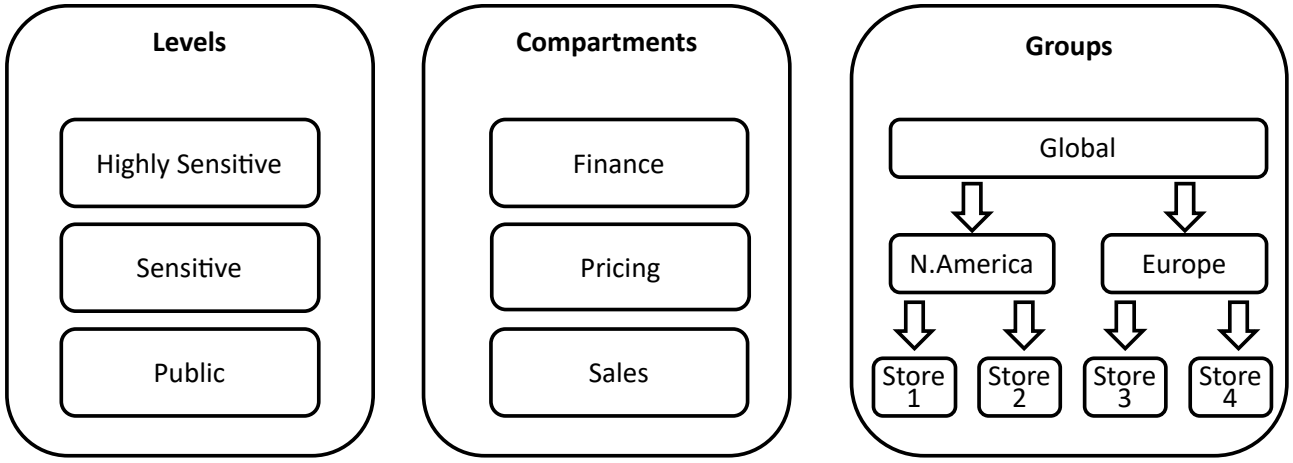


Figure 9-3: Example OLS policy showing Label Security components: Levels, Compartments, Groups

Labels are represented by the three components separated by a colon (:) as shown in Figure 9-4.

```
LEVEL:COMPARTMENTS:GROUPS
```

Figure 9-4: Label component organization

A data label representing sensitive finance data for Store 1 would be represented as `SENSITIVE:FINANCE:STORE1`. Similarly, pricing data for Store 1 would be labeled `SENSITIVE:PRICING:STORE1`. In this model, a Store 1 marketing

Copyright © 2026, Oracle and/or its affiliates

manager who requires pricing information for creating weekly sales newspaper inserts and promotional emails would be assigned a user label that includes the PRICING compartment.

Evaluating labels: Groups, Compartments, and Levels

To illustrate how labels are evaluated, the following examples begin with the group component and then introduce compartments and levels.

In the label design shown in Figure 9-5: Using the Group component in Label Security, users in each store can only see data regarding their store due to the restrictions imposed by the group label component. Users assigned the STORE1 group label can only view data labeled for STORE1. A North America regional director, however, may be assigned the N. AMERICA group label. Because groups are defined hierarchically, this allows access to all subordinate groups, such as STORE1 and STORE2. This hierarchical group structure enables centralized data storage while ensuring users only see data relevant to their organizational scope.

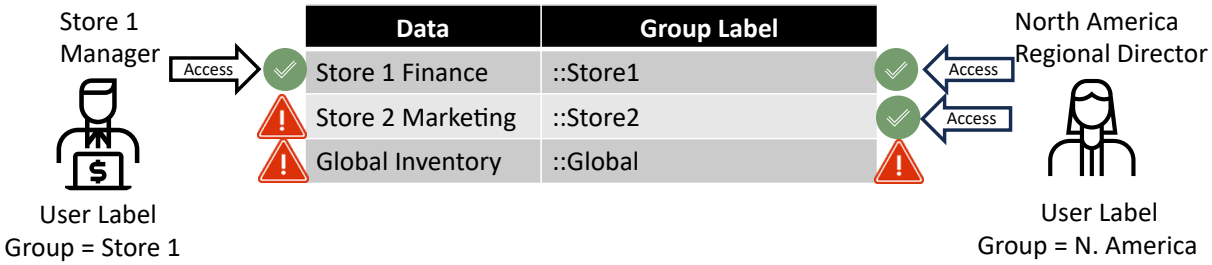


Figure 9-5: Using the Group component in Label Security

The retail company required that financial and pricing data be visible only to staff with a legitimate business need. To enforce this, FINANCE and PRICING compartments were created, and the corresponding data records were labeled accordingly. User access was then reviewed and aligned with these compartments. The Store 1 marketing manager was granted access to pricing data to produce weekly sales materials, while the Store 1 manager was granted access to both pricing and financial data for Store 1. However, neither role was permitted to access financial data belonging to the N. AMERICA group.

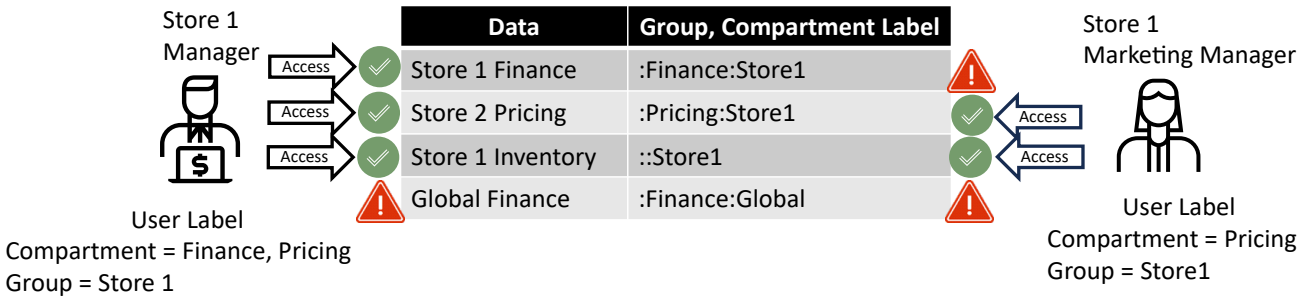


Figure 9-6: Using the Compartment and Group components in Label Security

As a public company, the global financial data and reports are classified as highly sensitive. All data is considered sensitive, but global financial data is specifically highly sensitive within the FINANCE compartment and GLOBAL group. Only users whose labels reflect the level HIGHLY SENSITIVE with the FINANCE compartment and the GLOBAL group may see this information.

Copyright © 2026, Oracle and/or its affiliates

User access conditions

User access to data requires that their user label meets the following conditions in all three components:

- **Level**: The user's label must be at the same or a higher level than the data's label. For instance, a user labeled `SENSITIVE` can access `SENSITIVE` or `PUBLIC` data, but not `HIGHLY SENSITIVE` data.
- **Compartment**: The user's label must contain every compartment listed on the data. Data labeled with the `FINANCE` compartment requires the user's label to include FINANCE. If the data label specifies both FINANCE and PRICING, the user's label must include both compartments.
- **Group**: The user's group label must match at least one group from the data label or encompass it hierarchically. For example, if the data is labeled `STORE1`, users with `STORE1`, `N. AMERICA`, or `GLOBAL` group labels can view the data.

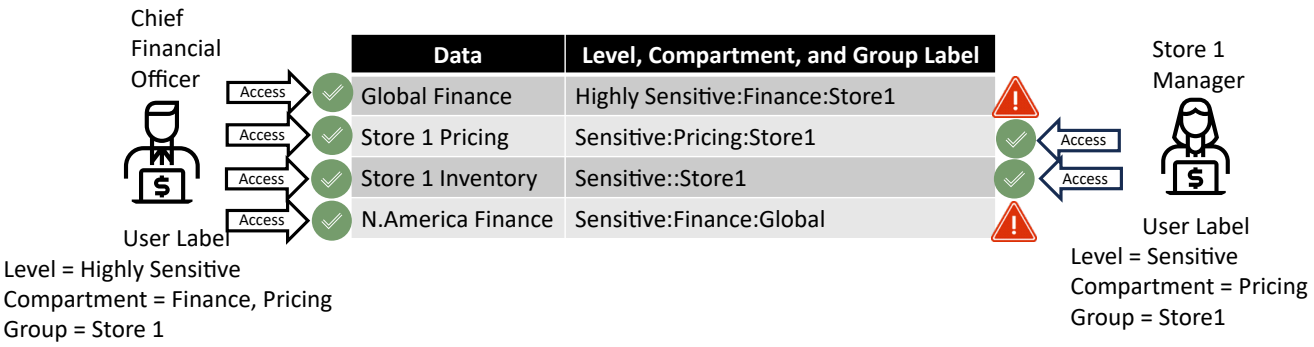


Figure 9-7: Using the Level, Compartment, and Group components in Label Security

Label Security use cases

Label Security was originally developed to meet the stringent data labeling and mandatory access control requirements of government, military, and intelligence agencies. However, commercial organizations have also found that Label Security simplifies data access authorization. For example, in the commercial scenario described above, each retail store previously operated its own applications and databases, making consolidated reporting challenging and maintaining parallel systems costly. By consolidating data into a single database, Label Security enables robust access controls while reducing operational complexity.

Another common commercial application relates to data protection laws, such as the European Union's General Data Protection Regulation (GDPR), which require restricting data processing based on user consent. In this case, groups and compartments can record the types of processing a data subject has consented to, which is reflected in the data label. Applications then set a user session label, ensuring that records the application should not process remain inaccessible.

A major advantage of Label Security over VPD or RAS is a complete authorization infrastructure that provides automatic mediation, enforcement, and simplified administration. Label Security uses hierarchical user and data labels comprised of levels, compartments, and groups that parallel real-world scenarios without requiring administrators to write PL/SQL policy functions. Once the model is configured for user and data labels, access control is consistently enforced across all interfaces, eliminating the need for explicit configuration on a per-table basis. Label Security is recommended when access control logic can be defined using user and data labels.

Copyright © 2026, Oracle and/or its affiliates

Real Application Security

Previously, we saw that VPD manages data access at the row and column level. However, in most modern applications, end users do not interact directly with the database and typically do not have individual database accounts. Instead, a single database account represents the application, and it queries and updates data on behalf of all end users, who are managed by the application itself.

Because the database is not aware of individual end-user identities, the application must handle per-user access control. This approach requires additional development effort and often results in inconsistent policy enforcement, particularly if other tools, clients, or applications can access the database directly and bypass the application layer.

To address these challenges, Oracle introduced Real Application Security in Oracle Database 12c as a next-generation access control framework within the database. RAS allows two-tier and three-tier applications to define, provision, and enforce access requirements declaratively in the database layer, providing a simpler, more maintainable framework that can replace VPD in nearly all scenarios.

RAS is built on a policy-based authorization model that recognizes application users, privileges, and roles in the database, controlling access to both static and dynamic groups of records that represent business objects. This approach resolves the scalability, maintainability, and security limitations associated with complex VPD policies and supports common application patterns such as master-detail tables and temporary assignments.

Application user identity propagation

RAS securely propagates the identity of the application end user to the database. Like VPD or OLS, RAS enforces access control policies at the database level, regardless of which application accesses the data. The application can initiate multiple application user sessions and switch between them using a single database connection from a connection pool. Note that the application user does not have its own schema to store data objects or a dedicated connection to the database as a regular database user would. RAS can also enforce access control policies on traditional database users.

Data Realms and access control lists

RAS supports fine-grained column and row access restrictions, similar to VPD, but uses a more declarative syntax. Administrators define data realms for table rows, each using a WHERE clause-like syntax to identify subsets of data. Access control lists (ACLs) are then attached to each data realm, specifying which users or roles can perform which operations on data within that realm.

As illustrated in Figure 9-8: Example illustrating the data access controls with Real Application Security, three example data realms might be created for the same table: a public data realm covering all employee records, a self-data realm (a dynamic realm that protects each user's own record), and a manager data realm (another dynamic realm that covers employees reporting to the user).

Copyright © 2026, Oracle and/or its affiliates

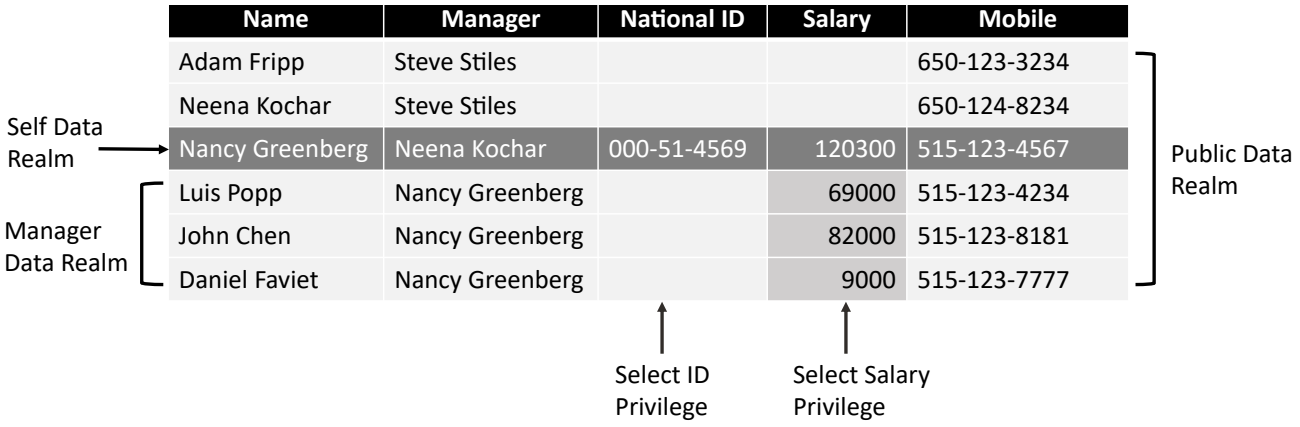


Figure 9-8: Example illustrating the data access controls with Real Application Security

For each data realm, an access control list defines who can access that data realm and under what conditions. For the public data realm, all employees have only the `SELECT` privilege.

All non-protected data should be visible to any employee. Because `SSN` and `Salary` data are protected and not included in the ACL for the public data realm, data from those columns does not appear when queried by someone with only public access. The self-data realm is associated with an ACL to view `SSN` (Authorize SSN) and Salary (Authorize Salary) data so that an employee can view their own sensitive data.

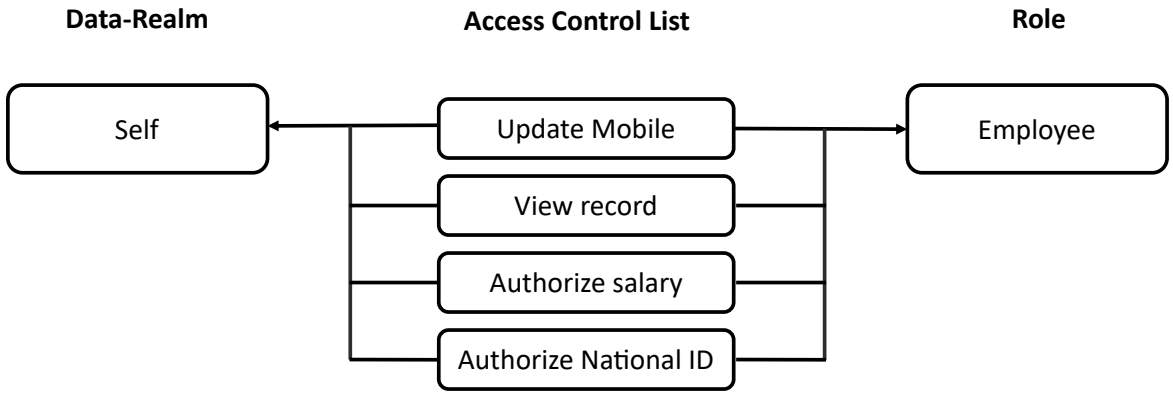


Figure 9-9: Example of a RAS Access Control List (ACL)

Managers can view their direct reports' salary information via the manager data realm, which is governed by an ACL permitting access to salary data. Both self and manager data realms are dynamic, returning the relevant rows and columns according to the application user's session identity.

Application-defined privileges

RAS also allows the database to enforce additional security policies unique to each application. Applications can define privileges beyond standard SELECT, INSERT, UPDATE, and DELETE operations, such as those for approving vacation requests or processing invoices. Administrators may assign access to sensitive columns such as salary or SSN only for users holding specific application-defined privileges (for example, UPDATE_SALARY or VIEW_SSN).

Minimal privilege design

Connecting through the RAS Java interface offers more security than traditional methods, which often require application accounts to have all privileges necessary for every user operation, including those needed for

Copyright © 2026, Oracle and/or its affiliates

maintenance and installation. These highly privileged accounts become prime targets for attackers. RAS enables the use of minimally privileged application accounts, so even if a compromise occurs, access to sensitive data remains protected. RAS realms, ACLs, and policies can be fully managed using the RAS PL/SQL API.

With built-in support for propagating application user sessions to the database, RAS allows security policies on data to be expressed directly in terms of the application-defined users' roles and data operations. The RAS security model enables uniform specification and enforcement of access control policies on business objects irrespective of the access path. Using declarative access control policies on application data and operations, RAS enforces security close to the data and enables end-to-end security for both three-tier and two-tier applications. The declarative model for RAS is much simpler to maintain and extend than VPD.

Choosing the right technology

Virtual Private Database is the most straightforward access control feature to begin using. However, when VPD PL/SQL policies become complex, the code to create the WHERE predicate clause becomes difficult to manage. If the access control policy is simple and will remain that way, VPD is an appropriate choice.

Real Application Security is a more powerful data-driven access control technology. It is enterprise-ready with many enterprise patterns built in. It can serve as a direct replacement for VPD, working directly with database users, roles, and objects. If an organization is developing a new application or has the ability to modify existing application code, RAS provides a robust data security framework that separates access control policies from the applications. The application simply invokes the RAS APIs to connect application users to the database session and set any needed context variables.

Label Security is an excellent choice because the logic to evaluate label precedence is handled automatically. However, work is required to manage and maintain user and data labels. Both VPD and RAS provide row and column-level control, while Label Security operates only at the row level.

Organizations should evaluate these technologies based on several factors: the complexity and expected evolution of access control requirements, the ability to modify application code, the need for application user identity propagation, and whether the access control model can be represented using hierarchical labels. For new applications with complex authorization needs, RAS is typically the best choice. For scenarios where data classification naturally maps to hierarchical labels, Label Security offers the simplest administration. For simple, stable access control requirements on existing applications that cannot be modified, VPD remains a viable option.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Protecting row and column data with Virtual Private Database
- DB Security - Label Security

Copyright © 2026, Oracle and/or its affiliates

Summary

Organizations increasingly need access control that reaches beyond whole tables down to individual rows and columns, especially as AI agents in business applications repeatedly try to fetch whatever data might help complete a task. When authorization lives only in application code, changes get expensive, enforcement varies from app to app, and controls can be sidestepped through direct database paths like SQL clients, analytics tools, AI agents, or Model Context Protocol (MCP) connections. Centralizing fine-grained authorization in the database keeps policies consistent, no matter how the data is accessed.

Data-driven authorization is presented as a more grounded way to make access decisions: evaluate privileges and session context, but also evaluate the data values themselves. In an employee table, for example, a manager should only see direct reports because the policy checks the manager column to decide which rows qualify. The same idea maps cleanly to other common expectations: customers should see only their own service records, help desk technicians only the tickets assigned to them, and sales managers only opportunities tied to their teams. A detailed HR example makes the point vivid in a mixed-sensitivity table: Nancy Greenberg can see general contact and manager information for everyone, can see her own national identification number, and, as a manager, can see salaries for her direct reports, while other sensitive values remain out of view.

Three Oracle Database technologies form the core toolbox for enforcing this kind of fine-grained control: Virtual Private Database (VPD), Oracle Label Security (OLS), and Real Application Security (RAS).

- **VPD** attaches a PL/SQL function, written by an administrator or security team, as a policy. The database rewrites queries by adding a predicate that filters rows and can return `NULL` for unauthorized columns, enforcing controls even when the query itself is unrestricted. It's described as the most straightforward on ramp, with the tradeoff that policy logic can become harder to manage over time because it lives in PL/SQL rather than declarative SQL.
- **OLS** labels each row with a classification and evaluates access using levels, compartments, and groups. The retail example uses levels like `HIGHLY SENSITIVE, SENSITIVE`, and `PUBLIC`; compartments such as `FINANCE` and `PRICING`; and hierarchical groups like `STORE1`, `N. AMERICA`, and `GLOBAL`, with explicit access conditions for each dimension.
- **RAS**, introduced in Oracle Database 12c, propagates application end-user identity into the database, defines data realms and access control lists declaratively, supports application-defined privileges, and enables minimal-privilege application accounts to reduce the blast radius of compromise.

Selection guidance comes down to fit and friction: consider policy complexity and expected change, constraints on modifying application code, whether end-user identity propagation is required, and whether the authorization problem naturally maps to a label model. In general, RAS is positioned as a strong match for new applications with complex authorization, OLS fits clean classification models, and VPD works well for simpler, stable policies, especially when applications can't be modified.

With data driven authorization in place, the narrative naturally widens from who can see which rows and columns to how sensitive values should look when they do appear on screen. The next chapter on Data Masking is where policy meets presentation, turning raw secrets into carefully curated shadows so teams can keep moving fast without letting regulated details spill into the wrong hands. After that, Encryption takes the stage as the final vault door, protecting data at rest and in motion so that even if something is copied, intercepted, or mishandled, what remains is unreadable noise to anyone without the keys. Keep reading, because these two chapters complete the arc from access control to safe exposure, to durable protection, and together they form a practical playbook for building systems that stay trustworthy under real world pressure.

Copyright © 2026, Oracle and/or its affiliates

Chapter 10

Masking Sensitive Data

Why masking sensitive data matters

Applications handle large amounts of sensitive information, including personally identifiable information (PII), financial records, healthcare data, and proprietary business details. Organizations face a fundamental challenge: extracting maximum value from data while protecting it from unauthorized access or misuse. Limiting the exposure of sensitive data to users who do not need to see it is a challenge that grows more acute as data proliferates across development, testing, analytics, and operational systems.

The business case for using actual production data is compelling. Development teams need realistic datasets to build and test applications that function correctly under real-world conditions. Marketing and customer service teams deliver better outcomes when they can analyze customer behavior patterns and interaction history. Call center representatives resolve issues faster when they can review purchase histories and account details. Fraud analysts require actual transaction patterns to train detection models. Each of these scenarios delivers clear business value, but each also increases the risk of exposing sensitive information to unauthorized parties or unintended uses.

Most organizations create multiple copies of production databases for use in testing, development, user acceptance testing, and training. These non-production copies may also be shared with partners or third-party vendors for specialized analytics or fraud detection. Each additional copy raises the probability that sensitive information could be compromised through misconfigured access controls, insider threats, lost devices, or external attacks. As organizations embrace artificial intelligence and machine learning, these same environments are increasingly used to train models and power retrieval-augmented generation (RAG) pipelines. Without proper controls, real PII can proliferate into systems that were never designed to protect it.

Data masking addresses this challenge directly. By transforming sensitive data into realistic but fictitious values, organizations protect customer privacy while preserving the utility of data for legitimate business needs. Masked data holds little or no value to attackers, yet it remains valuable for development, testing, analytics, training AI models, and operational support. When testing a point-of-sale application, development teams can use actual inventory levels and store configurations while substituting realistic but fictitious values for customer names, email addresses, and payment details. This approach maintains referential integrity and application functionality while substantially reducing the risk of data breaches.

Data masking, combined with data redaction and data subsetting, forms an integral part of a comprehensive data privacy strategy. These techniques help organizations enforce the principle of least privilege, satisfy anonymization and pseudonymization requirements, and meet compliance obligations under regulations such as PCI-DSS, EU GDPR, CCPA, and India's DPDPA. In an environment where data privacy and security are both regulatory requirements and competitive differentiators, these strategies enable organizations to manage sensitive data confidently while supporting business objectives.

Understanding data masking

Data masking protects sensitive information by replacing real data with realistic but fictitious values. The fictitious data preserves the appearance and structure of the original, making it suitable for development, testing, and analytics, but it has little or no value to an attacker. Data masking is sometimes referred to as redaction, scrambling, or obfuscation—and is sometimes (informally) labeled anonymization—though important technical distinctions exist among these terms.

Unlike encryption, data masking is not designed to be reversed. Once data has been masked, the original values are typically not recoverable through standard administrative processes. This irreversibility is a strength when the goal is to permanently remove sensitive information from non-production environments. Masked data should remain consistent across related tables and respect referential integrity rules. For example, if a customer ID appears in both a customer master table and an orders table, the masked value must be identical in both locations to preserve the relationships that applications depend on.

Copyright © 2026, Oracle and/or its affiliates

Data masking strategies fall into two broad categories: static and dynamic.

Static data masking modifies the stored data itself, permanently replacing sensitive values with masked alternatives. This approach is used when sensitive data must be removed from non-production environments such as development, testing, training, or analytics systems. Static masking may also be applied when sharing data with external organizations for purposes like fraud analysis or clinical research. After static masking is complete, all users who access the data see only the masked values because the original unmasked data no longer exists in that environment.

Dynamic data masking—referred to in Oracle Database Security documentation as Data Redaction—prevents the proliferation the proliferation of sensitive data by masking values as they are retrieved, while leaving the stored data unchanged. With dynamic data masking, the results returned to a user depend on the masking policy in effect. Some users may see masked values, while others with appropriate authorization see the original data. Dynamic masking is typically applied in production environments where different users require different levels of access to sensitive information.

Oracle Database provides comprehensive support for both static and dynamic masking. Organizations can use Oracle Data Safe or Oracle Data Masking and Subsetting for static masking. For dynamic data masking, Oracle offers Data Redaction (included with Oracle Advanced Security), Virtual Private Database, and Real Application Security.

Static data masking

Static data masking creates datasets in which the original sensitive data no longer exists. After masking is applied, it does not matter who attempts to access the data. Everyone sees only the masked version, and the original values are gone. This permanence makes static masking the appropriate choice for non-production environments where developers, testers, and analysts need realistic data but have no legitimate need to access actual customer, employee, or financial details.

Databases masked in this way can be provisioned to development, testing, and business analytics teams while minimizing the exposure of sensitive information. Organizations can also use statically masked data when sharing datasets with external partners, vendors, or research institutions.

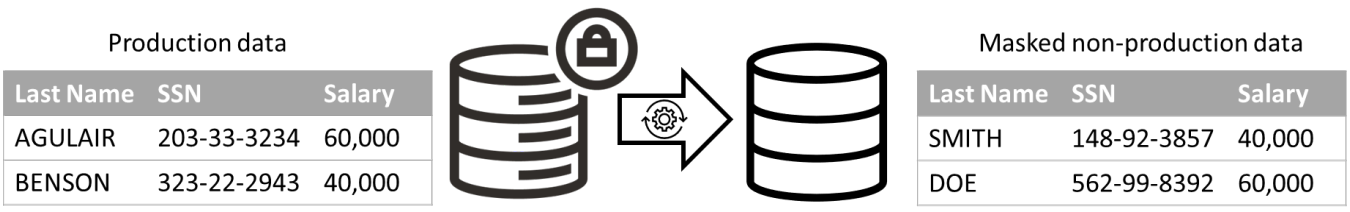
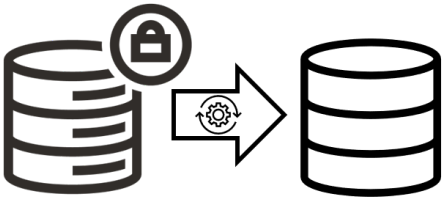


Figure 10-1: Example static data masking illustrates static data masking. In this example, actual employee names from a production database are replaced with realistic but fictitious names. Real Social Security numbers are transformed into random values that follow the correct format. Salary information is shuffled so that the overall payroll total remains accurate for aggregate analysis, but no individual record reveals actual compensation data. The production database retains sensitive information and requires full security controls. The non-production database, containing masked data, carries less inherent risk and can be managed with reduced security overhead.

Copyright © 2026, Oracle and/or its affiliates

Production data

Last Name	SSN	Salary
AGULAIR	203-33-3234	60,000
BENSON	323-22-2943	40,000



Masked non-production data

Last Name	SSN	Salary
SMITH	148-92-3857	40,000
DOE	562-99-8392	60,000

Figure 10-1: Example static data masking policy

Organizations can perform static data masking using services external to Oracle Database, such as Oracle Data Safe or Oracle Data Masking and Subsetting.

Use cases for static data masking

Static data masking supports several important scenarios:

Reduce risk and simplify compliance: By masking sensitive personal data in non-production environments, organizations lower the probability and potential impact of a data breach. Less sensitive data in the environment means less to steal and less damage if a breach occurs. In some cases, masking can remove entire systems from the scope of privacy regulations such as GDPR and CCPA, simplifying compliance audits and reducing the administrative burden of managing regulated data.

Protect non-production environments: Developers and testers need realistic data to validate that applications perform correctly under production-like conditions. Data masking enables teams to work with real-world data structures, volumes, and relationships without exposing sensitive information. This approach also helps organizations meet regulations that explicitly prohibit the use of actual sensitive data in non-production environments, such as certain healthcare and financial industry standards.

Save time and reduce costs: Cloning a production database and applying masking produces a realistic test or development environment far more quickly than building synthetic datasets from scratch. Teams avoid the expense and complexity of creating and maintaining artificial data that may not accurately reflect production patterns. Masked production clones provide high-fidelity test data with minimal ongoing effort.

Secure data sharing and third-party analytics: When organizations need to share data with vendors, partners, or research institutions, static masking protects privacy and security. In specialized scenarios such as healthcare research or financial fraud analytics, reversible masking techniques can provide useful data for analysis without exposing sensitive details. Analysts can derive insights from masked datasets without ever accessing the original information.

Support training and education: Data masking allows students, trainees, and new employees to work with realistic data volumes and structures in educational settings or training programs without exposing private information. This use case is particularly important in industries like healthcare and finance, where professionals must learn to work with production systems before handling actual customer data.

Data Redaction

Dynamic data masking, referred to in Oracle Database Security as Data Redaction, prevents the proliferation of sensitive data by masking values as they are retrieved from the database, while leaving the stored data unchanged. Unlike static masking, which permanently transforms the data on disk, dynamic masking applies transformations at query time based on policy conditions. The result delivered to a user depends on the masking policy in effect: some users may see masked values, while others with appropriate permissions see the original data.

Copyright © 2026, Oracle and/or its affiliates

Dynamic masking is typically used in production environments where different users require different levels of access to sensitive information. The database enforces masking policies transparently, without requiring changes to application code or queries. This approach provides consistent protection across all applications and tools that access the database, eliminating gaps that can arise when security logic is scattered across application layers.

Figure 10-2: Example Data Redaction illustrates Data Redaction in action. A data analyst and a human resources partner both query the same employee table. The data analyst needs to see employee last names for reporting purposes but does not require access to full Social Security numbers or actual salary details. The database applies a redaction policy that hides part of the Social Security number and replaces the salary with a fixed value when responding to the analyst's query. The HR partner, who is authorized by policy to view full employee data, sees the complete Social Security number and actual salary. The data stored in the table remains unchanged. Only the query results are transformed based on policy.



Figure 10-2: Example Data Redaction policy

Oracle Database enables dynamic masking through several mechanisms: Data Redaction (part of Oracle Advanced Security), Virtual Private Database, and Real Application Security. Each approach offers different capabilities and trade-offs.

Use cases for Data Redaction

Data Redaction addresses several important scenarios:

Control access to specific data elements: Data redaction enables organizations to hide sensitive columns as data is delivered to database clients, ensuring that only users with a legitimate business need can view certain information. For example, business intelligence tools and SQL-based reporting systems rarely need to display credit card numbers, birth dates, or taxpayer identifiers in full. Redaction policies can mask these columns while allowing aggregate analysis and reporting to proceed normally. This approach is also valuable in artificial intelligence applications that use retrieval-augmented generation with relational data, where sensitive details should not be fed into language models or training pipelines.

Minimize access based on policy conditions: Organizations can configure dynamic masking to redact data when access conditions fail to meet policy requirements. For instance, a policy might prevent users from viewing sensitive columns when connecting from outside approved network segments, accessing the database outside normal business hours, or querying from applications that have not been explicitly authorized. This conditional masking provides an additional layer of defense in depth.

Protect sensitive data in legacy applications: Many legacy applications were designed before modern privacy standards took effect and may display sensitive information even when users do not need to see it. Updating application code to implement fine-grained access controls is often difficult, expensive, or impossible when source code is unavailable or the application is provided by a third-party vendor. Data redaction allows organizations to hide

Copyright © 2026, Oracle and/or its affiliates

or transform sensitive data at the database layer, reducing risk without modifying application logic. This ensures consistent protection regardless of which application or tool accesses the data.

Support training and education: Data redaction allows trainees and students to work with production or production-like systems in educational or training programs without exposing private information. Redaction policies can be configured to mask sensitive columns for training accounts while allowing experienced staff to see full details.

Data Subsetting

Data Subsetting helps organizations meet data minimization requirements and reduce risk in non-production environments. When the full production dataset is not necessary for testing, development, or analytics, subsetting extracts a smaller, representative portion of the data. This approach reduces the amount of sensitive information at risk, lowers storage costs for database copies, and accelerates operations like backups, restores, and environment refreshes.

Subsetting allows organizations to extract only the data needed for a specific purpose. An analytics team might require only 20 percent of the production dataset, or only records collected within the last year, or only data specific to a particular geographic region. This targeted extraction provides teams with relevant, realistic data while keeping security risks and operational costs in check.

Data Subsetting is often combined with Data Masking in a single workflow. Organizations first extract the required subset from production, then apply masking to any sensitive information within that subset before provisioning it to non-production environments. This combined approach maximizes the business value of data, protects sensitive information, and avoids wasting resources on unnecessary storage and processing.

Together, Data Masking, Data Redaction, and Data Subsetting limit the opportunities for attackers to steal sensitive data, reduce the impact of a database compromise, and help organizations maintain compliance with privacy regulations such as PCI-DSS and EU GDPR.

Masking formats

Masking formats define how original data is transformed into anonymized values during the masking process. Oracle Database provides a comprehensive library of masking formats that range from simple transformations to sophisticated predefined patterns designed for common types of sensitive data. Simple masking formats are basic transformations that can be applied to data. Table 10-1: Simple masking formats describes several commonly used simple formats:

Masking format	Description
Fixed value	Replaces original values with a fixed value. For example, all email addresses might be replaced with "fake@no.co".
Random values	Replaces original values with random values, usually within a specified range or constrained by data type.
Random list	Replaces values in a column using a predefined list or array of values. For example, first names might be replaced with names randomly selected from a list of common names.
Substitution	Replaces values in a column with values selected from another table. This format is useful when masked values must maintain certain properties, such as valid postal codes for actual cities.
Encryption	Applies a cryptographic algorithm to mask the data. This type of masking is used when there is a need to reverse the masking later, such as when sharing data with a third party for analysis and subsequently recovering the original data after processing is complete.

Copyright © 2026, Oracle and/or its affiliates

Format preserving randomization	Randomizes data while preserving its structure and format. Replaces letters with letters and digits with digits, but keeps the data length, special characters, and the case (upper or lower) of characters unchanged. This technique is often used for data elements like national identifiers, postal codes, and license plate numbers.
Shuffle	Randomly reorders the values within a column so that statistical properties such as distributions and totals are maintained. Often applied to columns like age, gender, or salary.
SQL expression	Uses a custom SQL expression to generate masked values. Example: email addresses can be generated from columns containing first and last names. SQL expression masking is applied after other columns have been masked allowing masked email addresses to be constructed in the form "first_name.last_name@company.com" while maintaining consistency between the name and email columns. Also, can be used to mask data contained in large objects (LOBs), including BLOBs, CLOBs, and NCLOBs.

Table 10-1: Simple masking formats

Organizations commonly combine multiple simple transformations to create realistic masked data that follows the formatting and syntax rules appropriate for specific data elements. For example, a masked street address might be constructed by concatenating a four-digit random number between 1 and 9999, a random selection from a list of common road names such as Main, Elm, or Broadway, and a random selection from a list of street types such as Street, Road, Avenue, Boulevard, or Circle. The result is a fictitious but plausible street address.

In addition to simple formats, Oracle Database provides predefined masking formats for common types of sensitive data, including credit card numbers, telephone numbers, Social Security numbers, and other national identifiers. These predefined formats combine several simple transformations and often include additional processing to ensure that masked values appear realistic and follow the correct syntax rules. Figure 10-3: Predefined masking formats in Oracle Data Safe shows a selection of more than 50 predefined formats available in Oracle Data Safe. Predefined formats simplify the design and implementation of masking policies.

Copyright © 2026, Oracle and/or its affiliates

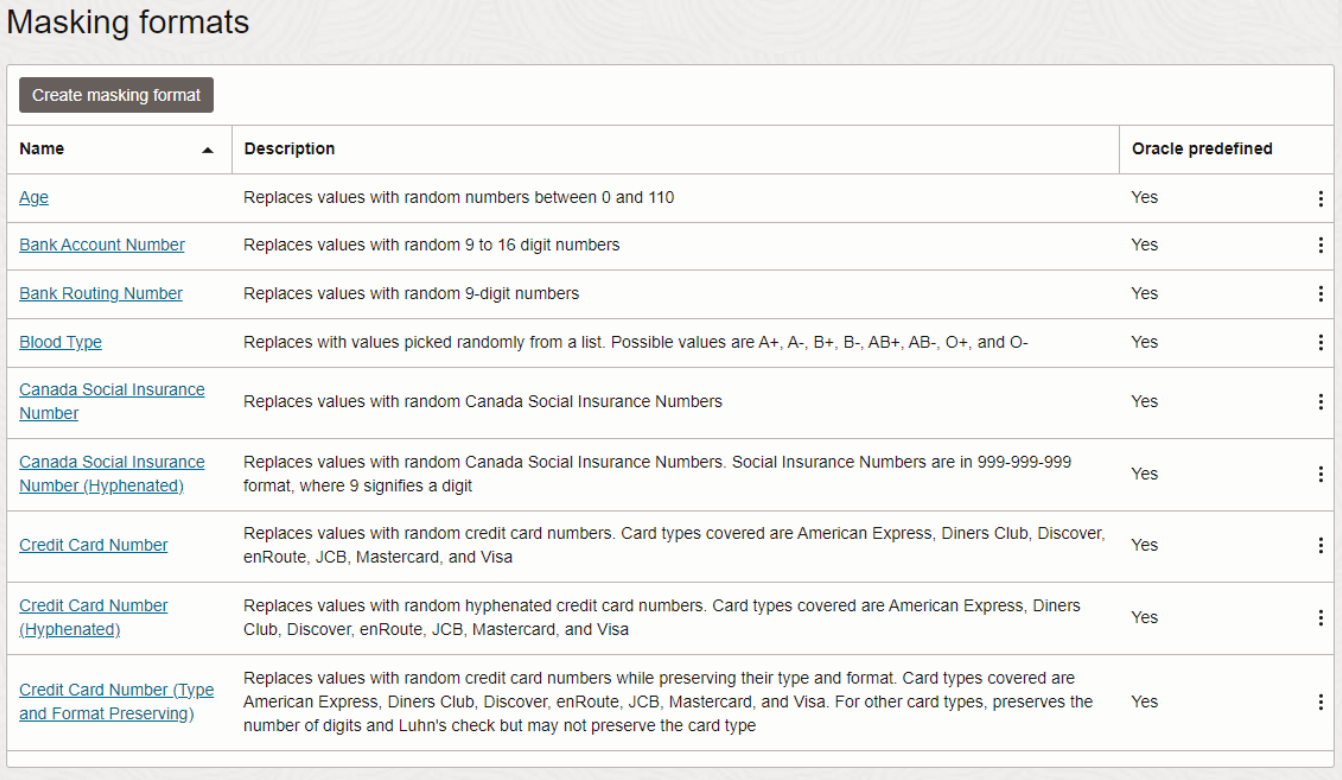


Figure 10-3: Predefined masking formats in Oracle Data Safe

Organizations are not limited to the formats Oracle provides. Custom masking formats can be created by combining existing formats or by defining entirely new transformation logic. If an organization uses unique data elements, such as proprietary account numbers or custom identifiers, it can create masking formats tailored to those requirements. This flexibility ensures that masked datasets are both secure and useful for their intended purposes.

Masking techniques

Masking techniques describe how transformations are applied to datasets in ways that preserve application functionality and data relationships. These techniques make it possible for complex applications to work correctly with masked data in non-production environments. Table 10-2 describes several important masking techniques.

Copyright © 2026, Oracle and/or its affiliates

Masking technique	Description
Conditional masking	Applies different masking formats to different rows based on specific conditions. For example, data about citizens from multiple countries can have unique masked national identifiers based on their country of origin. Similarly, credit card numbers can be masked while preserving their original card type and format.
Compound masking	Masks related data contained in multiple columns as a group. For example, city, state, and postal code can be masked together to ensure the values remain consistent. This prevents nonsensical results such as New York City located in Texas.
Deterministic masking	Produces the same masked output value for a given input value consistently across multiple databases, applications, or masking runs. Deterministic masking allows organizations to maintain referential integrity in a multi-database test environment. When the same customer ID appears in multiple databases, deterministic masking ensures it is masked to the same value everywhere.
Reversible masking	Encrypts sensitive data using a cryptographic key so that the original values can be recovered when necessary. This technique is useful when organizations must share masked data with third parties for business processing and later recover the original data after processing is complete. For example, healthcare data might be shared with researchers in masked form and then unmasked for follow-up patient care.
Shuffle masking	Randomly reorders values within a column, breaking the one-to-one mapping between related data elements. For example, a column containing salaries can be shuffled to break the association between employees and their compensation while preserving the overall distribution of salaries for statistical analysis.

Table 10-2: Summary table of available Masking techniques

Figure 10-4 demonstrates these techniques in action. Conditional masking produces different masked formats for national identifiers based on country codes. Random strings replace license plate numbers while preserving the original format. Deterministic masking ensures that employee names are masked consistently across all databases and masking operations. Shuffle masking disassociates health records from patient identities while preserving complex data types like diagnostic images.

Copyright © 2026, Oracle and/or its affiliates

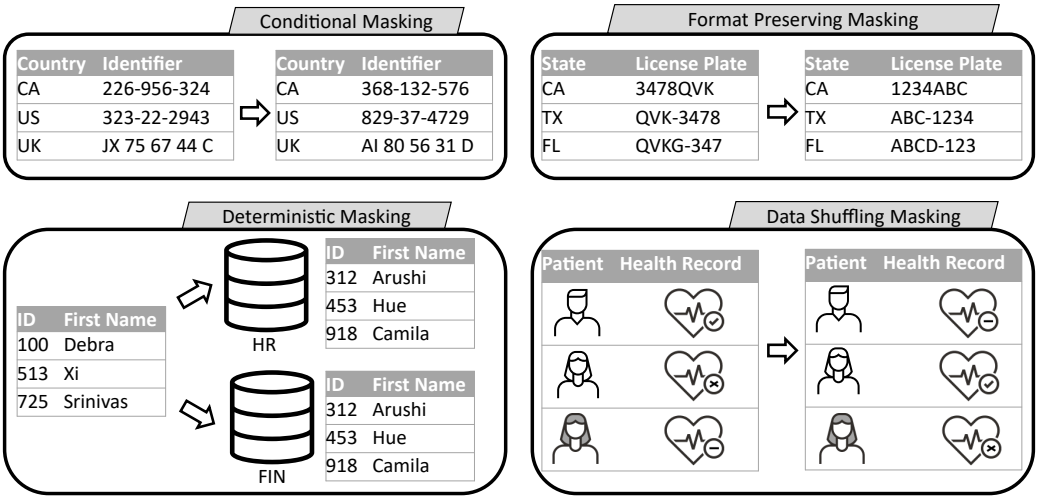


Figure 10-4: Masking formats and techniques in action

Implementing static data masking

Oracle Database provides two primary tools for static data masking: Oracle Data Safe and Oracle Data Masking and Subsetting. Both tools support comprehensive masking workflows, but they differ in deployment model, feature set, and operational flexibility.

Masking sensitive data using Oracle Data Safe

Oracle Data Safe provides cloud-based static data masking with a rich set of simple masking formats and a library of over 50 predefined formats. To mask a data column, users select the desired masking format from a dropdown list. If Oracle Data Safe's data discovery feature has already been used to identify sensitive data, the tool automatically suggests default masking formats based on the data types it detected. Organizations can also create and add custom masking formats to the library, providing the flexibility to meet unique requirements.

Masking reports

The Oracle Data Safe console summarizes the results of masking operations in detailed reports. Each report contains information about the database and masking policy applied, the timestamp of the masking job, and statistics showing which columns were masked, the masking formats used, and the volume of data transformed. These reports can be downloaded as PDF files to support audit and compliance requirements, helping organizations demonstrate that they are managing data privacy obligations effectively.

Copyright © 2026, Oracle and/or its affiliates

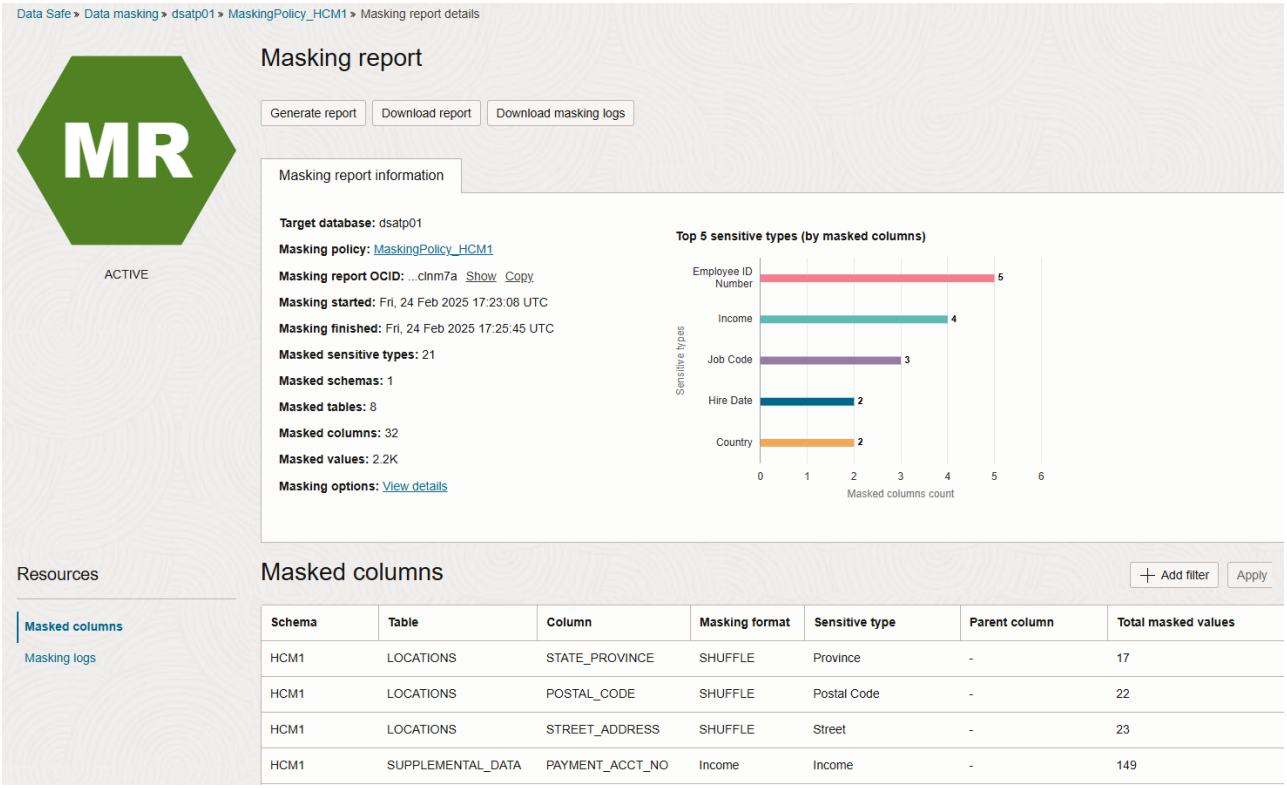


Figure 10-5: Example Data Safe masking report

Masking sensitive data using Oracle Data Masking and Subsetting

Oracle Data Masking and Subsetting, available as an Oracle Enterprise Manager management pack, provides static masking and subsetting capabilities for on-premises and cloud deployments.

Data masking workflow

The workflow for Oracle Data Masking and Subsetting is similar to Oracle Data Safe but includes some important differences. Organizations begin by running data discovery to build an application data model that identifies sensitive columns and their referential relationships. Data discovery is a mandatory step with Oracle Data Masking and Subsetting. Only sensitive columns identified through data discovery are eligible for masking.

After sensitive columns have been identified, organizations create a masking definition by assigning masking formats to those columns. Like Oracle Data Safe, Oracle Data Masking and Subsetting supports both simple and predefined masking formats and allows the creation of custom formats.

To apply the masking definition, organizations generate a masking script and execute it within the Enterprise Manager framework or independently. Figure 10-6, taken from the Oracle Data Masking and Subsetting console, provides an overview of the workflow for both masking and subsetting.

Copyright © 2026, Oracle and/or its affiliates

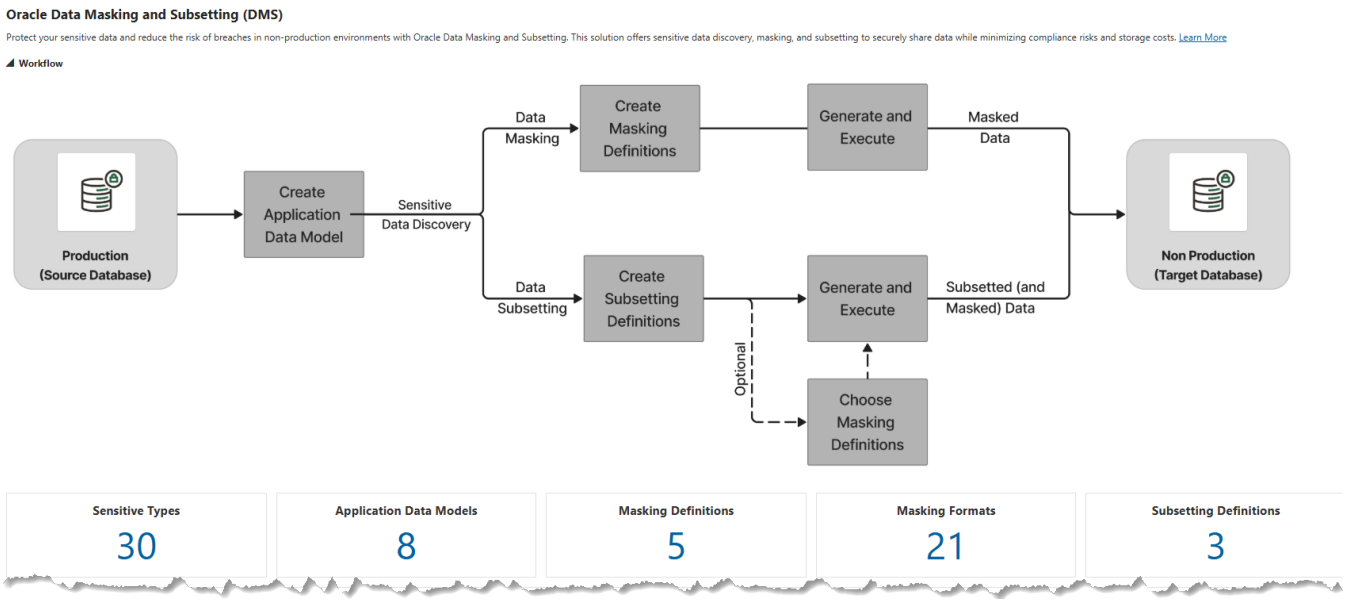


Figure 10-6: Overview of Data Masking in Enterprise Manager

In-database and in-export masking modes

Oracle Data Masking and Subsetting offers two modes for masking: in-database and in-export. The different modes provide flexibility in how masking operations are executed. Figure 10-7 illustrates the difference between the two modes.

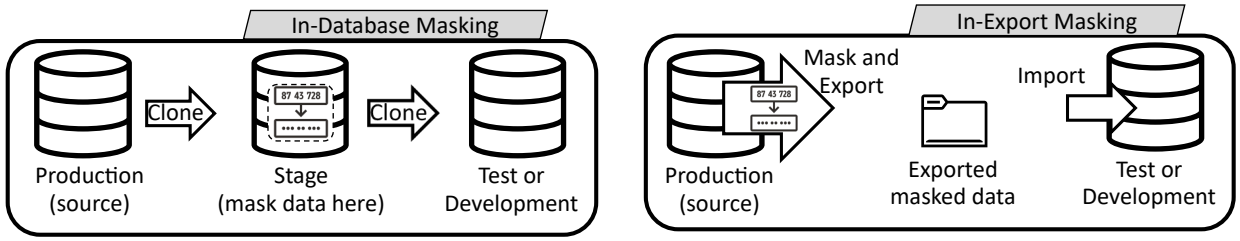


Figure 10-7: Masking modes for Data Masking and Subsetting

The in-database mode applies masking and subsetting within a non-production database, typically a clone of the production database, with minimal impact on the production environment. Organizations start by creating an unmasked staging copy of the production database, then apply masking and subsetting within that staging copy. From this masked staging copy, additional database clones can be created as needed. Data masking irreversibly modifies the stored data and should not normally be performed on a production database.

The in-export masking mode applies masking and subsetting during a Data Pump export operation. In this mode, Oracle Database masks data as it is exported from the production database but leaves the original values untouched. The masked data is written to standard Data Pump files, which can then be imported into non-production databases or shared for analytics. This approach eliminates the need to create a staging copy of the production database and ensures that sensitive data remains only within the production perimeter. When using in-export masking, the masking formats applied should be reasonably simple, and the export process should be scheduled outside peak usage hours to avoid performance impacts.

Copyright © 2026, Oracle and/or its affiliates

Choosing between Oracle Data Safe and Oracle Data Masking & Subsetting

Oracle Data Safe and Oracle Data Masking and Subsetting both provide robust static masking capabilities, but they are designed for different deployment scenarios and operational preferences.

Oracle Data Masking and Subsetting works exclusively with Oracle Database Enterprise Edition, while Oracle Data Safe can mask any edition of Oracle Database, including Standard Edition. This makes Oracle Data Safe a more versatile option for organizations with mixed database environments.

Oracle Data Safe streamlines the masking workflow by automatically matching masking formats to sensitive data types identified during data discovery. This automation allows organizations to create and execute masking policies quickly. Oracle Data Masking and Subsetting requires administrators to manually associate each column with the desired masking format, providing greater control but requiring more manual effort.

Oracle Data Masking and Subsetting offers the unique capability to create a masked Data Pump export directly from the production database. While this in-export masking mode limits the complexity of the formats that can be applied, it is a convenient way to generate pre-masked datasets for distribution to remote developers or external analysts.

Oracle Data Masking and Subsetting generates masking scripts that can be downloaded and executed independently of the Enterprise Manager framework. This capability simplifies the orchestration of test and development database provisioning workflows. With Oracle Data Safe, masking execution is always performed by the service and cannot be separated from the platform.

Organizations should evaluate these differences in the context of their specific requirements, including database edition, deployment model, operational complexity, and integration with existing management tools.

Implementing dynamic data masking with Data Redaction

Oracle Database enables dynamic masking of sensitive columns through several mechanisms, including database views, Virtual Private Database (VPD), Real Application Security (RAS), and Data Redaction. Data Redaction, included as part of Oracle Advanced Security, offers the most flexible and powerful approach to dynamic masking.

Data Redaction in Oracle Database

Data Redaction provides one of the most flexible approaches to dynamically masking sensitive data. Rather than altering the data stored in the database, it masks information dynamically as query results are returned. Data Redaction works with the most common column data types and applies across tables, views, and materialized views. Redacted values preserve the key characteristics of the original data, including data type and optional formatting, so downstream applications and processes continue to function as expected.

This seamless behavior extends to both applications and the database itself. Applications receive data in the format they expect, even when sensitive values have been masked. Within the database, redaction leaves buffers, caches, and storage untouched, so it does not interfere with operations such as exporting data with Oracle Data Pump, performing backups and restores with Oracle Recovery Manager, or running advanced configurations such as Oracle Real Application Clusters, Oracle Active Data Guard, or Oracle GoldenGate.

What makes Data Redaction particularly powerful is where and how it operates. Transformation happens inside the database kernel, driven by declarative policy conditions that are transparent and easy to manage. Unlike masking approaches that depend on application code, Data Redaction policies are enforced directly within Oracle Database, ensuring consistent protection across all application modules and tools. This eliminates gaps that can arise when security logic is scattered across multiple codebases or when organizations lack access to application source code.

Copyright © 2026, Oracle and/or its affiliates

Oracle AI Database 26ai introduces significant enhancements to Data Redaction, designed for environments where business applications increasingly depend on real-time analytics and complex SQL. Redaction must remain effective even as queries grow more sophisticated. These new capabilities include:

Support for mathematical and set functions. Data Redaction now supports aggregate functions such as SUM, COUNT, MIN, and MAX, as well as set operations such as UNION, INTERSECT, and MINUS. It also extends to advanced SQL constructs such as WITH clauses and OUTER JOINs. Queries execute against the original underlying values in the database, but when query results are returned, only the displayed output is redacted. For example, if a JOIN or UNION includes a redacted column, the operations proceed using the actual values, but the final result is redacted. This change enables complex queries to execute without exposing sensitive data or producing errors.

Support for sorting and grouping on redacted columns. Operations such as GROUP BY, ORDER BY, and DISTINCT can now be applied to columns with redaction policies. The query uses the actual values to group, sort, or remove duplicates, but when the results are displayed, the sensitive column values are redacted. For example, in a hospital setting, patient records may be grouped by insurance policy number. The grouping happens using the actual policy numbers, but the report displays them in redacted form such as XXXXX123.

Views with expressions. Redacted columns can now be used in SQL expressions on both regular and inline views. In addition, functions such as CONCAT, MIN, MAX, COALESCE, and TRIM are supported when applied to redacted columns in views. This ensures consistent redaction across complex view definitions and improves usability in reporting scenarios.

Support for redacting virtual columns in function-based indexes. Virtual columns used in function-based indexes are now automatically redacted when their underlying base columns are redacted. This ensures consistent redaction behavior even when sensitive information is accessed through virtual columns.

Data Redaction is included as part of Oracle Advanced Security. There is nothing additional to install or configure at the database level. Organizations simply create redaction policies and begin protecting sensitive data immediately.

Data Redaction policies

Organizations create Data Redaction policies at the table and column level using the DBMS_REDACT PL/SQL package. Each redaction policy specifies:

- The schema, object, and column to be redacted
- The format of redaction and any parameters associated with that format
- The enforcement expression that determines when redaction is applied

Figure 10-8 shows an example of creating a Data Redaction policy. In this example, the policy specifies that the first five characters of the SSN column in the CUST_INFO table within the HR schema will be redacted whenever the enforcement expression evaluates to true (in this case, always).

```
DBMS_REDACT.ADD_POLICY(
   object_schema       => 'HR',
   object_name         => 'CUST_INFO',
   column_name         => 'SSN',
   policy_name         => 'REDACT_CUST_SSN',
   function_type       => DBMS_REDACT.PARTIAL,
   function_parameters => DBMS_REDACT.REDACT_US_SSN_F5,
   expression          => '1=1')
```

Figure 10-8: Example Data Redaction policy creation for partial redaction

The enforcement expression provides precise control over when to redact data. Organizations can base redaction policies on a wide range of factors, including usernames, client identifiers, database roles, session information such as

Copyright © 2026, Oracle and/or its affiliates

client IP addresses and program modules, and context from Oracle Application Express (APEX), Oracle Real Application Security (RAS), or Oracle Label Security (OLS). Multiple conditions can be combined for more granular control.

If organizations use Oracle Enterprise Manager, they can take advantage of the Redaction Policy Expression Builder. This tool guides administrators through creating policy conditions using supported context attributes and provides an intuitive interface for specifying protected columns, transformation types, and enforcement conditions.

Data Redaction formats

Data Redaction policies allow organizations to choose different masking techniques based on business requirements:

- **Full masking**: Masks all sensitive data in a column.
- **Partial masking**: Masks only a portion of the data, such as the first five digits of a Social Security number.
- **Rule-based masking**: Uses regular expressions to define precisely which parts of the data to mask.
- **Random masking**: Replaces sensitive data with random values
- **Nullify**: Replaces sensitive data with a NULL value

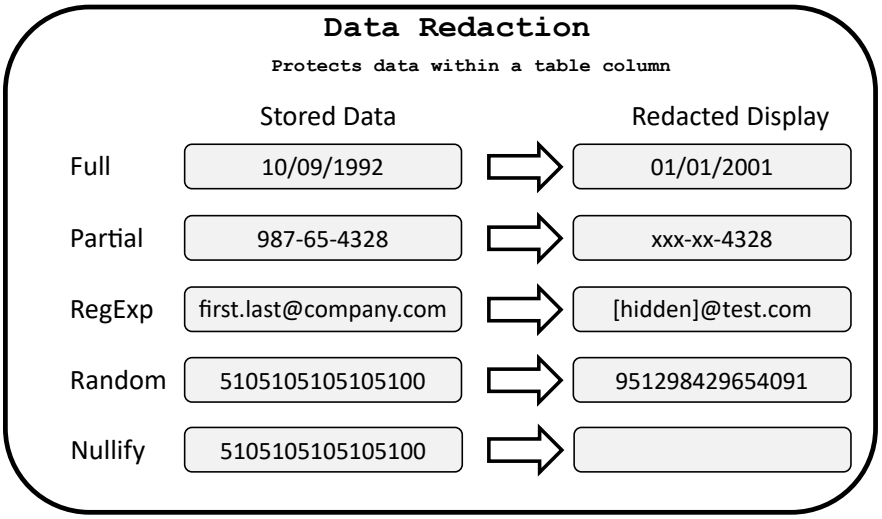


Figure 10-9: Example showing results using various Data Redaction formats

Figure 10-10 and Figure 10-11 demonstrate how application users see data before and after Data Redaction policies are applied.

Copyright © 2026, Oracle and/or its affiliates

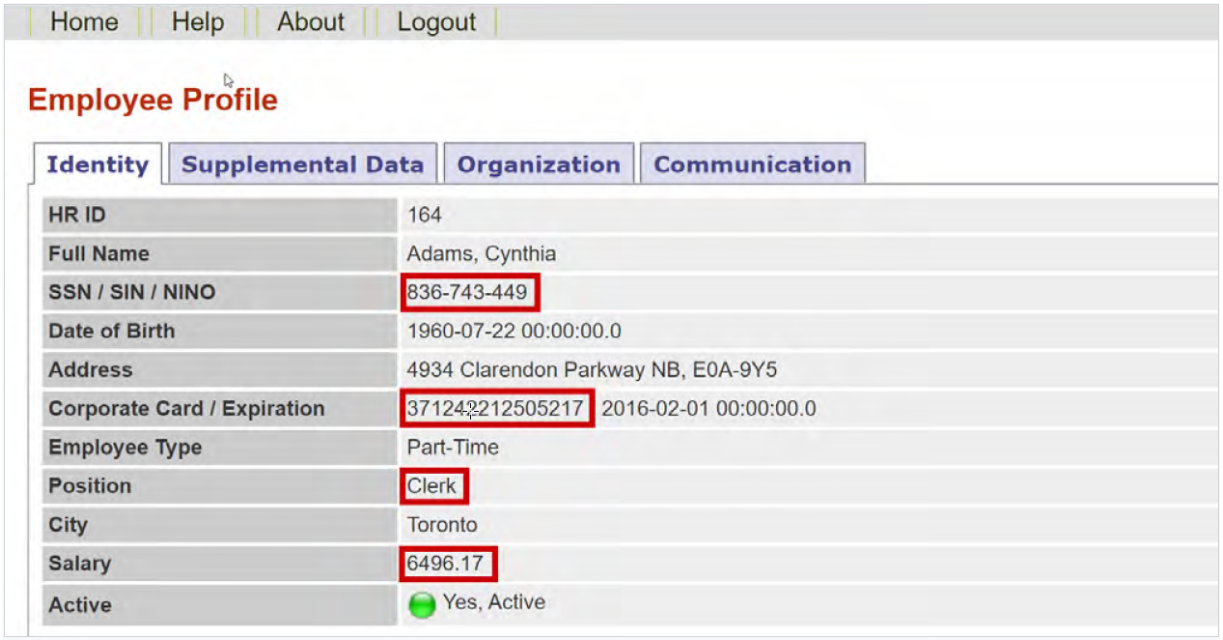


Figure 10-10: Example showing unredacted application data

Figure 10-11 shows the same data with redaction policies applied to the `SSN`, `Corporate Card, Position`, and `Salary` fields using partial, regular expression, random, and full transformations.

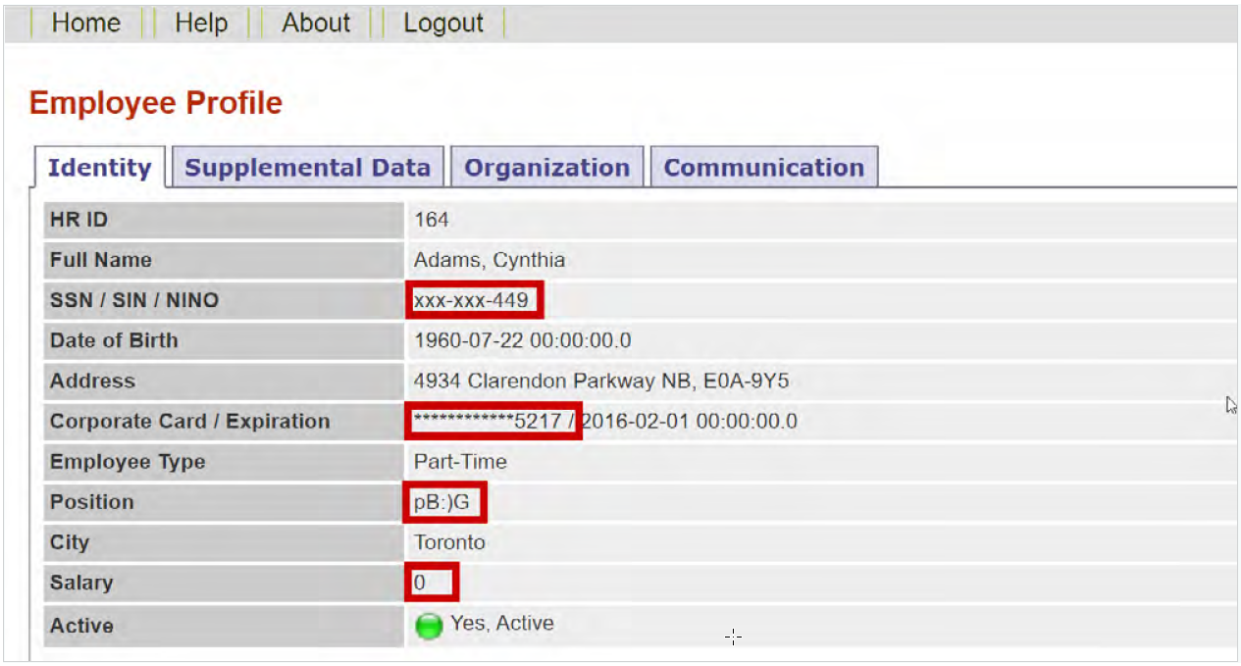


Figure 10-11: Example showing redacted application data

In addition to these simple formats, Data Redaction supports predefined redaction formats for common types of sensitive data, including credit card numbers, Social Security numbers, postal codes, email addresses, and phone numbers. Table 10-3 shows a selection of the predefined redaction formats. Organizations can also create custom redaction formats to match unique requirements.

Copyright © 2026, Oracle and/or its affiliates

Format	Description
US SSN first five (formatted)	Redacts the first 5 numbers of a Social Security number when the column is a VARCHAR2 data type. For example, the string 987-65-4320 becomes XXX-XX-4320.
US SSN last four (formatted)	Redacts the last 4 numbers of a Social Security number when the column is a VARCHAR2 data type. For example, the string 987-65-4320 becomes 987-65-XXXX.
US SSN total (formatted)	Redacts the entire Social Security number when the column is a VARCHAR2 data type. For example, the number 987-65-4320 becomes XXX-XX-XXXX.
US SSN first five	Redacts the first 5 numbers of a Social Security number when the column is a NUMBER data type. For example, the number 987654320 becomes XXXXX4320.
US SSN last four	Redacts the last 4 numbers of a Social Security number when the column is a NUMBER data type. For example, the number 987654320 becomes 98765XXXX.
US SSN total	Redacts the entire Social Security number when the column is a NUMBER data type. For example, the number 987654320 becomes XXXXXXXXX.
CA SIN six nines plus last three	Redacts the Canadian Social Insurance number by replacing the first 6 digits by 9 (a number). For example, 123456789 becomes 999999789.
CA SIN last three	Redacts the Canadian Social Insurance number by replacing the first 6 digits by X (string). For example, 123456789 becomes XXXXXX789.
CA SIN last three (formatted)	Redacts the Canadian Social Insurance Number by replacing the first 6 digits by X (string). For example, 123-456-789 becomes XXX-XXX-789.
UK NIN alpha (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 becomes ET XX XX XX D.
UK NIN alpha	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 becomes ETXXXXXXD.
Credit Card last four (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 becomes ****–****–****–5100.
Credit card last four	Redacts the credit card number (other than American Express) by replacing everything but the last 4 digits by *. For example, the credit card number 5105105105105100 becomes ************5100.
AMEX last five (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 becomes ****-******-10005.
AMEX zeros plus last five	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 becomes 0000 000000 10005.
US Zip code (varchar)	Redacts a 5-digit postal code when the column is a VARCHAR2 data type. For example, 95476 becomes XXXXX.

Copyright © 2026, Oracle and/or its affiliates

US Zip code (number)	Redacts a 5-digit postal code when the column is a NUMBER data type. For example, 95476 becomes XXXXX.
Date to 1970	Redacts dates to 01-JAN-1970.
Date to millennium	Redacts dates to 01-JAN-2000.
North America phone number (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.
North America phone number (zeros)	Redacts the North American phone number by leaving the area code but replacing everything else with 0. For example, 6505550100 gets redacted to 650000000.
North America phone number	Redacts the North American phone number by leaving the area code, but replacing everything else with X. For example, 6505550100 is redacted to 650XXXXXXX.

Table 10-3: Summary of predefined redaction formats

Implementing Data Subsetting

Oracle Data Masking and Subsetting provides comprehensive data subsetting capabilities that allow organizations to create smaller datasets derived from production data while maintaining referential integrity. For example, a large 2 TB production dataset can be subset down to 100 GB for development or analytics. This smaller dataset enables more efficient operations while using fewer resources.

As shown in Figure 10-12, data subsetting supports several strategies for reducing table size. Organizations can extract a percentage of a large table for application development, pull data from a specific geographic region for localized analysis, or select rows from a particular partition or sub-partition.

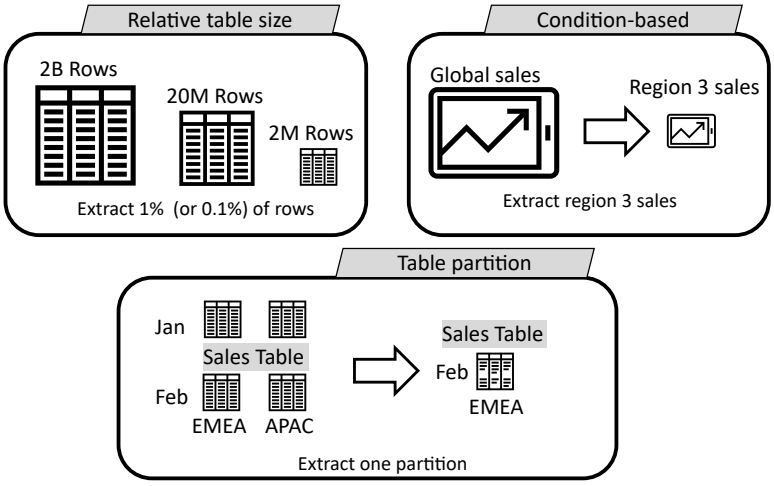


Figure 10-12: Data Subsetting strategies

Copyright © 2026, Oracle and/or its affiliates

Table 10-4 describes different methods for structuring the conditions used to extract data for subsetting.

Term	Description
Random Sampling	Randomly selects a subset of data from a dataset. This approach is useful when a sample that represents the entire dataset is needed, such as for development or testing
Stratified Sampling	Use this method when the dataset contains different groups or strata. Stratified sampling ensures that each group is adequately represented, and organizations can define rules for sampling each group.
Time-based Subsetting	Subsets data based on time criteria, such as selecting records from a particular date range. For example, organizations may want the oldest or the most recent data for analytics.
Conditional Subsetting	Selects data based on certain conditions or criteria. For example, organizations might extract records where an attribute meets a certain threshold.
Top-N Queries	Selects a subset of data based on the top or bottom N records for a particular attribute. For example, organizations might extract the top 1000 customers by sales volume.

Table 10-4: Subsetting conditions

Organizations can combine subsetting and masking in a single workflow. First, extract only the data required from production, then mask any sensitive information within that subset before provisioning it to non-production environments. This combined approach maximizes the business value of data, protects sensitive information, and avoids wasting resources on unnecessary storage and processing.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Get Started with Oracle Data Safe Fundamentals
- Data Redaction
- Redaction for Autonomous Databases
- Data Masking and Subsetting

Copyright © 2026, Oracle and/or its affiliates

Summary

Modern applications generate and consume huge volumes of sensitive information, from personally identifiable information and financial records to healthcare data and proprietary business details. As that data spreads into development, testing, analytics, training, and daily operations, organizations run into a familiar tension: they want the realism of production data without the exposure that comes from unnecessary visibility.

Realistic datasets unlock real value. Developers validate behavior against believable edge cases, marketing and customer service teams find patterns faster, call centers resolve issues with less back-and-forth, and fraud models learn from meaningful signals. But every new copy of production data widens the attack surface, increasing the odds of misconfigurations, insider misuse, lost devices, and external attacks. The risk climbs even higher when AI and machine learning teams reuse these environments for model training and retrieval augmented generation pipelines, because real PII can end up in places that were never designed to protect it.

Data masking is presented as the pragmatic middle path: convert sensitive data into realistic but fictitious values that preserve utility while stripping away the value of theft. A key point is that masking is not designed to be reversed, and that irreversibility is a feature when sensitive information must be permanently removed from non-production environments. Masking also preserves table relationships through consistent replacements that maintain referential integrity. Alongside masking, the discussion positions data redaction and data subsetting as complementary privacy pillars that support least privilege, anonymization and pseudonymization expectations, and compliance obligations such as PCI DSS, EU GDPR, CCPA, and India DPDPA. A practical example grounds the idea: point-of-sale testing can keep real inventory and store configurations while swapping in plausible names, email addresses, and payment details.

Two primary approaches shape how organizations apply these protections:

- Static **Data Masking** permanently replaces stored values so the target environment contains only masked data. That makes it a strong fit for development, testing, training, analytics, and external data sharing where real customer, employee, or financial details are not needed.
- **Data Redaction**, or dynamic data masking, leaves stored data untouched but transforms query results based on policy. This enables production systems to show different users different versions of sensitive fields without application changes.

Examples reinforce the difference: static masking can replace employee names with realistic substitutes, randomize Social Security numbers into the correct format, and shuffle salaries to preserve totals for aggregate analysis while protecting individual compensation.

Implementation is described as a series of tradeoffs across speed, control, and deployment preferences. For data masking, the chapter compares Oracle Data Safe and Oracle Data Masking and Subsetting: Data Safe emphasizes cloud service execution and automated format suggestions, while Data Masking and Subsetting emphasizes mandatory discovery and more manual format assignment. Masking can also be done in-database using a staging clone, or in-export during Data Pump export. Data Redaction is described as a database kernel capability driven by declarative policies that can incorporate identities, roles, session attributes, and application context, with Oracle AI Database 26ai enhancements to keep redaction effective for complex SQL, grouping, sorting, and views with expressions. Data Subsetting rounds out the toolbox by extracting only the needed slice of production data, cutting risk and cost, speeding refreshes, and pairing naturally with masking in a combined workflow.

The next chapter shifts to protection to encryption, positioning it as a way to keep sensitive data authentic while allowing it to move through production systems with confidence. Continue reading to see how encryption layers onto masking and redaction strategies, and how it can serve as a bright, dependable shield that helps organizations move faster without losing control of what matters most.

Copyright © 2026, Oracle and/or its affiliates

Chapter 11

Data Encryption and Key Management

Why encryption matters for database security

Strong authentication and authorization controls protect the front door to the database, but attackers often attempt to bypass these normal access paths. Encryption ensures that data remains unreadable even when attackers circumvent access controls, because encrypted data is useless without the corresponding encryption key. Encryption shifts the security challenge from protecting large volumes of data to protecting a much smaller and more manageable set of keys.

Attackers can bypass database controls in several ways. They may intercept data traveling across the network between database clients and servers; on many networks, capturing traffic and reading clear-text information is straightforward. They can also exploit operating system privileges to read database files directly. Unencrypted backups present another high-value target, whether stored on disk, tape, offsite media, or in cloud storage. Because bypass attacks do not create a database session, traditional access controls offer no protection. Encrypting data at rest and in transit ensures that intercepted traffic, copied files, and stolen media do not expose sensitive information.

Organizations that implement encryption gain several strategic advantages:

- Persistent protection: Data stays protected at rest, in transit, and in backups when encryption is applied appropriately.
- Simpler control surface: Security teams can focus on key security and key lifecycle management instead of tracking every copy of the data.
- Compliance support: Many regulations and security standards require-encryption for sensitive or personally identifiable information (PII). Encrypting data helps organizations meet these requirements and reduce audit findings.

Encryption should be used as part of a defense-in-depth strategy. Organizations must ensure strong key management practices, limit key access according to least privilege principles, and monitor key usage to detect and respond to anomalies quickly. This chapter explains how encryption protects data in motion and at rest for Oracle Database. It also discusses considerations for securely storing and managing the encryption keys that ultimately safeguard the encrypted data.

Encrypting data in motion

Securing data in motion is essential for preventing eavesdropping and tampering. Network encryption is built into Oracle Database across Enterprise and Standard editions, enabling organizations to protect data as it moves between clients and the database. Two primary options are available for encrypting data in motion: industry-standard Transport Layer Security (TLS) and Oracle Native Network Encryption (NNE). Both options provide confidentiality to prevent eavesdropping and integrity protection to prevent tampering or replay attacks. TLS also adds server authentication and can optionally authenticate the client.

Most Oracle client software supports both NNE and TLS, including thin and thick JDBC drivers, Oracle Instant Client, SQL*Plus, SQL Developer, and the Visual Studio Code SQL Developer extension.

Transport Layer Security (TLS)

TLS provides both authentication and encryption for database connections. Organizations can configure server-authenticated TLS, in which the server presents an identifying certificate, or mutual TLS (mTLS), in which both the server and the client present certificates. This approach protects data in motion and can authenticate the client to the database. TLS certificates are stored in the system certificate store or in an Oracle Wallet.

Organizations that use PKI certificate-based authentication for database access depend on TLS to establish the secure channel over which authentication occurs.

Copyright © 2026, Oracle and/or its affiliates

TLS cipher negotiation and protocol versions

With TLS, the server and client negotiate and select the strongest mutually supported cipher suite. Organizations can explicitly require specific algorithms if desired.

Oracle AI Database 26ai and newer versions support TLS 1.3, while older versions support TLS 1.2. Organizations whose clients and applications both support TLS 1.3 should use it. TLS 1.3 streamlines the handshake process and uses modern cipher suites, improving security posture and reducing connection setup time compared to TLS 1.2.

Starting with Oracle AI Database 26ai (release 23.26.0), organizations can secure TLS 1.3 connections with Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) or the quantum-resistant ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism) algorithms. This flexibility allows organizations to choose well-established or post-quantum security based on risk profile and compliance needs. The January 2026 release update (23.26.1) added a hybrid key exchange option that combines ECDHE and ML-KEM. In hybrid mode, the database performs key exchange using both algorithms, producing two independent keys that are combined to form the final session key. This dual approach helps ensure that neither classical nor post-quantum vulnerabilities alone compromise the security of encrypted data. As long as at least one of these algorithms remains uncompromised, the encrypted connection is protected.

Native Network Encryption

Native Network Encryption (NNE) encrypts each connection with a unique shared secret key. When a client connects to an NNE-enabled database, the client and server negotiate a session-specific key used only for that connection. Neither side needs a certificate. Database administrators can enable NNE by adding a simple entry to the database network configuration file, and clients require no changes. Because NNE does not rely on external infrastructure or certificates, many organizations choose it for its straightforward setup.

Federal Information Processing Standard (FIPS)

Optionally, TLS and NNE can run in FIPS mode to enforce stricter control over approved cryptographic libraries. Oracle AI Database 26ai allows organizations to choose between FIPS 140-2 and FIPS 140-3 modes.

Security and performance considerations

Stronger ciphers typically increase compute overhead. However, TLS 1.3 often performs better than TLS 1.2. The updated handshake removes an extra round trip, which speeds up connection establishment, and the standardized cipher suites are modern and efficient, delivering slightly better performance while strengthening security.

For more information on configuring TLS encryption, refer to the Oracle AI Database Security Guide.

Encrypting data at rest

Organizations can encrypt data at multiple points in the technology stack: the application layer, the database layer, or the storage layer. In general, the lower encryption is implemented in the stack, the less intrusive and better performing it is. However, encrypting lower in the stack addresses fewer threat types.

Encryption strategies: application, database, and storage layers

Application-layer encryption sits at the top of the stack. Applications encrypt data before writing it to the database and decrypt it when needed. Each application must manage encryption and decryption throughout its workflows. This approach can affect database performance and reduce the types of queries that can run on encrypted columns. For example, common analytic queries that search ranges or compute values may be slow or fail because indexes are often less useful with application-encrypted data. If multiple applications share the same database, they must also share access to the same encryption keys to encrypt and decrypt data.

Copyright © 2026, Oracle and/or its affiliates

Application-layer encryption provides strong protection but adds significant development and operational overhead, and it limits meaningful relational operations on the data. Because the database may no longer use indexes effectively, performance can degrade. Application-layer encryption can make sense for narrow cases, such as storing a secret string or a blob that must not be visible to any other database user, including administrators, and that does not require database operations.

Database-layer encryption encrypts whatever data the application sends and stores the result in encrypted form. Multiple applications can use the same database without changes because the encryption is transparent to them. This protects against attacks that bypass the database and attempt to read data directly from storage. It does not protect against someone who connects through the database with authorization to read the data. Access controls remain critical: access controls guard the front door; encryption protects the windows and back doors.

Storage-layer encryption (file or volume encryption) protects data at rest against someone who bypasses the operating system and accesses storage directly. This provides limited mitigation for database threats. The main problem is that file or volume encryption lacks database user context. If Oracle Database runs as the operating system user "oracle," any other user logged in as "oracle" or "root" can access decrypted data. Because the goal is to prevent database bypass, a control that automatically decrypts data for operating system utilities that read or copy files (for example, file-search and file-copy tools) is not effective.

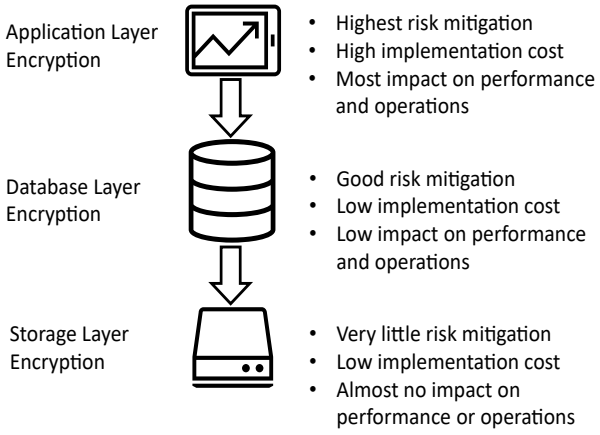


Figure 11-1: Encryption offers trade-offs at each layer

Storage-layer encryption usually has minimal impact on stability and operations, but results depend on the solution. Encrypting file systems such as ZFS or Microsoft BitLocker, or secure-wipe encryption built into high-end media such as that used in Exadata storage systems, tend to be stable. Third-party file system encryption that inserts itself between the operating system and the database can introduce instability and complicate upgrades. It can also delay patching because teams must wait for the file system vendor's dependent patch before applying database or operating system patches. Because of these complications, Oracle cannot support database issues that arise from the use of third-party file system encryption.

Evaluating encryption solutions

Regardless of where-encryption is implemented, organizations should evaluate solutions by the following criteria:

- Which risks does the encryption mitigate?
- What is the performance impact?
- How simple is installation and use?
- How transparent is it to existing applications?
- What resources are needed to convert data between plaintext and encrypted forms?

Copyright © 2026, Oracle and/or its affiliates

- What are the patching considerations?
- How robust is key management?

Neither application-layer nor storage-layer encryption alone is optimal for protecting data in Oracle Database. Organizations need a mechanism that improves security and reliability while minimizing implementation effort and performance overhead.

Combining encryption strategies

One size rarely fits all. Many organizations blend multiple encryption strategies. For example, a retailer that retains financial data, including customers' credit cards for reuse, and personal data such as date of birth, physical address, email address, phone number, and purchase history might implement:

1. **Application-layer encryption for credit card numbers**: Data is encrypted at the point of collection and remains encrypted except when used to support purchases.
2. **Database-layer encryption for all other application data:** This protects against bypass attacks such as ransomware and limits data theft exposure while allowing analytics and, where customers have opted in, marketing use.
3. **Storage-layer encryption**: The SAN uses media configured for secure wipe when failed hardware must return to the manufacturer.

Transparent Data Encryption

When attackers try to bypass database controls and read files directly, data must remain protected. Transparent Data Encryption (TDE), a feature of Oracle Advanced Security, encrypts data within the database to ensure it remains unreadable outside an authenticated session.

How TDE works

TDE encrypts data before writing it to storage and decrypts it when reading it back. Authorized users and applications see decrypted (plaintext) data, but anyone who tries to read data through the operating system or from storage sees only encrypted data. Because database files are encrypted, backups that include that data are also encrypted, so stolen backup copies do not expose sensitive information.

Nothing changes for applications. Organizations keep using the database as they do today. The database enforces the access controls that have been defined and denies access to anyone who is not authorized, while TDE protects data at rest from bypass attacks.

This transparency is equally valuable for emerging AI workloads. Data derived from sensitive sources, including vector embeddings, training datasets, and inference caches, deserves the same protection as the original production data, and TDE provides this protection at the tablespace level with no application modifications.

Copyright © 2026, Oracle and/or its affiliates

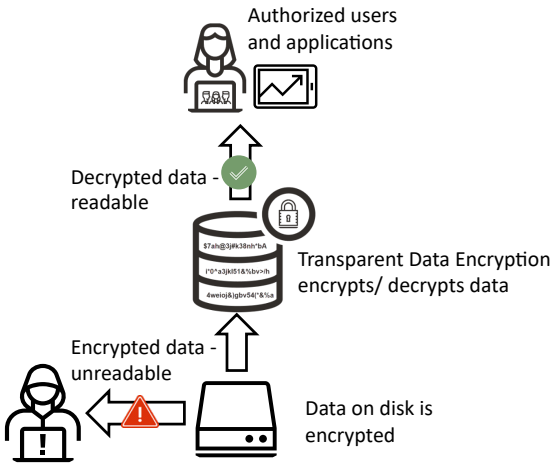


Figure 11-2: Users and applications access the decrypted data in database memory

Tablespace encryption and column encryption

TDE can encrypt whole tablespaces or individual columns. In almost all cases, tablespace encryption will be the optimal choice. With tablespace encryption, database administrators do not need to track which columns to encrypt or worry about someone inserting sensitive data into a column that would not usually be considered sensitive. With tablespace encryption, administrators also do not need to worry about column characteristics such as indexes and constraints.

Creating encrypted tablespaces

Creating an encrypted tablespace is straightforward:

```
SQL> CREATE TABLESPACE investigations
     DATAFILE '$ORACLE_HOME/dbs/investigations.dbf' SIZE 50M
     ENCRYPTION [USING 'AES256' MODE 'XTS'] ENCRYPT;
```

Figure 11-3: Example syntax creating an encrypted tablespace

Note: The default encryption algorithm in Oracle AI Database 26ai is AES256 using XEX-based tweaked-code book mode with ciphertext stealing (XTS) encryption. Shorter key lengths (AES192 and AES128) are available, but not advised due to concerns about quantum computing resistance. Pre-26ai databases implemented AES using cipher feedback (CFB) mode. Oracle AI Database 26ai still offers CFB as an option for backward compatibility.

Organizations may configure their databases to encrypt new tablespaces automatically, even if an application does not create encrypted tablespaces as part of the application installation process.

When TDE is used to encrypt tablespaces, sensitive data remains encrypted throughout the database, whether in tablespace storage files, temporary or undo tablespaces, or other files such as redo logs. In addition, TDE can encrypt RMAN database backups and Data Pump exports.

TDE automatically leverages special instructions in Intel/AMD (AES-NI), SPARC, ARM, and IBM Power CPUs to accelerate cryptographic operations. In addition, TDE tablespace encryption integrates with the performance optimizations built into Oracle Engineered Systems such as Exadata and the Oracle Database Appliance. For example, TDE tablespace encryption works seamlessly with Exadata Hybrid Columnar Compression (EHCC) and Smart Scan technology.

Copyright © 2026, Oracle and/or its affiliates

Migrating existing data to encrypted tablespaces

When encrypting tablespaces in an existing database, there are two choices – online and offline.

Online tablespace encryption, introduced with Oracle Database 12.2.0.1 enables zero-downtime migration from plaintext to encrypted data or converts from one encryption algorithm to another. Online encryption temporarily uses additional storage of the same size as the largest processed tablespace. In the sample command shown in Figure 11-4, datafile hr_data01.dbf will be copied and written in encrypted form into datafile hr_data01_enc.dbf. When the encryption operation completes, the original datafile hr_data01.dbf will be overwritten with zeroes and deleted.

```
ALTER TABLESPACE hr_data ENCRYPTION ONLINE [using ‘algorithm’] ENCRYPT
[FILE_NAME_CONVERT=('hr_data01.dbf','hr_data01_enc.dbf')];
```

Figure 11-4: Example syntax for online tablespace encryption

TDE also supports two offline encryption methods that do not require extra storage but do require the tablespace to either be offline or the database to be in mount mode. Both are "standby-first" migration methods if the database is protected by a standby database. The procedure is to turn off the managed recovery process, encrypt the standby database, resume managed recovery and wait for the recovery process to catch up, then perform a switchover and encrypt the new standby (formerly primary) database. In this scenario, the downtime for encrypting any size database is as short as the duration of one or two switchovers.

The first offline encryption method simply has administrators take the tablespace offline and issue an ALTER TABLESPACE command to encrypt all data files in the tablespace. The command can be as straightforward as shown below if the default encryption algorithm and mode are acceptable (AES256 in XTS mode in Oracle AI Database 26ai, AES128 in CFB mode in Oracle Database 19c). If desired, administrators can override the defaults by specifying the algorithm and mode of choice.

```
ALTER TABLESPACE hr_data ENCRYPTION OFFLINE [using 'algorithm' mode 'XTS'] ENCRYPT;
```

Figure 11-5: Example syntax for offline tablespace encryption

The second offline encryption method allows administrators to encrypt the data files, potentially multiple in parallel:

```
ALTER DATABASE datafile '$ORACLE_HOME/dbs/investigations.dbf' ENCRYPT;
```

Figure 11-6: Example syntax for offline data file encryption

This command encrypts the data files with the default encryption algorithm.

TDE and the Database Configuration Assistant (DBCA)

Starting with Oracle Database 21c, DBCA can create encrypted databases with a wallet-based TDE setup, so the database is encrypted from the very start. In addition to creating a new encrypted database, DBCA in 26ai allows organizations to:

- Create new encrypted databases with their keys managed in Oracle Key Vault (OKV).
- Encrypt an unencrypted database.
- Migrate an encrypted database from wallet-based TDE setup to centralized key management with OKV.

Key management

Encryption is only as strong as key management. Organizations must store keys securely, restrict access according to least privilege principles, and validate backup and recovery processes. This section describes options for key management, beginning with local key management using Oracle Wallets and progressing to centralized key management with Oracle Key Vault.

Copyright © 2026, Oracle and/or its affiliates

TDE two-tier key architecture

At the start of this chapter, the message was clear: "*Encryption reduces the problem of securing a large amount of data to a more straightforward problem of securing a much smaller key used to encrypt the data.*" That idea sits at the heart of encryption in the real-world. Encrypted data stays protected only as long as the encryption keys stay secret, which is why disciplined key management is so important.

With the right key management approach, attackers are far less likely to discover an encryption key and unlock access to sensitive data. It also helps keep active keys from being misplaced, ensures keys are rotated periodically, and makes sure older keys are securely archived so encrypted backup data sets remain usable over time. Many regulations, including PCI-DSS, require physical separation between encrypted data and encryption keys and recommend periodic key rotation to shrink the exposure period if a key is ever compromised.

TDE uses a two-tier key architecture to create and manage the keys used for encryption. The first tier is a master encryption key (MEK) used throughout the database. The second tier is a set of data encrypting keys (DEK) that are unique for tablespaces or tables, depending on whether column or tablespace encryption is used.

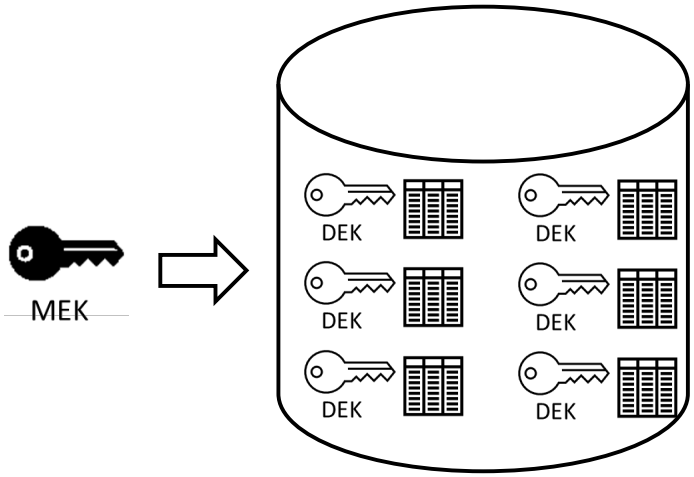


Figure 11-7: TDE uses a two-tier key architecture to improve key management security

As shown in Figure 11-7, a single MEK is kept outside of the database and is used to encrypt and decrypt the internally generated DEK. Those DEKs, in turn, encrypt tablespace data files or table columns. The database automatically manages DEKs behind the scenes.

This two-tier key architecture streamlines key rotation. When rotating the MEK, there is no need to decrypt and re-encrypt all data, only the much smaller set of DEKs. Online tablespace encryption can be used if there is ever a need to rotate the DEKs or switch to another encryption algorithm.

Administrators perform key management operations, such as create, open, rotate, and backup, through SQL commands or, if the database is an Oracle Cloud database service, through the OCI console. Oracle Database 12c introduced the SYSKM administrative privilege, enabling key administrators to carry out key management operations without requiring the powerful SYSDBA privilege.

The MEK is always stored outside of the database, separated from encrypted data, and managed by the key administrator in the Oracle Wallet, Oracle Key Vault, or OCI Vault as shown below in Figure 11-8.

Copyright © 2026, Oracle and/or its affiliates

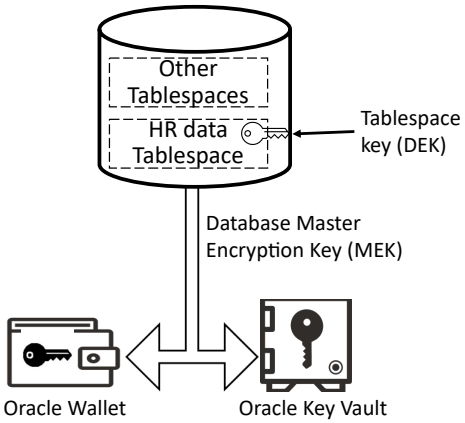


Figure 11-8: Example TDE encryption keys managed by either Oracle Wallet or Oracle Key Vault

The database master encryption key is created when TDE is first enabled, and it can be rotated using an `ADMINISTER KEY MANAGEMENT` command. Tablespace keys are automatically created when a tablespace is encrypted, and a tablespace key can be rotated using an `ALTER TABLESPACE` command. Each encrypted tablespace has its own tablespace key, which is stored in the tablespace header and encrypted using the master key.

With Oracle AI Database 18c and later, externally generated MEKs can be imported into the key store, enabling bring your own key (BYOK) functionality for those who prefer to use an external key generator.

Local key management with Oracle Wallets

By default, Transparent Data Encryption (TDE) stores the master encryption key in an Oracle Wallet that follows the PKCS#12 standard. The wallet's contents are encrypted with an AES256 key derived from the wallet passphrase. If the wallet is lost or corrupted, or the passphrase is forgotten, the encrypted data cannot be recovered. TDE has no built-in recovery 'back doors' for lost passwords. Organizations must securely back up the password-protected wallet and test recovery regularly.

To streamline daily operations, organizations can use local auto-login wallets so the database can open the wallet without manual passphrase entry. A local auto-login wallet is tied to the server where it is created, which reduces the risk of misuse on another system. Some operational guidance when using wallets includes:

- Limit wallet and key access to authorized administrators only.
- Protect wallet files with strong operating system permissions and encrypted backups.
- Document and test wallet backup and recovery steps on a defined schedule.
- Rotate keys according to policy and coordinate rotations with maintenance windows.
- Monitor key usage and alert on anomalies.

Centralized key management with Oracle Key Vault

As more databases are encrypted, managing hundreds or thousands of wallets becomes risky and time-consuming. Lost or corrupted wallets and forgotten passphrases can put data and entire databases at risk. Many regulations and security best practices also prohibit storing encryption keys on the same server as the encrypted data, and many organizations require clear separation of duties between database administrators and key management administrators.

Copyright © 2026, Oracle and/or its affiliates

Oracle Key Vault (OKV) helps organizations solve these challenges. With OKV, organizations centralize and secure master keys, wallets, and other secrets, reduce the operational burden of wallet sprawl, and lower the risk of losing access to encrypted data. Key Vault provides:

- **Central control**: Store, manage, and distribute TDE master keys, wallets, and secrets from a single secure service.
- **Strong security**: Keep keys off database servers and enforce least-privilege access with separation of duties.
- **Operational resilience**: Automate backup of wallets and keys, streamline rotations, and reduce the chance of lockouts.
- **Compliance support**: Align with regulations and security best practices that require externalized key management and auditable controls.

With Key Vault, organizations gain an enterprise-grade key manager that secures and reliably distributes the keys and secrets their databases depend on, while simplifying day-to-day operations.

Oracle AI Database connects to the Key Vault (or, more typically, the Key Vault cluster) through a PKCS#11 library installed on the database server as shown in Figure 11-9 below.

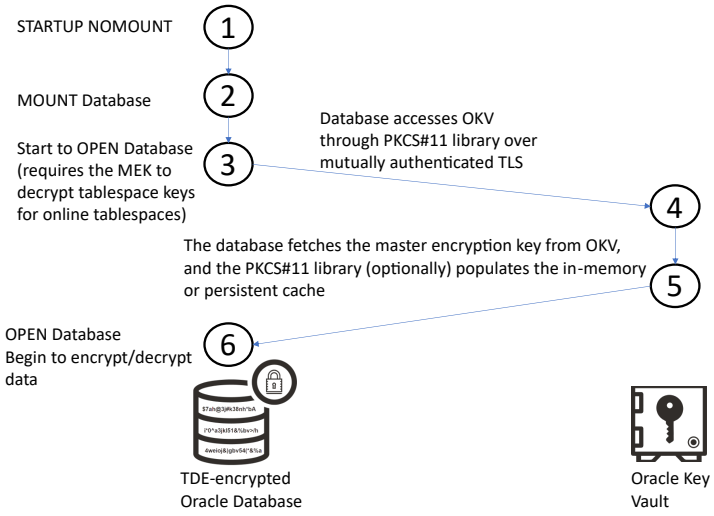


Figure 11-9: Oracle Key Vault securely delivers MEKs to Oracle AI Databases on demand.

Oracle Database connects to the Key Vault (or, more typically, the Key Vault cluster) through a PKCS#11 library installed on the database server. The master key is used for many database operations. Operations that require the database to consult the MEK include:

- When the database is started (opened)
- When a new tablespace is created
- When a tablespace or data file is brought online
- When a redo log group switch occurs
- When a new archive log is created
- When a parallel query process is begun
- During direct-path access to the database
- Initial access to an encrypted column

Copyright © 2026, Oracle and/or its affiliates

- GoldenGate and Data Guard operations

There is also a "heartbeat" check of the external keystore to ensure it remains available. If the database cannot access the master encryption key, it immediately closes the keystore to preserve data and prevent corruption. To avoid downtime, organizations should keep Key Vault (OKV) highly available so it can respond to key requests quickly. OKV delivers continuous availability and high scalability for read and write key operations using a multi-master cluster architecture. Organizations can create clusters with up to 16 nodes across geographically distributed data centers. As long as the database can reach any node in the cluster, operations continue. Organizations can place nodes on-premises and in major cloud providers (Oracle, Amazon, Microsoft, and Google). Because all nodes are active, each node holds a copy of all keys and shares the key management and distribution load.

Databases can connect to any OKV node to retrieve master encryption keys. OKV synchronously replicates key updates and authorization changes to a partner node to protect against data loss if a node fails. If a connection fails or a node goes down, database servers transparently fail over to a nearby active node for read and write operations without downtime or administrator intervention.

To guard against transient network issues, the default configuration uses a local persistent cache. OKV endpoints cache keys in an encrypted software keystore on the database server and serve them if the network link to OKV is unavailable. Organizations can protect the persistent cache with a random password that differs from the keystore password. Most key operations are satisfied from this cache, which improves performance and reduces dependency on the network.

Organizations can configure keys as non-extractable so they never leave the OKV cluster boundary. When the database must decrypt a data encryption key (DEK), it requests the decrypted DEK over an encrypted channel from OKV.

Multi-cloud and hybrid use cases

Oracle Key Vault allows organizations to retain ownership and portability of encryption keys, enabling seamless database movement across OCI, Azure, AWS, Google Cloud, and on-premises deployments. This empowers enterprises to meet regulatory requirements for key control while ensuring operational flexibility across diverse infrastructure environments.

Unlike native cloud KMS solutions, which lock encryption keys to a single cloud provider or environment, Oracle Key Vault enables several strategic use cases:

- **Multi-cloud disaster recovery**: Architecting a primary database in one cloud (for example, AWS) with a standby in another (for example, Azure) is unworkable with native KMS, as keys cannot be accessed cross-cloud. OKV provides an independent key management fabric, allowing both sites to operate and failover seamlessly.
- **Hybrid and multi-cloud operations**: Relying solely on native KMS creates a fragmented security posture, requiring management of multiple incompatible key management systems across clouds or between on-premises and cloud. OKV consolidates key management for all environments into one consistent platform.
- **Backup portability and repatriation**: Database backups are tied to the KMS that encrypted them, making cross-cloud or on-premises restoration impossible with native KMS. OKV ensures keys are independent and accessible from any environment, enabling true backup portability.
- **Intra-cloud regional migration**: Even within a single cloud, migrating databases between regions can result in dependencies and risk if keys remain in the original region's KMS. OKV's distributed architecture provides local key access in every location, ensuring regional independence and resilience.

Copyright © 2026, Oracle and/or its affiliates

- **On-premises clean room recovery**: Security best practices often require restoring cloud backups to isolated on-premises environments for forensics or breach recovery. Native KMS cannot support this, but OKV can be deployed within the clean room, enabling secure recovery.
- **Multi-cloud sharding**: For globally distributed sharded databases that span multiple clouds, native KMS solutions create "encryption islands." OKV's cross-cloud key management allows sharded architectures to function as a single logical database.

By delivering these capabilities, Oracle Key Vault unifies and simplifies key management to support modern data mobility, disaster recovery, and security requirements across hybrid and multi-cloud environments.

OKV supports TDE keys across Oracle Database deployments that use Multitenant, Oracle Real Application Clusters (RAC), Data Guard, globally distributed (sharded) databases, GoldenGate, and other Oracle technologies. Beyond TDE, OKV provides SSH key governance and manages:

- Java Keystores
- MySQL and MongoDB encryption keys
- Solaris Crypto keys
- Oracle ASM Cluster File System (ACFS) volume encryption keys

If organizations prefer local wallets, OKV can back them up on a schedule as a key escrow service. OKV can also back up Java Keystores and PKI certificate wallets.

Key Vault with DBMS_CRYPTO

Separating keys and secrets from application code increases security and simplifies development. Key Vault offers an integration accelerator that allows developers, after adapting a few configuration files to their environment, to implement key management for custom cryptographic applications using Oracle Database's DBMS_CRYPTO package.

Key Vault for SSH key governance and secrets management

Beyond encryption, Oracle Key Vault (OKV) helps organizations govern SSH keys and control remote server access across their environment. Centralized SSH key management protects private and public keys from loss and misuse and allows organizations to define how keys are created, distributed, rotated, and revoked. During a security incident, organizations can disable SSH access across servers with a single action or API call. Because OKV is the central point for SSH key access, it also provides a clear audit trail that shows who accessed a server, when access occurred, and the source of the request.

OKV is also a comprehensive secrets manager for credentials used in daily IT tasks. Many teams rely on locally managed Secure External Password Store (SEPS) wallets for jobs such as nightly RMAN backups, batch data loads, or refreshing materialized views. Because those wallets reside on the database server, they can be stolen. Instead, organizations can store database account passwords in OKV and retrieve them securely at runtime via an API call. Centralizing credentials in OKV strengthens security and removes the overhead of protecting hundreds or thousands of SEPS wallets used for automation. It also simplifies tightly controlled sharing of those passwords between databases.

C and Java SDKs let developers integrate Java, C, or C++ applications with OKV to manage encryption keys, passwords, and other secrets and retrieve them on demand.

OKV protects both public keys and non-extractable private SSH keys from loss and misuse. With centralized SSH key governance, organizations can enforce consistent policies fleet-wide and, when needed, shut off SSH access quickly during an incident while retaining a complete audit trail.

Copyright © 2026, Oracle and/or its affiliates

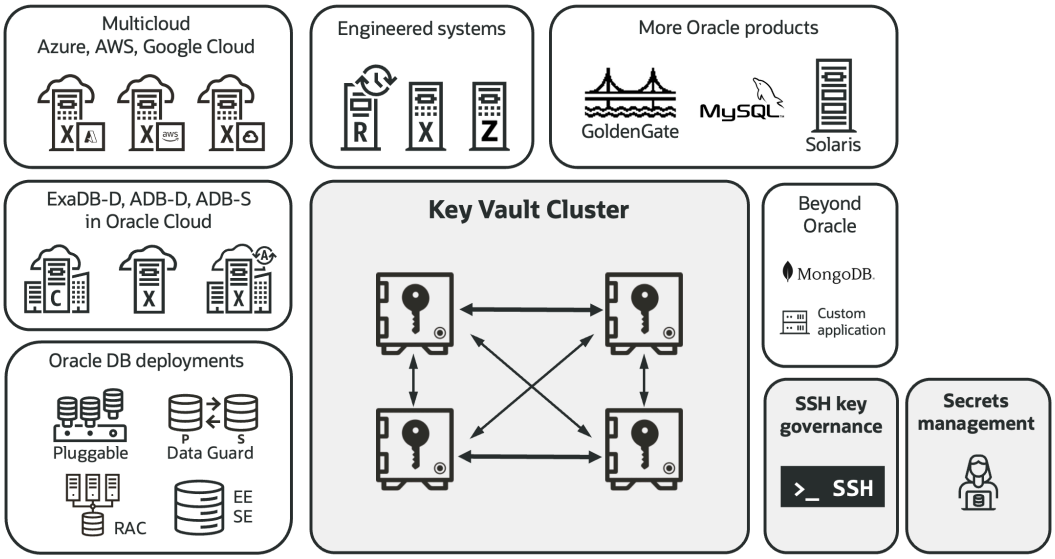


Figure 11-10: Oracle Key Vault manages Oracle TDE keys as well as other keys and security objects.

Key Vault APIs for automation

Key Vault's RESTful APIs make it easy to monitor a Key Vault cluster, add or remove nodes from an existing cluster, or even create a new cluster. Onboarding new servers and managing encryption keys may also be automated through the APIs, with the process of distributing and enrolling Key Vault endpoints fully automated. Automation allows database administrators to encrypt their databases while the security team securely manages the master keys within Oracle Key Vault.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- DB Security – Native Network Encryption
- Protect database communications using 1-way TLS
- DB Security – ASO (Transparent Data Encryption)
- Pluggables, Clones, and Containers – Securing Oracle Multitenant with TDE
- DB Security – Key Vault
- DB Security – Key Vault (SSH Key Management)

Copyright © 2026, Oracle and/or its affiliates

Summary

Encryption is the safety net that keeps working even when other defenses get sidestepped, especially in bypass scenarios where no database session is ever created. Real-world attack paths make the stakes concrete: sniffing data on the wire, using operating system privileges to read database files, or simply walking off with unencrypted backups sitting on disk, tape, offsite media, or in cloud storage. In practical terms, encryption turns exposed data into gibberish unless the attacker also gets the key. That shift naturally moves the spotlight to strong key custody, disciplined key lifecycle practices, and sharp-eyed detection of suspicious key activity.

On the network side, protecting data in motion is best handled with TLS or Oracle Native Network Encryption. Both provide confidentiality and integrity protections that help block tampering and replay. TLS stands out as the more identity-rich option because it authenticates the server and can also authenticate the client through mutual TLS, using certificates stored in an Oracle Wallet or the system certificate store. Version and algorithm choices matter too: Oracle AI Database 26ai and newer support TLS 1.3, which improves handshake efficiency compared to TLS 1.2, with additional options such as ECDHE, ML KEM, and a hybrid exchange combining the two. Oracle Native Network Encryption, by contrast, is often the smoother operational path. It avoids certificates, can frequently be enabled with straightforward network configuration, and still generates a unique key for each connection.

For data at rest, it helps to think in layers: application layer, database layer, and storage layer encryption. Lower layers can be less intrusive, but they may cover fewer threat types. Application layer encryption can raise the cost of development and operations, and it can restrict the database's ability to use indexes and run certain analytic queries. Storage layer encryption can also come up short against bypass attacks when operating system users like `oracle` or `root` can access decrypted data. For Oracle Database, Transparent Data Encryption (TDE) is the purpose-built database layer choice: it encrypts data before it is written to storage, keeps it unreadable outside an authenticated session, and encrypts backups that include encrypted files, all without requiring application changes. In almost all cases, tablespace encryption is the recommended approach over column encryption. It's also important to plan how encryption gets applied: encrypted tablespaces can be created directly, and existing data can be encrypted through online conversion or offline methods, including standby-first patterns that keep downtime largely limited to switchovers.

The unifying theme is simple and unavoidable: encryption is only as strong as the keys behind it. TDE's two-tier key architecture helps operationally because an externally stored master encryption key protects internal data encrypting keys, and rotating the master key does not require decrypting and re-encrypting all data. Key storage, however, comes with real operational risk. Wallet-based storage using PKCS#12 Oracle Wallets works, but losing the wallet or forgetting the passphrase can make encrypted data unrecoverable. That's why disciplined backups and regular recovery testing are essential. Auto-login wallets can reduce day-to-day friction, but they're tied to a single server. At enterprise scale, Oracle Key Vault offers centralized key management that reduces wallet sprawl, supports separation of duties, provides high availability via multi-master clustering, adds local persistent caching to reduce network sensitivity, and can keep keys within the cluster boundary through non-extractable configuration while still meeting database needs over secure channels.

The next chapter turns from locking data down to lighting the environment up. Database Auditing and Database Activity Monitoring focuses on seeing what is happening as it happens, where the database tells its story in crisp, actionable signals that help teams validate controls, spot anomalies, and transform security from a hope into a measured, watchful advantage.

Copyright © 2026, Oracle and/or its affiliates

Chapter 12

Using and Managing Audit Data

Why auditing matters

Auditing is often viewed as one of the less exciting areas of security, something everyone knows they should do but too often postpones. There are many reasons for this. Some don't know what to audit; some don't know how to audit; some don't know what to do with the audit records once they have audited; and some believe the audit performance overhead is a penalty that doesn't justify the process.

Simply collecting audit records is only the first part of the equation. That collected data is worthless unless it is used. Database activity should be actively reviewed and monitored because accounts are always at risk of being compromised or misused, and early detection can help reduce the damage from a compromise.

As AI agents and MCP-driven pipelines connect to databases alongside human users, an audit strategy must account for both. Machine-generated queries can arrive at high volume, at unusual hours, and through service accounts that can blur traditional accountability. Capturing and distinguishing this activity is essential to maintaining visibility and control.

The value of database auditing

Database auditing helps organizations catch malicious behavior sooner, satisfy regulatory obligations, and move faster during incident investigations. It also creates an authoritative audit trail that shows who did what, when it happened, and where the request came from, bringing clarity, accountability, and stronger security. The data collected through auditing is useful for all of the following:

- **Detecting and understanding malicious activity**: Spot suspicious behavior quickly, uncover root causes, and reduce dwell time.
- **Assisting investigations**: After a breach or other suspicious activity, replay the timeline to see exactly what happened and when.
- **Deterring inappropriate behavior**: When people know activity is audited, they are less likely to take risky or malicious actions.
- **Non-repudiation**: Provide proof that an event occurred, including who performed it, when it happened, and the source of the request.
- **Supporting regulatory compliance audits**: Show a verifiable record of activity, especially access to or changes in financial and other sensitive data.
- **Monitoring privileged user** activity: Keep visibility into actions by database administrators (DBAs) and other privileged accounts that access critical systems and data. These high-value accounts are frequent targets for attackers.
- **Assisting troubleshooting**: For example, if hundreds of application servers connect to a database and one keeps trying to log in with the wrong credentials, failed login auditing helps you find the culprit and fix the issue quickly.

When monitoring and auditing become a standard part of day-to-day operations, organizations reduce risk, strengthen compliance, and give their teams the evidence they need to act quickly and accurately.

Oracle Database unified auditing

Oracle Database auditing is designed to be both robust and highly customizable, so organizations can fine tune auditing to match their specific security requirements. Auditing in Oracle has matured over decades, from its debut in Oracle Database 7 (1992) through the unified audit framework introduced in Oracle Database 12.1 (2012). Starting with Oracle AI Database 26ai, unified auditing becomes the only supported framework. By bringing previously

Copyright © 2026, Oracle and/or its affiliates

separate audit trails together into one normalized audit trail, teams can search once and get precise answers faster. Key advantages include:

- **Improved performance**: Overhead is usually 3% or less for audit generation, even under heavy load of up to 200 audit records per second.
- **Stronger security**: The unified audit trail lives in the dedicated AUDSYS schema and is protected from DML operations, including update, delete, and truncate.
- **Greater flexibility**: Conditional auditing and an extensible audit trail make it easier to capture exactly what is needed and nothing more.

Note: Legacy auditing is desupported starting in Oracle AI Database 26ai. If an organization upgrades while still using legacy audit, the existing traditional audit settings will continue to be honored, and audit records will keep collecting in their respective trails to help bridge the gap to unified auditing. In Oracle AI Database 26ai, traditional audit settings cannot be created or updated and can only be deleted.

Transitioning from legacy to unified auditing

Legacy auditing (traditional auditing) is desupported in Oracle AI Database 26ai, which means new 26ai databases simply do not include the legacy audit facility. When you upgrade an existing database to 26ai, what happens next depends on where you are starting from.

If the source database is using unified auditing with no traditional audit settings, the upgraded database keeps running on unified auditing. If the source database still relies on legacy auditing, those traditional audit settings remain in effect after the upgrade, and audit records continue to land in the traditional audit trail until you turn them off with NOAUDIT. Also note that in 26ai you cannot create or update legacy audit policies; you can only delete the ones that already exist.

The practical takeaway is to plan the move to unified auditing either before the upgrade or immediately after, using the guidance below.

- Oracle Database provides predefined unified audit policies to help you get started. If you upgraded from release 11g, at a minimum enable these policies, which address common security and compliance needs:
    - Security-relevant activity (ORA_SECURECONFIG), such as use of the CREATE USER system privilege
    - Logon failures (ORA_LOGON_FAILURES)

  All new databases created from release 12.2 and later automatically enable ORA_SECURECONFIG and ORA_LOGON_FAILURES.

If you have highly customized legacy audit settings, you have the following choices to transition them - listed in order of preference:

- Create custom unified audit policies that use the richer features of unified auditing to make your policies more conditional, selective, and focused. The unified audit facility supports fine-grained conditions and an extensible audit trail.
- If you're unfamiliar with unified audit syntax, use the syntax converter script in My Oracle Support note 2909718.1. It converts legacy audit configurations into syntactically correct unified audit policies and writes them to a script for your review. Oracle strongly recommends reviewing and enhancing the generated policies-add conditions or audit application context values-before you enable them.

After you convert legacy settings to unified audit policies, run NOAUDIT to disable traditional audit policies and remove traditional audit parameter settings. Update any purge jobs created with the DBMS_AUDIT_MGMT PL/SQL package to align with unified audit configuration.

Copyright © 2026, Oracle and/or its affiliates

Utilizing audit data

Organizations often operate hundreds or even thousands of databases where both user and administrator activity must be audited and monitored to satisfy regulations and help detect security breaches. While monitoring database activity is essential, it can produce audit trails with substantial amounts of data. Because these audit trails are spread across many databases, analysis and reporting quickly become challenging. Once an organization moves beyond a handful of databases, audit data should be moved into a consolidated repository to support safekeeping, analysis, alerting, and report generation.

Database auditing is frequently strengthened with solutions that collect and store audit data from an entire fleet of databases. Oracle Data Safe and Oracle Audit Vault and Database Firewall (AVDF) both provide this capability. In both cases, audit data is collected from the target databases and stored in a centralized audit repository. Both solutions include built-in information lifecycle management to manage retention and archiving of audit data, a wide set of audit reports with flexible filtering to meet requirements, and alerting features that highlight specific user actions or unusual behavior.

Both Data Safe and AVDF take a comprehensive approach, offering numerous benefits, including the following highlights:

- Centralize audit data in a dedicated repository to preserve data integrity and security
- Monitor database activity in near-real-time to capture audit records, analyze them against policies, and alert on violations
- Use built-in and customizable reports to meet security and compliance needs
- Aggregate and search audit records at scale to support forensic analysis
- Investigate breaches and suspicious activity with high-volume collection and efficient search across audit and network events
- Store and retrieve historical audit data
- Track user privilege changes to quickly identify risky escalation
- Demonstrate monitoring of critical systems and access to sensitive data with audit-trail analysis and auditor-ready reports to support regulatory compliance.

Audit data best practices

Collecting audit data is only the beginning. The real value comes from turning that data into action. As noted earlier, database activity is audited to do the following:

- Identify anomalous or suspicious activity
- Support regulatory compliance audits
- Assist in troubleshooting operations

This is where analysis, reporting, and alerting become essential. The centralized repository created for audit data effectively serves as a data warehouse for activity across the entire database fleet. It can be used to do the following:

- Mine audit records for post incident forensic analysis
- Generate alerts for unusual or policy violating behavior
- Produce reports for the stakeholders who need them
- Accelerate troubleshooting by correlating events across systems

As AI agents increasingly generate queries autonomously, audit data becomes the primary mechanism for confirming that those agents remain within authorized boundaries. Agent driven activity should be correlated against established

Copyright © 2026, Oracle and/or its affiliates

policies to verify that MCP servers and AI pipelines accessed only the data they were entitled to, and to flag any deviation for investigation.

Organizations that process sensitive data are often subject to regulations such as the General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), Payment Card Industry Data Security Standard (PCI-DSS), India's Digital Personal Data Protection Act (DPDPA), Health Insurance Portability and Accountability Act (HIPAA), IRS Publication 1075, Sarbanes-Oxley Act, and the UK Data Protection Act. Most regulations require visibility into who accessed sensitive data and a verifiable record of that access, often retained for several years.

Centralizing audit management

Database auditing is a foundational control for database security and regulatory compliance. In addition to defining audit policies, organizations also need to manage the audit data those policies produce. Audit data management typically includes the following:

- Controlling how much space audit data consumes in your databases
- Protecting audit records from tampering or deletion
- Retaining audit data to meet security goals and regulatory requirements

The strongest approach is to collect audit data into a centralized repository. Centralization makes it practical to move audit data out of source databases, freeing space and easing day-to-day operational burden.

A dedicated repository also provides a more controlled environment to protect audit records from accidental changes and deliberate manipulation. It supports separation of duties by preventing users who administer or access the originating databases from altering the audit trail.

In addition, a centralized repository makes retention and lifecycle policies easier to apply consistently. Organizations can archive audit data to low-cost storage when it is no longer needed for reporting or analysis and delete older records when retention obligations no longer apply.

What to audit

Auditing every single action may sound like the safest option, but it often brings unintended costs: more storage, slower performance, and a flood of records that can drown out the signals that matter. In practice, most audit records are routine and add little security value, which can make malicious behavior harder to spot. A smarter approach is selective auditing that targets significant events. Strong audit policies capture the details that matter and deliberately skip routine actions from known sources. With less noise, teams can surface threats faster and respond with greater confidence.

In most cases, effective auditing keeps policies focused on the following types of database events while ignoring routine activity generated from known sources.

1. Security-relevant events
2. Privileged user activity
3. Sensitive data access

Oracle Database includes predefined audit policies for common requirements, making it easy to get started. Once the right policy is selected, the next step is to run the appropriate command to enable it. Here is an example of enabling one of the most commonly used policies:

Copyright © 2026, Oracle and/or its affiliates

```
AUDIT POLICY ORA_SECURECONFIG;
```

Figure 12-112: Enabling a predefined audit policy

The following sections review the three recommended event types and the policies that enable organizations to begin auditing quickly.

Auditing security-relevant events

Some database actions deserve extra attention because their impact can ripple well beyond a single table or schema, and they can hint at abuse or an attempt to cover tracks. Security-relevant events that should be audited include:

- Failed login attempts – consider auditing all logins, with exceptions for allowlisted service accounts connecting from known locations and running known programs
- Schema structural changes – create/alter/drop table, view, stored procedures, and similar DDL
- Account changes – create/alter/drop of users, privilege and role grants, password changes, etc.
- Security policy changes, especially for schemas that contain sensitive data
- Activities using system privileges
- Large-scale data extraction using components such as Oracle Data Pump
- Activities from dormant database accounts

Three predefined unified audit policies help capture most of the activity above:

- ORA_SECURECONFIG
- ORA_LOGIN_LOGOUT (named ORA_LOGON_FAILURES in older database versions)
- ORA_ACCOUNT_MGMT

These policies cover the security-relevant events most organizations care about first. They can be enabled with a simple PL/SQL call (as shown in Figure 12-1) or through the audit management consoles in Oracle Data Safe or Oracle Audit Vault and Database Firewall (AVDF). The *Database Schema Changes* policy can also be provisioned from either the Data Safe or AVDF management console.

In addition, many other security-relevant events are always audited as part of the *Mandatory Audit*. These are activities that can potentially cause significant issues. For the complete list of auditable events covered by the mandatory audit, see the product documentation here.

Privileged user activity auditing

Because privileged accounts carry the highest risk, monitoring them is one of the fastest ways to spot unusual behavior and detect sensitive data leaks if an account is compromised or misused. From an auditing perspective, any database user granted system privileges such as SELECT ANY TABLE or ALTER USER should be treated as privileged, with activity monitored for potential abuse or misuse.

It is also important to note that top-level statements by administrative users (for example, SYSDBA, SYSKM) issued while the database is in the closed or mount state are included in the mandatory audit. Mandatory auditing is enabled by default and cannot be disabled.

Note: In this context, “top-level” refers to the statement the user actually enters. Auditing top-level statements captures only those statements, not the internal or recursive SQL generated as a result. For example, if top-level statements are audited and a DBA runs DBMS_STATS.GATHER_DATABASE_STATS, only the call to execute the PL/SQL package is audited. This approach avoids capturing the SQL executed inside named procedures and

Copyright © 2026, Oracle and/or its affiliates

functions, which could otherwise generate thousands of audit records, and it significantly reduces routine administrator noise.

To capture top-level actions by administrative accounts such as SYS during normal database operations, enable additional audit policies. If Data Safe or AVDF is used to monitor database activity, provision the *Admin Activity Auditing* policy to audit activities by privileged administrators, including SYS.

Database connections can be made through the database listener (remote) or locally from an account logged into the database server. Auditing local sessions is critical because they can bypass network monitoring and may connect to the database without any authentication beyond the database server operating system account that initiates the session. Oracle recommends auditing all top-level operations that originate from a local session. Create an audit policy as given below to capture all top-level actions with appropriate audit conditions and consider enabling it for interactive users who have direct access to the database.

```
CREATE AUDIT POLICY local_db_access
ACTIONS ALL
WHEN '(SYS_CONTEXT (''USERENV'', ''IP_ADDRESS'') IS NULL)'
EVALUATE PER SESSION
ONLY TOPLEVEL;
```

Figure 12-2: Auditing local access to the database

When there is a need to monitor every action taken by privileged accounts that access sensitive data, the most effective approach is to target those users directly. With Data Safe or AVDF, the *User Activity Auditing* policy can be provisioned to audit all activity for a defined set of users. Another option is to create a custom unified audit policy that captures everything a user does. Figure 12-3 provides an example: the first statement creates a policy that audits all actions, and the second statement enables that policy for the user DBA_DEBRA.

```
CREATE AUDIT POLICY actions_all_pol
ACTIONS ALL
ONLY TOPLEVEL;


AUDIT POLICY actions_all_pol BY DBA_DEBRA;
```

Figure 12-3: Example syntax for auditing all activity by a user

It is important to keep in mind that auditing everything a user does can dramatically increase the volume of audit records, which can drive up storage costs and add performance overhead.

Sensitive data access auditing

Auditing access to, and updates of, sensitive data helps teams detect inappropriate activity sooner and helps discourage people from accessing or changing data when they do not have a clear business reason.

A practical approach is selective auditing of sensitive data access, which provides visibility into access and updates without generating unnecessary noise. The first step is to scan the database for sensitive data using Data Safe, AVDF, Oracle Database Security Assessment Tool (DBSAT), or Oracle Enterprise Manager. Chapter Three provides more detail on sensitive data discovery.

Once sensitive data has been found, a custom audit policy can be created to audit access outside normal application behavior. Figure 12-4 shows an example that audits all activity against the employees table in the HR schema (which is presumed to contain sensitive data). The first statement creates the policy and identifies the object to monitor. The second statement enables the policy and applies it to all user.

Copyright © 2026, Oracle and/or its affiliates

```
CREATE AUDIT POLICY actions_on_hr_emp_pol
ACTIONS ALL
ON hr.employees;


AUDIT POLICY actions_on_hr_emp_pol;
```

Figure 12-4: Example syntax for auditing activity on the EMPLOYEES table containing PII

The policy above will likely capture routine access from the HR application, which can create a large volume of low value audit records and introduce additional performance and storage overhead. A better option is to audit only access to the table that happens outside the approved application. The updated example in Figure 12-5 filters most routine activity and focuses the audit trail on meaningful events.

```
CREATE AUDIT POLICY actions_on_hr_emp_pol
ACTIONS ALL
ON hr.employees
WHEN 'SYS_CONTEXT (''USERENV'', ''HOST'') NOT IN (''APPSERVER1'',''APPSERVER2'')'
EVALUATE PER SESSION;


AUDIT POLICY actions_on_hr_emp_pol;
```

Figure 12-512: Example syntax for auditing activity from outside of an application

In the modified policy, a WHEN condition is added to treat access through the application server as the trusted path. The policy checks the session context value HOST to confirm the connection came from the application server. To keep overhead low, the condition is evaluated once per session using EVALUATE PER SESSION, and the result is reused for the rest of that session.

As a result, activity on the table that does not come from the application server generates audit records, while routine application server activity does not. The WHEN condition can use complex AND/OR logic, but it is limited to 4,000 bytes. If more logic is required, a custom PL/SQL function can be used. For more complex conditions, the EVALUATE PER SESSION clause can greatly reduce overhead.

A practical strategy is to focus audit policies on privileged user activity, security relevant events, and sensitive data access. This keeps auditing centered on what matters, cuts down unnecessary audit records, and still supports audit goals.

Because audit policies are flexible by design, teams can target exactly what is needed while filtering out routine activity that drags down performance, increases storage costs, and makes unusual or malicious behavior harder to spot.

Effective audit policies

Here are a few general guidelines to follow when designing effective audit policies:

1. Do not duplicate mandatory audits - creating policies that duplicate the audit information already captured with the mandatory audit has a slight but measurable impact on performance since all audit policies must be evaluated for every command. Refer to the Oracle Database Security Guide for a complete list of mandatory auditable events corresponding to your database version. For instance, avoid duplicating mandatory audit policies in your own custom policies.
2. Leverage predefined audit policies - Oracle Database includes several predefined "best practice" unified audit policies that cover common security-relevant audit settings. You only need to enable them to start collecting audit data. Available policies vary by version, so check the Oracle Database Security Guide for your release.

Copyright © 2026, Oracle and/or its affiliates

For instance, you'll find a description of the predefined audit policies that ship with Oracle AI Database 26ai here. If you use Data Safe or AVDF, you can enable many of the recommended audit configurations and predefined Oracle Database audit policies using a single click or via a REST API call.

3. Create custom audit policies for context-dependent scenarios - some configurations are unique to your environment-for example, sensitive data access-so create custom unified audit policies in Oracle Database for those cases.

Auditing in the age of AI

As agentic AI systems become more autonomous and start executing complex tasks at machine speed, auditing moves from a nice to have to a practical necessity. By watching how agents invoke tools, share context, and run tasks through MCP pipelines that interact with Oracle Database, organizations can illuminate the data access patterns that sit behind autonomous actions.

Unified auditing helps deliver end-to-end visibility into how agentic AI accesses the database through MCP pipelines. When an agent sends MCP orchestrated SQL into Oracle Database, unified auditing captures the executed statements along with context, user identity, and source details. That makes it easier to detect unusual access patterns, unauthorized privilege usage, or abnormal data extraction that automation may trigger.

Oracle SQLcl MCP Server also makes tracing easier by identifying itself in the database session using the PROGRAM, MODULE, and ACTION properties in V$SESSION. MCP Server interactions can then be logged to UNIFIED_AUDIT_TRAIL using the statement shown in Figure 12-6.

```
AUDIT CONTEXT NAMESPACE userenv ATTRIBUTES module, action;
```

Figure 12-612: Example syntax for auditing attributes from USERENV

The sample unified audit record below captures an interaction between Oracle SQLcl MCP Server and a GitHub Copilot agent, in which the agent queries data in Oracle Database using SQL generated by Claude.

	CLIENT_PROGRAM_NAME	SQL_TEXT	APPLICATION_CONTEXTS
1	SQLcl-MCP	SELECT /* LLM in use is claude-3.5-sonnet */ EMPLOYEE_ID, …	(USERENV,ACTION=run-sql); (USERENV,MODULE=claude-3.5-sonnet)

Figure 12-712: Example syntax for auditing Agentic AI interactions

Managing and using audit data with Oracle Data Safe

An audit data warehouse supports these obligations by automating retention policies and producing auditor ready reports, helping organizations demonstrate consistent monitoring and timely analysis.

Two proven options are available to centralize, analyze, and act on audit data at scale:

1. **Oracle Data Safe –** A database security cloud service running in Oracle Cloud Infrastructure (OCI). Data Safe helps manage day-to-day security and compliance for Oracle AI Databases, both on premises and in the cloud. It includes policy management and retention, centralized audit collection, a broad library of reports with flexible filters, and alerting for specific activities or unusual behavior, helping teams identify issues faster and streamline reporting.
2. **Oracle Audit Vault and Database Firewall** (AVDF) – Deploy on premises or in the cloud to unify database auditing with network level SQL traffic monitoring. AVDF combines native audit data collection with network-based capture for commonly used databases, including Oracle, Microsoft, IBM, and SAP. It provides a highly scalable, enterprise grade data warehouse, collection agents, Database Firewall and SQL Firewall, robust

Copyright © 2026, Oracle and/or its affiliates

reporting and analytics, and a flexible alert framework. Recommended auditing policies can be applied with a single click to accelerate time to value.

The rest of this chapter explains how to manage audit data volume and build a centralized repository that delivers the analysis, reporting, and alerting needed for security and compliance.

With Oracle Data Safe, audit records are centralized from Oracle AI Databases, including SQL Firewall violations (see Chapter Eight). After audit data is collected, it can be safely deleted from the source databases to reclaim space and reduce overhead. In Figure 12-8 below, a fleet wide retention policy is configured in Data Safe. Audit records can be retained for up to seven years, up to one year online and up to six years in archive storage, helping meet compliance requirements without overloading databases.

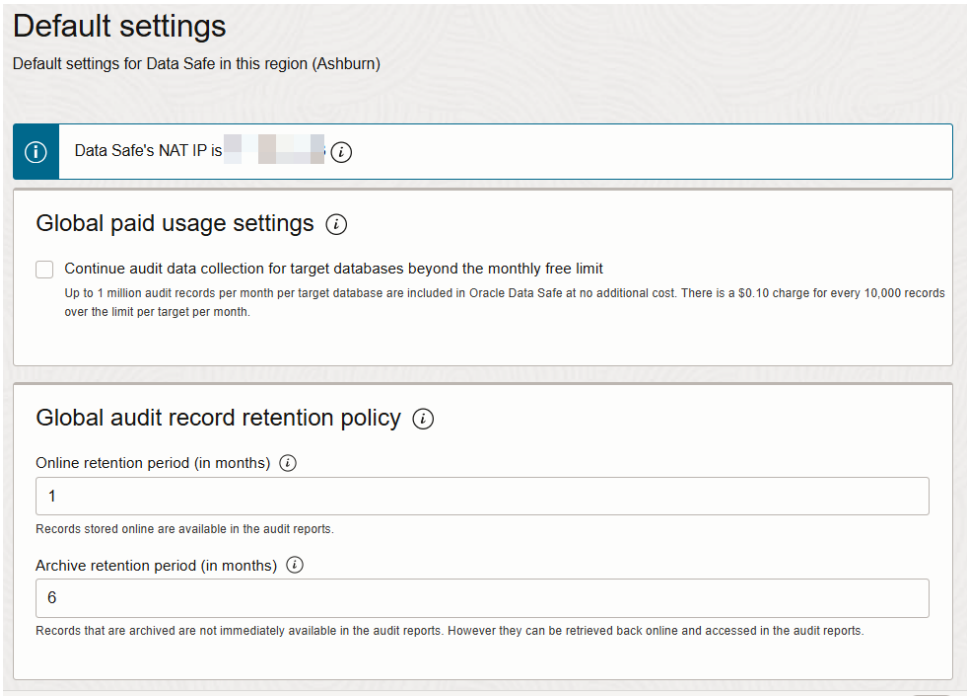


Figure 12-8: Data Safe audit data global retention settings

As shown in the following figure, the fleet setting can be overridden for individual databases that require different retention periods.

Copyright © 2026, Oracle and/or its affiliates

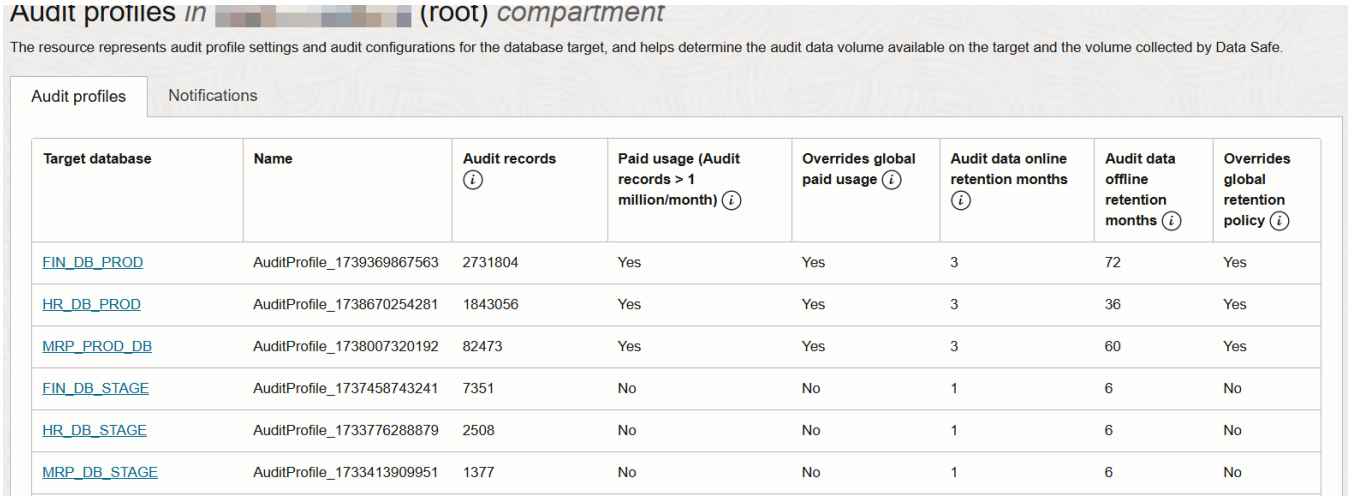


Figure 12-9: Data Safe individual database retention settings

Data Safe provides interactive reports that help track specific users, objects, or actions. It also supports forensic investigations and delivers summary reports for compliance and activity reporting. Reports can be scheduled, downloaded as PDFs, and run on recurring schedules to document periodic compliance reviews without manual effort. To stay ahead of risk, teams can choose predefined alert policies that flag unusual activity that may indicate compromise or create custom alert policies tailored to the environment.

Data Safe supports Oracle AI Database across deployments in Oracle Cloud Infrastructure, Cloud@Customer, on-premises, and other clouds such as AWS or Azure.

After Data Safe provisions and enables audit policies for a target database, it collects audit records for the activities being audited. To keep source databases from retaining too much audit data, the optional auto-purge feature (disabled by default) can be enabled to delete audit data from the target after collection.

Audit insights

Audit Insights provide a fast way to surface patterns and outliers across the fleet, helping teams focus on investigations, tune policies, and report with confidence. With Audit Insights, organizations can do the following:

- Spot spikes in Data Definition Language (DDL) and Data Manipulation Language (DML) activity
- Track trends in failed logins and user or entitlement changes
- See which databases, schemas, and objects contribute most to audit volume
- Find the audit policies that generate the most records
- Identify client connections and database users that generate the most audit volume
- Verify that audit policies capture the intended activity on target schemas and objects
- Filter results by time period and target scope to keep analysis focused on investigations

Copyright © 2026, Oracle and/or its affiliates

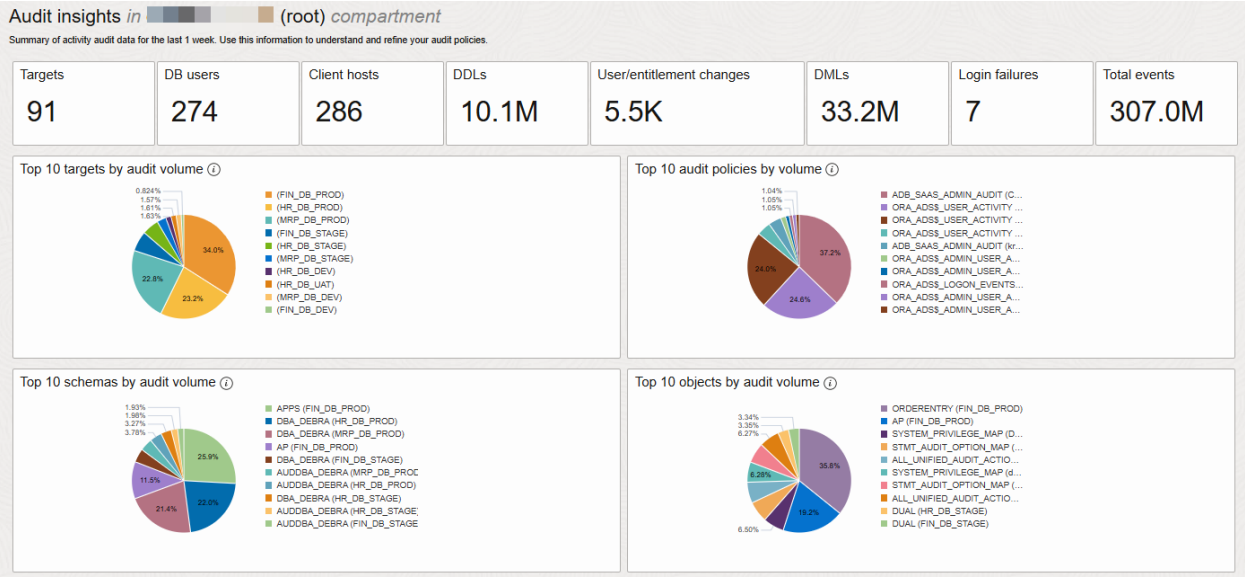


Figure 12-10: Data Safe audit insights

Audit Insights metrics and charts can be used to examine audit record volume and quickly spot outliers. By analyzing the top contributors by audit volume, including databases, schemas, objects, users, client connections, and policies, teams can pinpoint where auditing should be tightened or streamlined. Adjusting noisy or incomplete policies reduces overhead, keeps attention on high value activity, and strengthens the overall security of target databases.

Activity auditing dashboard & reports

The activity auditing dashboard presents charts and tables that summarize audit events from the past week across all target databases. Filters can be used to adjust the time range, focus on all databases within a specific compartment, or drill into a single target database for deeper analysis.

Audit records should be centralized in the Data Safe repository to support confident analysis and reporting. Once data is in the repository, interactive reports are available for user activity tracking and forensic investigation, along with summary reports for routine compliance and operational reporting.

Data Safe includes a library of predefined audit reports that surface detailed database activity, including the following:

- Privileged user activity
- Login activity
- Schema changes
- Entitlement changes
- Access to sensitive data

These reports help teams quickly spot issues, investigate incidents, and document oversight. See Table 12-1 for examples of predefined audit reports available in Data Safe.

All activity	**All audit activities**
Admin activity	Database activities of admin users (as identified in Data Safe's user assessment feature)
User/entitlement changes	User creation and deletion, privilege and role changes

Copyright © 2026, Oracle and/or its affiliates

Audit policy changes	All changes in audit policies
Login activity	Database login attempts
Data access	Database query operations
Data modification	Data modification activities (DMLs)
Database schema changes	Database schema changes (DDLs)
Common user activity	Database activities of common users in pluggable databases
Database errors	Errors reported in the database for audited activities
Data extraction	Data Pump and RMAN activities in the database
Sensitive data activity	Database activity on sensitive objects (as identified in Data Safe's data discovery feature)

Table 12-1: Data Safe audit reports

Filters can be set and report columns can be adjusted to meet specific needs. To reuse a report with these customizations, the modified report can be saved as a custom audit report. Predefined or custom reports can also be scheduled to run at specific times, with results available for download as PDF or XLS files.

Audit retention

Audit retention plays a central role in meeting organizational policies and regulatory requirements, which commonly require audit records to be preserved for a defined period, often several years. This ensures compliance while maintaining a dependable history for future investigations, audits, or reviews.

Data Safe uses a two-tier repository model: online storage for timely reporting and analysis, and offline archive storage for long term retention. Retention settings can be applied globally across all databases or adjusted for specific databases that require a different policy. When set, individual database retention settings override global settings.

Online retention ranges from one month to 12 months. Offline retention ranges from zero months to 72 months. Together, these settings allow audit records to be retained for up to 84 months (seven years). For retention beyond seven years, a request can be submitted through Oracle Support.

Data Safe automatically transitions audit data from online to offline storage when the online retention window ends and permanently deletes the data once the offline retention period expires.

When older records are needed, archives can be retrieved from offline storage into the Data Safe online repository for analysis and reporting. Retrieved archives remain available for up to one month before returning to the archive automatically, and they can be released back to the archive manually at any time.

Copyright © 2026, Oracle and/or its affiliates

Alerts

Alerts can be used on target databases to monitor high priority activities and notify the right stakeholders when events occur. Data Safe generates an alert when a specified audit event occurs on a target database, according to the alert policies that have been enabled. Teams can enable multiple preconfigured alert policies, or create custom policies tailored to their environment. See Table 12-2 for the list of preconfigured alerts.

Audit policy	Severity level
Failed logins by admin users	Critical
Profile changes	Critical
Audit policy changes	High
Database parameter changes	High
Database schema changes	Medium
User creation/modification	Medium
User entitlement changes	Medium

Table 12-2: Alert policies in Oracle Data Safe

Create custom alerts to scan incoming audit data and apply filters that detect specific conditions. For example, as shown in Figure 12-11, an alert can be triggered when an audit record involves the EMPLOYEES table and the actor is not the application account.

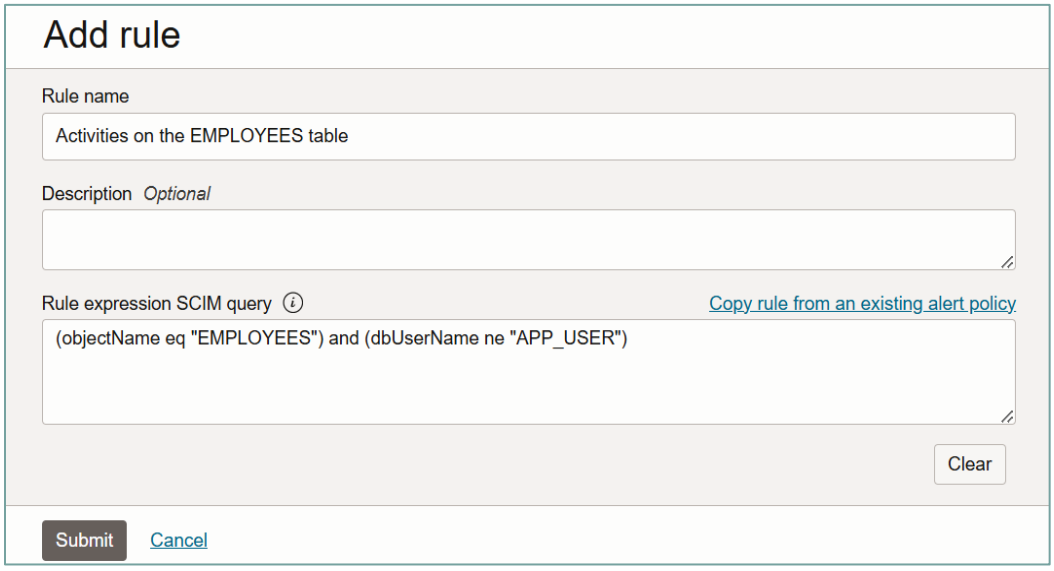


Figure 12-11: Data Safe example of a custom alert policy

Use System for Cross-domain Identity Management (SCIM) style filters to define alert conditions and report filters with a consistent, flexible query syntax.

Data Safe also provides a built-in alert dashboard for viewing and managing alerts. For notifications, Data Safe alert events can be connected to Oracle Cloud Infrastructure (OCI) Events and Notifications to route messages to selected channels (for example, email).

Copyright © 2026, Oracle and/or its affiliates

Auditing with Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall (AVDF) is a scalable, flexible solution for database auditing and network activity monitoring that helps centralize audit data and surface risk sooner. It consolidates audit records from databases, operating systems, directories, file systems, and applications into a single repository for analysis, alerting, and reporting. AVDF is deployed as a full stack software appliance on dedicated hardware or on virtual machines.

AVDF supports the most common database types, including Oracle AI Database, MySQL, Microsoft SQL Server, SAP Sybase, IBM Db2 LUW, PostgreSQL, and MongoDB. It can also ingest almost any standard format audit source using audit collection plugins that read from database tables and CSV, XML, or JSON files. For targets that are not accessible through the quick collector framework, the included software development kit (SDK) can be used to build custom collectors.

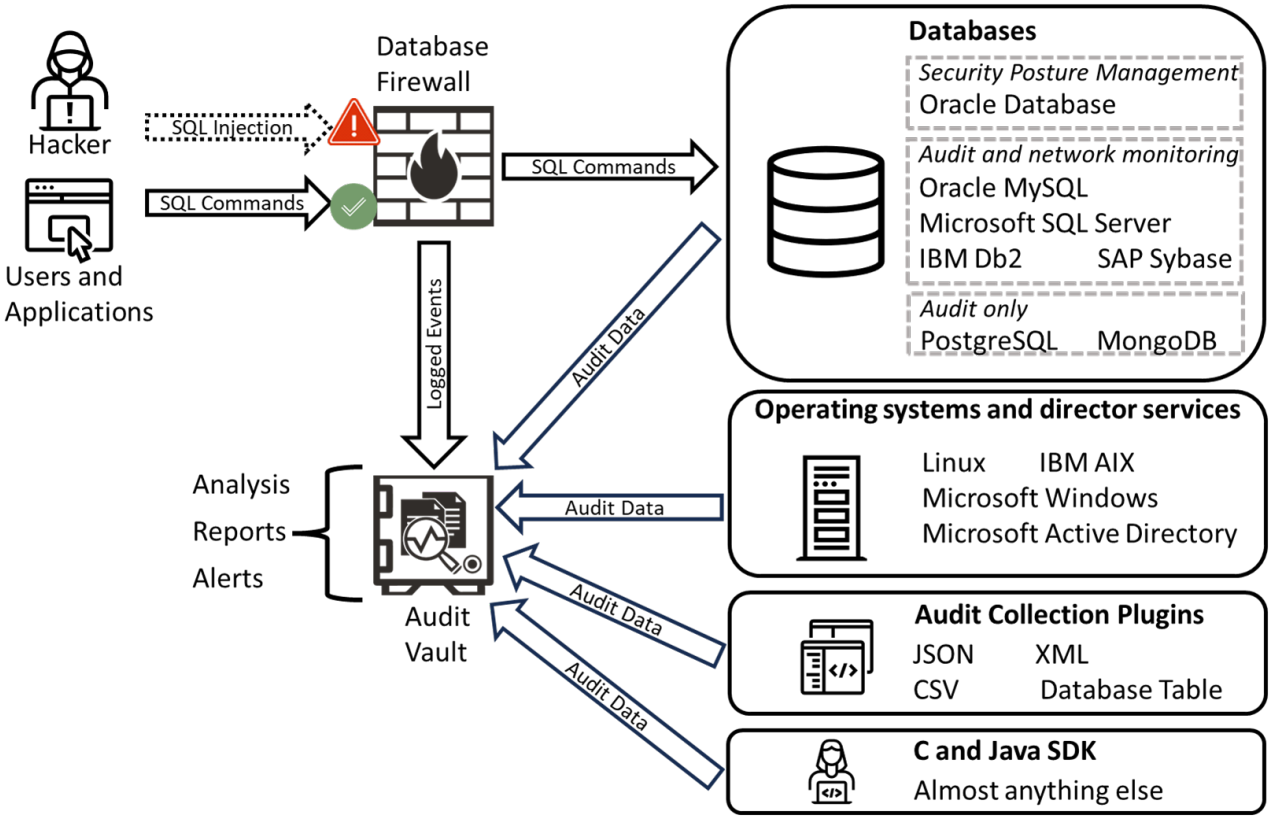


Figure 12-12: Functional architecture showing Oracle Audit Vault and Database Firewall data flows

External monitoring

Native auditing captures only the activity that is explicitly defined in audit policy. That precision keeps noise down, but it can also miss first time clients or unexpected SQL. External monitoring addresses that blind spot. AI agents and MCP servers often introduce novel SQL, either by generating queries that were not anticipated in existing policies or by reaching beyond their intended table scope. Database Firewall captures this traffic even when no audit policy exists for it, creating an added layer of visibility for agent driven activity that does not match established patterns.

Database Firewall (see Chapter Eight) can log traffic that falls outside the policy allowlist, including new clients, unseen SQL, and other outliers. Although these records may not contain full session context, they are valuable for anomaly detection and efficient triage.

Copyright © 2026, Oracle and/or its affiliates

Activity data collection

Oracle Audit Vault Server collects audit data from databases, operating systems, and application tables or files, and it ingests events from Database Firewall and policy violations from SQL Firewall. It then maps all collected data to a normalized schema and stores it in the Audit Vault repository, a specially configured Oracle AI Database. By normalizing the data, AVDF enables a unified reporting view across heterogeneous sources, from databases and operating systems to applications and directories.

For Oracle and Microsoft SQL Server, AVDF can capture before and after values, preserving a clear history of who changed data and the previous and updated values, which supports data governance. For Oracle AI Database, AVDF also evaluates security configuration, including drift detection, tracks user entitlements, and monitors stored procedure changes, as discussed in Chapter Three. When teams are unsure what to audit, AVDF offers recommended audit policies for Oracle AI Database that can be provisioned in a few clicks.

The Audit Vault Agents collect audit data from target audit trails, including databases and operating systems, and forward it to Audit Vault Server. Audit Vault Agents poll targets on a schedule, collecting only records added since the last run and sending only the deltas. One agent can collect from multiple target types. AVDF also supports agentless collection for Oracle AI Database and Microsoft SQL Server, pulling audit data directly from targets without installing an agent on target host machines.

Audit reports

Audit Vault offers dozens of out of the box reports that illuminate database activity, privileged user activity, schema changes, entitlement changes, stored procedure modifications, before and after data value changes, and login failures. Predefined compliance reports are also available for GDPR, PCI-DSS, SOX, HIPAA, Gramm-Leach-Bliley Act (GLBA), DPA, and IRS Publication 1075.

AVDF reports provide detailed insight into database activity events collected from both audit data and the Database Firewall (see Chapter Eight for more details on the Database Firewall). Coverage includes all audit and network activity, privileged user activity, schema changes, entitlement changes, stored procedure modifications, before and after data value changes, sensitive data access, and login failures.

Copyright © 2026, Oracle and/or its affiliates

Error! Reference source not found. summarizes the different types of reports in AVDF.

Report type	Description
Activity	These reports track database access activity, including audited SQL statements, application access, and user logins. Typical activities include: • All activity • Data modifications • Before and after values of modified data • DDL activity • Failed logins
Alerts	Alert reports display the raised alerts and their status.
Stored Procedure Audit	This report tracks stored procedure changes, including create, replace, and drop.
Oracle AI Database Firewall	These reports detail SQL traffic monitored by Database Firewall, including statements that generated warnings or were blocked by the database firewall policy.
User Entitlements	These reports show user privilege information for Oracle Databases. AVDF supports establishing a baseline and reporting drift from it, making it easier to monitor and investigate changes.
User Correlation	For Oracle Database targets on Linux, these reports correlate database events with the original Linux OS user. This is especially useful when users log in with named accounts and then switch to special purpose users, such as the database owner, usually ORACLE, using *su* or *sudo*.
Database Vault Activity	The Database Vault Activity report shows Database Vault events, which capture policy or rule violations and unauthorized access attempts.
Anomalous Activity	These reports pertain to new or previously dormant users accessing the system, or users accessing the system from IP addresses not seen before.
Assessment Reports	These reports bring together security assessment data from Oracle AI Databases and provide recommendations designed to strengthen the security of the Oracle AI Database system. They also include drift reporting against an established security baseline.

Table 12-3: Built-in reports in AVDF

Auditors can interact with AVDF reports in a web interface and export results as PDF or XLS files. AVDF also supports scheduling reports to run automatically and distributing them by email. Because the AVDF repository is built on an Oracle AI Database and exposes an open, documented schema, organizations can query it using their preferred reporting tools and integrate it with platforms such as Splunk or their SIEM.

Audit insights

The Audit Insights dashboard offers a fast, bird's eye view of activity captured from auditing and network-based events across one or many databases. It is designed to start with a high-level overview and then support deeper analysis through clickable charts and tables. Events can be explored by context, such as database target, user, privileged user, event type, and more, to sharpen investigations and streamline reporting as shown below.

Copyright © 2026, Oracle and/or its affiliates

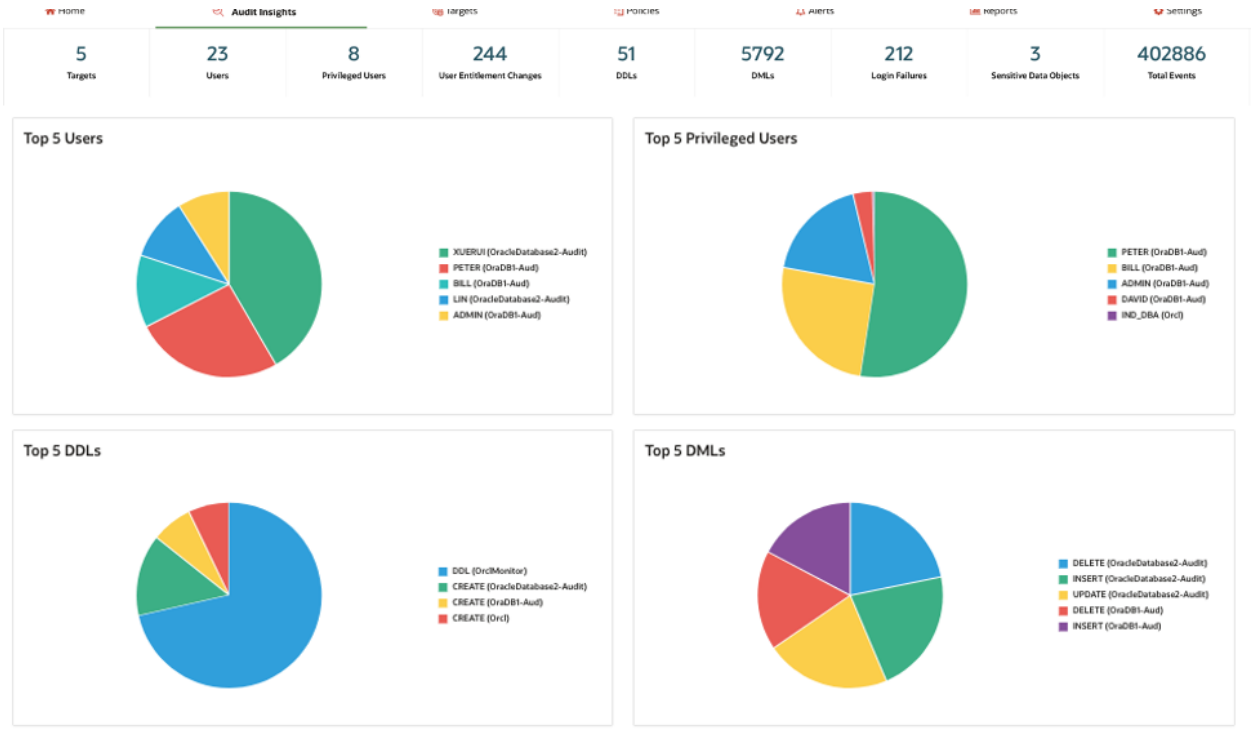


Figure 12-13: AVDF Audit Insights

Before/After data value changes

A key strength of AVDF is its ability to follow data changes end to end, making it much easier to show who changed what and when. Each time an update occurs, AVDF records both the original value (before) and the new value (after), along with the actor and timestamp.

That level of detail can be a real advantage for organizations aiming to meet strict data governance expectations, especially in heavily regulated sectors like healthcare and financial services. In fact, some regulators, including India's Ministry of Corporate Affairs, require before and after tracking for financial data. Auditors routinely lean on these records to trace the lifecycle of specific attribute changes and to confirm that the right controls are in place and working properly.

Before and after value tracking is available for Oracle and Microsoft SQL Server databases.

AVDF special support for Oracle AI Databases

Audit collection works across many sources, and network event collection covers most major databases. AVDF also includes capabilities designed specifically for Oracle AI Database as part of its database security posture management (DSPM) features.

Provisioning audit policies

Audit scope should align with an organization's policies and risk profile. In practice, this typically includes privileged activity, security relevant events such as user creation and privilege or role grants, changes to database structures, and access to sensitive data, especially access that occurs outside the application.

Audit policy design requires balance: capture enough detail to satisfy regulations and support investigations, while keeping policies targeted to limit noise, performance impact, and storage costs. AVDF supports this by offering a baseline set of recommended policies for Oracle AI Database that can be enabled with a single click. It also provides preconfigured compliance policies aligned to standards such as CIS and STIG.

Copyright © 2026, Oracle and/or its affiliates

Monitor database security configuration and detect configuration drift

As discussed in Chapter Three, AVDF assesses database security and detects configuration drift. Fleet level reporting helps teams quickly spot configuration issues and decide which databases should be addressed first.

Chapter Three also covered how AVDF monitors user entitlements and tracks drift in privilege and role assignments. AVDF reports list each user's privileges and roles and can compare results to a baseline or an earlier snapshot to highlight what changed. This supports a certify once, then monitor approach that reduces manual effort and lowers the risk of human error in periodic reviews.

Drift detection for stored procedures

Stored procedures often encapsulate sensitive business logic and data access, which makes them an attractive target for attackers seeking to introduce backdoors. AVDF periodically checks Oracle AI Databases for new or modified stored procedures and generates reports for review, helping teams spot potential tampering quickly when compromise is suspected.

Sensitive data discovery

AVDF helps organizations find and classify sensitive data so they can see where risk is concentrated and quantify how much sensitive data they have. As discussed in Chapter Four, it can discover sensitive columns, label them by data type (for example, personal, financial, or health), and report volumes by database, schema, table, and column.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Get started with Oracle Data Safe fundamentals → Lab 6: Audit database activity
- Integrate Oracle Data Safe with Applications and Services → Labs 3, 4 and 5
- DB Security – Audit Vault and DB Firewall

Copyright © 2026, Oracle and/or its affiliates

Summary

Database auditing is a practical control that only delivers value when teams actively review the data and turn findings into action. It is the difference between guessing and knowing. Audit records become truly authoritative only when someone examines them, mines them for meaning, and uses them to drive decisions.

Auditing is worth doing because it helps organizations catch suspicious behavior earlier, build a defensible timeline during investigations, discourage risky actions, support non-repudiation, demonstrate compliance, keep a close watch on privileged accounts, troubleshoot issues such as repeated failed logins, and reconstruct prior states after an incident. A newer factor makes this even more important: AI agents and MCP-driven pipelines can generate floods of unexpected SQL, often through service accounts. If that activity is not captured and clearly separated, accountability can quickly blur.

On the technical side, standardize on Oracle Database unified auditing, which is the only supported framework starting with Oracle AI Database 26ai. Legacy auditing is desupported; in upgraded 26ai environments, traditional settings are limited to deletion. To move quickly, start with predefined unified audit policies such as ORA_SECURECONFIG and ORA_LOGON_FAILURES. If you have legacy customizations, convert those settings into custom unified audit policies, and then review and enhance the converted output before enabling it. Keep the design grounded: do not duplicate mandatory audits, prioritize event types that genuinely matter, and use conditions to prevent normal application noise from flooding the audit trail. For example, audit activity on hr.employees while filtering approved application hosts, and evaluate conditions once per session to keep overhead down.

Operationally, plan for centralized audit management. Scattered trails across a fleet quickly become difficult to report on, retain, and investigate. Oracle Data Safe can centralize audit records, support retention with online and offline tiers for up to seven years, optionally auto-purge on targets after collection, and provide Audit Insights alongside predefined reports and configurable alerts. Those alerts can use SCIM-style filters for custom conditions, and notifications can be routed through OCI Events and Notifications. For a broader approach, AVDF centralizes both audit and network activity monitoring via a software appliance deployment model. It maps diverse sources into a normalized repository schema, provides dozens of out-of-the-box and compliance reports, and adds Database Firewall for deeper visibility. The firewall captures traffic that can slip past native audit policy scope, including first-time clients and the unexpected SQL that may accompany agent-driven activity.

With the audit foundation in place, the narrative naturally turns from visibility to restraint. The next chapter picks up where auditing leaves off by confronting how agentic workflows, vibe coding shortcuts, and shadow AI can reshape risk, especially when automation starts acting at machine speed with valid credentials. The proper approach is toward a layered, database-centric strategy that tightens the “who” and the “how” through least-privilege and runtime guardrails, keeps data unreadable by default through encryption and external key management, and constrains what AI can see through controls such as Database Vault, VPD, OLS, and RAS, backed by monitoring and blocking capabilities like AVDF and SQL Firewall. For teams that want to move fast without losing the plot, the next chapter is the map from bright ideas to hardened practice, turning AI ambition into governed execution with Oracle Database security controls as the steady hand on the wheel.

Copyright © 2026, Oracle and/or its affiliates

Chapter 13

Securing Agentic AI

Security risks in the age of AI

Artificial intelligence is transforming industries, driving innovation and reshaping day-to-day business operations. What differentiates this moment is not only AI's ability to generate text or code, but its growing ability to act by planning work, executing tasks, and interacting with enterprise systems. As Agentic AI systems expand across enterprises, they bring major productivity gains while extending the security perimeter in ways many organizations are still working to map, measure, and manage.

AI-generated coding accelerates delivery, but speed can create security risk if it is not carefully managed. Quickly produced code can introduce vulnerabilities that are hard to detect, reproduce, or trace, especially when developers accept suggestions without full validation. In parallel, AI assistants embedded in development pipelines, ticketing systems, and deployment tooling can introduce new opportunities for abuse, including prompt injection, insecure tool use, accidental exposure of secrets, and the quiet spread of flawed patterns across multiple services. The shift is clear: the AI revolution is not just another tool in the stack. It changes the pace, scale, and shape of the security problem.

Risk increases when AI goes beyond generating code and starts working directly with enterprise data and systems. AI that can access data, process it, and make changes is powerful, but it can be risky without proper security. Autonomy without guardrails can lead to data overreach, unintended disclosure, or destructive actions, particularly when permissions are overly broad for convenience, access is not well segmented, or auditability and policy enforcement fall behind the technology's capabilities. A sustainable approach pairs autonomy with strong data safeguards so innovation remains under control.

Enterprise security strategy therefore needs to evolve from protecting static applications to governing dynamic, AI mediated actions. Organizations benefit from security built for AI-era workflows, including least-privilege data access, strong identity controls, reliable auditing, and governance that can keep pace with automation and orchestration.

Organizations can tap into AI's momentum while keeping the business protected with proven, enterprise-grade database security. With Oracle AI Database 26ai, orchestration layers like MCP servers, and Private Agent Factory, a no-code platform that lets teams build and deploy AI agents entirely within their own environment, AI can move fast while governance stays firmly in place.

Modern risks: vibe coding and shadow AI

The era of "move fast and break things" has shifted into something far more insidious: moving fast and accidentally breaking things without even realizing what has been built. This is the backdrop for "vibe coding," a growing practice in which developers lean on AI assistants to generate entire functions, classes, or even application architectures from natural language prompts, rather than from deep technical understanding.

While this democratization of coding can unlock remarkable productivity gains, it also creates a serious blind spot: code that appears to work flawlessly while its underlying mechanics remain opaque to the humans responsible for operating it.

The security implications are stark. Unvetted AI-generated code can introduce injection vulnerabilities, expose API keys through improper error handling, implement flawed authentication logic, or inadvertently create data exfiltration pathways. When AI-generated solutions touch sensitive customer data, financial transactions, or healthcare records, organizations can end up trusting black boxes built by other black boxes, a recursive uncertainty that should concern any organization committed to strong data governance.

This vibe coding trend also feeds directly into the broader challenge of Shadow AI: the unauthorized deployment of AI tools and agents across enterprises by well-meaning employees pursuing efficiency gains. Unlike traditional shadow IT, where workers might install unapproved software, Shadow AI carries a fundamentally different risk profile. These tools do not merely process data, they make autonomous decisions about it, learn from it, and may leak it across organizational boundaries.

Copyright © 2026, Oracle and/or its affiliates

As a result, the security risks multiply: prompt injection attacks that manipulate AI behavior, training data poisoning from uploaded proprietary information, unauthorized data transmission to third-party AI services without encryption or audit trails, and the complete bypass of data loss prevention systems. A marketing analyst using an unapproved AI assistant to "quickly clean up" a customer database could expose personally identifiable information to external servers. A developer allowing an AI agent to refactor authentication code could introduce privilege escalation vulnerabilities. A financial team feeding proprietary models into a chatbot for "analysis" could effectively hand competitive intelligence directly to competitors who use the same AI service. The question is not whether an organization has Shadow AI; it is how much damage remains hidden from view.

Vibe coding

Software development has stepped into a new era, one where momentum and creativity can sometimes race ahead of safety. Vibe coding, a term introduced by OpenAI co-founder Andrej Karpathy in 2025, is changing how applications get built. The flow is simple: describe what you want to an AI model, let it generate the code, run a quick test, and if it works, keep moving. The catch, as mentioned above, is that teams do not always stop to review every line of code, every database query, or every potential security gap. It feels fast and efficient. But when deep reviews are skipped, the risk does not disappear, it just goes quiet until a high stakes moment forces it into the open.

Consider what can happen when essential safeguards are missed:

Privacy breach: Picture an AI-powered dashboard wired directly into your Oracle Database. It runs smoothly, the charts look great, and then the AI summary helpfully prints customer names, credit card numbers, and personal contacts because there is no protection layer. Suddenly, sensitive data is on screen in front of people who were never supposed to see it. This is not malicious; it is a lack of protection. Oracle Database Vault, Fine-Grained Access Control (FGAC), Row-Level Security (RLS), Virtual Private Database (VPD), Oracle Label Security (OLS), or Real Application Security (RAS) could have filtered and limited access to confidential data, keeping private information secure even before it reached the LLM.

Accidental deletion: In a quick demo, it is surprisingly easy to give an AI agent broader permissions than it truly needs. Someone asks, "Can you remove John Smith's record?" and the AI takes the instruction at face value, deleting every record with that name. In this example, three John Smiths and their entire histories vanish. With a read-only account, this could not have happened. There is no malicious intent here. It is simply what happens when proper controls are missing. The very Oracle Database security features mentioned above are built to prevent exactly these kinds of mistakes.

Copyright © 2026, Oracle and/or its affiliates

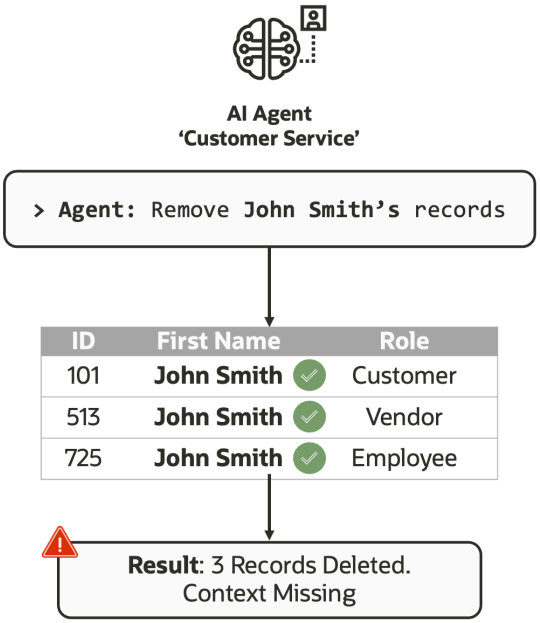


Figure 13-1: The danger of valid access at machine speed

Vibe-coded applications also bring a higher risk of SQL injection. When AI-generated SQL reaches your database without solid validation, attackers get a familiar and dangerous foothold. Oracle SQL Firewall helps by blocking unauthorized SQL statements before they execute.

Shadow AI

Vibe coding can accelerate development, but it can also encourage shortcuts that weaken security. A larger threat appears when AI tools operate entirely outside IT's control. Picture a marketing analyst, not a developer, assembling a dashboard with a free AI tool or an unauthorized MCP server. Or consider a browser extension quietly establishing a connection to a corporate database. This is Shadow AI: the use of artificial intelligence tools without security or IT oversight.

Shadow AI pushes Shadow IT into more complex and higher risk territory. Shadow IT typically involves employees installing unapproved software. Shadow AI tools, however, do not merely store information. They process data, generate insights, and make decisions. That shift introduces fresh compliance and privacy exposure because unsanctioned AI solutions can now build apps, automate work, and directly access sensitive data. The scale is staggering. It's predicted that by 2027, more than half of a typical organization's employees will acquire, modify or create technology outside IT's visibility. The driver is straightforward: AI tools are easier to access than ever.

Example: An analyst, rushing to finish a report, links an external language model to a corporate database through an unauthorized MCP server using real credentials. Sensitive customer records and transactions are rapidly analyzed and stored beyond organizational control. That single shortcut can trigger compliance violations and serious data residency risks.

Undocumented, unmonitored access steadily undermines data protection. At the same time, attempting to ban every tool is not realistic. Restrictions can be difficult to enforce, can hurt morale, and often invite even more workarounds.

Securing data at the source

The core lesson is clear: even as AI tools and access patterns evolve, the data point of origin remains the strongest, and sometimes only, line of defense. Oracle supports data protection with solutions such as Database Vault, FGAC, RLS, VPD, OLS, and RAS. These security controls help safeguard information regardless of the tool, user, or workaround attempting to connect.

Copyright © 2026, Oracle and/or its affiliates

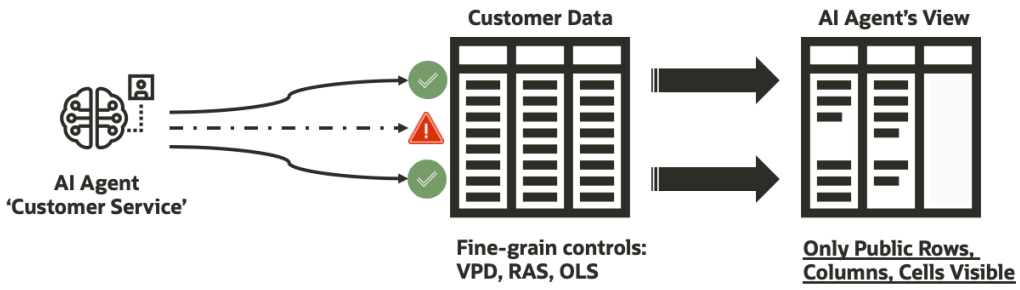


Figure 13-2: AI Agents accessing only authorized rows, columns, and cells

While it is not possible to anticipate every new AI trend, it is possible to control who can see what inside an Oracle database. Securing data at the source remains the strongest defense. Even when every browser extension or SaaS signup cannot be fully monitored, database controls can still limit the data those pathways can reach. Database level security serves as a practical shield against threats that may arrive without warning.

Building a future-proof AI security strategy

Unchecked AI adoption brings real risks, but stepping away from AI isn't realistic when the speed, automation, and competitive upside are compelling. The smarter path is to build a future-proof AI security strategy where the Oracle Database becomes the control plane for what AI can see, do, and change no matter how fast your developers ship or how many agents you plug in.

The essential question is simple: how can AI and automation be embraced without triggering data loss or compliance nightmares? A durable strategy pairs proven Oracle Database controls with modern, AI-aware enforcement so you can move quickly without turning compliance into a constant fire drill.

How to tame AI-induced chaos without killing innovation

- Policies and guardrails are a necessary start, but they're only paper shields unless the database itself enforces them at runtime. Oracle's advantage is that you can put security at the data source, so every access pattern (apps, humans, scripts, and AI agents) runs through the same hardened gate.
- Build controls where the data lives: database-native enforcement reduces "workarounds" and brittle middleware dependencies.
- Design for prevention and response: assume misconfigurations, leaked credentials, prompt-injected SQL attempts, and over-permissioned service accounts will happen, then engineer to contain blast radius.
- Protect the protectors: keys, policies, and audit trails must be isolated and tamper-resistant, not sitting in the same place an attacker-or a runaway agent-can reach.
- Context aware policies help ensure only the right users, whether human or AI, can see the data they are intended to access.

AI is changing the rules across the board. Oracle helps teams keep innovating, while staying secure. That is the difference with Oracle AI Database 26ai.

Below are the key Oracle Database Security controls covered throughout this book, along with how they work together to reduce risk from agentic AI, vibe coding, and shadow AI.

Secure the 'who' and the 'how'

Start by making identity and privilege boring again: tightly scoped, measurable, and continuously improved. Privilege Analysis helps you find and remove unused or excessive grants, so accounts adhere to least-privilege-especially

Copyright © 2026, Oracle and/or its affiliates

important when MCP-driven AI or automation needs database access and you want to prevent privilege creep and escalation.

Then lock down how access happens, not just who has a password. Controls like SQL Firewall in Oracle AI Database 26ai are designed to inspect incoming SQL and restrict database activity to explicitly authorized SQL (and can log or block violations), which helps contain risks such as SQL injection and compromised accounts-even when the request originates from automated systems.

Keep data unreadable by default and usable by design

Future-proofing means assuming data files, backups, or storage layers may be exposed, and planning so exposure doesn't equal disclosure. Oracle Transparent Data Encryption (TDE) encrypts sensitive data in tables and tablespaces and supports encrypting backups, while transparently decrypting for authorized access.

Keys are the crown jewels, so they need their own castle. TDE's master encryption key is stored in an external keystore (such as an Oracle wallet or Oracle Key Vault), keeping keys out of the database and reducing the chance that backend access translates into full plaintext access.

Constrain Agents with Unmovable Guardrails

Modern AI risk often looks like "valid credentials doing the wrong thing at machine speed," so you need enforcement that watches behavior and stops bad moves early. Oracle Audit Vault and Database Firewall (AVDF) inspects SQL traffic and can allow, log, alert, substitute, or block SQL, while storing firewall events in Audit Vault Server alongside audit data for a unified activity view.

For Oracle AI Database 26ai specifically, SQL Firewall provides allowlisting-style control, inspecting SQL and connections (including local/network traffic and PL/SQL-originated statements) and enabling logging or blocking of statements outside policy, using session context signals to reduce the risk of stolen or misused service credentials.

Key Oracle Database security controls mapped to AI reality

Below is a summary of how Oracle Security and Database Security controls work together as a layered strategy. This approach helps ensure that shadow AI, vibe coding shortcuts, and agentic workflows do not quietly become pathways for data exfiltration or trigger unexpected compliance issues.

Privilege analysis – Privilege Analysis helps identify and remove unused or excessive privileges, so every account aligns with the principle of least-privilege. When MCP-driven AI connects to databases, only the required rights are granted, helping prevent privilege escalation and limiting what AI can do.

Oracle Transparent Data Encryption (TDE) – TDE automatically encrypts data at rest. This means that even if someone gains backend access to data files, they cannot view raw, unencrypted information.

Oracle Key Vault (OKV) – OKV centralizes key management for encryption and credentials. This helps prevent AI systems with backend data access from reaching keystores that contain master encryption keys.

Oracle Database Vault (DBV) – Database Vault restricts access to protected data, including for privileged users and automated tools. It helps control which applications and agents can connect using a trusted path and which can interact with sensitive information, keeping AI within defined operational boundaries.

Oracle Data Redaction – Data Redaction dynamically masks sensitive data at query time. AI systems see only the fields or details required, while everything else is automatically redacted.

Oracle Audit Vault and Database Firewall (AVDF) – Audit Vault and Database Firewall monitor and audit database activity and can block unauthorized actions. AVDF can detect abnormal AI-driven queries and maintain a complete audit trail. Database Firewall and the SQL Firewall in Oracle AI Database 26ai can block malicious or injected SQL before it runs, providing runtime protection even if a misconfigured MCP server or rogue AI issues a harmful command.

Copyright © 2026, Oracle and/or its affiliates

Oracle Data Safe (OCI Cloud Service) – Oracle Data Safe, delivered as a cloud service, assesses configuration, reviews user entitlements, discovers sensitive data, and manages masking and SQL Firewall policies. It also collects and reports on auditing activity, helping keep AI-accessible datasets governed and compliant.

Data Masking and Subsetting Pack (DMS) – Before training or testing AI models, Data Masking and Subsetting Pack anonymizes and reduces datasets. It enables masking of sensitive or personally identifiable data to support safe model development and reduce privacy risk. Masking can be managed with Data Safe or DMS in Enterprise Manager.

Oracle Virtual Private Database (VPD) – VPD enforces fine-grained access control at query time. For AI agents, it ensures queries return only the rows they are entitled to view, and nothing more.

Oracle Label Security (OLS) – OLS enables data labeling by classification, such as Public, Confidential, or Restricted, and enforces access through clearance-based policies.

Oracle Real Application Security (RAS) – RAS brings user and session security into the database access layer. AI agents can be treated as application users with distinct privileges, with access controlled at the row, column, and cell level.

Copyright © 2026, Oracle and/or its affiliates

Summary

AI-enabled development and operations move at a new tempo, and that tempo can outstrip traditional review and governance. When AI is treated as a fast path to working code, teams can unknowingly accept opaque implementations that carry injection vulnerabilities, flawed authentication logic, or data exfiltration pathways. The risk profile intensifies when the AI system can directly access enterprise data and systems, because valid credentials paired with machine speed can produce rapid overreach, unintended disclosure, or destructive actions before humans can intervene.

A consultative strategy centers on governing dynamic, AI-mediated actions with controls that hold even as tools and access patterns shift. Oracle Database is framed as the control plane for enforcement, with strong identity and least-privilege as the baseline. Privilege Analysis helps remove unused or excessive grants, while SQL Firewall and AVDF provide policy enforcement and visibility over SQL activity, including the ability to log or block unauthorized statements. TDE keeps data at rest encrypted, and OKV centralizes key management, so a compromise of storage or backend access does not automatically translate into plaintext exposure.

The narrative also highlights how Shadow AI expands Shadow IT into a higher risk domain because unsanctioned AI tools can process data, generate insights, and make decisions without oversight. Examples include a marketing analyst using an unapproved assistant to clean a customer database, a developer letting an AI agent refactor authentication code, and a financial team submitting proprietary models for analysis, each creating pathways to privacy exposure, privilege escalation vulnerabilities, or competitive intelligence loss. Another example describes an analyst linking an external language model to a corporate database through an unauthorized MCP server, resulting in sensitive records and transactions analyzed and stored beyond organizational control.

The consultative advice is to resist tool-by-tool bans that invite workarounds, and instead secure data at the source using database controls such as Database Vault, FGAC, RLS, VPD, Oracle Label Security (OLS), and Real Application Security (RAS), so governance persists regardless of how AI connects.

With AI agents, dashboards, and automation increasingly reaching across environments, the next pressure point is not just what connects to Oracle Database, but where everything runs and which identity system is trusted when clouds intersect. The discussion becomes more vivid when identities live in one place, applications run in another, and data sits somewhere else, all while security teams must preserve consistent controls and auditability. Continue into “Multicloud Database Security” for a clear, practical path through token-based access with OCI IAM and Microsoft Entra ID, cloud agnostic security services such as Data Safe, and patterns for secure service to service connectivity that reduce password sprawl while keeping governance intact across OCI, AWS, Azure, GCP, and on premises footprints.

Copyright © 2026, Oracle and/or its affiliates

Chapter 14

Database Security in a Multicloud Environment

Understanding multicloud strategies

Multicloud environments are now both pervasive and essential to modern enterprise IT. They shape how organizations build resilience, manage risk, and deliver services at speed. Understanding how multicloud works, and why it is adopted, helps leaders make sharper infrastructure decisions and extract more value from cloud investments.

A multicloud strategy leverages cloud computing services from two or more cloud providers. Organizations may also encounter the term "hybrid cloud," which describes environments that blend on-premises resources with cloud services. Many enterprises combine both approaches—running some services on-premises while distributing others across different cloud providers. Although the technically precise term for this combination is "hybrid multicloud," industry practitioners typically shorten it to simply "multicloud".

Most organizations already operate in multicloud environments. While some have migrated entirely to the cloud, most large enterprises retain on-premises systems with no immediate migration plans. Organizations adopt multicloud strategies for several compelling reasons: reducing costs, leveraging best-in-class services, avoiding vendor lock-in, improving redundancy, optimizing performance, and meeting data sovereignty requirements. By combining the strengths of multiple cloud providers, organizations can customize their infrastructure to meet specific business and technical requirements while maximizing cloud computing benefits.

This chapter reviews the security and management of Oracle AI Database in a multicloud environment.

The multicloud security challenge

Multicloud architectures create scenarios where identity management resides in one cloud, applications run in another, security tools operate in a third, and data lives in yet another. Think of this as a distributed orchestra—each section plays in a different concert hall, yet all must perform in harmony. Organizations select different cloud providers based on how effectively and economically each delivers specific services.

Consider a common architecture: databases running in Oracle Cloud Infrastructure (OCI), applications hosted in Amazon Web Services (AWS), and business intelligence services operating in Microsoft Azure. Figure 14-1: High-level illustration of services in a multicloud environment illustrates this complexity—identity management in one cloud, applications in a second, data storage in a third, with security and analytics tools potentially scattered across fourth and fifth providers.

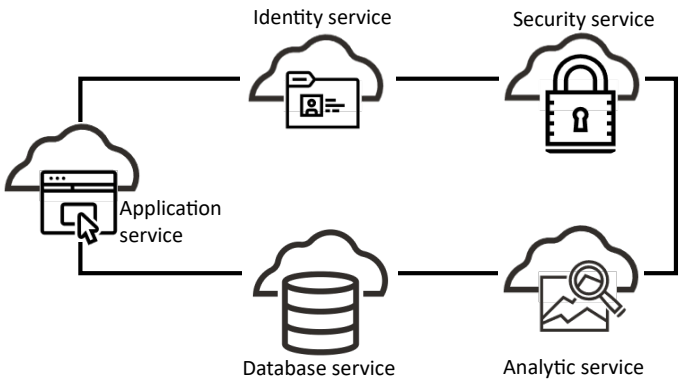


Figure 14-1: High-level illustration of services in a multicloud environment

This distributed landscape poses significant security challenges. However, Oracle AI Database security offers a distinct advantage: security features function consistently across any cloud environment, and the tools and services supporting Oracle AI Database security remain cloud-agnostic.

Copyright © 2026, Oracle and/or its affiliates

Establishing identity as the foundation

Although services may span multiple clouds, organizations typically designate one cloud provider to manage enterprise identities. Oracle AI Database includes native support for Microsoft Entra ID and Oracle Cloud Infrastructure IAM. Other cloud identity services generally integrate with Oracle AI Database through supported authentication protocols. For instance, identity service provider Okta connects to Oracle AI Database using RADIUS.

Multicloud applications and tools either accept identity tokens issued by the chosen identity provider or enable federation between the organization's identity provider and their cloud identity service. Establishing a single source of truth for identity enables centralized management of authentication and authorization for cloud users, extending the same capabilities already applied to on-premises applications and databases. Users expect single sign-on (SSO) functionality and resist managing multiple identities with different credentials across applications.

Identity integration approaches

Organizations can implement cloud identity integration through several strategies:

1. **Network extension model**: Design the cloud network as an extension of the on-premises network and continue using existing on-premises identity solutions (such as Microsoft Active Directory) for cloud resources. While this approach leverages existing skills and knowledge, it forgoes cloud identity features like forwardable tokens and cross-platform authentication.
2. **Synchronization model**: Synchronize the on-premises identity solution with cloud identity services. Users access cloud resources through the cloud identity solution and on-premises resources through the on-premises identity solution. User management typically occurs in one environment and synchronizes to the other.
3. **Migration model**: Migrate the on-premises identity solution entirely to the cloud. As part of this migration, on-premises resources connect with the cloud identity solution alongside new cloud resources.

When selecting an identity cloud service, organizations must evaluate whether users and resources will interact directly with the identity service, federate with another identity service, or federate with multiple identity services, depending on application, tool, and data locations.

Managing Oracle AI Database 26ai in multicloud

Managing Oracle AI Database 26ai security was more straightforward when all databases resided within on-premises data center boundaries. In multicloud environments, Oracle AI Databases can exist in OCI, non-OCI clouds (Azure, Google Cloud, AWS), and on-premises locations simultaneously. Managing database security across cloud boundaries should not require working with multiple tools and skill sets simply because databases run in different environments.

Consistent Database Security features

Security features familiar from on-premises Oracle AI Databases function identically in cloud deployments. These capabilities include:

- Transparent Data Encryption (TDE)
- Database Vault
- Label Security
- Data Masking and Subsetting
- Real Application Security
- Virtual Private Database

Copyright © 2026, Oracle and/or its affiliates

Note: Managed database services like Amazon RDS may restrict certain features. For example, RDS does not support Database Vault.

Cloud agnostic security services

Multicloud operations make it possible to use data security services designed to work across cloud environments. Several of these have been covered throughout this book:

- **Data Safe** – runs in Oracle Cloud Infrastructure but can be used with Oracle AI Databases anywhere, including in OCI, AWS, Azure, Google Cloud, and on-premises. Data Safe is included with all Oracle AI Database cloud services, including both Autonomous Database and Exadata Cloud Services running in Azure, AWS, and GCP.
- **Key Vault** – available in the OCI marketplace. May also be installed on-premises, in AWS, Google Cloud, or Azure. With reliable network connectivity, a Key Vault node installed in one location can service databases in other locations.
- **Audit Vault and Database Firewall** – available in the OCI marketplace. May also be installed on-premises or in AWS. With network connectivity in place, Audit Vault and Database Firewall deployed in one location can service databases in others.

Leveraging native Oracle services across clouds

Managing resources separately in each cloud environment is operationally demanding, potentially expensive, and can increase security risk as complexity grows across disparate tools and capabilities. Oracle addresses this challenge through strategic partnerships with major cloud service providers.

Oracle partners with most major cloud platforms, including AWS, Azure, and Google Cloud, offering Oracle Autonomous Database and Oracle Exadata Database Service. Regardless of where those databases run, organizations can register them with Oracle Data Safe, extending consistent security coverage across multicloud environments. This architecture places Oracle AI Database Services on OCI inside hyperscaler data centers, enabling organizations to run entire application stacks in one cloud while eliminating the network latency and costs associated with moving large data volumes across clouds.

Oracle manages private network connections between OCI child sites running on hardware installed within hyperscaler data centers and Oracle data centers in the region. This configuration allows other OCI services direct access to Oracle AI Database Services in the respective hyperscalers for seamless integration.

Enabling secure multicloud service access

Services require the ability to connect securely across clouds without the overhead of managing, storing, and rotating passwords. Creating trust relationships between clouds and using service principals to connect one cloud service to another provides high security without the administrative burden of managing credentials separately in each cloud.

Databases also need connections to external data sources like object storage to represent data in structured formats and enable database tools to analyze the data. If a database cannot connect to object storage in another cloud, organizations must transfer data between clouds to access it—adding cost and complexity.

Oracle Autonomous Database addresses this challenge by accepting identity tokens from cloud authentication services offered by Azure, Google Cloud, and Amazon Web Services to connect to resources like object storage in those clouds. This capability eliminates the need to transfer data from one cloud to another.

Multicloud service access capabilities represent one of the fastest evolving areas in database security.

Implementing multicloud Identity Management

Enterprise users value the convenience of single sign-on across their web applications, and that same smooth experience should carry over when they access services in other clouds, including Oracle AI Databases. To support

Copyright © 2026, Oracle and/or its affiliates

this, Oracle AI Database can be configured to accept tokens from Microsoft Entra ID or Oracle Cloud Infrastructure Identity and Access Management (IAM).

Oracle Cloud Infrastructure Identity and Access Management (IAM)

IAM users, OCI services, resources, and applications hosted on the OCI platform can use OCI Identity and Access Management (IAM) credentials to obtain an IAM database token that enables access to an Oracle AI Database in OCI. IAM tokens provide stronger security than basic database passwords and can take advantage of IAM multifactor authentication when a user requests a token.

For applications or services that rely on a different cloud identity provider, that provider can be federated with IAM. This allows users, applications, or services to continue authenticating with the non-OCI identity service, then use IAM database tokens to access OCI databases.

Oracle AI Database clients can accept IAM tokens through the client API, from the file system, or by retrieving the token directly from IAM. Users only need to sign in to their preferred tool, and a properly configured client will obtain the database token from IAM on the user's behalf. If the user does not have a current session token, IAM prompts the user to sign in securely.

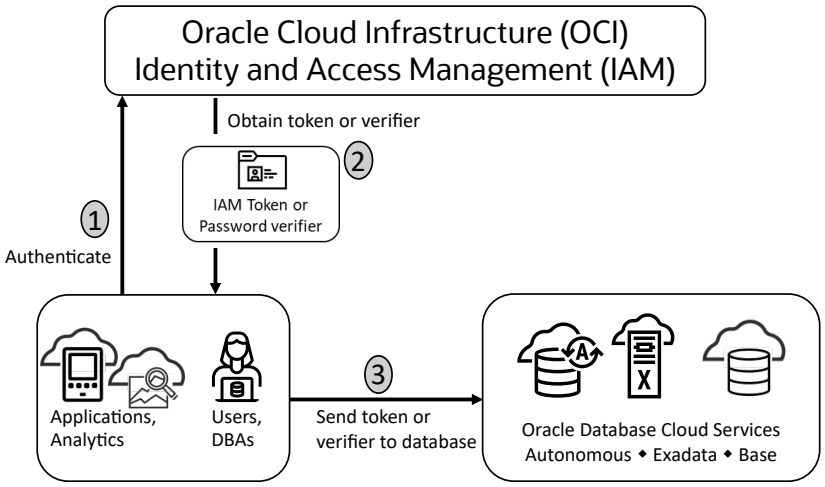


Figure 14-2: OCI database integration with IAM

Copyright © 2026, Oracle and/or its affiliates

Microsoft Entra ID integration

Microsoft Entra ID OAuth2 access tokens can be used to access Oracle AI Databases in OCI, AWS, Azure, Google Cloud, or on-premises from applications hosted on Azure or from Azure services. These services and applications can present Entra ID managed identities (service principals) to the database through an OAuth2 token, allowing the service or application to connect to the database as the service or application. Azure services and applications running in the Azure cloud can also request on behalf of (OBO) tokens for the logged in user and then access Oracle AI Database as the application user. This end-to-end authentication supports enforcement of user access controls and activity auditing within the database.

Oracle AI Database clients obtain Entra ID tokens in much the same way they obtain OCI IAM tokens. Entra ID access tokens can be passed through the client API, retrieved from the file system, or requested directly from Entra ID using the Entra ID user's session token. If no session token exists, the client can direct the user to authenticate with Entra ID.

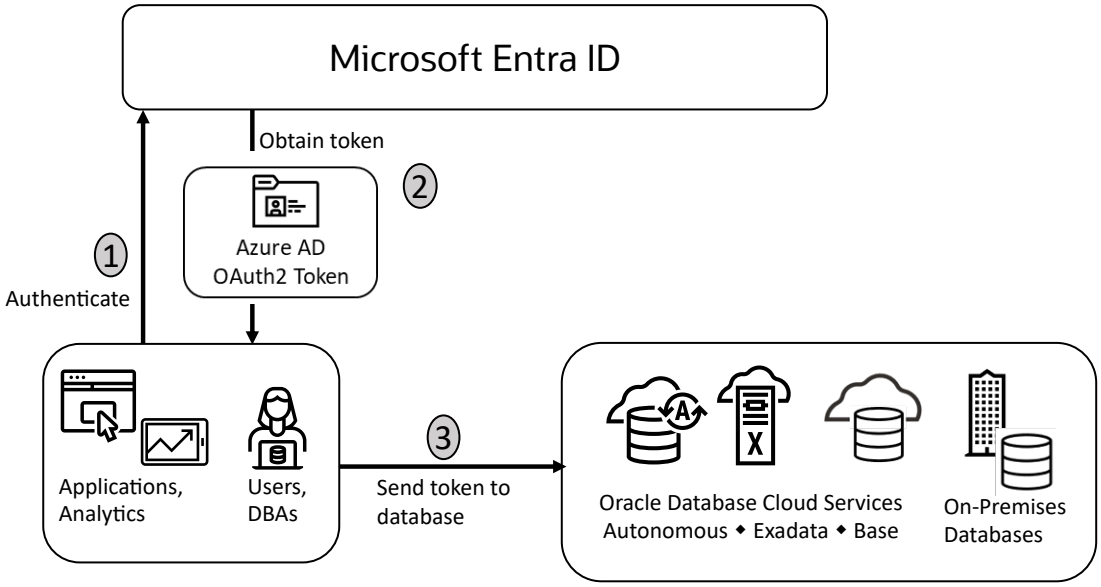


Figure 14-3: Oracle AI Database integration with Entra ID

Microsoft Power BI single sign-on

Many organizations want to use Microsoft Power BI to analyze application data stored in Oracle AI Databases in OCI, and they typically want single sign-on access to both Power BI and the data itself. Before multicloud identity integration, that often meant migrating data into another cloud environment just to make analysis possible. With Oracle's integration with Power BI, organizations can use Entra ID single sign-on and Power BI to access Oracle AI Database, whether the data is stored in OCI or on-premises.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Simplify User Management by integrating Oracle Autonomous Database and Oracle Cloud Identity

Copyright © 2026, Oracle and/or its affiliates

Summary

Multicloud environments have become the norm for enterprise IT, increasing the need for thoughtful strategies that blend the strongest capabilities of multiple providers. With Oracle Database services available across major cloud environments, organizations can place workloads close to applications and users to drive very high throughput and low latency performance while keeping operations consistent across clouds and on-premises. High availability capabilities help protect business critical systems from outages and disruptions by supporting robust continuity designs across environments. Most importantly, Oracle Database brings the strongest database security capabilities into multicloud deployments, with consistent controls and enforcement across locations, helping organizations protect sensitive data, reduce risk, and maintain trust as they scale and innovate.

Managing security across this fragmented landscape calls for a disciplined focus on identity as the foundation of trust. A common recommendation is to designate a single provider to manage enterprise identities, so users benefit from one consistent login experience across all applications. Oracle AI Database supports this approach through native integration with major identity services such as Microsoft Entra ID and OCI IAM. Whether organizations synchronize directories or migrate fully to the cloud, well-designed authentication protocols extend enterprise security standards across every part of the multicloud ecosystem.

Operational consistency becomes essential when databases run at the same time in local data centers and multiple public clouds. Oracle helps ensure that core security capabilities, including data masking and real application security, operate consistently in every environment. As a result, an administrator managing a database in Azure can rely on the same tools and techniques used to manage a database in a local facility. Strategic partnerships reinforce this model by enabling Oracle Database services to run natively inside hyperscaler data centers, reducing latency and simplifying the technical landscape for IT teams.

To enable secure communication among these varied services, modern architectures should lean on established trust relationships rather than cumbersome password management. Service principals and identity tokens allow applications to connect securely to databases and object storage across cloud boundaries. This shift in service access reduces administrative overhead and lowers the risk of credential theft. With these advanced security measures in place, organizations can expand their multicloud footprint and pursue innovation without introducing new vulnerabilities into critical data systems.

By leaning into these multicloud strategies, Oracle AI Database and its proven security capabilities become more than a hybrid deployment model. It becomes a dynamic asset that helps the enterprise innovate with confidence. The result is the best data platform embedded in any Cloud delivering the ultimate flexibility and deployment choice. The next step is to apply this agility deliberately and convert infrastructure complexity into a clear competitive advantage.

Copyright © 2026, Oracle and/or its affiliates

Chapter 15

Ransomware and Zero Trust

Ransomware: When data gets taken hostage

Ransomware has grown from an occasional disruption into a disciplined, profit-driven business model that uses data as bargaining power. Today's campaigns rarely stop at encrypting files and demanding payment; they often blend credential theft, lateral movement, privilege escalation, and "double extortion," where attackers threaten to leak sensitive data if ransoms are not paid. For enterprises, the fallout is seldom contained: downtime spills into customer services, regulatory exposure rises, and recovery becomes a tense, time-constrained call. In that landscape, the database stands out as a natural target because it holds the crown jewels and keeps the applications that run the business alive.

A practical, consultative way to build ransomware resilience is to assume compromise will happen somewhere, then ensure database security helps mitigate the risk of that compromise turning into a full-scale disaster. The focus is on limiting how far stolen credentials can reach, reducing the chances that attackers can tamper with or encrypt critical data at the source, and making detection and recovery swift, trustworthy, and auditable. Oracle Database security capabilities support this layered posture through strong identity and access controls that curb privilege misuse, encryption and key management that reduce the value of exfiltrated data, auditing and activity monitoring that highlight suspicious behavior early, and configuration hardening that blocks common routes to escalation.

This chapter presents a phased framework for ransomware resilience that prioritizes immediate, high-impact defensive actions before longer-term architectural overhauls. The guiding premise is clear: organizations should secure quick, measurable wins that blunt ransomware's most damaging tactics, specifically data destruction and data theft, before embarking on a broader Zero Trust transformation.

Phase 1: Immediate high impact defense zeroes in on two critical capabilities that can be deployed quickly and yield immediate protection:

- **Encryption and key management:** Transparent Data Encryption (TDE) paired with Oracle Key Vault (OKV) renders stolen data unintelligible, defusing double and triple extortion tactics where attackers threaten to leak exfiltrated information.
- **Immutable recovery infrastructure:** Zero Data Loss Recovery Appliance (ZDLRA) and Oracle AI Database Zero Data Loss Autonomous Recovery Service (ZRCV) deliver synchronized, immutable backups that ransomware cannot destroy, enabling recovery without paying a single ransom dollar.

Phase 2: The Zero Trust journey represents a sustained and continuous improvement. Zero Trust principles, spanning configuration monitoring, attack surface minimization, authentication hardening, granular access controls, and comprehensive activity monitoring, create overlapping layers of defense that shrink both the likelihood and blast radius of successful attacks. These initiatives demand ongoing organizational commitment and are best approached as iterative improvements rather than one-time projects.

By deploying Phase 1 encryption and recovery capabilities first, organizations immediately strip ransomware of its primary leverage while buying valuable time to execute the Zero Trust transformation. This pragmatic approach delivers measurable risk reduction up front and lays the groundwork for enduring resilience.

Understanding the ransomware imperative

Major cybersecurity agencies worldwide, including the European Union's Agency for Cybersecurity (ENISA), the United States' Cybersecurity and Infrastructure Security Agency (CISA), and the United Kingdom's National Cyber Security Centre (NCSC), consistently rate ransomware as the most immediate and disruptive threat organizations face today. Criminal groups leverage ransomware to generate revenue, while state-backed actors increasingly deploy it to

Copyright © 2026, Oracle and/or its affiliates

cause widespread damage, distract defenders from parallel operations, and funnel funds into economies operating under international sanctions.

Defining ransomware and its evolution

Ransomware is a type of cyber-attack where criminals seize control of your systems or data and demand payment to restore access or to keep stolen information from going public. In many campaigns, attackers encrypt files to block access. Others skip encryption entirely and rely purely on data theft and extortion to compel rapid payment.

The origins of ransomware stretch back decades. The first widely recognized attack surfaced in 1989, when Dr. Joseph Popp distributed malware-infected floppy disks to roughly 20,000 researchers at a World Health Organization AIDS conference. Dubbed the AIDS Trojan, it encrypted files and demanded US$189 for decryption. This predated the dark web; victims sent payments by mail to a post office box in Panama.

Ransomware activity exploded around 2015. Cryptocurrencies such as Bitcoin made ransom payments easier to route and harder to trace, while Ransomware-as-a-Service (RaaS) platforms lowered the barrier to entry, enabling a wider pool of actors to build and distribute ransomware at scale.

Modern campaigns target servers as aggressively as user workstations. Most ransomware variants both encrypt files and exfiltrate them to remote servers. Even if victims pay, there is no guarantee attackers will delete the data or refrain from selling it in underground markets. Payment also offers no assurance of successful file recovery.

AI-Powered ransomware attacks

The integration of artificial intelligence into ransomware attacks marks a significant escalation in cyber threats. AI-enhanced reconnaissance tools automatically scan vast amounts of public and leaked data to identify vulnerable systems, map database architectures, and extract credentials in hours rather than weeks. Once inside a network, AI-powered malware adapts to defensive responses in real time, modifies its own code to evade detection, and optimizes lateral movement toward high-value targets — in one documented case achieving domain dominance on a corporate network in under an hour with no human intervention.

Beyond initial compromise, AI transforms post-breach exploitation: automatically categorizing exfiltrated data, generating tailored extortion notes, and deploying chatbots to negotiate ransoms around the clock. These capabilities allow threat actors to operate with greater speed, precision, and scale across every phase of their campaigns, from crafting convincing phishing emails to executing final stage encryption.

Organizations should recognize that traditional signature-based defenses and static security rules fall short against AI-augmented threats. This reality makes the pairing of encryption (to guard against data theft) and immutable recovery (to guarantee restoration) even more vital, since these defenses hold firm regardless of how sophisticated the attack becomes.

A two-phased strategy for database ransomware defense

In an ideal world, organizations would have solutions that block ransomware with 100 percent effectiveness, especially since it has been a problem for more than 35 years. In reality, ransomware continues to evolve in ways that improve outcomes for ransomware gangs. As American football coach Vince Lombardi famously observed, “Offense only needs to succeed once. Defense must succeed every time.” Even highly prepared organizations can face ransomware attacks that cannot be blocked.

Organizations need both immediate tactical wins and long-term strategic depth. Below is an example of a practical, two-phase approach.

Phase 1: Immediate high impact defense (encryption & recovery)

Copyright © 2026, Oracle and/or its affiliates

These are comparatively low-effort, high-impact measures that directly counter the two primary database-level ransomware threats:

- **Data loss through encryption**: Ransomware encrypts data files and renders databases unusable
- **Data theft for extortion**: Attackers exfiltrate data and threaten public disclosure unless ransom is paid

Phase 1 solutions neutralize both threats: encryption makes stolen data worthless, and immutable backups ensure recovery without paying ransoms. These capabilities can often be quickly implemented and deployed, and deliver immediate, measurable risk reduction.

Phase 2: The Zero Trust journey of continuous improvement

Zero Trust is a broad architectural transformation built on the assumption that breaches are inevitable, and it removes implicit trust across the environment. Although these practices can sharply reduce the likelihood of successful attacks, they require sustained organizational commitment, cross-functional coordination, and ongoing refinement. As a result, most organizations approach Zero Trust as an enduring program rather than a finite project, steadily strengthening their security posture over months and years.

The following sections describe both phases, starting with immediate, high-impact measures that deliver quick wins, then shifting to the longer-term Zero Trust journey.

Phase 1: Immediate high impact defense

Phase 1 focuses on deploying two critical defensive capabilities that can be implemented relatively quickly and provide immediate protection against ransomware's most damaging tactics. By prioritizing encryption and recovery infrastructure, organizations reduce an attacker's primary leverage: the ability to destroy data and extort payment by threatening to leak stolen information.

These measures function as the "final line of defense," helping ensure organizational continuity when preventive controls fall short. Even if attackers bypass perimeter defenses and compromise accounts, well-implemented encryption and recovery capabilities can prevent catastrophic data loss and reduce the pressure to pay ransoms.

Neutralizing data theft and extortion through encryption

Data theft has become a hallmark of modern ransomware. Many campaigns forgo encrypting data altogether. Instead, attackers steal it, show proof, and monetize the theft by threatening public disclosure.

Ransomware gangs employ increasingly inventive monetization strategies:

- **Simple extortion**: Encrypting data and demanding payment for decryption, without threatening exposure.
- **Double extortion**: Encrypting AND exfiltrating data, then threatening to publish or leak it unless ransom is paid.
- **Triple extortion**: Layering additional pressure tactics on top of encryption and data leak threats—for example, launching DDoS attacks, directly contacting customers, employees, or partners, or employing "high-pressure publicity" where data subjects are notified to intensify pressure on the victim organization.
- **Traditional sales**: Selling stolen data on the dark web regardless of whether ransom is paid, particularly identity related information that enables fraud.

These tactics are not merely theoretical. In November 2023, the ALPHV/BlackCat ransomware group broke new ground by filing a formal complaint with the U.S. Securities and Exchange Commission (SEC) against its own victim, a publicly traded digital lending company. After exfiltrating data without encrypting any systems, ALPHV reportedly grew frustrated with the company's refusal to negotiate and escalated the pressure by reporting it to the SEC for allegedly failing to disclose the breach within four business days under the agency's new cybersecurity incident

Copyright © 2026, Oracle and/or its affiliates

disclosure rules. The group even published screenshots of the SEC complaint form and the acknowledgment it received.

Notably, the rule ALPHV cited had not yet taken effect. Even so, the episode marked a troubling new chapter in ransomware pressure tactics. Security researchers have since observed a growing pattern of gangs leveraging regulatory frameworks, including GDPR in Europe, to coerce victims into paying. This approach, often described as "compliance extortion," adds another layer of pressure for organizations already navigating the operational, legal, and reputational fallout of a breach.

Transparent Data Encryption: The primary defense

The key to mitigating data theft lies in understanding how ransomware harvests data. Most ransomware mounts "out-of-band" attacks that circumvent database session controls by directly accessing underlying data files on storage or sniffing network traffic. These attacks bypass application-level security entirely.

Transparent Data Encryption (TDE), part of Oracle Advanced Security, is the primary defense against out-of-band attacks. When attackers steal encrypted database files, they obtain nothing but useless encrypted blocks that cannot be decrypted without the encryption keys. This neutralizes double and triple extortion tactics. Attackers have nothing of value to leak.

There are two ways to utilize TDE to encrypt data:

- **Data at rest**: TDE encrypts all data stored in database files using AES-256 encryption. This is the default in Oracle AI Database 26ai and should always be enforced.
- **Data in motion**: Data traveling between clients and database servers should be encrypted using either TLS (Transport Layer Security) or Oracle Native Network Encryption (NNE), both of which are included with the Oracle Database license.

Critical requirement: Off-server key management

WARNING: Encryption becomes a liability if attackers can access both encrypted data files AND encryption keys during an attack. Storing keys on the same server as the encrypted data creates a critical vulnerability. Attackers can simply delete or corrupt the keys, permanently locking the organization out of its own encrypted data.

The requirement to store encryption keys separately from the data they protect is well established across major regulatory frameworks. NIST Special Publication 800-57, the foundational key management guidance that most regulations defer to, explicitly requires documented key management plans and separation of keys from encrypted data. PCI DSS is among the most prescriptive, directly mandating that keys be stored separately from the data they encrypt. HIPAA, aligned with NIST guidelines, calls for keys to be stored and managed independently from protected health information. Storing keys off-server is not merely a security best practice; for many organizations, it is a compliance obligation.

Oracle Key Vault (OKV) addresses this requirement directly, moving encryption keys off server and managing the complete lifecycle of TDE master encryption keys. OKV provides:

- Centralized Key Lifecycle Management: Manages TDE keys for all Oracle AI Database instances plus encryption keys for backup infrastructure like ZDLRA, delivering unified key management across the entire data protection ecosystem.
- Compliance Support: Provides audit trails for key access and automated key rotation policies to satisfy requirements under PCI DSS, HIPAA, and other regulatory frameworks.
- SSH Key Management: Centralizes storage, management, and rotation of SSH keys used for server access, eliminating unmanaged "orphaned" keys that attackers exploit for lateral movement.

Copyright © 2026, Oracle and/or its affiliates

- Recovery Clean Room Support: In ransomware recovery scenarios, OKV can be restored from a known good backup, or keys can be exported and imported into an isolated recovery environment, ensuring data remains decryptable even if the primary data center is compromised.

Futureproofing: Quantum-safe cryptography

As organizations map out their encryption strategies, they must also reckon with emerging threats from quantum computing. “Harvest now, decrypt later” attacks involve adversaries collecting encrypted data today with the intention of decrypting it once quantum computers gain the power to break current cryptographic algorithms.

Oracle AI Database 26ai addresses this risk by supporting quantum-safe cryptographic algorithms for both data at rest and in motion. The release integrates NIST approved post quantum algorithms including ML-KEM (for key exchange) and ML-DSA (for digital signing) and enables quantum safe TLS to protect data traffic between clients and database servers.

For organizations safeguarding highly sensitive or long-lived data such as financial records, healthcare information, and government data, implementing quantum-safe cryptography now forestalls future compromise of today's encrypted backups.

Eliminating data loss through immutable recovery

While encryption neutralizes data theft, organizations still face the threat of data destruction. When ransomware encrypts production database files, only two options remain: pay the ransom and hope the decryption key works, or restore from backup. The second option is viable only if backups exist and have not been destroyed by the ransomware.

The problem: Ransomware targets backups

Modern ransomware has grown fiendishly effective at hunting down and destroying traditional backups. Ransomware developers understand that eliminating recovery options forces victims to pay, so they actively seek out and corrupt backup files, delete backup catalogs, and compromise backup management systems.

Traditional backup architectures may fail because of the following reasons:

- **Reachable through standard protocols**: Traditional backup systems rely on well-known protocols such as CIFS/SMB, NFS, and SSH/SFTP that ransomware is specifically engineered to discover and exploit. Any process running with sufficient privileges can enumerate network shares, locate backup repositories, and delete or corrupt files using the same standard interfaces that legitimate backup software depends on.
- **Backups that can be modified**: Standard backup files can be deleted by any account with sufficient privileges, including compromised administrator accounts under attacker control.
- **Inconsistent recovery points**: When multiple interconnected databases are backed up at different times, restoring them produces data inconsistencies where transactions appear in one system but are missing from others.
- **Slow recovery**: Traditional restore processes can stretch across days or weeks, leaving business operations paralyzed throughout.

The solution: Immutable, synchronized backup architecture

Immutable backups consists of (read-only) backup files that cannot be deleted, altered, or corrupted, even by compromised administrator accounts wielding high level privileges. This architectural requirement has become mandatory for ransomware resilience.

Oracle delivers this capability through specialized recovery infrastructure:

Copyright © 2026, Oracle and/or its affiliates

- **Zero Data Loss Recovery Appliance (ZDLRA)**: An on-premises appliance providing real time, synchronized database protection with built in immutable architecture.
- **Oracle AI Database Zero Data Loss Autonomous Recovery Service (ZRCV)**: A cloud native autonomous recovery service in Oracle Cloud Infrastructure with managed immutability and air gap separation options.
- **ZFS Storage Appliance**: High performance storage with immutable snapshots via retention locks, suitable for snapshot based recovery strategies.

Critical feature: Synchronized recovery point objectives

When ransomware strikes multiple interconnected databases, restoring them to different points in time creates data inconsistency gaps where some transactions exist in one system but vanish from related systems. The consequences can be severe:

- Broken referential integrity (orphaned foreign keys)
- Financial discrepancies between orders, invoices, and payments
- Weeks or months of painstaking manual reconciliation to track down and repair missing transactions
- Potential regulatory violations if financial records fail to balance

ZDLRA and ZRCV solve this problem through real time synchronization. Every committed transaction is captured as part of the transaction itself, ensuring all databases can be restored to precisely the same point in time. This "zero-data loss" capability means no committed transactions are lost and recovery achieves perfect consistency across the enterprise.

Infrastructure requirement: Recovery clean rooms

Ransomware recovery should be treated as *disaster* recovery, not merely *database* recovery. Organizations need dedicated infrastructure to host recovered systems while production environments undergo malware remediation. Traditional hot, warm, and cold site strategies still apply, but cloud providers now offer enhanced options:

- **Secure Enclaves**: Oracle Cloud Infrastructure provides isolated environments separated from production by network air gaps, preventing ransomware from jumping to recovery infrastructure
- **Recovery Clean Rooms**: Highly isolated environments purpose built for data and application recovery, designed to bring up a minimally viable IT environment rapidly after an incident
- **Continuous Incremental Restore**: Cloud environments can execute continuous incremental restores in standby mode, dramatically compressing Recovery Time Objectives (RTO) from days to hours

Phase 1 Implementation roadmap

Organizations should prioritize these immediate, high impact actions:

1. **Assess encryption coverage**: Verify that Transparent Data Encryption (TDE) is enabled for all production databases containing sensitive data, and identify any gaps that could leave data exposed or unrecoverable in a ransomware attack
2. **Implement off-server key management**: Deploy Oracle Key Vault (OKV) and migrate all TDE master encryption keys off database servers to centralized key management
3. **Validate backup immutability**: Confirm that ZDLRA, ZRCV, or ZFS Storage Appliance are configured with immutable retention locks that prevent backup deletion or tampering by ransomware
4. **Review recovery point synchronization**: For multi-database applications, ensure ZDLRA or ZRCV provides synchronized, zero data loss recovery points across all interconnected systems
5. **Test recovery procedures**: Conduct tabletop exercises and actual recovery tests to validate both technical capabilities and organizational readiness

Copyright © 2026, Oracle and/or its affiliates

These Phase 1 measures can typically be implemented in weeks to months, depending on organization size, and can quickly blunt ransomware's ability to destroy data or drive payment demands through data theft. With these fundamentals in place, organizations can then begin the longer term Zero Trust journey with less urgency and lower risk.

Phase 2: The Zero Trust journey

While Phase 1 encryption and recovery capabilities deliver immediate protection, Phase 2 shifts the focus to comprehensive architectural transformation, reducing the likelihood that ransomware ever reaches database servers in the first place. Zero Trust is an ongoing program of continuous improvement rather than a finite project.

Zero Trust: The foundation for long term resilience

Zero Trust architecture is rooted in the principle of "never trust, always verify." Rather than relying on perimeter defenses and implicit trust, Zero Trust verifies every request based on identity, device or workload posture, and context—regardless of where the request originates.

Zero Trust assumes breaches will occur and that any account, credential, or system can be compromised. Organizations reduce reliance on any single point of trust by:

- Enforcing strong, phishing-resistant authentication and least-privilege access
- Continuously verifying identity and session risk, not just at login
- Segmenting networks and applications to contain blast radius
- Granting just-in-time, time-bound access to limit standing privileges
- Monitoring events and enforcing policy in real time

By removing implicit trust and continuously validating each action, organizations limit lateral movement, prevent privilege abuse, and make ransomware campaigns significantly harder to execute.

Typical Zero Trust implementations

The US Cybersecurity and Infrastructure Agency (CISA) offers a graphical model of the technology components involved in Zero Trust:

Copyright © 2026, Oracle and/or its affiliates

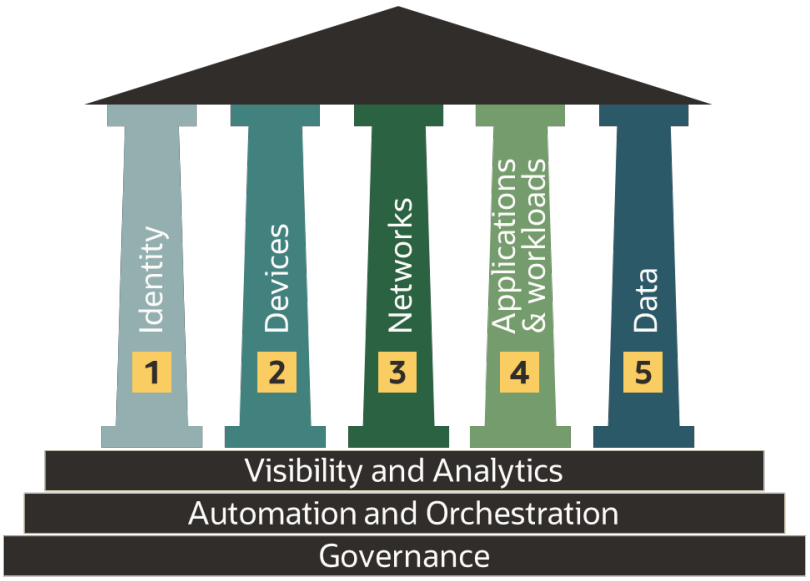


Figure 15-1: CISA Zero Trust maturity model pillars

As the model illustrates, Zero Trust influences nearly every area of IT. Implementations typically bring renewed attention to fundamental security practices and raise expectations for consistent, measurable controls. Common Zero Trust initiatives that affect database operations include:

- Increased network segmentation enforced by newer and more capable firewalls
- Mandates to route administrator connections to databases and database servers through bastion hosts or jump servers
- Adoption of privileged account managers, particularly for operating system accounts like ROOT or ORACLE and database accounts like SYS and SYSTEM.

Four core Zero Trust projects for databases

Applying Zero Trust principles to database security can be organized into four major project areas. Each represents an ongoing initiative that organizations continuously refine and strengthen.

Project 1: Configuration assessment and drift monitoring

Zero Trust Principle: Never assume systems remain in their intended secure state. Instead, continuously verify configuration compliance. Databases are intricate platforms with dozens of security relevant parameters. Misconfigurations are among the most common entry points for ransomware. A Zero Trust approach calls for ongoing validation that databases remain correctly configured, rather than relying on the assumption that an initially secure setup will stay secure over time.

Implementation Approach:

1. **Establish security baselines**: Use Oracle Data Safe, Audit Vault and Database Firewall, or Database Security Assessment Tool (DBSAT) to perform initial security assessments and document compliant configurations
2. **Enable continuous drift detection**: Configure Data Safe or Audit Vault and Database Firewall to automatically flag deviations from defined security standards

Copyright © 2026, Oracle and/or its affiliates

3. **Treat changes as potential IoCs**: Investigate every unexpected configuration change as a potential indicator of compromise

Relevant Oracle utilities, features, services, and products:

- Data Safe for security assessment and drift detection
- Audit Vault and Database Firewall for security assessment and drift detection
- Database Security Assessment Tool for one-time assessment
- Enterprise Manager Database Lifecycle Management Pack

Project 2: Attack surface minimization

Zero Trust principle: Assume attackers will reach the databases. Limit the damage they can cause by removing unnecessary data, privileges, and access points. Most database breaches hinge on compromised accounts that attackers can simply use to log in. Every unneeded privilege, role, or extra copy of sensitive data expands the blast radius of a successful compromise.

Implementation Approach:

1. **Data minimization**:
    a. Remove sensitive data from test and development databases using data masking, so breaches of non-production systems reveal nothing of value.
    b. Use data subsetting to reduce the volume of data copied into non-production environments
2. **Privilege minimization**:
    a. Deploy Database Vault to prevent even DBAs from viewing application data, and enforce Trusted Paths that help stop application accounts from being misused.
    b. Conduct privilege analysis to identify and remove unused privileges and roles.
    c. Implement data driven controls (Virtual Private Database, Real Application Security, Label Security) to limit what each user can see

Relevant Oracle utilities, features, services, and products:

- Data Safe Data Masking and User Assessment
- Enterprise Manager Data Masking and Subsetting Pack
- Audit Vault and Database Firewall entitlement monitoring/reporting
- Oracle AI Database Privilege Analysis
- Database Vault
- Data driven controls like Virtual Private Database, Real Application Security, and Label Security
- Database Security Assessment Tool

Project 3: Authentication strengthening

Zero Trust Principle: Never trust passwords or any single authentication factor. Identity should be verified using the strongest available methods. Authentication is the bedrock of access control. Every authorization decision depends on correctly confirming who is making the request. A Zero Trust model calls for a deliberate shift away from password only access toward multi factor authentication and strong, cryptography-based methods.

Implementation approach:

Tiered Authentication Strategy

Copyright © 2026, Oracle and/or its affiliates

- **Superuser accounts (e.g., SYSDBA, SYSKM)**: These accounts should be protected with Privileged Account Manager (PAM) and used only when absolutely necessary. SYSDBA is rarely required for routine, day to day database administration.
- **Administrator accounts (DBAs, security administrators, application administrators)**: Integrate with Active Directory/Entra ID/Oracle IAM and enforce multi factor-authentication where possible. It may also be appropriate to secure these accounts with a PAM.
- **User accounts (data analysts, developers, testers, business intelligence users):** Require strong authentication (Kerberos, certificates, MFA) for users who connect directly to the database.
- **Application service accounts:** Use the strongest authentication the application supports. When strong authentication is not feasible, implement multi factor _authorization_ via Database Vault or SQL Firewall. Use Gradual Password Rollover to enable password changes without downtime. Monitor logins for unusual patterns.

Note: Starting with Oracle AI Database 26ai and Oracle Database 19c (July 2025 update), MFA can be enabled for local database accounts using Oracle Mobile Authenticator or Cisco Duo without any client-side changes.

For accounts that must continue to rely on passwords, organizations should also audit login activity and watch for indicators of compromise. Common signals include new IP addresses connecting to the database, unfamiliar programs suddenly in use (e.g., AI agents, vibe coded apps), connections at unusual times of day or during non-working hours, or multiple sessions originating from different geographical locations.

Organizations may also want to enforce a trusted path for those password dependent accounts. This involves defining the specific conditions under which they can connect or access sensitive data. By restricting access to approved network paths, client programs, and time windows, even a compromised password is far less likely to be usable outside tightly controlled boundaries.

Relevant Oracle utilities, features, services, and products:

- Database **strong authentication** utilizing Kerberos, certificate, RADIUS, OCI IAM token, and Entra ID OAuth2 tokens
- Database **gradual database password rollover** to safely change an application's password without requiring application downtime
- Database **centrally managed users** via a directory e.g., Microsoft Entra ID
- Database **unified audit** to audit logins by users authenticated via password
- **Database Vault** *Trusted Path* to lockdown accounts with weaker authentication so they can only be used under certain conditions

Project 4: Comprehensive database activity monitoring

Zero Trust principle: Never assume preventive controls are 100% effective. Continuously monitor activity to detect threats early and respond quickly. Perfect security is not achievable, especially in environments where some level of data access must be allowed to keep the business moving. Detective controls provide the safety net by revealing when actions drift beyond approved boundaries, giving incident responders a chance to contain issues before they cascade into broader damage.

Organizations should deploy a layered mix of auditing and network-based monitoring to surface anomalies that may indicate malicious or unauthorized activity, and to support investigations when incidents occur. At a minimum, the database audit trail should capture all data definition and control language activity. If someone creates a new user, grants a role or privilege, replaces a stored procedure, or copies an existing table into a new one, there should be a record. If someone accesses sensitive data outside an authorized application, there should be a record. And if

Copyright © 2026, Oracle and/or its affiliates

privileged users touch data directly, there should absolutely be a record. This kind of visibility turns detective controls into something practical and actionable, rather than an aspirational checkbox.

Implementation approach:

- **Audit security relevant activities**: Capture all data definition and control language operations (creating users, granting privileges, altering procedures)
- **Monitor sensitive data access**: Log instances where sensitive data is accessed outside normal application flows
- **Track privileged user activity**: Record all actions performed by DBAs and other high privilege accounts
- **Identify anomalies**: Watch for indicators of compromise, including new IP addresses, unusual connection times, unexpected programs (AI agents, vibe coded apps), and geographic anomalies

Relevant Oracle utilities, features, services, and products:

- Database **Unified Auditing** to capture security relevant activity
- **Audit Vault and Database Firewall/SQL Firewall** to examine ALL database commands and identify anomalies (note SQL Firewall is a new capability of Oracle AI Database 26ai.)
- **Data Safe** Activity Auditing or **Audit Vault and Database Firewall** for audit analysis, reporting, and alerting

Copyright © 2026, Oracle and/or its affiliates

Zeroing in on Zero Trust

No single set of steps makes a system perfectly secure, and there is always more that can be done. Implementing the practices outlined above will materially strengthen the security posture and reduce implicit trust to a level that is practical for most environments.

The “to learn more about…” pointers included after each topic are intended to help navigate related sections. Zero Trust touches nearly everything discussed because it brings these concepts together into a single, consistent approach.

For organizations planning a Zero Trust initiative, the scope should remain realistic. Zero Trust is not implemented in a few weeks and then declared complete. It is a long running program, and many of the most mature organizations never “finish” Zero Trust. They iterate, measure, and continually adjust controls to improve security over time. This discipline matters because attacks continue to evolve, especially in an AI-driven world.

Defending against ransomware takes ongoing vigilance and diligence, yet the outlook keeps improving as technology and infrastructure advance each day. Zero Trust remains the most effective approach for blocking ransomware. And if an attack does succeed, law enforcement agencies advise victims not to pay the ransom. With encryption and an immutable, zero data loss recovery plan in place, it becomes far more practical to follow that advice.

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Tales from the Dark Side: Hacking the Database

Copyright © 2026, Oracle and/or its affiliates

Summary

Ransomware is widely treated by major cybersecurity agencies as the most immediate and disruptive threat facing organizations today, with criminal groups using it for revenue and state-backed actors increasingly using it to cause broader disruption. Modern campaigns often go beyond simple encryption to combine credential theft, lateral movement, privilege escalation, and double extortion, where attackers steal data and threaten public disclosure. Some operations skip encryption entirely and rely purely on theft and extortion. Even when victims pay, there is no guarantee the data will be deleted, not sold, or successfully recovered, which makes prevention, containment, and trustworthy recovery central to resilience planning.

A consultative two-phase strategy is presented to balance urgent risk reduction with long-term architectural improvement. Phase 1 focuses on immediate, high-impact defenses that blunt ransomware's two primary database-level threats: data loss through encryption and data theft for extortion. The central message is to deploy encryption and recovery capabilities first, because they function as a final line of defense when perimeter controls fail and accounts are compromised. This approach aims to reduce pressure to pay ransoms by making stolen data less useful and recovery more dependable.

On the encryption front, Transparent Data Encryption (TDE) stands out as the go-to protection against out-of-band attacks that slip around application-level controls. It's important to treat "data at rest" and "data in motion" as two different battlefields. Data at rest is protected with AES-256 for database files, while data in motion relies on TLS or Oracle Native Network Encryption (NNE). One operational rule is nonnegotiable: encryption keys must be kept separate from the data they protect. Oracle Key Vault (OKV) supports that separation with centralized key lifecycle management for TDE keys, audit trails, automated key rotation policies, and the ability to restore keys into isolated recovery environments.

On the recovery side, the emphasis shifts to immutable backups, because modern ransomware is not content to attack production systems. It actively searches for and destroys traditional backups using standard protocols and privileged access. ZDLRA and ZRCV are positioned as a way to deliver immutable, synchronized protection, including recovery points aligned across interconnected databases to avoid transaction gaps and inconsistency. Planning should also include recovery clean rooms and isolated environments, allowing restored systems to operate while production environments go through malware remediation. Options like air gap separation and continuous incremental restore can further compress recovery timelines when every minute matters.

With Phase 1 protections in place and a clear view of the Zero Trust journey ahead, the next practical question is where security becomes the default rather than the ongoing fight for consistency. That is where Oracle Autonomous AI Database enters the story. Chapter 16 shifts from "how to withstand ransomware" to "how to start strong and stay strong," highlighting the service's hardened baseline configuration, always-on encryption for data at rest and in motion, and default auditing, along with the operational advantage of automated patching and maintenance. It also draws a crisp line around the shared responsibility model, clarifying what the service secures automatically and what the organization must still govern, such as user privileges, data sensitivity, and access policy. For readers who want defenses that hold up not only in a clean architecture diagram but also at 2:00 a.m. during an incident, the next chapter translates the principles from this ransomware discussion into concrete, repeatable security posture in an autonomous environment.

Copyright © 2026, Oracle and/or its affiliates

Chapter 16

Securing the Autonomous Database

Database Security in Autonomous AI Database

As more organizations adopt the Oracle Autonomous AI Database (ADB), it is worth understanding what distinguishes it from other Oracle AI Database services and editions. Knowing those differences makes it easier to prioritize risk mitigation and protect data, particularly as AI agents and automated tools become regular database users.

ADB comes with standardized, hardened security configurations that help reduce the time and cost spent managing settings across databases. It also applies security patches and updates automatically, which means teams spend less time chasing maintenance tasks and more time focusing on outcomes. Together, these capabilities help protect databases and data from costly and potentially disastrous vulnerabilities and breaches. Oracle Autonomous AI Database also encrypts data at rest and in motion using industry-standard cryptographic algorithms, and it provides tools to further restrict or isolate sensitive data from privileged users, developers, data analysts, application administrators, and AI agents.

Even with built-in protections, some responsibilities remain firmly with the organization. Autonomous AI Database automates many infrastructure and security tasks, but it has no visibility into whether granted users behave according to your organization's policies, nor does it inherently understand the sensitivity of the data it stores. These remain the responsibility of the database owner. That is why this chapter matters. It helps clarify how responsibilities are shared between the organization and the database operations team, and it highlights the tools available to control risk and secure the system with greater precision.

The value of Autonomous AI Database

ADB minimizes unnecessary human labor, and the errors that accompany it, so your DBAs and developers can focus on higher-value activities instead of time-consuming routine tasks.

Broadly speaking, Autonomous AI Databases handle many routine tasks and provide these features:

- Automates database and infrastructure provisioning, management, monitoring, backup, recovery, and tuning.
- Provides preventative protection against many unplanned and planned downtime, and enables rapid, automatic recovery from outages without downtime. Autonomous AI Database availability and performance management are taken to the next level using AI-based autonomy that integrates multiple areas of diagnostics and enables analysis and action to be taken at runtime to minimize or eliminate operational disruption.
- Protects itself from many vulnerabilities and attacks. Oracle Cloud Infrastructure provides continuous threat detection, while Autonomous AI Database automatically applies all security updates online and provides "always on" end-to-end encryption. This preventative approach is critical because most security breaches that leverage a known vulnerability occur after the patch that mitigates that vulnerability is already available!

Security benefits of Autonomous AI Databases

By this point in the book, readers have likely developed a strong appreciation for data risk. Reducing that risk is one of the top concerns in application development and deployment. As earlier chapters in this book show, there is a lot to consider when securing a database. The practical goal is for the Autonomous AI Database to do as much of that work as possible, freeing teams to focus on security tasks that are not practical to automate.

Autonomous AI Database reflects years of security development and hands-on experience operating critical workloads for some of the world's largest and most security-conscious customers. Oracle has been a leader in database security for decades, and that expertise has been built into Autonomous AI Database by tightening areas of risk that customers commonly agree should be handled consistently. Very few principles apply universally across

Copyright © 2026, Oracle and/or its affiliates

more than 400,000 customers in 175 countries and every industry vertical. Still, there are many areas where broad agreement exists, and everything in that category that can be automated has been automated.

Security capabilities of Autonomous AI Databases

The security capabilities discussed so far are built into Oracle Autonomous AI Database. Together, they establish a baseline security posture that is already stronger than most on-premises environments and flexible enough to meet highly stringent security requirements. Most Oracle Maximum Security Architecture (MSA) technologies are widely adopted, de facto industry standards for protecting and monitoring Oracle AI Database environments. These security capabilities include:

- **Assessment and configuration –** Earlier chapters highlight how critical secure configuration is. Autonomous AI Database begins from a secure configuration baseline. Application data is isolated into well-defined tablespaces, sensible password policies are applied to meet most customer requirements, and security-relevant initialization parameters are set with the CIS Benchmarks and DISA STIG in mind. Because Data Safe is included with the Autonomous AI Database service, the database can be monitored for configuration drift.
- **Encryption for data in motion –** Autonomous AI Database is automatically configured to use industry-standard Transport Layer Security to encrypt data in transit between the database service and clients or applications. Required client certificates and networking information are automatically packaged for the service consumer when the service is provisioned**.**
- **Encryption for data at rest –** Data in the Autonomous AI Database is automatically encrypted using Oracle Transparent Data Encryption**.** Because the data is encrypted, backups of the data are also encrypted.
- **Flexible key management –** Key management for Autonomous AI Database can be Oracle-managed (Autonomous AI Database creates and manages the encryption keys) or customer-managed. With customer-managed keys, the encryption keys are stored in OCI Vault for Autonomous AI Databases running in OCI, or in Key Vault for Autonomous AI Database running in OCI, all third-party clouds, or on Exadata Cloud@Customer.
- **Automated separation of duties –** The Autonomous AI Database removes direct access to the database node and local file system. Additional separation between service administrators and service consumers is provided through Operations Control, a feature of Oracle Database Vault. This separation of duties, a key Oracle Cloud differentiator, reduces the risk of administrator misconduct and eliminates the ability of service administrators to view or modify data stored in Autonomous AI Database. As with Transparent Data Encryption, Database Vault has been continuously enhanced, with features such as Operations Control added explicitly to support Autonomous AI Database.
- **Database auditing configured by default and customizable on demand –** Autonomous AI Database is preconfigured with audit policies that monitor privileged user activity and logon failures, along with optional preconfigured policies aligned with the Center for Internet Security (CIS) audit benchmarks, account management, and more.
- **Assessing the database security and its data –** Data Safe is included with Autonomous AI Database to help assess and secure databases. Data Safe helps teams understand data sensitivity, evaluate data risks, mask sensitive data, implement and monitor security controls, assess user security, monitor user activity, and address data security compliance requirements. Sensitive and regulated data in Autonomous AI Database can be identified and protected by registering the database with Data Safe.
- **Masking sensitive data for non-production use -** Production data is often well protected in production, then copied into mostly insecure nonproduction environments for development. In practice, if a database contains production data, it still requires production level protection. To provide developers a sanitized environment, Data Safe supports identifying sensitive data, creating masking templates, and masking data for non-production Oracle AI Databases, including Oracle Autonomous AI Database.

Copyright © 2026, Oracle and/or its affiliates

- **Reduced opportunity for human error –** Human error contributes to many data breaches and is one of the most difficult risks to mitigate. Autonomous AI Database reduces this risk by automating a significant portion of database administration. It further reduces error by limiting the range of commands an Autonomous AI Database user is permitted to run.
- **Automated patching, upgrades, and maintenance –** A major advantage of Autonomous AI Database is the ability to apply security patches and upgrades without downtime. Autonomous AI Database is patched automatically when new security patches become available, with a patch frequency of at least monthly, which is more than most customers can reasonably achieve on their own. Much of this capability builds on well-tested, mature Oracle AI Database technologies such as Real Application Clusters for rolling online RAC patches and cloud service process automation.

These are only the fundamentals. Autonomous AI Database also locks down access to sensitive packages such as UTL_HTTP, UTL_TCP, and UTL_SMTP, tightly controlling interactions between the database and the surrounding network. OCI administrator access to the databases is restricted by Database Vault Operations Control.

Taken together, these Autonomous AI Database capabilities provide a security framework that addresses core security requirements for most organizations out of the box, enabling operations and security teams to focus on elevating enterprise security posture further.

Shared responsibility

Some security responsibilities remain outside the reach of Autonomous AI Database self-securing capabilities. Oracle implements data privacy controls to ensure that no one at Oracle can view the data placed into an Autonomous AI Database. As a result, Oracle does not know whether the data in a customer database is sensitive, nor whether it is subject to specific regulatory requirements.

Oracle also has no visibility into a customer's database users or AI agents, their job functions, or the policies that determine which data they should or should not access. Oracle similarly does not know the contextual signals that often inform access decisions, such as working hours, client programs, and normal work locations. This limits Oracle's ability to determine whether additional protections are being applied to satisfy regulations, laws, or the risk profile of the data stored in the database. It also means Oracle cannot determine whether users have been granted appropriate privileges or whether access aligns with customer policy.

To support the security requirements that cannot be handled autonomously, Oracle provides a full suite of security tools and services with Autonomous AI Database.

A primary tool is Oracle Data Safe. Data Safe reports on users, their privileges, and their activity to help identify over privileged users or users who are not acting according to policy. Data Safe also scans the database for sensitive data, helping teams understand where data risk resides and potential exposure based on the types and amount of data stored. It also helps address the proliferation of sensitive data by enabling data masking for non-production environments.

In addition, Autonomous AI Database includes the same security capabilities Oracle AI Database provides in other deployments, including Database Vault, Data Redaction, and Label Security. These capabilities are available to use and can be configured to meet customer requirements.

Copyright © 2026, Oracle and/or its affiliates

Hands-On learning with Oracle LiveLabs

Oracle LiveLabs provides free, self-paced workshops for hands-on experience with Oracle technologies. These resources let you put the concepts from this chapter into practice, safely and at your own pace. To get started, check out the following labs:

- Autonomous AI Database for Security Administrators
- Prevent unauthorized access in Autonomous AI Database with Database Vault
- Securing a legacy application using Database Vault on Autonomous AI Database
- Simplify user management by integrating Autonomous AI Database and Oracle Cloud Identity
- How do I create a database user using Data Tools

Summary

Oracle Autonomous AI Database is designed to start strong and stay strong, secure from day one and steadily reinforced through proven operational discipline and Oracle Maximum Security Architecture. It combines a hardened baseline configuration, always-on encryption, automated patching and maintenance, and default auditing so the safeguards most teams want are applied consistently, not left to chance or calendar reminders. Just as importantly, it reduces the likelihood of human error by automating large portions of routine administration and by limiting the range of commands users are allowed to run.

At the same time, this is not a set-it-and-forget-it approach. Oracle implements privacy controls so that Oracle personnel cannot view customer data. As a result, Oracle does not know whether that data is sensitive or subject to specific regulatory requirements, and Oracle has no visibility into customer users, AI agents, their roles, or the contextual signals that shape access decisions.

That shared responsibility reality is exactly why the surrounding security portfolio matters. Tools like Oracle Data Safe help teams discover sensitive data, assess user privileges and activity, identify drift, and mask data for non-production environments, while Database Vault, Data Redaction, and Label Security offer additional ways to restrict, isolate, and protect sensitive information across a range of enterprise needs.

The outcome is a security foundation built to help organizations move forward with AI confidently by automating what can be automated and enabling teams to precisely govern what cannot be automated.

At this point, there is a strong, secure foundation to build on: an Autonomous AI Database environment built to begin strong and stay strong. With a hardened baseline, always on encryption, automated patching, and default auditing, much of the routine day-to-day risk is already taken off the table. Still, the real impact in database security rarely comes from any single control. It comes from how the layers work together, how the right priorities are chosen, and how an organization turns solid capabilities into a repeatable security posture, especially as AI agents and automated tools become everyday database users.

That is where the concluding chapter (“Putting it all together”) takes the lead. Instead of focusing on individual mechanisms, it shifts to practical orchestration: how to apply defense in depth as one coherent strategy, how to decide where to start, and how to align controls with what matters most, namely the sensitivity of the data, its business criticality, and the surrounding threat environment. If this chapter helped make the shared responsibility model clear, the next chapter provides a playbook for turning that clarity into crisp priorities and durable outcomes.

Copyright © 2026, Oracle and/or its affiliates

Chapter 17

Putting It All Together

Oracle Maximum Security Architecture

The goal throughout this journey into securing the Oracle AI Database has been to offer useful insight into the many defense-in-depth strategies that can help keep data secure. A wide landscape has been covered, from locking down privileges and encrypting data to auditing activity and stepping into the newest frontier: Securing Agentic AI. When these layers are brought together, see Figure 17-1, they create the *Maximum Security Architecture*, a clear blueprint for securing data and minimizing risk.

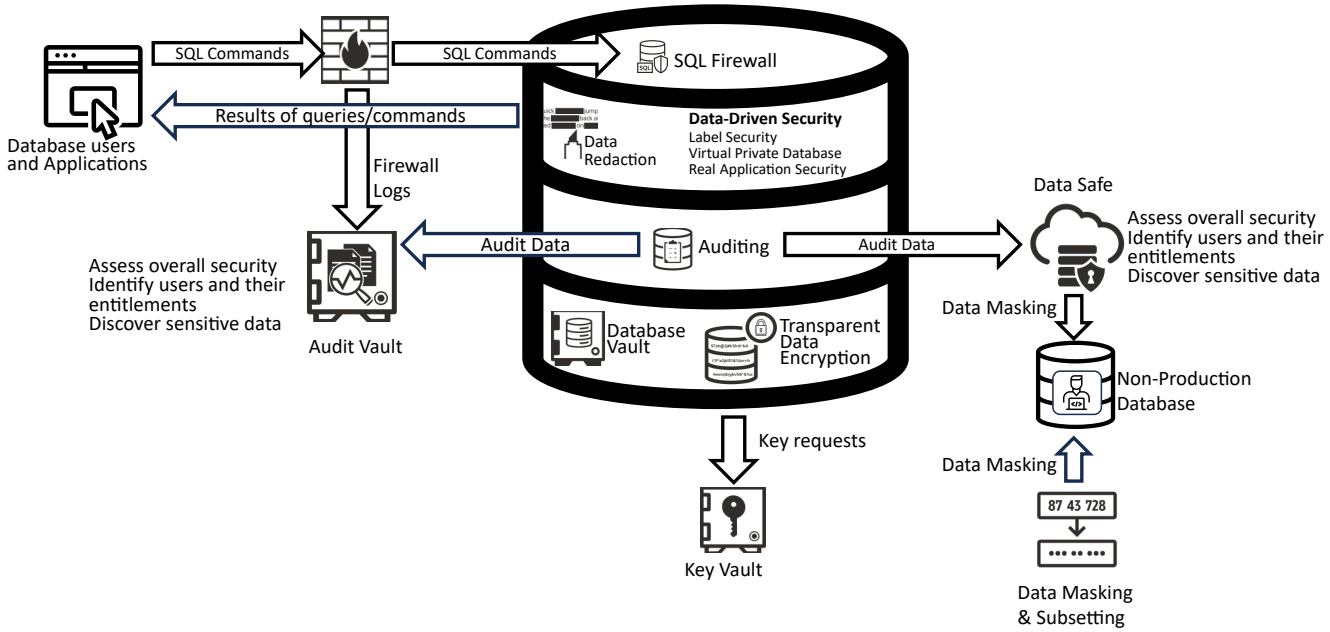


Figure 17-1: Oracle Database Maximum Security Architecture

Database security can feel like a sprawling discipline, especially when the scope includes hundreds of configuration parameters, diverse authentication methods, layered access controls, multiple encryption options, comprehensive auditing, and tools that span assessment, monitoring, and enforcement. This book has walked through each of these areas in depth, presenting the technical capabilities Oracle Database offers to help organizations protect their most sensitive assets. Now it is time to step back and see how these pieces fit together into a coherent, actionable security strategy.

This chapter synthesizes the main themes from the preceding chapters, highlights the connections between different controls, and provides practical guidance on how to move from understanding to implementation. The goal is to leave readers with a clear roadmap, not just a list of features, and to reinforce that database security is not a one-time project but an ongoing discipline that adapts as threats evolve, regulations change, and business needs shift.

The defense-in-depth framework revisited

Throughout this book, defense-in-depth has served as the organizing principle. The idea is straightforward: no single control will stop every attack, so organizations layer multiple controls across different dimensions. If one layer fails or is bypassed, others remain in place to detect, contain, or block the threat.

Oracle Database security capabilities naturally align with four widely shared objectives across the broader digital security landscape:

1. **Assess security posture** to identify misconfigurations, vulnerabilities, and drift before attackers exploit them
2. **Control access to data** through strong authentication, least-privilege authorization, and fine-grained policies

Copyright © 2026, Oracle and/or its affiliates

3. **Monitor user activity** with comprehensive auditing and real-time anomaly detection
4. **Protect data against theft** using encryption, key management, and secure recovery

These four pillars are not independent silos. Assessment informs which access controls to prioritize. Access controls depend on knowing where sensitive data resides. Monitoring validates that access controls work as intended. Encryption ensures that if all other layers fail, stolen data remains unreadable. When implemented together, these capabilities create a security posture that is far more resilient than any single feature alone.

Database security applied medallion framework

The recommended approach is to implement security controls that align with the sensitivity of the data, its importance to the business, and the surrounding threat environment. One practical way to start is to inventory systems and group them into levels based on sensitivity. These levels could be labeled Bronze, Silver, Gold, and Platinum. Bronze systems might include internal portals, employee directories, and Wikis. Silver systems might include business transaction systems, supplier information, and parts catalogs. Gold systems might include data subject to regulatory compliance, whether General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), Personally Identifiable Information (PII), Payment Card Industry (PCI), Health Insurance Portability and Accountability Act (HIPAA), or Sarbanes-Oxley Act (SOX). Platinum systems might include highly sensitive and restricted data, such as quarterly sales numbers, sales forecasts, M&A activities, and intellectual property such as source code.

The next step is to define and implement controls appropriate for each category. For example, an organization might identify the controls shown in Figure 17-2 and apply them as layered protections that build on one another as systems move up through the levels. Databases at the Bronze level and above must be securely configured and kept current with security patches. Otherwise, attackers can exploit an unpatched system, then use it as a command-and-control base for further attacks or as a staging area for the data they uncover. In addition, privileged user activity on these machines should be monitored and audited to capture meaningful changes in configuration or user access privileges.

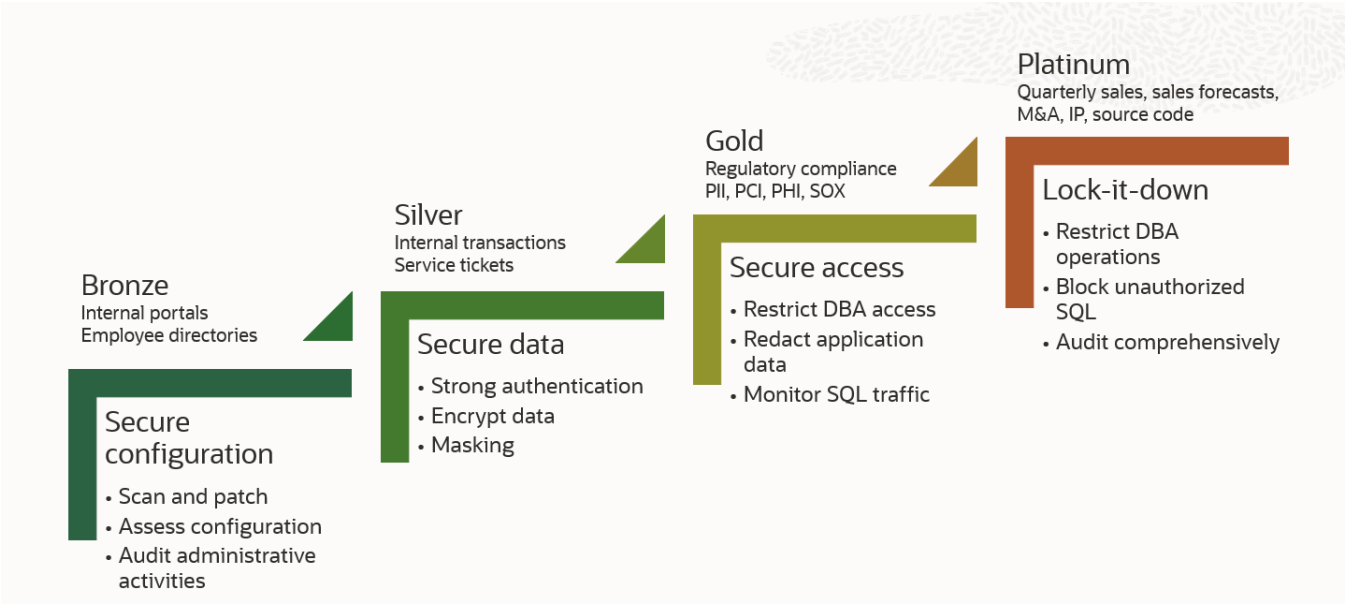


Figure 17-2: Example of implementing security controls based on system sensitivity.

Copyright © 2026, Oracle and/or its affiliates

Additional measures can be implemented for Silver and above databases to protect data from unauthorized users. These measures can include encrypting data in transit, preventing it from being viewed directly from the operating system, and keeping it out of test and development machines. These databases may also require privileged users to authenticate with strong passwords or to use PKI or Kerberos based authentication. The primary security controls for Silver typically include encryption, both in motion and at rest, masking, and strong authentication.

For Gold and above databases, additional controls should shield data from privileged users, autonomous AI agents, and anyone else without a legitimate business need. Privileged user access must be restricted while still enabling administrators to perform their duties, and SQL activity over the network should be monitored to quickly identify malicious attempts to exploit application vulnerabilities. To address basic compliance requirements, all PII, PCI, and PHI data must be protected.

Platinum databases should include all the security controls used for Bronze, Silver, and Gold. In addition, Platinum databases should be further locked down because they contain the most sensitive data in the organization. Additional measures can include restricting who can log on to the database server, monitoring SQL operations in real time, guarding against potential SQL injection attacks, and implementing comprehensive audit policies so forensics teams can determine the impact and method of attack in the event of a breach. When AI agents touch Platinum systems, the stakes are highest:

- Deploy SQL Firewall allowlists to govern AI generated SQL
- Audit every agentic session
- Enforce Database Vault command rules, realms, and trusted paths so autonomous agents never stray beyond their boundaries.

Depending on priorities and the organization's security strategy, deployment can begin from either end of the spectrum. One option is a controls first approach, starting by assessing and securing database configurations, then moving on to configuring transparent data encryption, and continuing by implementing one control at a time across the environment. Another option is to start with the most critical and sensitive systems and implement all necessary controls on those systems one at a time. Both approaches are valid, and in practice organizations often use a combination of the two.

AI shifts the access pattern from human-driven to tool-generated SQL, and if AI agents are already touching production data, those systems should move to the front of the line. A vibe-coded dashboard hitting a Bronze database with real customer is a Platinum-level risk.

An effective strategy accounts for business objectives, people, resources, and time, protecting "toothbrushes" and "diamonds" with the appropriate level of security, maximizing return on security investment. Then the layers fall into place: hardened, patched foundations; encryption and masking to protect data in motion and at rest; strong authentication; and monitoring and auditing that transform uncertainty into evidence.

As agentic AI changes how databases are accessed by generating SQL, orchestrating workflows, and touching production data, Platinum environments require the most rigorous controls: allowlisting AI-generated SQL, auditing every agentic session, and enforcing strict boundaries so autonomy never becomes overreach.

The takeaway is clear: Prioritize wisely, layer the right controls, and organizations can move fast, embrace AI, and still protect the one asset they cannot replace, data.

Chapter-by-Chapter journey: Key takeaways

To understand how the pieces connect, it helps to revisit the journey taken through this book. Each chapter addressed a specific domain within database security, building on the foundations laid earlier.

Copyright © 2026, Oracle and/or its affiliates

Chapters 1 and 2: Foundation and threat landscape

The opening chapters established why database security matters now more than ever. Chapter 1 sets the context: database security is no longer a specialized concern. It is an enterprise risk driver shaped by regulations, ransomware, and rapidly evolving attacker tactics. It also argues for pragmatic focus over theory. Databases concentrate sensitive assets in one place, making them high-value targets. Regulatory requirements such as GDPR, PCI DSS, and breach notification laws create legal and financial consequences for organizations that fail to protect data. At the same time, threat actors have grown more sophisticated, deploying ransomware campaigns that combine encryption, exfiltration, and extortion, and leveraging AI to accelerate reconnaissance and exploitation.

With the motivation established, the next chapter defines the security outcomes to protect data in real environments. Chapter 2 introduces a goal-based approach, including the "Dirty Dozen" risk areas, and anchors regulatory alignment in three durable pillars: encrypt data, control access, and audit access. The Dirty Dozen framework cataloged the most common database security risks, from insecure configurations and unpatched systems to weak authentication and unprotected backups. The goal-based approach to security introduced in Chapter 2 provided a practical lens: align every task to a clear security outcome, choose the right mix of technical and process controls, and focus on reducing real risk rather than chasing compliance checkboxes.

These chapters also introduced a critical mindset: think like an attacker. Understanding how adversaries probe for default credentials, exploit SQL injection, intercept unencrypted traffic, and scrape database files without establishing sessions helps prioritize the controls that make attacks harder and costlier.

Key takeaways:

- Database security is driven by converging forces: regulatory mandates, sophisticated threats, and AI-amplified risks
- The Dirty Dozen provides a common vocabulary for database vulnerabilities
- A goal-based approach ensures security work delivers measurable risk reduction
- Defense-in-depth means layering controls so that no single failure leads to compromise
- Adopting an attacker mindset helps prioritize hardening efforts

Once outcomes are clear, the next practical question is how to measure the current state and detect drift over time.

Chapter 3: Database Security Posture Management

Security assessment is systematic reconnaissance of an organization's own defenses. Chapter 3 explains how Database Security Posture Management is best thought of as the discipline that turns database security from a hopeful checklist into a repeatable, defensible practice. Attackers automate reconnaissance to find these weaknesses, so defenders must do the same.

Four tools address assessment at different scales. DBSAT provides standalone, no-cost assessments for individual databases or small fleets. Data Safe offers cloud-based, fleet-wide assessment with automated scheduling, drift detection, and centralized reporting. AVDF delivers similar capabilities with on-premises deployment for organizations with data residency requirements. Oracle Enterprise Manager Database Lifecycle Management integrates assessment into broader infrastructure management workflows.

Assessment is not a one-time activity. Configuration drift is inevitable as databases are patched, applications are updated, and administrators make changes. Continuous or recurring assessment ensures that drift is detected and remediated before attackers exploit it. Establishing baselines, tracking changes over time, and generating compliance reports aligned with standards such as CIS Benchmarks, DISA STIG, and GDPR help demonstrate due diligence to auditors and regulators.

Copyright © 2026, Oracle and/or its affiliates

Privilege Analysis extends static assessment by tracking which privileges and roles accounts actually use. This data-driven approach enables confidently revoking unused permissions, shrinking the blast radius when credentials are compromised.

Key takeaways:

- Assessment tools identify misconfigurations before attackers exploit them
- DBSAT, Data Safe, AVDF, and DBLM serve different scales and deployment models
- Continuous assessment detects configuration drift and validates remediation
- Privilege Analysis reveals unused privileges, enabling least-privilege enforcement
- Patching keeps databases current with security updates and closes known vulnerabilities

Chapter 4: Sensitive Data Discovery

Effective protection starts with knowing where sensitive data resides. Chapter 4 introduced sensitive data discovery, the process of scanning database schemas to identify columns containing personally identifiable information (PII), financial data, health information, and other regulated and sensitive data types.

DBSAT Discoverer, Data Safe Data Discovery, and Oracle Data Masking and Subsetting all classify sensitive data using similar techniques that vary slightly in pattern matching, metadata lookups, and pattern-based data sampling. Discovery reports categorize findings by sensitivity level and recommend appropriate protections: encryption for high-risk data, access controls for medium-risk data, and auditing for all sensitive data.

Discovery is particularly valuable before migrating databases to new environments, provisioning data to non-production systems, or responding to regulatory inquiries. Organizations that cannot quickly answer "Where is customer PII stored?" struggle to demonstrate compliance and respond effectively to incidents.

Key takeaways:

- Sensitive data discovery identifies where PII, financial data, and health information reside
- Discovery tools classify data by sensitivity and recommend risk-appropriate protections
- Regular discovery scans track how sensitive data footprint changes as applications evolve
- Discovery informs encryption, access control, auditing, and masking decisions

Chapter 5: Strong Authentication

Authentication sits at the foundation of database security. Chapter 5 emphasized that if authentication is weak, all downstream access controls become less effective because attackers can impersonate legitimate users.

Passwords remain common but represent the weakest authentication method. When passwords are unavoidable, strong password profiles enforce complexity requirements, limit failed login attempts, and mandate periodic rotation. The password rollover feature introduced in Oracle Database 19c allows applications to accept both old and new passwords during transition periods, reducing downtime risk during password changes.

Stronger authentication methods reduce credential theft risk. PKI certificates, Kerberos, RADIUS, and cloud identity tokens provide phishing-resistant authentication. Local user multifactor authentication (MFA) adds a second factor for interactive logins, mitigating credential stuffing and brute force attacks.

Centralized user management through directory services like Microsoft Active Directory or cloud identity platforms such as Oracle Cloud Infrastructure IAM, Microsoft Entra ID, and Okta streamlines account lifecycle management. Centralization ensures that when employees leave or change roles, database access is automatically updated without relying on manual processes prone to error and delay.

Copyright © 2026, Oracle and/or its affiliates

Proxy authentication enables administrators and developers to work as application accounts without knowing application passwords. Audit trails capture both the proxied identity and the real user, preserving accountability.

Key takeaways:

- Authentication strength determines the effectiveness of downstream access controls
- Avoid passwords where possible; use PKI, Kerberos, RADIUS, or cloud identity tokens
- When passwords are required, enforce strong profiles and enable local user MFA
- Centralized user management reduces orphan accounts and streamlines role changes
- Proxy authentication preserves accountability while avoiding shared credentials
- Secrets managers like Oracle Key Vault centralize credential storage and distribution

Once identities are established, authorization becomes the next control point, defining what each identity is permitted to do within Oracle Database.

Chapter 6: Database Access Controls

Access control defines who can do what within the database. Chapter 6 introduced the four privilege types in Oracle Database: object privileges (fine-grained permissions on specific tables or procedures), schema privileges (scoped to a single schema, new in Oracle AI Database 26ai), system privileges (broad permissions across the database), and administrative privileges (focused on database administration tasks).

Roles aggregate privileges and simplify management. Instead of granting dozens of object privileges to each user individually, organizations create roles that bundle related permissions and assign roles to users. Secure application roles extend this model by requiring application context or custom logic before activation, ensuring roles are only enabled under approved conditions.

The principle of least-privilege guides access control design: grant only the permissions required for a user or application to perform its intended function, nothing more. Over-privileged accounts create risk because compromised credentials provide attackers with broader access. Custom DBA roles tailored to specific administrative functions (backup DBA, performance tuning DBA, security administrator) reduce the risk of granting the overly powerful default DBA role.

Centrally Managed Users (CMU) integrate database authorization with directory services, allowing roles and privileges to be managed externally. This approach works well for organizations with established identity management infrastructure and policies.

Key takeaways:

- Oracle Database provides four privilege types with different scope and power
- Schema privileges (new in Oracle AI Database 26ai) provide safer middle ground between object and system privileges
- Roles simplify privilege management and enable consistent policy enforcement
- Least privilege reduces blast radius when accounts are compromised
- Custom DBA roles replace the overly powerful default DBA role for specialized administrators
- Centrally Managed Users integrates database authorization with directory services

Least-privilege and roles help, but many programs fail when privileged users can override intent. The next chapter makes separation of duties enforceable.

Copyright © 2026, Oracle and/or its affiliates

Chapter 7: Database Vault and Separation of Duties

Database Vault addresses a critical challenge: how do you protect sensitive data from users who have powerful database privileges, including and especially DBAs? Chapter 7 explained that traditional privilege models grant DBAs nearly unlimited access, creating risk when DBA accounts are compromised or when separation of duties is required for compliance.

Database Vault introduces three core capabilities. Realms define protected zones within the database, restricting access to authorized users even if other users hold powerful privileges like SELECT ANY TABLE. Command rules control who can execute specific SQL statements under what conditions, enabling policies such as "only the application service account can connect from the application server" or "ALTER USER statements are only allowed during maintenance windows." Factors represent security-relevant context such as client IP address, program name, time of day, or database session attributes, enabling dynamic, context-aware authorization decisions.

Separation of duties ensures that no single individual has unchecked power. Database Vault enforces this by separating roles: the DV Owner manages realm and command rule definitions, the DV Account Manager creates and modifies users, and the DV Security Manager handles audit policies. This triad ensures that administrators cannot unilaterally disable controls, create backdoor accounts, or alter audit trails.

Key takeaways:

- Database Vault protects data from users with powerful privileges, including DBAs
- Realms define protected zones where access requires explicit authorization
- Command rules enforce conditional policies based on factors like client IP and program name
- Factors enable context-aware authorization decisions
- Separation of duties ensures no single administrator has unchecked power
- Trusted path restricts high-privilege accounts to approved connection conditions

With identity and authorization tightened, the book then addresses a persistent technical reality: applications are a common breach channel.

Chapter 8: SQL Firewall

SQL Firewall represents a proactive defense against SQL injection and unauthorized database commands. Chapter 8 explained that SQL injection remains a prevalent attack vector because many applications fail to properly validate user input. Attackers inject malicious SQL fragments into application inputs, tricking the database into executing unintended commands.

SQL Firewall operates by learning the SQL statements applications normally execute and then enforcing an allow list. Statements that do not match the allow list are blocked and logged. This approach provides strong protection without requiring application code changes.

Three enforcement modes provide flexibility. Capture mode records all SQL statements executed by a user or application. Analysis mode compares runtime statements against the allow list and logs violations without blocking. Enforce mode actively blocks statements that do not match the allow list, providing real-time protection.

SQL Firewall complements application-layer input validation. Even if application developers fail to sanitize inputs, SQL Firewall provides a database-native enforcement layer that blocks unexpected commands. This defense-in-depth approach reduces reliance on perfect application security.

Key takeaways:

- SQL Firewall blocks SQL injection and unauthorized database commands

Copyright © 2026, Oracle and/or its affiliates

- Allow lists define approved SQL statements for each user or application
- Capture mode learns normal behavior; enforce mode blocks deviations
- SQL Firewall protects without requiring application code changes
- Works alongside application input validation as defense-in-depth

Even when injection is addressed, many organizations still need to ensure users see only the rows and columns they are entitled to. That drives fine-grained access controls.

Chapters 9 and 10: Fine-Grained Access Controls

Standard privilege models operate at object-level granularity such as the ability to SELECT data from a table or view. Very commonly, many applications and use cases require finer control. Chapters 9 and 10 explored capabilities that restrict access at the row and column level. Chapter 9 explains why application-layer authorization is hard to keep consistent and why direct tools can bypass application checks. It introduces database-centric enforcement for row and column level control. Chapter 10 distinguishes static masking for non-production from dynamic redaction, or data redaction, for runtime display control, and it highlights data minimization through subsetting.

Virtual Private Database (VPD) enforces row-level security by dynamically appending WHERE clauses to SQL statements based on policy functions. VPD policies are transparent to applications and can implement complex business logic. Application contexts store security-relevant attributes in database session memory, enabling policy functions to make decisions based on current user identity, role, and application state.

Oracle Label Security (OLS) applies mandatory access control using hierarchical labels composed of levels, compartments, and groups. Data rows are labeled, users are labeled, and the database enforces access rules automatically. One of the interesting elements to OLS is that the data label that is always fixed to the row remains hidden by default. This transparency allows it to work with applications without requiring code changes. The Label Security model works well for government and military classifications (Top Secret, Secret, Confidential) and for commercial scenarios like multi-tenant SaaS applications where customers must not see each other's data. The ability to easily describe and validate the security policy makes OLS an attractive alternative to VPD especially when compliance auditors are trying to verify if the security controls are enforcing the security policy.

Real Application Security (RAS) provides a declarative, role-based authorization framework well-suited for mid-tier applications. RAS policies define which application users can perform which actions on which data, with enforcement occurring within the database. RAS simplifies authorization logic by moving it out of application code.

Data Redaction masks sensitive column values in query results based on user identity and context. Unlike masking tools that modify data, redaction occurs in real-time as queries execute. Predefined formats redact Social Security numbers, credit card numbers, and other common sensitive data types. Redaction protects production data when business users need access for analytics but should not see full sensitive values.

Key takeaways:

- Fine-grained access controls restrict data at row and column level
- VPD dynamically restricts rows based on policy functions and application context
- Label Security enforces mandatory access control using hierarchical labels
- Real Application Security provides declarative authorization for mid-tier applications
- Data Redaction masks sensitive values in query results without modifying stored data

After controlling what users can see and what data is copied, the next control layer typically addresses durability of protection at rest and in motion.

Copyright © 2026, Oracle and/or its affiliates

Chapter 11: Encryption and Key Management

Encryption transforms readable data into ciphertext, rendering it useless without the corresponding decryption key. Chapter 11 explained that encryption protects data when attackers bypass authentication and access controls, such as by stealing backup media, exploiting operating system privileges, or scraping database files directly from storage.

Oracle Database provides encryption for data in motion and data at rest. Native Network Encryption (NNE) and Transport Layer Security (TLS) protect data as it moves between servers, clients and the database, preventing eavesdropping and tampering. TLS 1.3, supported starting with Oracle AI Database 26ai, provides improved security and performance compared to TLS 1.2. Oracle AI Database 26ai supports the use of NIST-approved encryption algorithms and key exchange protocols, offering strong protection against both classical and quantum threats for its components.

Transparent Data Encryption (TDE) encrypts data at rest within the database. Tablespace encryption protects entire tablespaces, ensuring that data files, redo logs, and backups remain encrypted. TDE operates transparently to applications, requiring no code changes. Authorized users and applications see plaintext data, while attackers who bypass the database and access storage directly see only ciphertext.

TDE uses a two-tier key architecture. Data encryption keys (DEKs) encrypt tablespaces or columns. The master encryption key (MEK) encrypts the DEKs. This design simplifies key rotation: rotating the MEK only requires re-encrypting the much smaller set of DEKs, not the entire database.

Key management determines encryption effectiveness. Keys must be stored separately from encrypted data, rotated periodically, and backed up securely. Oracle Wallet provides local key management. Oracle Key Vault (OKV) centralizes key management, offering high availability, separation of duties, and operational resilience. OKV clusters support multi-master replication across geographically distributed nodes, ensuring key availability even during network or hardware failures.

For multicloud deployments, OKV enables true key portability. Unlike cloud-native key management services that lock keys to a single provider, OKV allows databases to move across Oracle Cloud Infrastructure, AWS, Azure, Google Cloud, and on-premises environments while maintaining consistent key access.

Key takeaways:

- Encryption protects data when attackers bypass authentication and access controls
- Oracle provides quantum-safe cryptography and protects both data at-rest and data in-motion using the latest NIST-approved algorithms
- Transparent Data Encryption encrypts data at rest transparently to applications
- Tablespace encryption protects data files, redo logs, and backups
- Two-tier key architecture simplifies key rotation
- Oracle Key Vault centralizes key management with high availability and multicloud support
- Proper key management is critical; without it, encryption provides false security

With preventive controls in place, security programs still need visibility into what happened, from privileged changes to sensitive data access. That leads to audit data strategy.

Chapter 12: Auditing and Monitoring

Prevention is important, but detection is essential. Chapter 12 explained that preventive controls should not be treated as perfect because business requirements often demand access to data. When prevention cannot be absolute, detective controls provide the next-best outcome: knowing when something improper occurred and having the evidence to prove it.

Copyright © 2026, Oracle and/or its affiliates

Unified Audit consolidates audit trail collection, replacing and updating the audit mechanisms of earlier Oracle Database versions. Audit policies define which actions to record: DDL statements (CREATE USER, GRANT ROLE), DCL statements (GRANT, REVOKE), DML statements (SELECT, INSERT, UPDATE, DELETE), and privileged operations. Predefined audit policies simplify initial configuration, while custom policies enable tailored monitoring.

Audit trails answer forensic questions: who performed the action, what did they do, when did it happen, where did they connect from, and which client program did they use. Audit records support incident investigation, compliance reporting, and behavioral analysis.

Oracle Audit Vault and Database Firewall (AVDF) and Oracle Data Safe centralize audit management across database fleets. AVDF collects audit data from multiple sources, stores it in a tamper-evident repository, and provides interactive reports and real-time alerting. Data Safe provides similar capabilities with cloud-based deployment. Both tools generate compliance reports aligned with regulatory requirements such as PCI DSS, HIPAA, and GDPR.

Database Firewall monitors all SQL traffic in real-time, identifying anomalies such as unusual query patterns, privilege escalation attempts, or SQL injection attacks. Unlike traditional network firewalls that operate at the packet level, Database Firewall understands SQL semantics and can detect application-layer attacks.

Key takeaways:

- Auditing provides detective controls when prevention cannot be absolute
- Unified Audit consolidates audit trail collection and simplifies policy management
- Audit trails support investigation, compliance reporting, and behavioral analysis
- Audit Vault and Database Firewall centralize audit management and provide real-time alerting
- Data Safe offers cloud-based audit management with fleet-wide visibility
- Database Firewall monitors SQL traffic and detects application-layer attacks

Chapter 13: Securing Agentic AI

Chapter 13 presents AI as a powerful productivity boost, while also noting that it can dramatically widen the security perimeter. It underscores a key risk: “valid credentials at machine speed” can quickly lead to overreach if guardrails are not strong.

AI-enabled development and operations move at a new tempo that can outstrip traditional review and governance. When AI is treated as a fast path to working code, teams can unknowingly accept opaque implementations that introduce injection vulnerabilities, flawed authentication, or data exfiltration pathways. Risk intensifies when AI can directly access enterprise data and systems, as valid credentials combined with machine speed can enable rapid overreach, unintended disclosure, or destructive actions before humans can intervene.

The narrative also warns that Shadow AI expands Shadow IT into a higher-risk domain, as unsanctioned tools can process data, generate insights, and make decisions without oversight. Examples include an analyst using an unapproved assistant to clean customer data, a developer letting an agent refactor authentication code, and a finance team submitting proprietary models for analysis—each creating pathways to privacy exposure, privilege escalation, or competitive intelligence loss. Another scenario describes an analyst linking an external language model to a corporate database via an unauthorized MCP server, pushing sensitive records and transactions beyond organizational control.

Rather than tool-by-tool bans that invite workarounds, the advice is to secure data at the source using database controls such as Database Vault, FGAC/RLS/VPD, OLS, and RAS so governance persists regardless of how AI connects.

Key takeaways:

Copyright © 2026, Oracle and/or its affiliates

- Agentic systems can increase the rate and volume of database actions, which raises the need for clear accountability
- Least-privilege and Privilege Analysis reduce overreach by limiting what automated accounts can do
- Runtime SQL controls (for example, SQL Firewall concepts) are positioned as guardrails that can log or block policy violations
- Data-layer controls such as VPD, OLS, RAS, and Data Redaction constrain what agents can retrieve

Chapter 14: Multicloud Security and Identity Management

Modern enterprise IT spans multiple clouds and on-premises environments. Chapter 14 addressed the challenge of maintaining consistent security controls when Oracle Databases run in Oracle Cloud Infrastructure, AWS, Azure, Google Cloud, and on-premises data centers simultaneously.

Managing security across this fragmented landscape calls for a disciplined focus on identity as the foundation of trust. A common recommendation is to designate a single provider to manage enterprise identities, so users benefit from one consistent login experience across all applications. Oracle AI Database supports this approach through native integration with major identity services such as Microsoft Entra ID and OCI IAM. Whether organizations synchronize directories or migrate fully to the cloud, well-designed authentication protocols extend enterprise security standards across every part of the multicloud ecosystem.

Operational consistency becomes essential when databases run at the same time in local data centers and multiple public clouds. Oracle Database features including TDE, Database Vault, Label Security, Data Masking, Unified Audit, and VPD operate identically regardless of where databases run. This consistency simplifies operations: security administrators do not need different skill sets or tools for databases in different clouds.

Oracle Database services such as Oracle Exadata Database Service and Oracle Autonomous Database are available across multiple clouds through Oracle Database@Multicloud. These services deliver the same capabilities whether deployed in Oracle Cloud Infrastructure or within hyperscaler data centers, with Oracle managing private network connections between environments.

Key takeaways:

- Multicloud deployments require consistent security controls across environments
- Oracle Database integrates with Oracle Cloud Infrastructure IAM and Microsoft Entra ID
- Identity federation enables single sign-on across clouds
- Core security features operate identically regardless of deployment location
- Oracle Database@Multicloud delivers consistent services across cloud providers

The next major step is to plan for ransomware and adopt a zero trust stance that assumes breach attempts will occur.

Chapter 15: Ransomware Resilience and Zero Trust

Chapter 15 brought together many threads from earlier chapters to address two contemporary challenges: Ransomware and Zero Trust architecture.

Ransomware has matured into a full-blown criminal enterprise, methodically combining credential theft, lateral movement, and multi-stage extortion to extract maximum financial damage from every breach. The emergence of Ransomware-as-a-Service (RaaS) platforms has lowered the barrier to entry, and AI-powered reconnaissance tools now allow attackers to map database architectures and achieve domain dominance in under an hour with no human intervention. Organizations that treat ransomware as a perimeter problem will find themselves dangerously under protected at exactly the wrong moment.

Copyright © 2026, Oracle and/or its affiliates

Phase 1 of a practical two-phase framework delivers immediate, measurable risk reduction. Transparent Data Encryption (TDE) renders exfiltrated data worthless, defusing double and triple extortion campaigns, while Oracle Key Vault (OKV) ensures encryption keys are managed separately from the data they protect, with full audit trails, automated key rotation, and portable key exports for isolated recovery environments. For organizations subject to PCI DSS, HIPAA, or frameworks aligned to NIST SP 800-57, off-server key management is not merely best practice; it is a compliance requirement.

Immutable backup infrastructure answers ransomware's deliberate campaign to eliminate recovery options. ZDLRA and ZRCV deploy retention locks that prevent any account, including compromised administrator accounts, from deleting or altering backups, and real-time transaction synchronization ensures all interconnected databases restore to exactly the same point in time. Dedicated recovery clean rooms, whether on-premises or in an isolated cloud environment, allow restored systems to serve the business while production infrastructure undergoes malware remediation.

With Phase 1 in place, Phase 2 builds lasting resilience through Zero Trust: continuous verification of identity and context, configuration drift monitoring via Oracle Data Safe and Audit Vault and Database Firewall, privilege minimization through Database Vault and privilege analysis, and activity monitoring that captures every privileged action and anomalous connection pattern. Zero Trust is best understood as a long-running program of iterative improvement, and with encryption and immutable recovery already operational, organizations can pursue that program from a far stronger starting position.

Key takeaways:

- Ransomware is rated the most disruptive cyber threat worldwide and has evolved to include double and triple extortion tactics that monetize stolen data even after victims pay.
- AI-powered ransomware now automates reconnaissance, adapts to defenses in real time, and can achieve domain dominance in under an hour, making static signature-based defenses insufficient.
- Transparent Data Encryption (TDE) neutralizes data theft as an extortion lever by rendering exfiltrated database files unintelligible to attackers who lack the encryption keys.
- Oracle Key Vault (OKV) must store TDE master encryption keys off-server, because co-locating keys with encrypted data allows attackers to permanently destroy an organization's access to its own data.
- Immutable backups through ZDLRA or ZRCV, with synchronized recovery points across interconnected databases, ensure ransomware cannot destroy recovery options or leave transaction gaps on restore.
- Zero Trust is a continuous program of iterative improvement, not a one-time deployment, and organizations that commit to it long-term build the most durable resilience against evolving ransomware tactics.

Chapter 16: Securing the Autonomous AI Database

Chapter 16 reviews security for the premium Oracle Autonomous AI Database (ADB) Cloud offering. For organizations evaluating or deploying ADB, security is built in from the start, not added after go-live. ADB launches with standardized, hardened configurations, patches without downtime, encrypts all data at rest and in transit by default, and continuously monitors for configuration drift through the included Oracle Data Safe service. This produces a baseline security posture that, by design, exceeds what most on-premises environments can sustain under routine operational pressures. It also reflects decades of Oracle experience running critical workloads for more than 400,000 customers across 175 countries and every major industry vertical.

From a risk management perspective, the architecture is strengthened by structural separation. ADB removes local file system and database node access entirely, and it deploys Database Vault Operations Control to prevent service administrators from accessing or altering customer data. It also reduces human error, a persistent driver of breaches, by automating large portions of database administration and limiting the commands any ADB user can run. These constraints are inherent to the service model, not compensating controls added later.

Copyright © 2026, Oracle and/or its affiliates

The shared responsibility model is central. Oracle maintains no visibility into customer data content, user roles, AI agent behaviors, or contextual signals (such as working hours, client programs, and normal work locations) used for access policy decisions. As a result, organizations must govern what ADB cannot govern on their behalf. Oracle Data Safe is the primary tool, enabling sensitive data discovery and classification, user privilege assessment, activity tracking, drift detection, and data masking for non-production use. Database Vault, Data Redaction, and Label Security further extend controls for more stringent regulatory or risk requirements.

Taken together, ADB shifts database security delivery by automating baseline configuration, patching, and encryption at scale, while supporting governance and policy through a flexible toolset. For security and IT leadership, the practical path is to use automation as the foundation and apply these tools where organizational context and human judgment are indispensable.

Key takeaways:

- Security is automated from day one
- Always-on encryption covers all data, everywhere
- Privileged-user risk is structurally addressed, not just monitored
- Human error is reduced through automation and command restrictions
- Shared responsibility requires active engagement of the surrounding toolset
- The security framework is built to scale with AI adoption

Connecting the Controls: How capabilities work together

The preceding summary highlights the breadth of Oracle Database security capabilities. Understanding how these capabilities interconnect transforms a collection of features into a coherent defense strategy.

Assessment informs everything else. Tools like DBSAT, Data Safe, AVDF, and DBLM identify which databases have weak password policies, overprivileged users, disabled auditing, or unencrypted tablespaces. These findings drive prioritized remediation. Without assessment, security efforts risk focusing on low-impact tasks while critical vulnerabilities remain unaddressed.

Sensitive data discovery guides protection decisions. Knowing where PII, financial data, and health information reside informs which tablespaces to encrypt, which columns to redact, which tables to protect with Database Vault realms, and which access to audit. Discovery ensures that protection investments align with data sensitivity and regulatory requirements.

Authentication establishes identity; authorization controls access. Strong authentication prevents attackers from impersonating legitimate users. Once identity is established, least-privilege authorization ensures users can only access data required for their job function. Fine-grained access controls (VPD, RAS, Label Security) further restrict access at the row and column level based on business logic and regulatory requirements.

Encryption provides defense-in-depth. Even if attackers operate outside the safety of the database, TDE ensures that stolen storage media containing database files, and backups remain unreadable. Key management with Oracle Key Vault ensures encryption remains effective by protecting keys, rotating them periodically, and preventing attackers from accessing keys alongside encrypted data.

Database Vault contains privilege misuse. Realms protect sensitive data from overprivileged users, including DBAs. Command rules and trusted path enforcement restrict high-privilege accounts so they can only be used under approved conditions. This limits the blast radius when powerful accounts are compromised.

Copyright © 2026, Oracle and/or its affiliates

SQL Firewall blocks unexpected commands. Even if attackers successfully authenticate and have legitimate privileges, SQL Firewall can block malicious statements that deviate from learned patterns. This provides application-layer protection without requiring code changes.

Auditing validates that controls work. Unified Audit records all security-relevant activity. Centralized audit management with AVDF or Data Safe enables fleet-wide visibility, real-time alerting, and compliance reporting. Audit trails support forensic investigation after incidents and demonstrate due diligence to auditors and regulators.

Privilege Analysis shrinks attack surface. By identifying unused privileges and roles, organizations can confidently revoke unnecessary access. This reduces the potential damage when accounts are compromised and simplifies compliance with least-privilege mandates.

These connections illustrate that Oracle Database security is not a menu of independent features but an integrated system. The most effective security postures deploy multiple capabilities together, creating overlapping layers of defense.

Copyright © 2026, Oracle and/or its affiliates

Closing thoughts: Security as an enabler

Security is often perceived as friction that slows down business operations, but the opposite is true when done well. Mature security programs enable organizations to move faster, take on new opportunities, and operate with confidence. As AI continues to accelerate at a remarkable pace, database security matters more than ever.

Customers increasingly factor security posture into purchasing decisions. A strong security track record reduces sales friction, accelerates due diligence, and strengthens win rates. Organizations that demonstrate robust data protection build trust with customers, partners, and regulators, creating competitive advantages in markets where trust is increasingly scarce.

Security also enables agility. Organizations with defensible security postures can migrate workloads to new clouds, adopt emerging technologies like Agentic AI, and respond to regulatory changes without scrambling to implement foundational controls at the last minute. Security becomes a strategic enabler rather than a blocking function.

The capabilities described in this book provide the tools needed to build resilient, defensible, and operationally efficient database security programs. Oracle Database security features integrate naturally into defense-in-depth strategies, align with Zero Trust principles, and scale from single databases to global fleets spanning multiple clouds.

The threat landscape will continue to evolve. Attackers will develop new techniques. Regulations will expand and tighten. Technologies like AI will introduce both new risks and new defensive capabilities. Organizations that treat database security as an ongoing discipline, continuously assessing posture, adapting controls, and learning from incidents, will be best positioned to protect their most valuable assets: data and the trust of those who depend on it.

The journey does not end here. It begins with a single assessment, a single remediation, a single improvement to configuration or policy. From there, the path forward becomes clearer, the benefits become tangible, and the organization builds security resilience that compounds over time.

Databases concentrate value, and attackers know it. Protecting that value is not optional. With the knowledge, tools, and roadmap provided in this book, organizations have what they need to secure their data, meet regulatory obligations, and enable the business to operate with confidence.

Copyright © 2026, Oracle and/or its affiliates

Appendix: Tools – Features, Options, Products, and Packs

Any discussion like this can quickly turn into a long list of products and features. We've tried to keep the main body focused on what you need to do and why it matters.

Below, you'll find the features, options, products, and packs mentioned in this guide, along with links to the relevant documentation. Additional features are listed here.

One more thing to keep in mind about these links: they were current when we wrote this section, but technical briefs can stay in circulation for a long time. Database releases and documentation change more frequently, so you should check docs.oracle.com for the latest version. The links here still help you identify the right manual (and often the specific chapter or appendix) to use as your starting point.

Database Security Assessment Tool (DBSAT)

DBSAT is a standalone utility included with your database support. DBSAT helps with security assessment, user assessment, and sensitive data discovery. There is no additional fee for the use of DBSAT. DBSAT works with all supported versions of the Oracle AI Database, on all supported operating systems, and can be run for databases on-premises or in the cloud – including non-Oracle clouds. The utility may be downloaded from My Oracle Support – check MOS note 2138254.1 for more information on downloading the tool. DBSAT documentation is available here.

Oracle Data Safe

Data Safe is an Oracle Cloud service, included with Oracle AI Database as a Service offerings and available for use with on-premises databases and Oracle AI Databases running on Oracle Cloud Compute. Data Safe includes several database security capabilities, including security assessment, user assessment, sensitive data discovery, sensitive data masking, unified audit policy control, audit data retrieval, reporting, and alert generation. Data Safe is Oracle's newest database security service and is rapidly evolving (two-week development sprints) new features and capabilities. Documentation for Data Safe is included here.

Enterprise Manager Database Lifecycle Management

Database Lifecycle Management is a management pack for Oracle Enterprise Manager Cloud Control. Database Lifecycle Management provides numerous functions for managing the lifecycle of your database, including configuration management, and can play a valuable part in security assessments. Information about configuration management using Enterprise Manager Database Lifecycle Management is available here.

Privilege Analysis

Privilege analysis is a database feature included with all databases and database services except for standard edition. It helps with user assessment, particularly with determining which privileges a user account has, but is not using. Privilege Analysis was introduced in Oracle AI Database 12c Release 1. Privilege Analysis documentation is available here.

Native Network Encryption

Native Network Encryption (NNE) is a database feature included with all databases and database services with the exception of Autonomous Database (Autonomous Database uses TLS instead of NNE). NNE encrypts data as it travels between the database and database client or application. NNE documentation is available here.

Copyright © 2026, Oracle and/or its affiliates

Transport Layer Security

Transport Layer Security (TLS) is a database feature included with all databases and database services. It is configured by default for Autonomous Databases. TLS encrypts data as it travels between database and database client or application. TLS documentation is available here.

Centrally Managed Users

Centrally Managed Users (CMU) is a database feature included with Oracle AI Database Enterprise Edition. CMU was introduced with Oracle AI Database 18c. CMU allows Oracle AI Databases to connect directly to Microsoft Active Directory. With CMU, users are created in Active Directory and mapped to database schemas. Optionally, database roles can be associated with Active Directory groups, and database role membership controlled by Active Directory group membership. CMU documentation is available here.

Enterprise User Security

Enterprise User Security (EUS) is a database feature included with Oracle AI Database Enterprise Edition. EUS was introduced with Oracle AI Database 8.1 (Oracle 8i). EUS allows Oracle AI Databases to connect to an Oracle directory service. With EUS, users are created in Internet Directory and database schemas are mapped to users or to collections of users within an LDAP organizational unit. Database roles can be associated with Internet Directory groups, and database role membership controlled by LDAP group membership. EUS documentation is available here.

Traditional Auditing

Traditional auditing was first introduced in Oracle AI Database 7 and was the primary auditing mechanism for Oracle AI Database until the release of Oracle AI Database 12c. Traditional auditing is being replaced by unified audit. Traditional auditing is desupported starting with Oracle AI Database 26ai. Traditional audit documentation is available here.

Fine-Grained Auditing

Fine-grained auditing was introduced in Oracle AI Database 9.0 (9i Release 1). As the name suggests, Fine-grained auditing allows audit policies to be focused more narrowly than traditional auditing, allowing audit policies based on columns. Fine-grained auditing also introduced the concept of conditional auditing, where audit records would only be generated if a certain condition evaluated to true. Documentation for fine-grained auditing is available here.

Unified Auditing

Unified auditing was introduced in Oracle AI Database 12.1 (12c Release 1). Unified auditing consolidates audit records into a single location, combining audit data from Database Vault, Label Security, Data Pump, SQL*Loader, Recovery Manager (RMAN), fine-grained auditing, and audit records generated from unified audit policies. Unlike traditional auditing, unified audit policies may be conditional, may choose to audit only top-level statements, and are extensible to include context information not in the default audit trail. Documentation for unified auditing is available here.

Enterprise Manager Data Discovery

Data Discovery is available in Enterprise Manager. Data discovery scans databases to locate sensitive data and helps drive Data Masking and Subsetting (DMS), Audit Vault and Database Firewall (AVDF) sensitive data audit reporting, and Transparent Sensitive Data Protection (TSDP) policies. Enterprise Manager Data Discovery stores the list of applications, tables, and relationships between table columns that are either declared in the data dictionary, imported from application metadata, or user specified. Data Discovery maintains sensitive data types and their associated

Copyright © 2026, Oracle and/or its affiliates

columns, and is used by test data operations, such as data subsetting and data masking, to securely produce test data. Like Data Safe, Data Discovery scans data contained within tables to find sensitive data.

Data Discovery is included at no additional cost with Oracle Advanced Security, Oracle AI Database Vault, Oracle Label Security, Oracle Data Masking and Subsetting, and Oracle Audit Vault and Database Firewall. Use Data Discovery to understand a database schema or schemas, including how columns relate to one another and which columns contain sensitive data. Documentation for Data Discovery is available here.

Oracle Advanced Security

Oracle Advanced Security (ASO) is a database option that includes Transparent Data Encryption, RMAN backup encryption, Data Pump export encryption, encrypted Database File System (DBFS), encrypted SecureFile LOBs, and Data Redaction. Advanced Security is one of the oldest database options, tracing its roots to the Advanced Networking Option introduced in Oracle 7. Documentation for Advanced Security's Transparent Data Encryption is available here. Documentation for Advanced Security's Data Redaction is available here.

Oracle Key Vault

Oracle Key Vault is a key management system supporting the Oracle infrastructure, optimized for use with Transparent Data Encryption. Key Vault provides continuous access to encryption keys with a multi-master fault-tolerant cluster architecture. In addition to protecting encryption keys, Key Vault also provides complete SSH key governance (an important capability for protecting all of your Linux servers, not just the database servers), secrets management, DBMS_CRYPTO key management, and much more. Documentation for Key Vault is available here.

SQL Firewall

SQL Firewall is a new component of Oracle AI Database 26ai. It is designed to detect and block SQL injection and other anomalies. SQL Firewall is included with both Oracle AI Database Vault and Oracle Audit Vault and Database Firewall. Documentation for SQL Firewall is available here.

Oracle AI Database Vault

Oracle AI Database Vault is a database option that provides advanced access control capabilities. Database Vault is commonly used to enforce separation of duties, block administrator access to sensitive data, and enforce trusted path access to data. Database Vault includes SQL Firewall. Documentation for Database Vault is available here.

Oracle Audit Vault and Database Firewall

Oracle Audit Vault and Database Firewall (AVDF) is a database activity monitor with heterogeneous capabilities – covering Oracle AI Database, Oracle MySQL, Microsoft SQL Server, PostgreSQL, MongoDB, IBM DB2, and SAP Sybase. For Oracle AI Databases, AVDF also provides security assessment, tracks user entitlements, performs drift detection, discovers sensitive data, and includes SQL Firewall. Documentation for Audit Vault and Database Firewall is available here.

Oracle Data Masking and Subsetting

Oracle Data Masking and Subsetting (DMS) is a management pack for Oracle Enterprise Manager. DMS removes risk from databases by replacing sensitive data with artificial values. DMS can also be used to create subsets of a database – smaller copies with only a portion of the original data. Documentation for Data Masking and Subsetting is available here.

Copyright © 2026, Oracle and/or its affiliates