Oracle Cloud Infrastructure Security

The Practice Workshop

Martin Berger
Stefan Oehrli

Martin Berger – Modern Data Platforms

Modern Data Platforms

stefan.oehrli@accenture.com

Oracle ACE Pro

Tech Architecture Manager

- Since 1998 working in IT
- More than 22 years of experience in Oracle databases
- Kestenholz / Jurasüdfuss / Switzerland
- 2 Junior-DBAs @ Home (8yrs & 10yrs)
- Firefighter & E-Biker
- Loves his companies' cultural values:
- curiosity, doers, network, space, and together
- Focus: Let’s go into the Cloud!
- Co-author of the book The Oracle DBA (Hanser, 2016/07)

martinberger.com

@martinberger_ch

Stefan Oehrli – Modern Data Platforms

Modern Data Platforms

stefan.oehrli@accenture.com

dbsym42

Oracle ACE Director

Terraform ASSOCIATE – HashiCorp CERTIFIED

Oracle ACE Director 2024

Tech Architecture Manager

- Since 1997 active in various IT areas
- More than 25 years of experience in Oracle databases
- Focus: Protecting data and operating databases securely
  - Security assessments and reviews
  - Database security concepts and their implementation
  - Oracle Backup & Recovery concepts and troubleshooting
  - Oracle Enterprise User and Advanced Security, DB Vault, …
  - Oracle Directory Services
- Co-author of the book The Oracle DBA (Hanser, 2016/07)

oradba.ch

@stefanoehrli

Oracle ACE

The Oracle ACE Program

400+ technical experts helping peers globally

- The Oracle ACE Program recognizes and rewards community members for their technical and community contributions to the Oracle community
- 3 membership levels: Director, Pro, and Associate
- Nominate yourself or a colleague at ace.oracle.com/nominate
- Learn more at ace.oracle.com

aceprogram_ww@oracle.com
Facebook.com/OracleACEs
@oracleace
Oracle ACE Program Group

Modern Data Platforms
VISON & MISSION

WHY? We are the game changer for our client's data platform projects

HOW? Maximum automation, maximum efficiency, maximum quality!

WHAT? We build innovative data platforms based on our blueprints and licensable assets and tools.

3 key benefits

1 Architecture expertise from hands-on projects

2 Delivery of tailor-made data platforms

3 Integrated Teams Like a rowing team, perfect alignment and interaction.

Tools and Blueprints

Key enabler for the implementation of modern data platforms at a high speed and quality.

Continuous Optimization

Tools and Blueprints are continuously optimized to the customer and project’s needs.

Expertise & Light Towers

Expert group for modern data platforms from technical implementation to project management and organization

OCI Security

Hands-On Practice Workshop

Course Schedule

Session Times and Breaks

from	to	Topic	
09:00	09:15	**Welcome**	Course Overview and Lab Setup Instructions
09:15	10:30	**Basic OCI Security**	In-Transit Encryption, Shielded Instances, Key Management
10:30	10:45	Coffee Break	
10:45	12:00	**Cloud Guard**	Overview, Recipes, CIS Scans, Remediation, Alerts, and Events
12:00	13:00	Lunch	
13:00	14:15	**Data Safe**	Overview, Setup, Auditing, Data Masking, Assessments,…
14:15	14:45	Coffee Break	
14:45	16:00	**Security Zones / WAF**	Overview, Zone Management, Rulesets, and CIS Compliance
16:00	16:30	**Wrap-Up**	Additional Resources, Next Steps, and Farewell

Security Service Products

Detection and Remediation


Cloud Guard


Maximum Security Zone


Security Advisor


Vulnerability Scanning

Data Protection


Vault Key Management


Vault Secrets Management


Data Safe

Certificates

OS and Workload Protection


Shielded Instances


Dedicated Host


Bastion


OS Management

IAM


IAM


MFA


Federation


Audit

Infrastructure Protection


DDoS Protection


Web Application Firewall


Security Lists / NSG


Network Firewall Appliance

Working Together

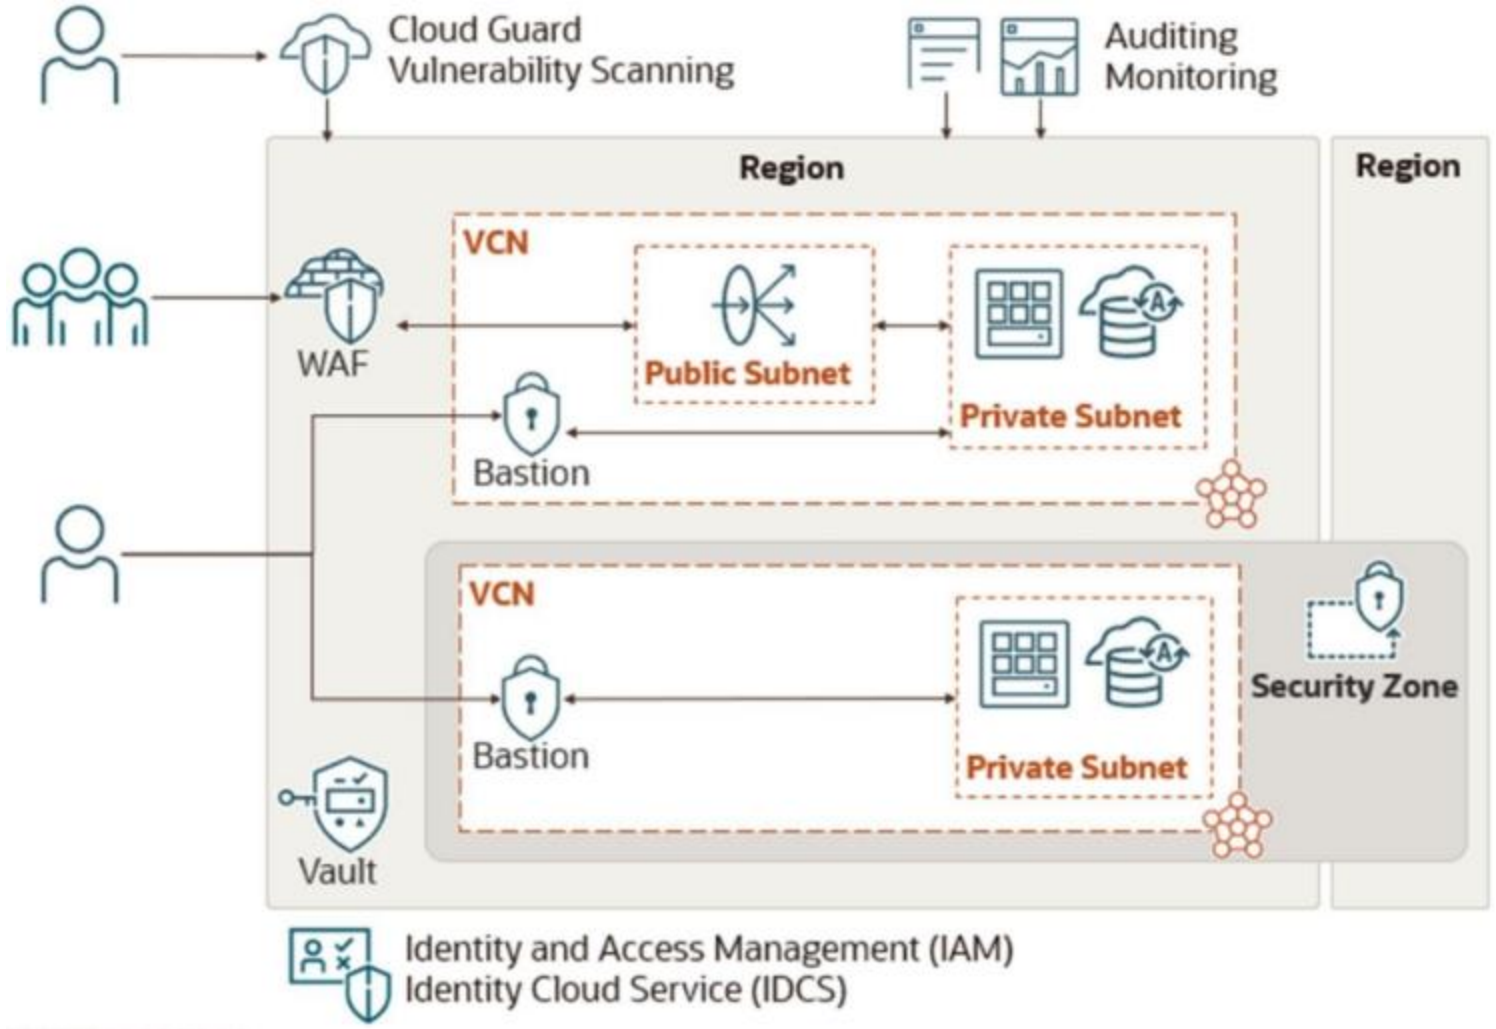

Source: oracle.com

OCI Lab Architecture Overview

A High-Level View of the Infrastructure

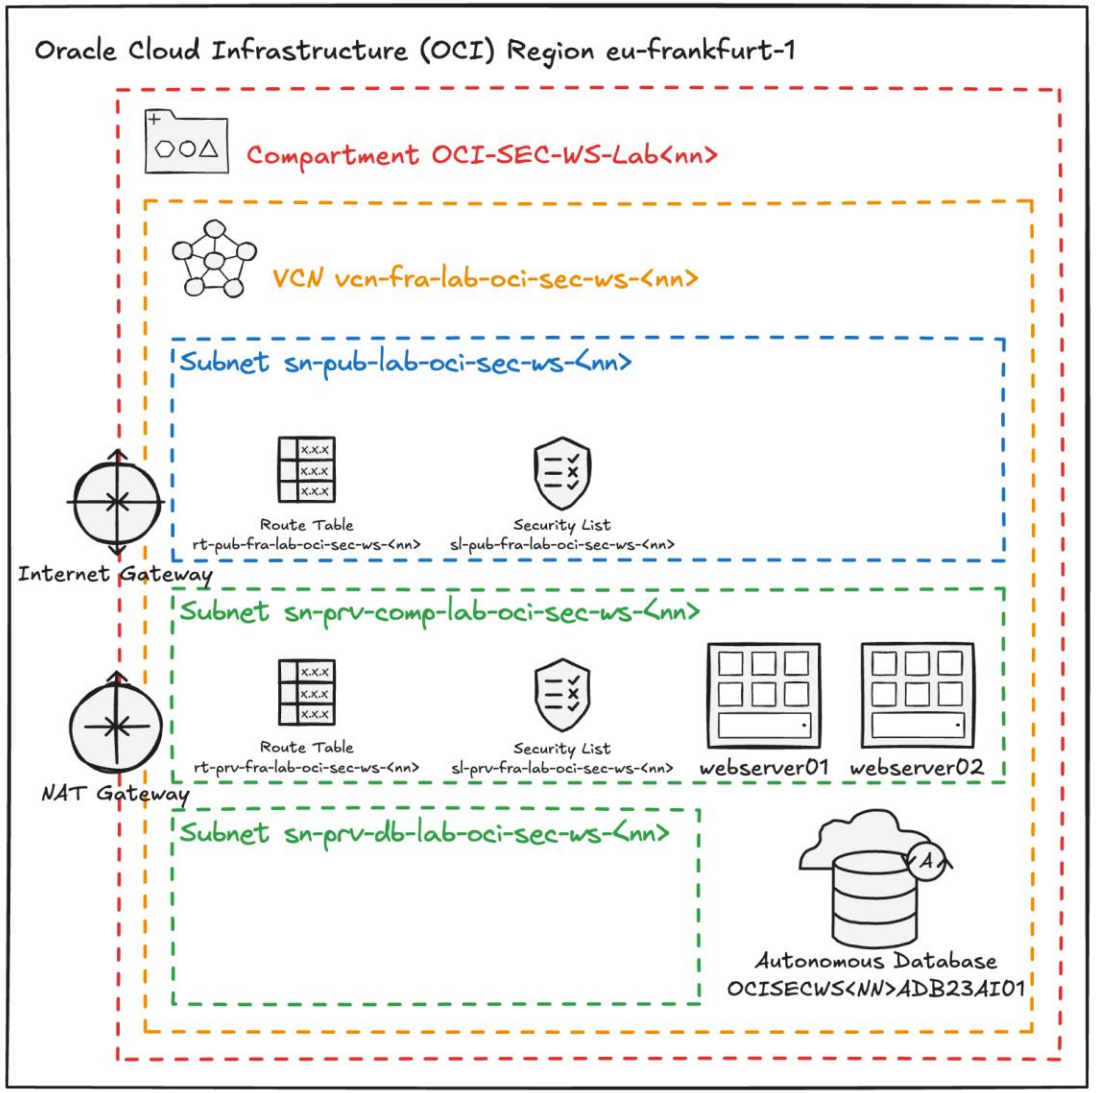

OCI Access Information

Resources and Cloud Environment Details

The course materials and exercises are available via the GitHub repository/website:

- OCI Walkthrough: https://code.oradba.ch/oci-sec-ws
- PDFs and Course Materials: https://code.oradba.ch/oci-sec-ws/others

For the workshop, each participant will have access to Oracle Cloud. Accenture is providing an environment with resources for the training day.

- **URL**: http://cloud.oracle.com
- **Tenant**: Provided by the instructor
- **User**: *lab-oci-sec-ws**NN***
- **Password**: Provided by the instructor
- **Compartment**: *OCI-SEC-WS-LAB-**NN***

Oracle Cloud Infrastructure Security

Basics

Martin Berger
Stefan Oehrli

Basics

Security for free.

1. Key Management
2. OS Management
3. Vulnerability Scanning
4. Shielded Instances
5. Hands-On Labs

1

Key Management

Security for free.

Encryption Basics

ENCRYPTION

Transform plain text into cipher text.

DECRYPTION

Transform cipher text into plain text.

Encryption Basics

Encryption at rest

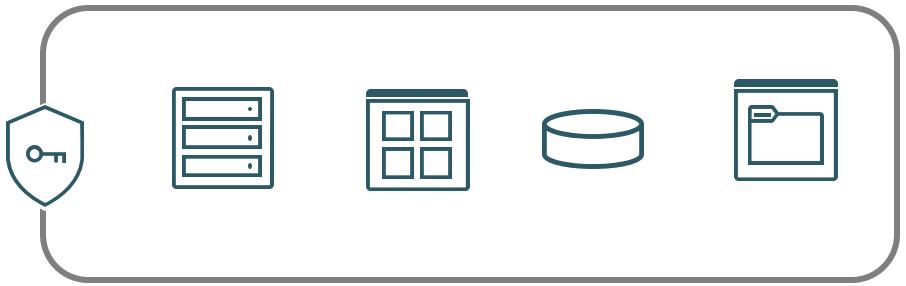

Encryption in transit

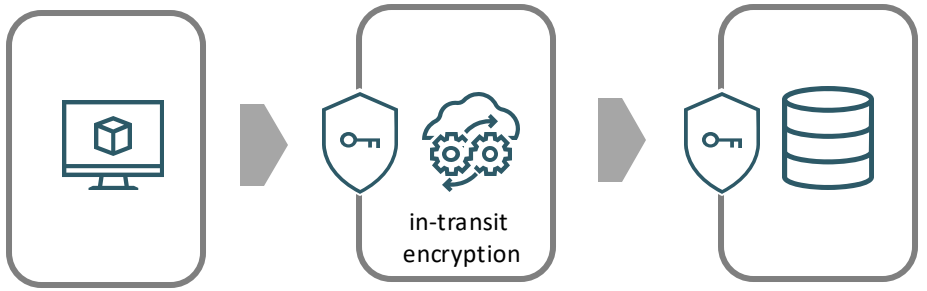

Protect stored, static data	**Purpose**	Protect data moving from one location to the other, through a private network.
Keeps data encrypted until a person or system provides a decryption key to access to the data	**How does it work**	Encrypts messages before transmission and decrypts upon arrival at the destination

Key Management Service

- The Oracle Cloud Infrastructure (OCI) Key Management Service (KMS) is a cloud-based service that provides centralized management and control of encryption keys for data stored in OCI. It simplifies key management by centrally storing and managing encryption keys, protecting data at rest and in transit by supporting various encryption key types, including symmetric and asymmetric keys.
- OCI KMS addresses security and compliance requirements by giving you more control over your encryption keys. You can bring your own keys (BYOK) to OCI, create them within OCI, or hold your own keys (HYOK) externally. Additionally, you can use FIPS 140-2 Level 3-certified hardware security modules (HSMs) to store and protect your encryption keys.

Example Use Case:
Block Volume Encryption

Key and Secret Management Concepts

- **OCI Vault**
  - A customer-managed encryption service that enables you to control the keys that are hosted in Oracle Cloud Infrastructure (OCI) hardware security modules (HSMs) while Oracle administers the Hardware Security Module HSM.
- **OCI Dedicated KMS**
  - A single-tenant HSM partition as a service that provides a fully isolated environment for storing and managing encryption keys. You can control and claim ownership of the HSM partitions and use standard interfaces, such as PKCS#11, to perform cryptographic operations.
- **OCI External KMS**
  - Enables you to use your own third-party key management system to protect data in OCI services. You control the keys and HSMs outside OCI, and you're responsible for the administration and manageability of those HSMs.

HSM: The OCI Vault Hardware Security Module (HSM) provides secure key management and cryptographic operations by using dedicated hardware devices to protect sensitive data and encryption keys.

Vault

- A managed service in Oracle Cloud Infrastructure that provides secure key management
- There are two different types available: **Virtual private vault** and **Vault in a shared Partition**.
- **Virtual private vault**:
  - creates the vault as a dedicated partition on the HSM
  - pricing based on the maximum usage against key limits
  - isolated cryptographic resources and enhanced security controls
  - greater performance

- **Vault in a shared partition:**
  - free for use
  - cryptographic resources are shared among multiple tenants

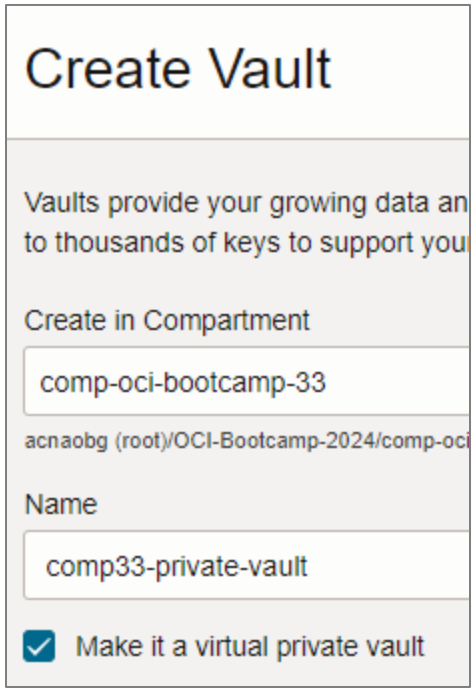

Master Encryption Key

- The Master Encryption Key in OCI Vault is the primary key used to encrypt and protect other keys and secrets, ensuring their secure management and storage within Oracle Cloud Infrastructure.

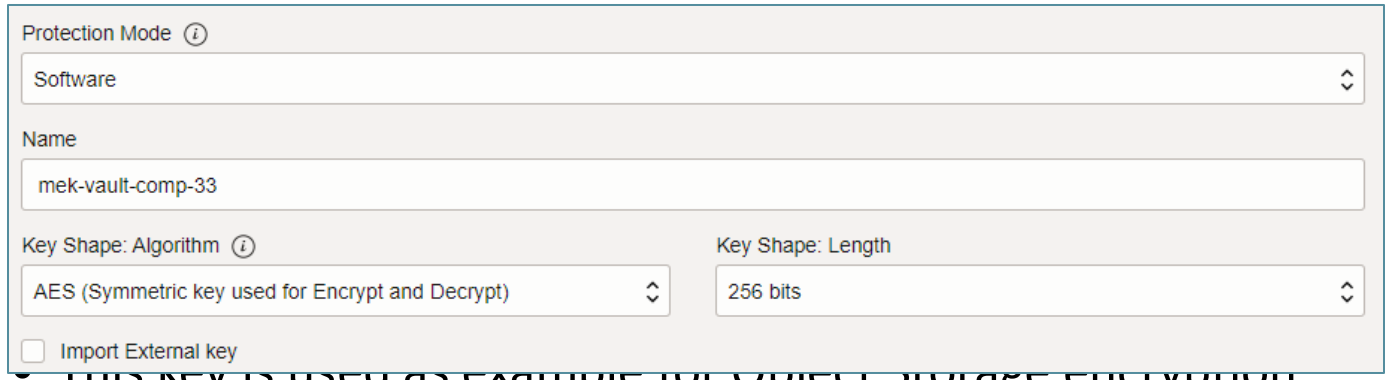

You can choose between **Software** and **HSM** protection mode, with different costs.

- This key is used as example for Object Storage encryption

CIS – Center of Internet Security OCI Benchmark

5.2.1 Ensure Block Volumes are encrypted with Customer Managed Keys (CMK). (Automated)

Description:

Oracle Cloud Infrastructure Block Volume service lets you dynamically provision and manage block storage volumes. By default, the Oracle service manages the keys that encrypt block volumes. Block Volumes can also be encrypted using a customer managed key.

Terminated Block Volumes cannot be recovered and any data on a terminated volume is permanently lost. However, Block Volumes can exist in a terminated state within the OCI Portal and CLI for some time after deleting. As such, any Block Volumes in this state should not be considered when assessing this policy.

Rationale:

Encryption of block volumes provides an additional level of security for your data. Management of encryption keys is critical to protecting and accessing protected data. Customers should identify block volumes encrypted with Oracle service managed keys in order to determine if they want to manage the keys for certain volumes and then apply their own key lifecycle management to the selected block volumes.

Impact:

Encrypting with a Customer Managed Key requires a Vault and a Customer Master Key. In addition, you must authorize the Block Volume service to use the keys you create. Required IAM Policy:

```
Allow service blockstorage to use keys in compartment <compartment-id> where
target.key.id = '<key_OCID>'
```

Can I rotate my keys?

Yes. You can regularly rotate your keys in alignment with your security governance and regulatory compliance needs or do it ad hoc in case of a security incident.

Regularly rotating your keys (for example, every 90 days) by using the console, API, or CLI, limits the amount of data protected by a single key.

Note: Rotating a key does not automatically re-encrypt data that was previously encrypted with the old key version; this data is re-encrypted the next time it's modified by the customer.

If you suspect that a key has been compromised, you should re-encrypt all data protected by that key and disable the prior key version.

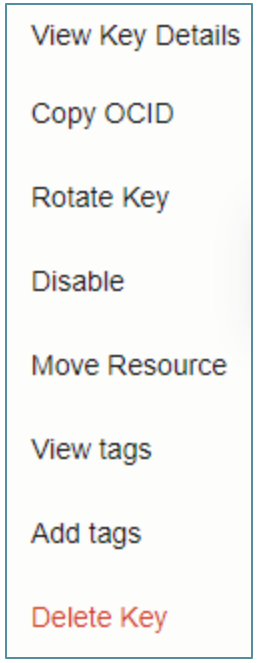

Backup and Replicate Vault and Keys

Backup:

- Backing up a private vault ensures that cryptographic keys and sensitive data can be restored in case of accidental deletion, corruption, or system failure.
- Regular backups help meet regulatory requirements for data protection and enable thorough audits of key management practices.
- Maintaining backups supports uninterrupted operations by allowing quick recovery of essential cryptographic resources, minimizing downtime in the event of an incident.

Replication:

- Replicate virtual private vaults across regions to meet compliance requirements or improve latency.
- Cross-region replication automatically syncs creation, deletion, updates, or moves of keys and key versions between source and destination vaults.
- The originating vault is called the source vault, and the receiving vault in the destination region is known as the vault replica.

Minimum Level: Private Vault

Secrets

- OCI Vault Secrets provides a secure location to store sensitive information such as passwords, API keys, and confidential data.
- It ensures that secrets are accessible only to authorized applications and users through fine-grained access policies.
- It facilitates automated secret management, including versioning and rotation, to enhance security and compliance.
- Example: Store the database credentials as a secret in OCI Vault. The application can securely retrieve these credentials at runtime, ensuring that sensitive information is not exposed in the codebase or configuration files.

```
curl -s -X GET \
-H "Authorization: Bearer <your_auth_token>" \
-H "Content-Type: application/json" \
"https://<region>.secrets.vaults.<domain>/20190301/secretBundles/<secret_ocid>" | jq -r
'.secret-bundle-content.content' | base64 --decode
```

Secret Creation

- Based on Master Encryption Key
- Different secret types like passphrase, SSH key etc.
- Auto-generation or manual input

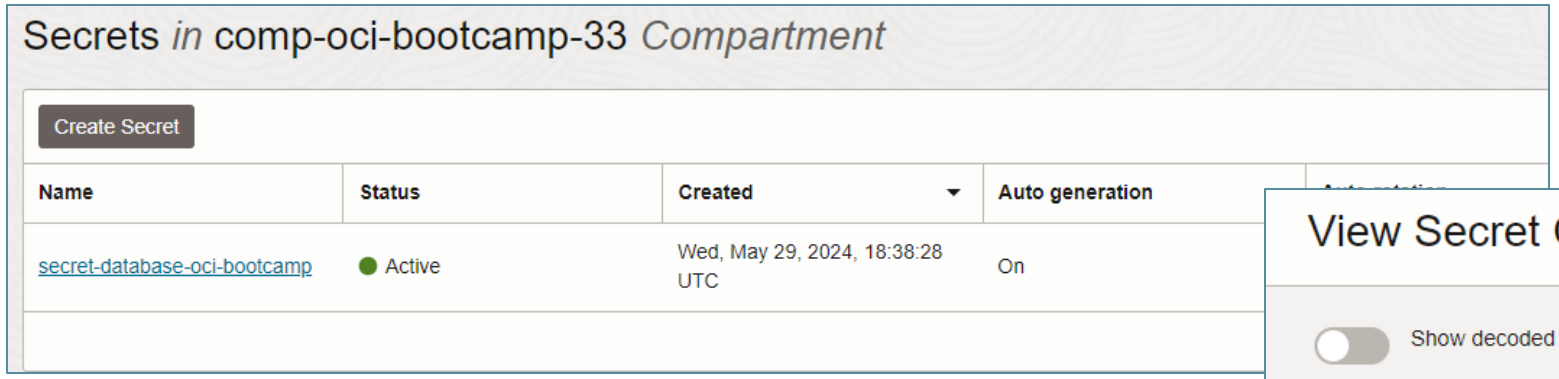

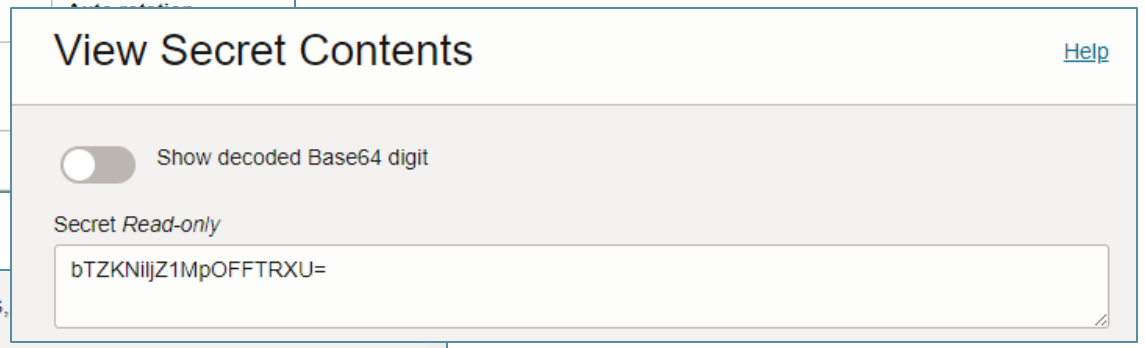

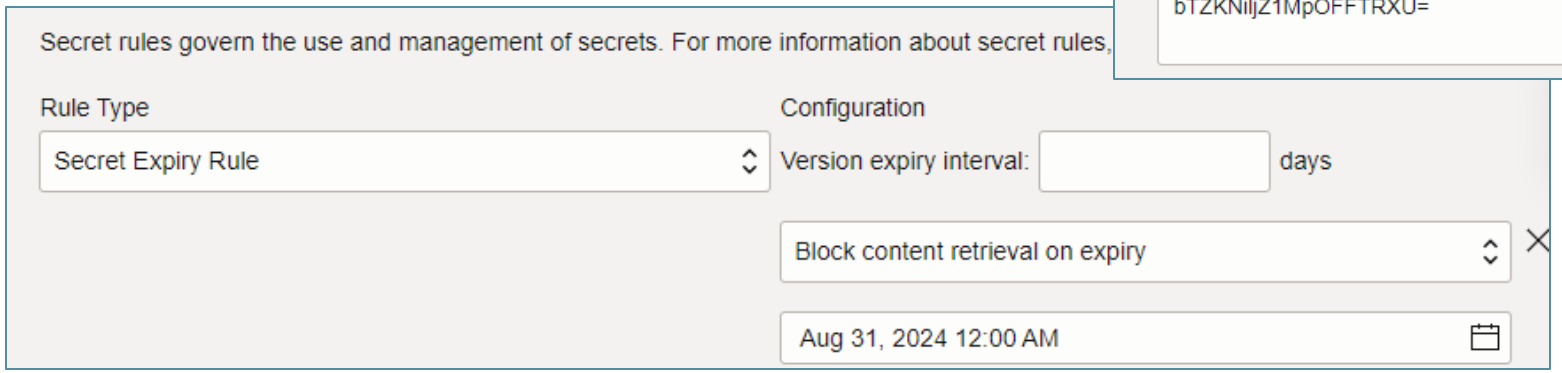

Deleting a Vault

- When you delete a vault, the vault and all its associated keys go into a pending deletion state until the waiting period expires.
- By default, the waiting period is set to 30 days, but it can be adjusted from a minimum of 7 days to a maximum of 30 days.
- When a vault is deleted, all its associated keys are also deleted.
- If replication is configured, deleting a vault in the source region also deletes the vault and any keys in the vault in the destination region.

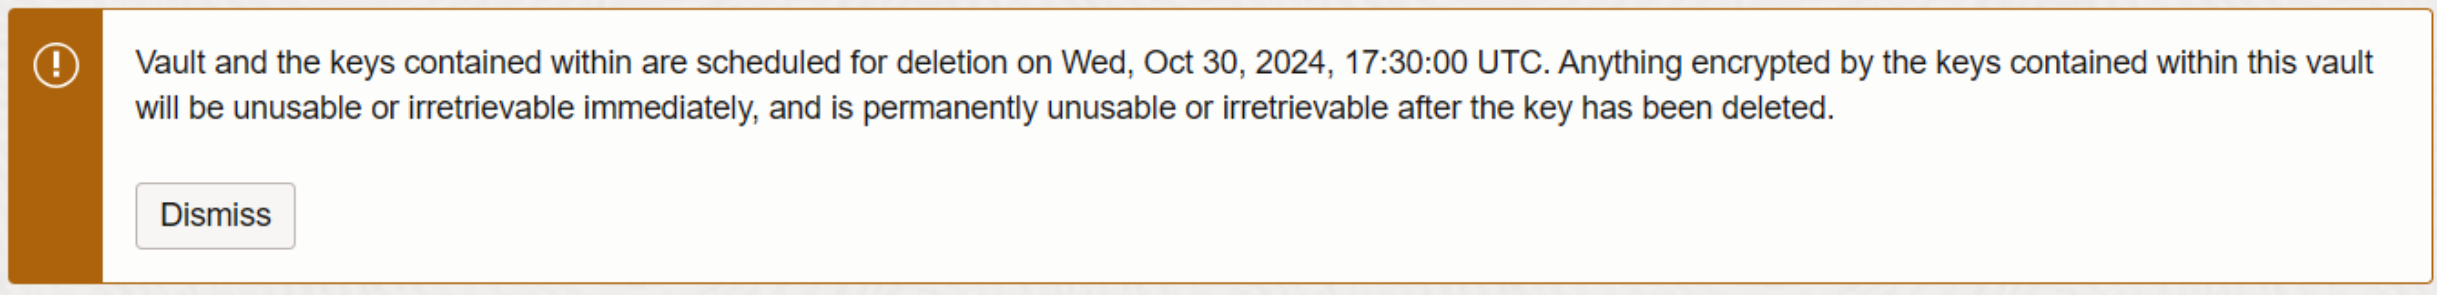

2

OS Management

Security for free.

OS Management Hub

- Next generation of OS management functionality
- OS Management Hub uses an agent plugin for managing and applying updates.
- Simplified OS Management:
  - Oracle OS Management Hub streamlines the management and monitoring of OS updates across on-premises, OCI, and supported third-party clouds.
- Inclusive Service:
  - OS Management Hub is included with Oracle Linux Support and OCI Compute subscriptions, providing comprehensive OS management at no additional cost.
- Auto-Patch capable:
  - Keep your OS up to date

OS Management Hub Architecture

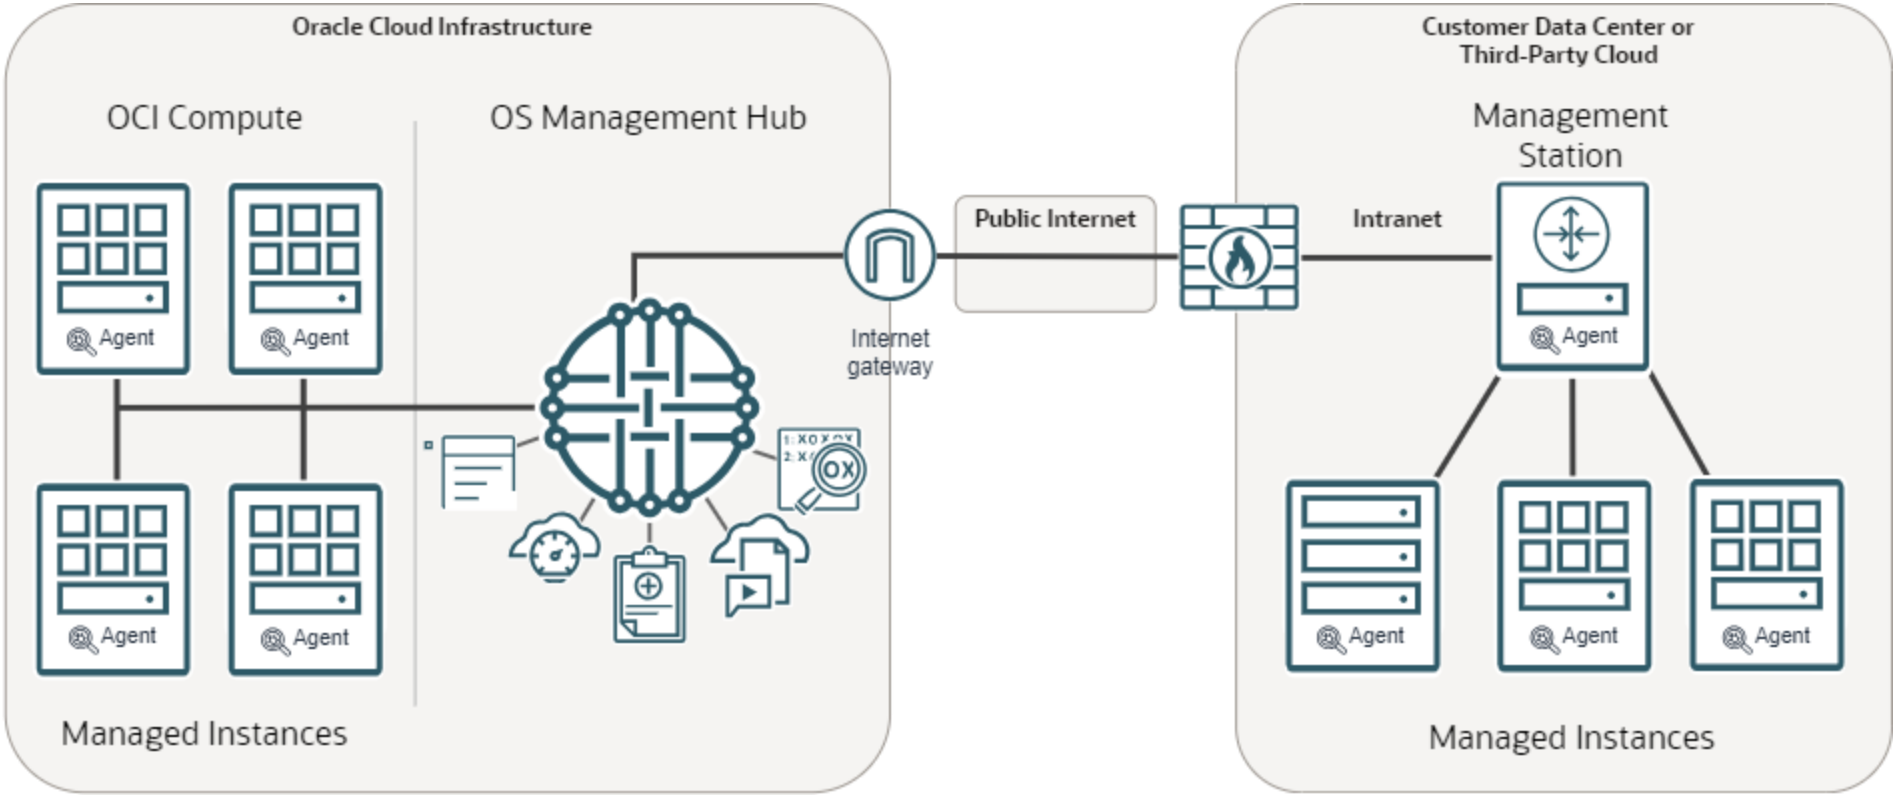

Source: oracle.com

OS Management Workflow 1/2

1. Adding Vendor Software Sources (create in root compartment first for later usage)
2. Create custom Software Source on compartment level
3. Create Service Profile
4. Verify OS Management Hub plugin is running on target instances

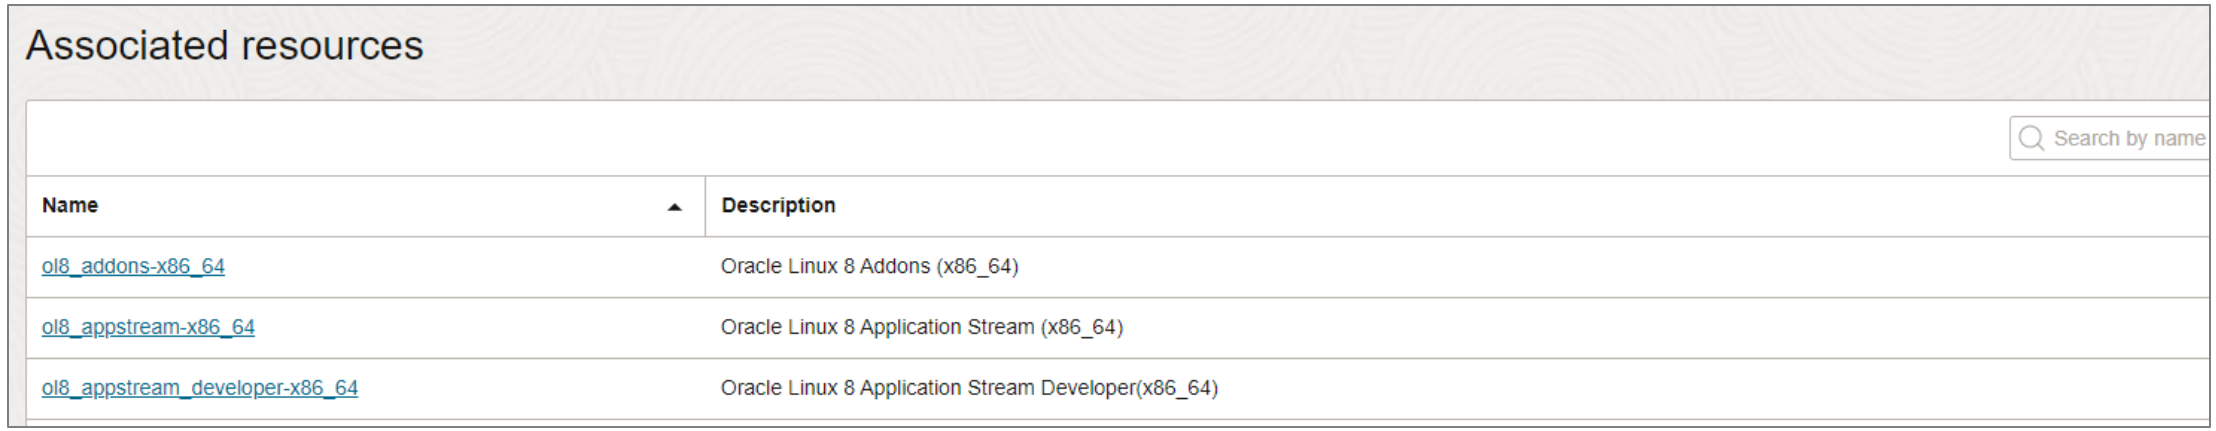

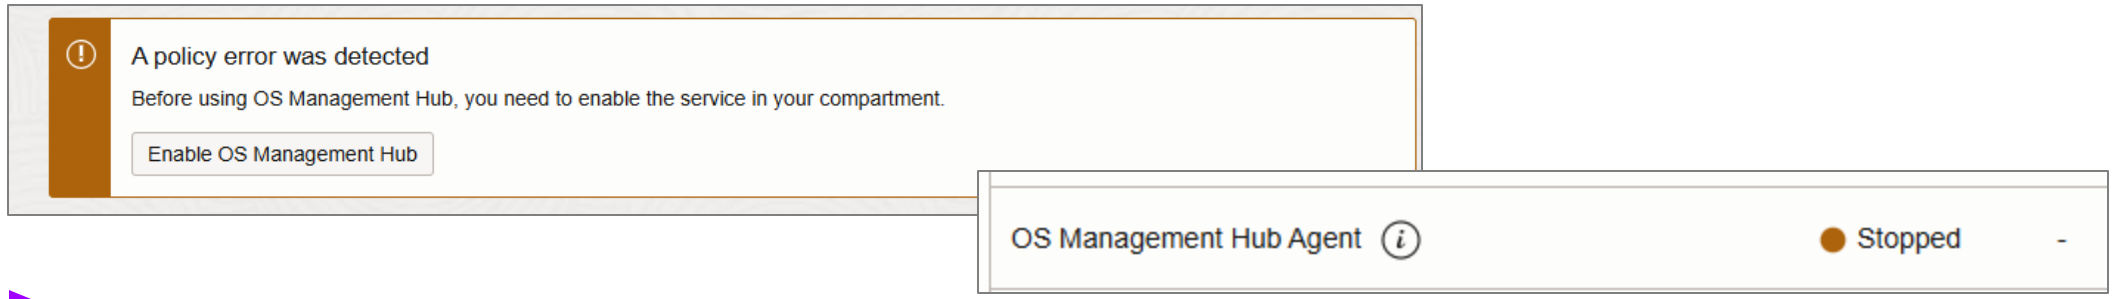

OS Management Workflow 2/2

5. Wait 10 mins..
6. (optional) create group with custom software source, attach instances

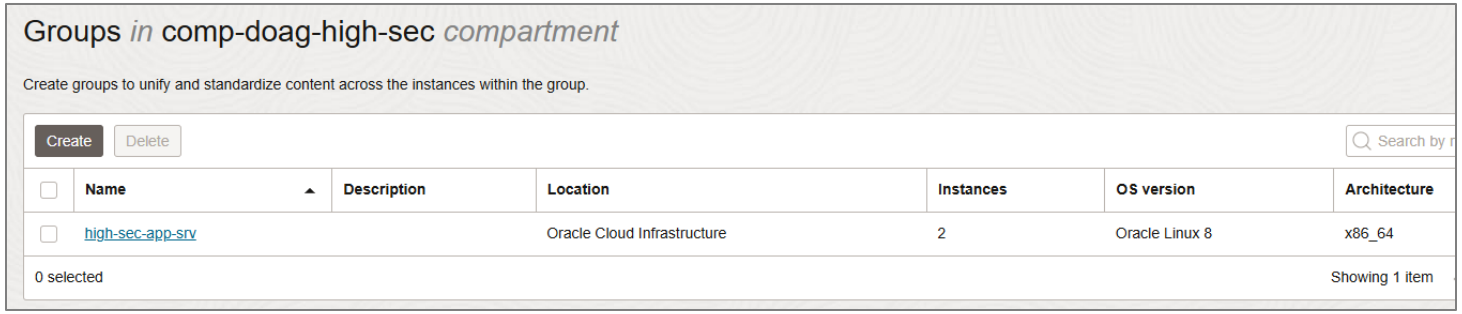

7. Verify Dashboard and Reports

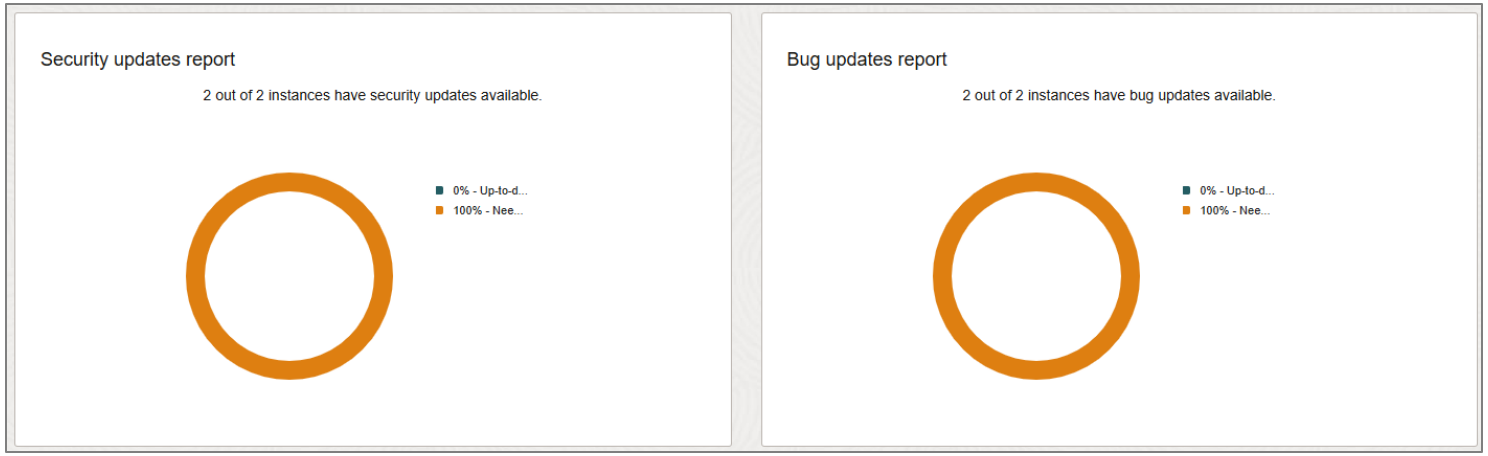

Verify OS Management & Categories

- On target hosts, any DNF/YUM command shows the enabled service

```
[opc@compute-high-sec-public ~]$ sudo dnf update
This system is receiving updates from OSMH.
Last metadata expiration check: 0:01:54 ago on Thu 17 Oct 2024 07:25:37 AM GMT.
Dependencies resolved.
```

OS Management Hub Category	Description
Security	An update that addresses security vulnerabilities found during development, testing, or reported by users. Security fixes usu ally have one or more associated CVE (Common Vulnerabilities and Exposure) names to identify the vulnerabilities.
Ksplice	An update used by Ksplice for installing zero-downtime security patches. The update job can include only Ksplice kernel updates, only Ksplice userspace updates, or both. See Using Ksplice for Oracle Linux for how to configure an instance to receive Ksplice updates.
Bug Fix	An update that fixes issues reported by users or discovered during development or testing.
Enhancement	An update that introduces new features, improved functionality, or enhanced performance in the package's software.
Other	An update that's not associated with any errata.

Reports

- You can apply the patches, fixes automatically or manually

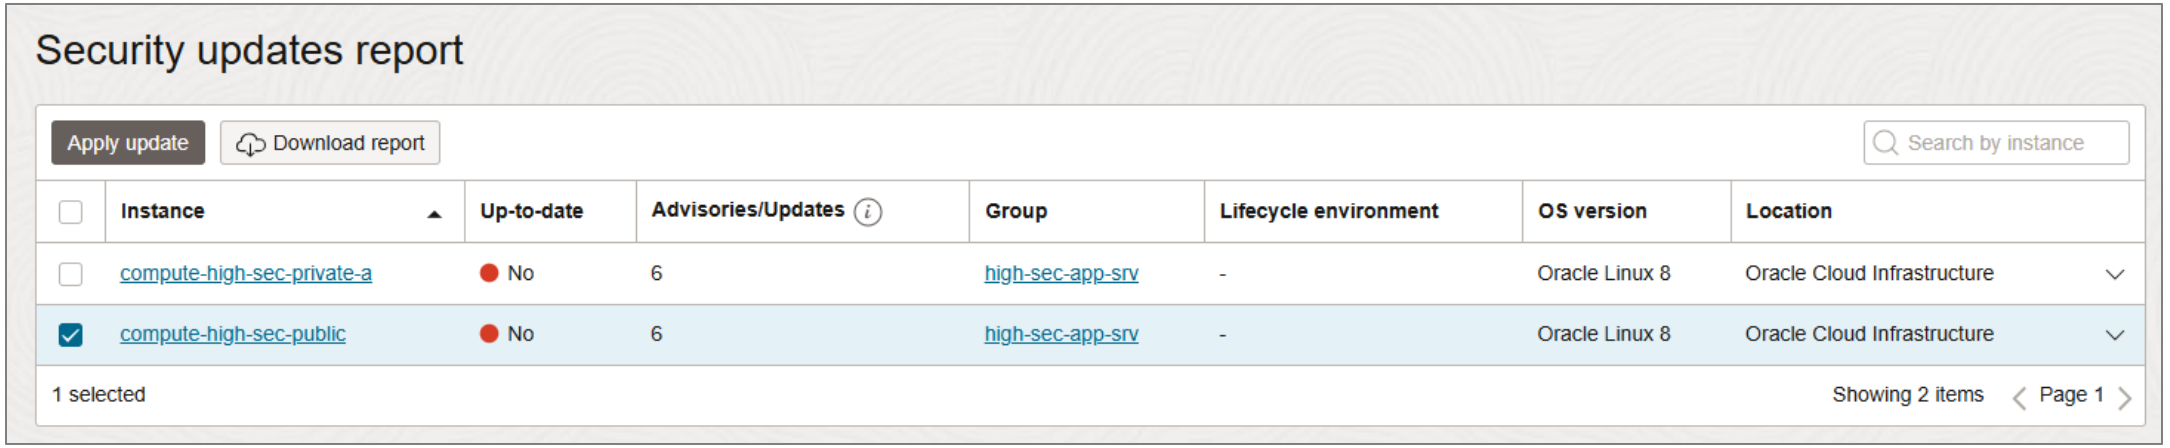

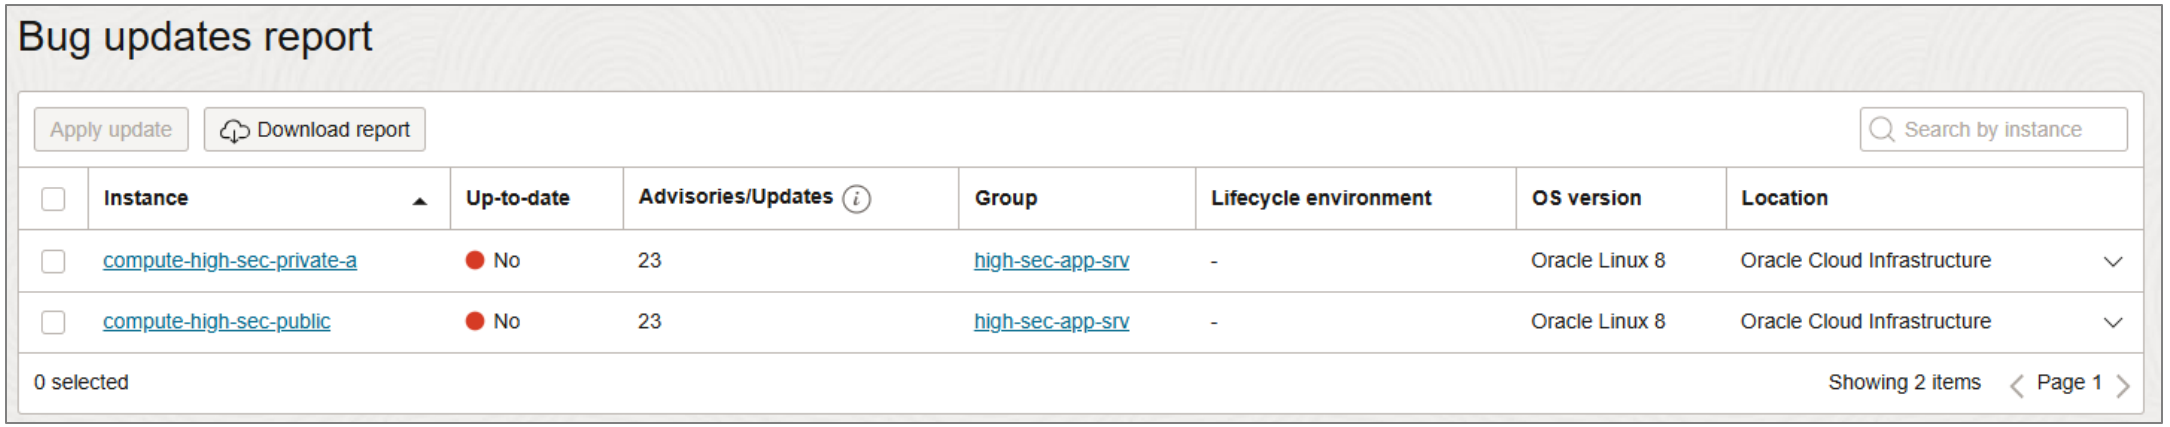

Update Schedule

- If you want to apply the patches and fixes automated and recurring, a job is required.
  - Can run once
  - Or scheduled, as example every Monday

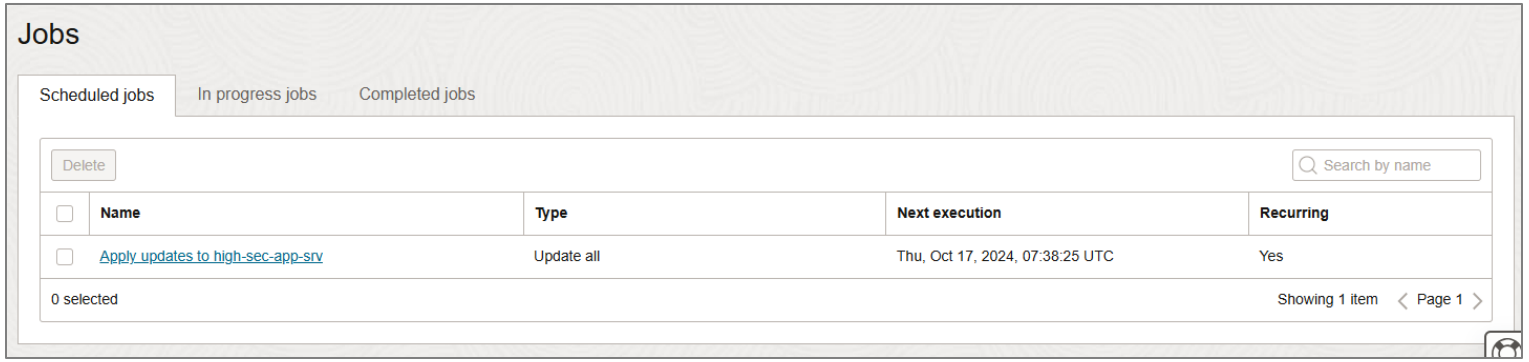

Scheduled

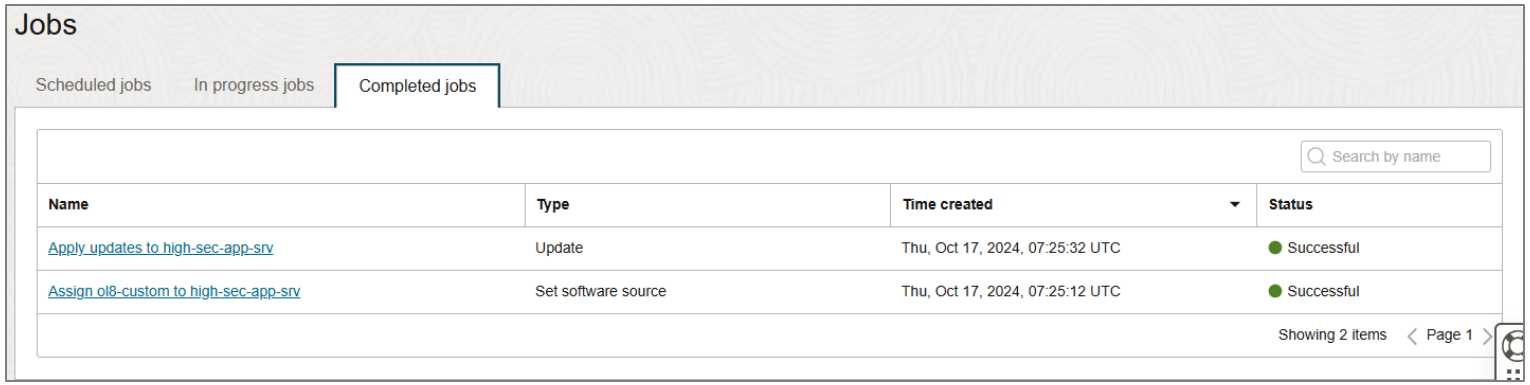

Completed

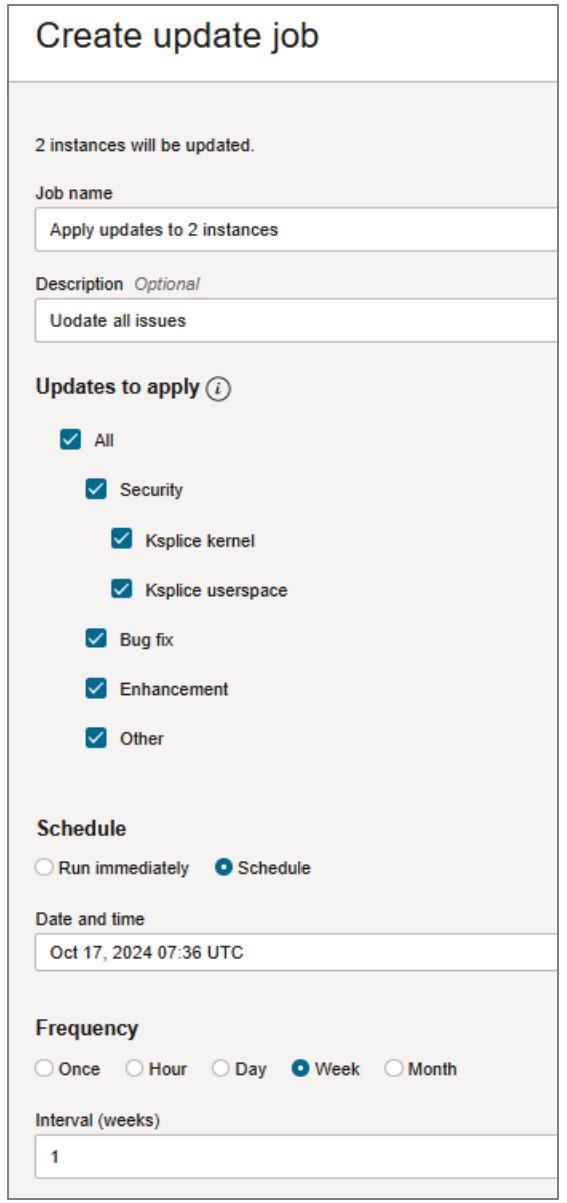

Finally

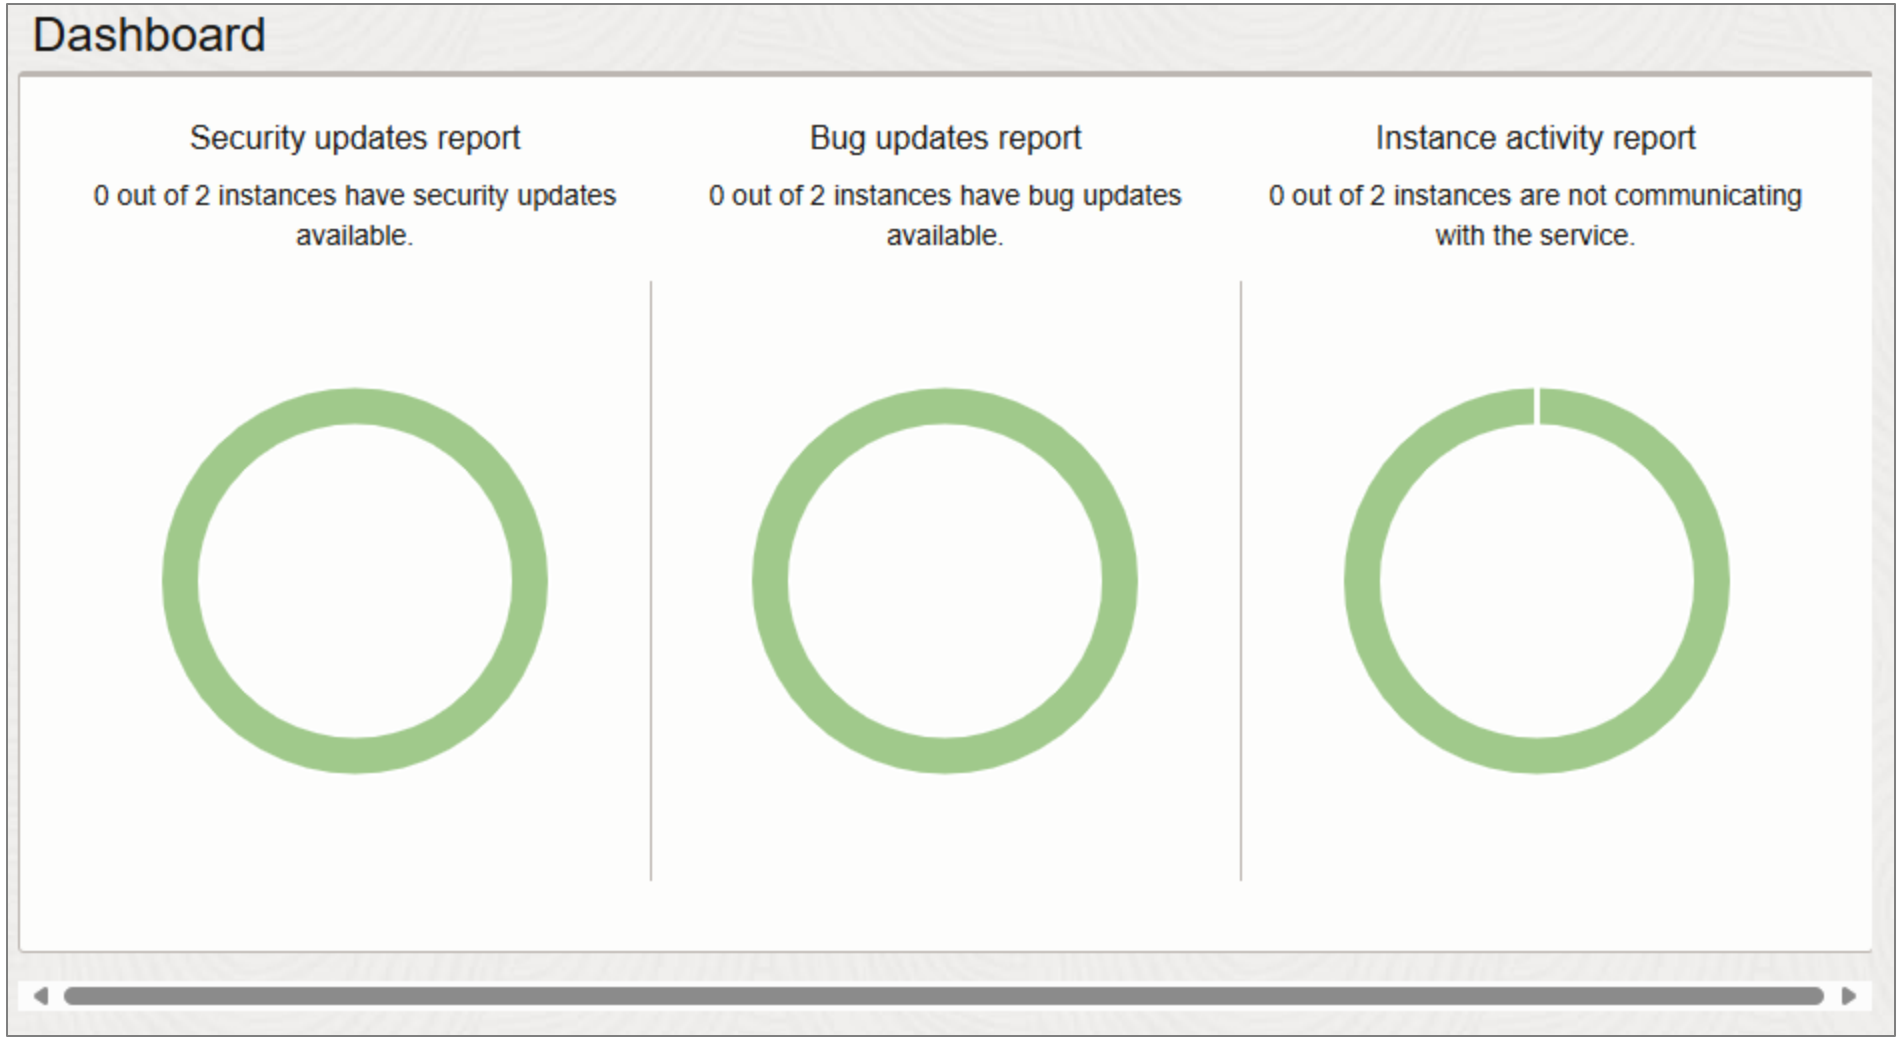

3

Vulnerability Scanning

Security for free.

Vulnerability Scanning Service

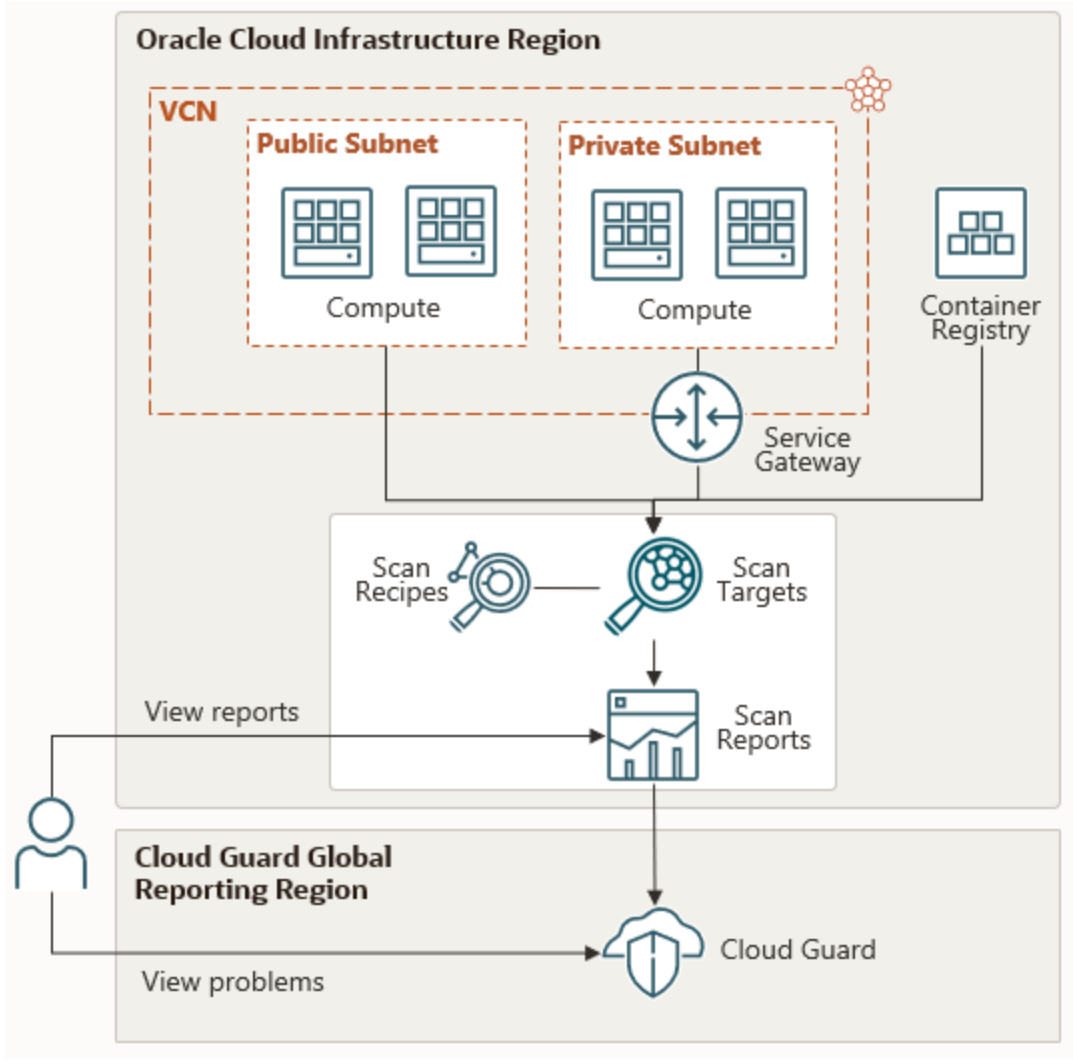

Source: oracle.com

Reports

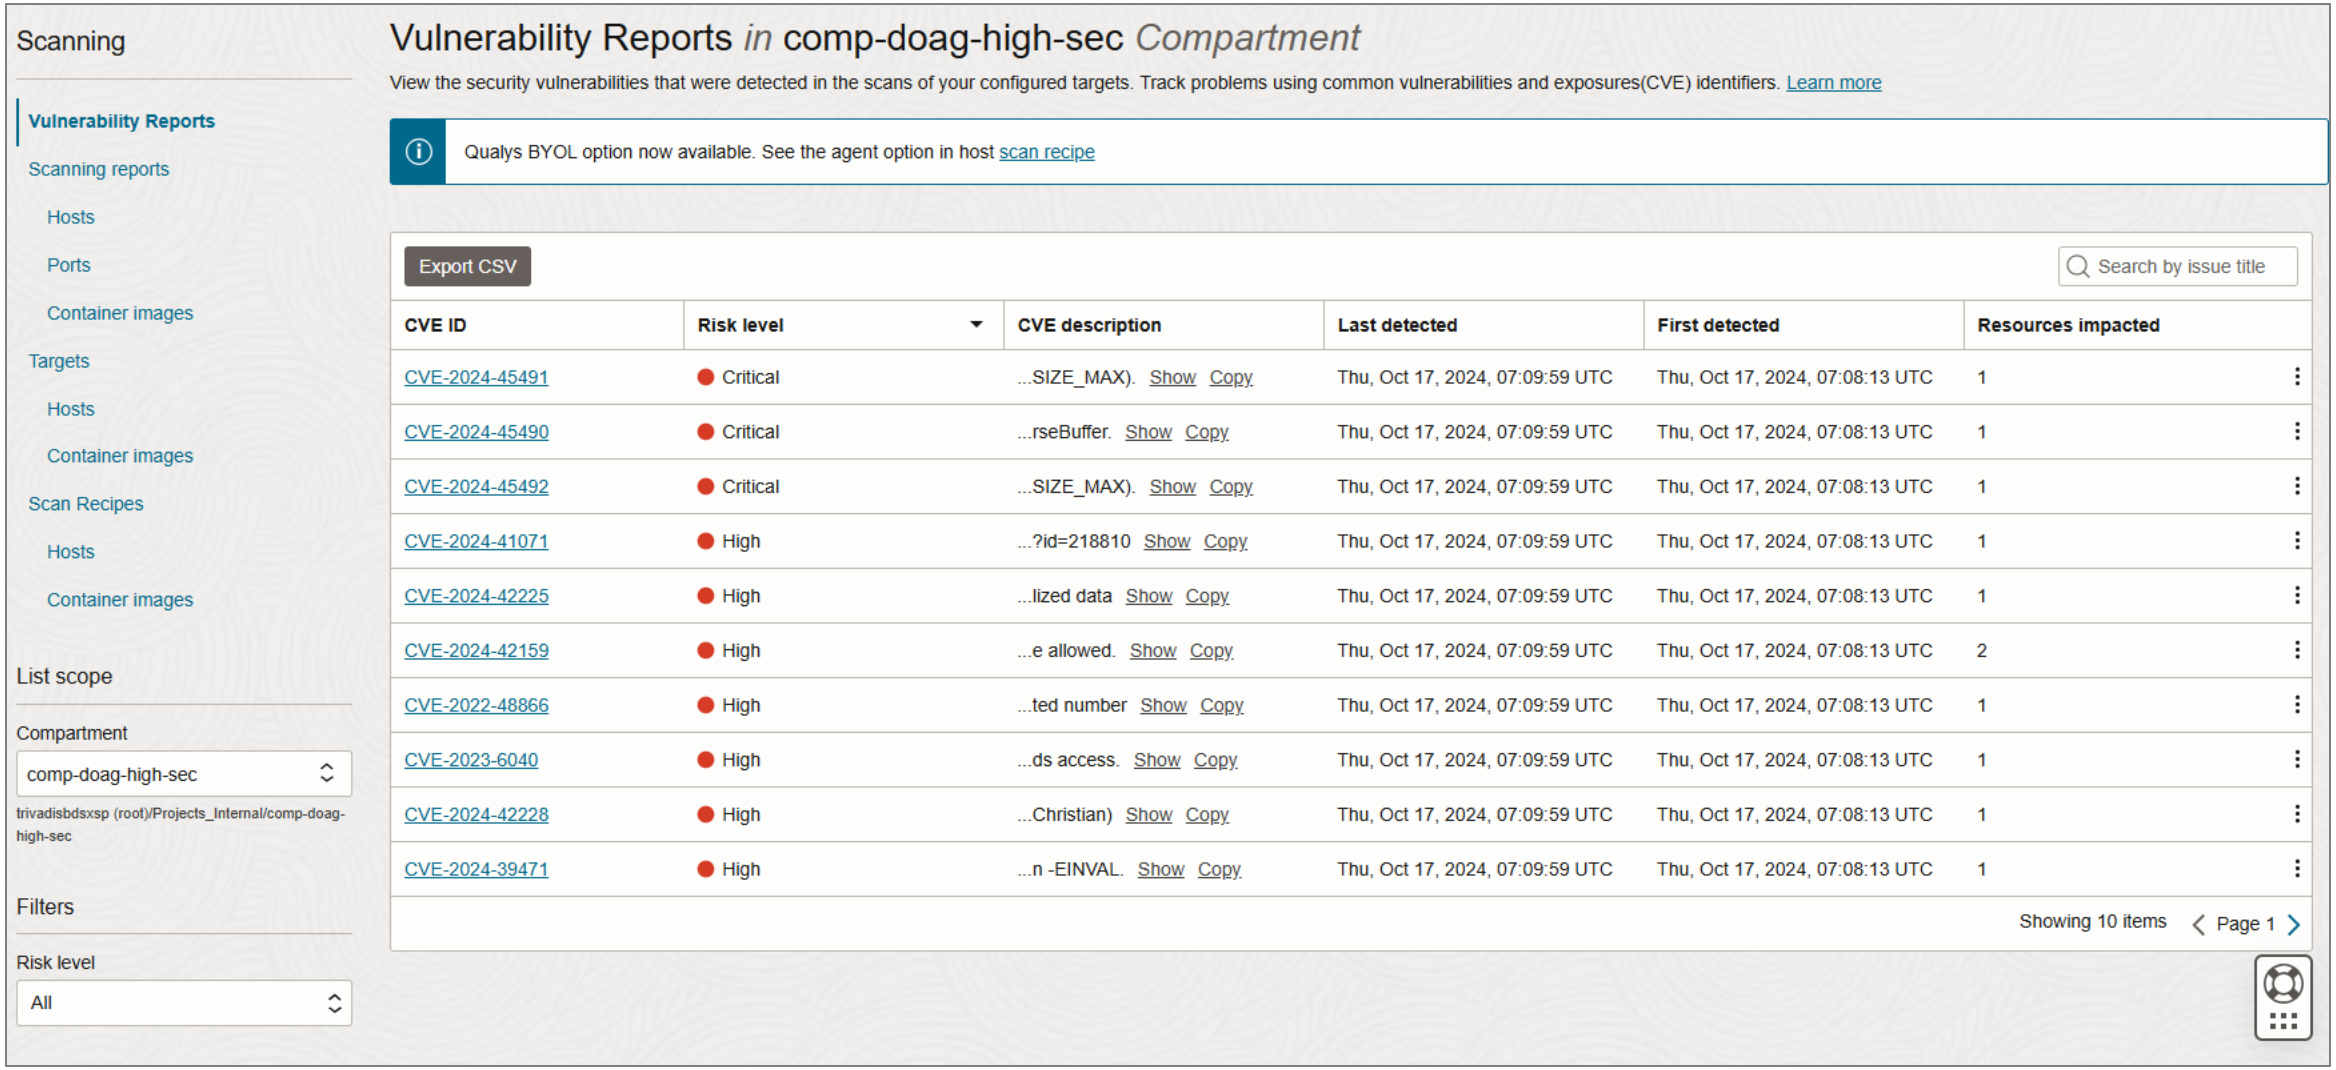

Vulnrability Sources

- The Vulnerability Scanning service detects vulnerabilities in the following platforms and using the following vulnerability sources.

Platform	National Vulnerability Database (NVD)	Open Vulnerability and Assessment Language (OVAL)	Center for Internet Security (CIS)
Oracle Linux	Yes	Yes	Yes
CentOS	Yes	Yes	Yes
Ubuntu	Yes	Yes	Yes
Windows	Yes	No	No

- Targets: Compute Instances and images with a Container Registry repository

Vulnerability Scanning

- Free or BYOL by Qualys

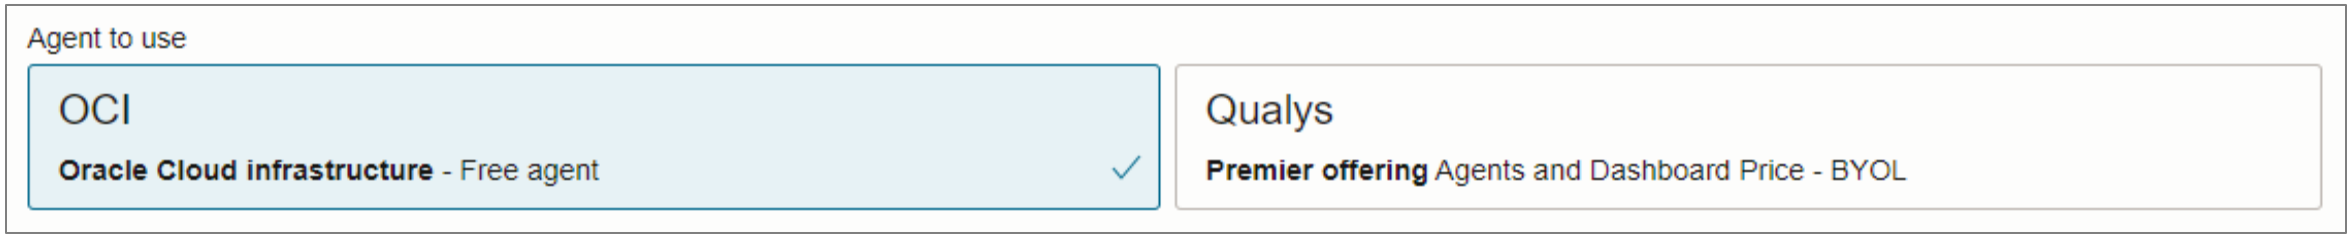

- Based on targets according Cloud Guard
- Pre-defined or adapted recipes
- Cloud Guard integration

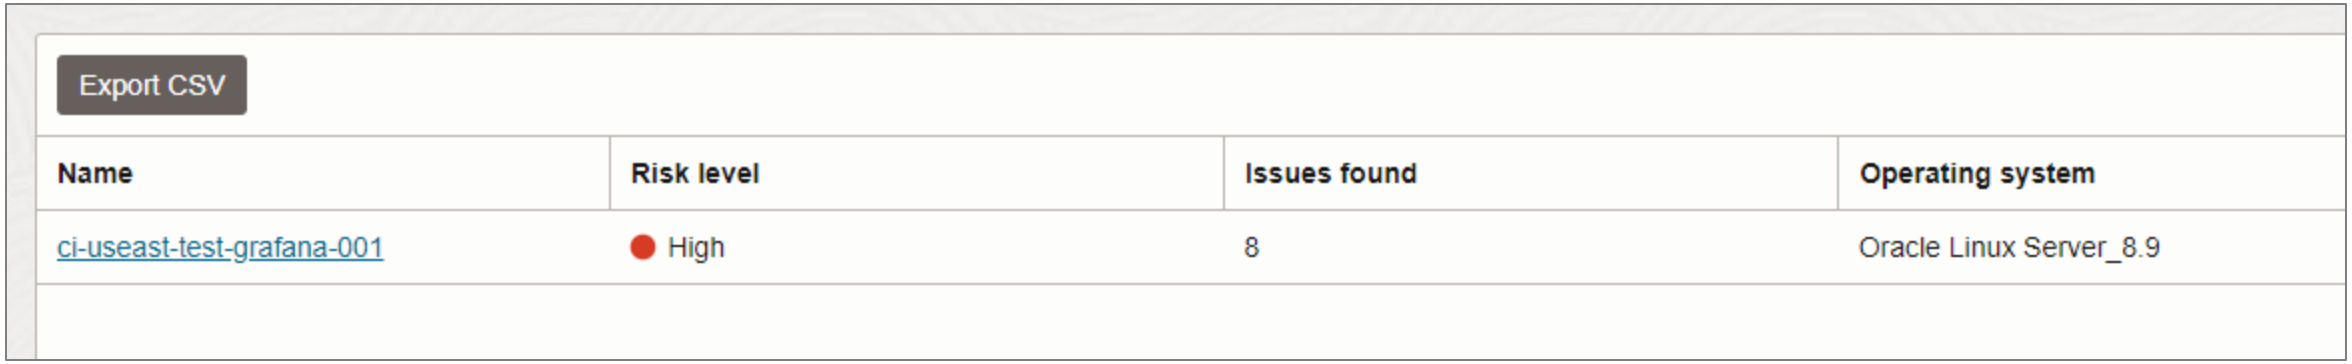

Recipes

- A CIS benchmark profile and OS folders can be selected
- Linux and Windows

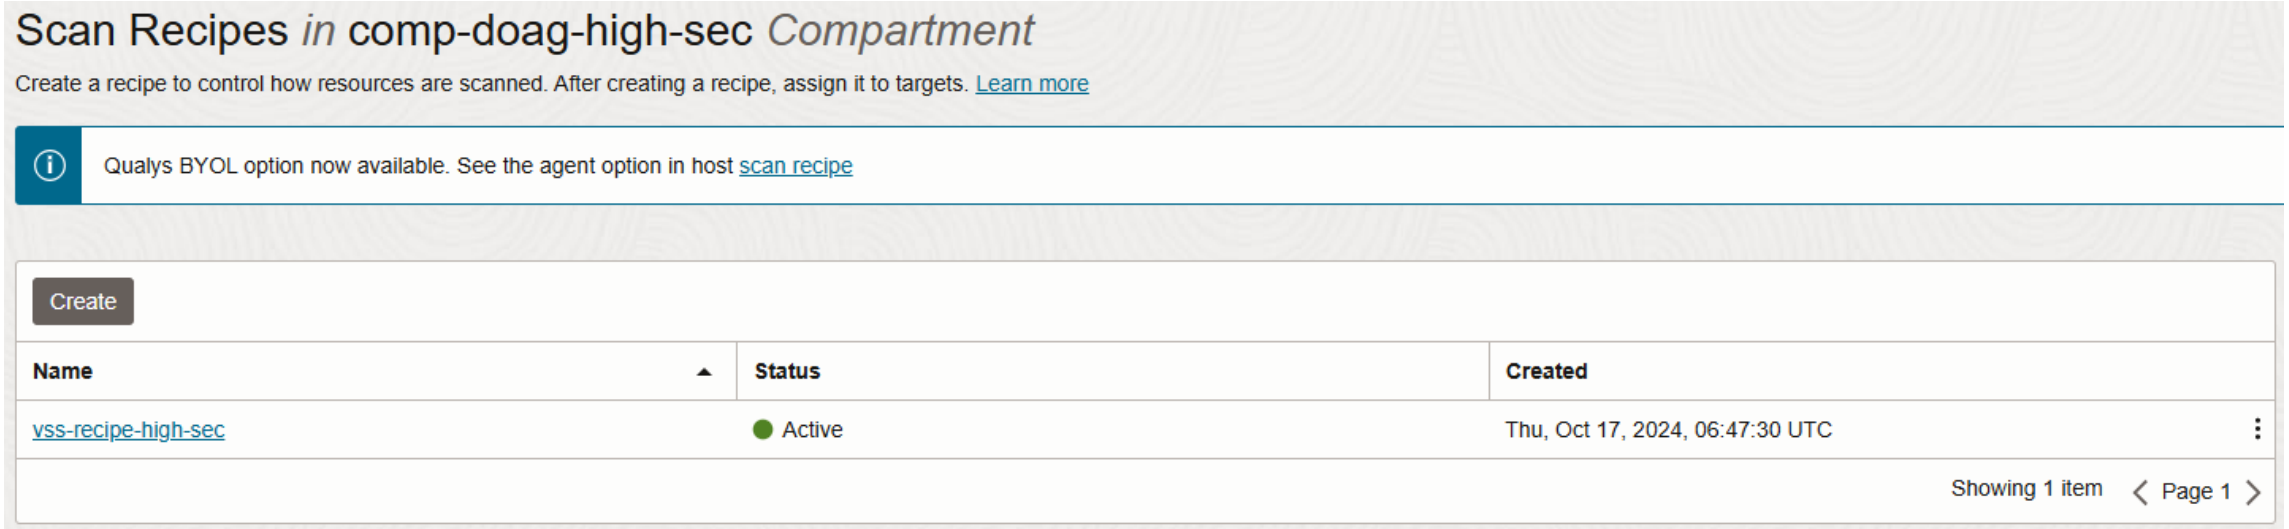

Targets

- Select all or defined Compute Instances only

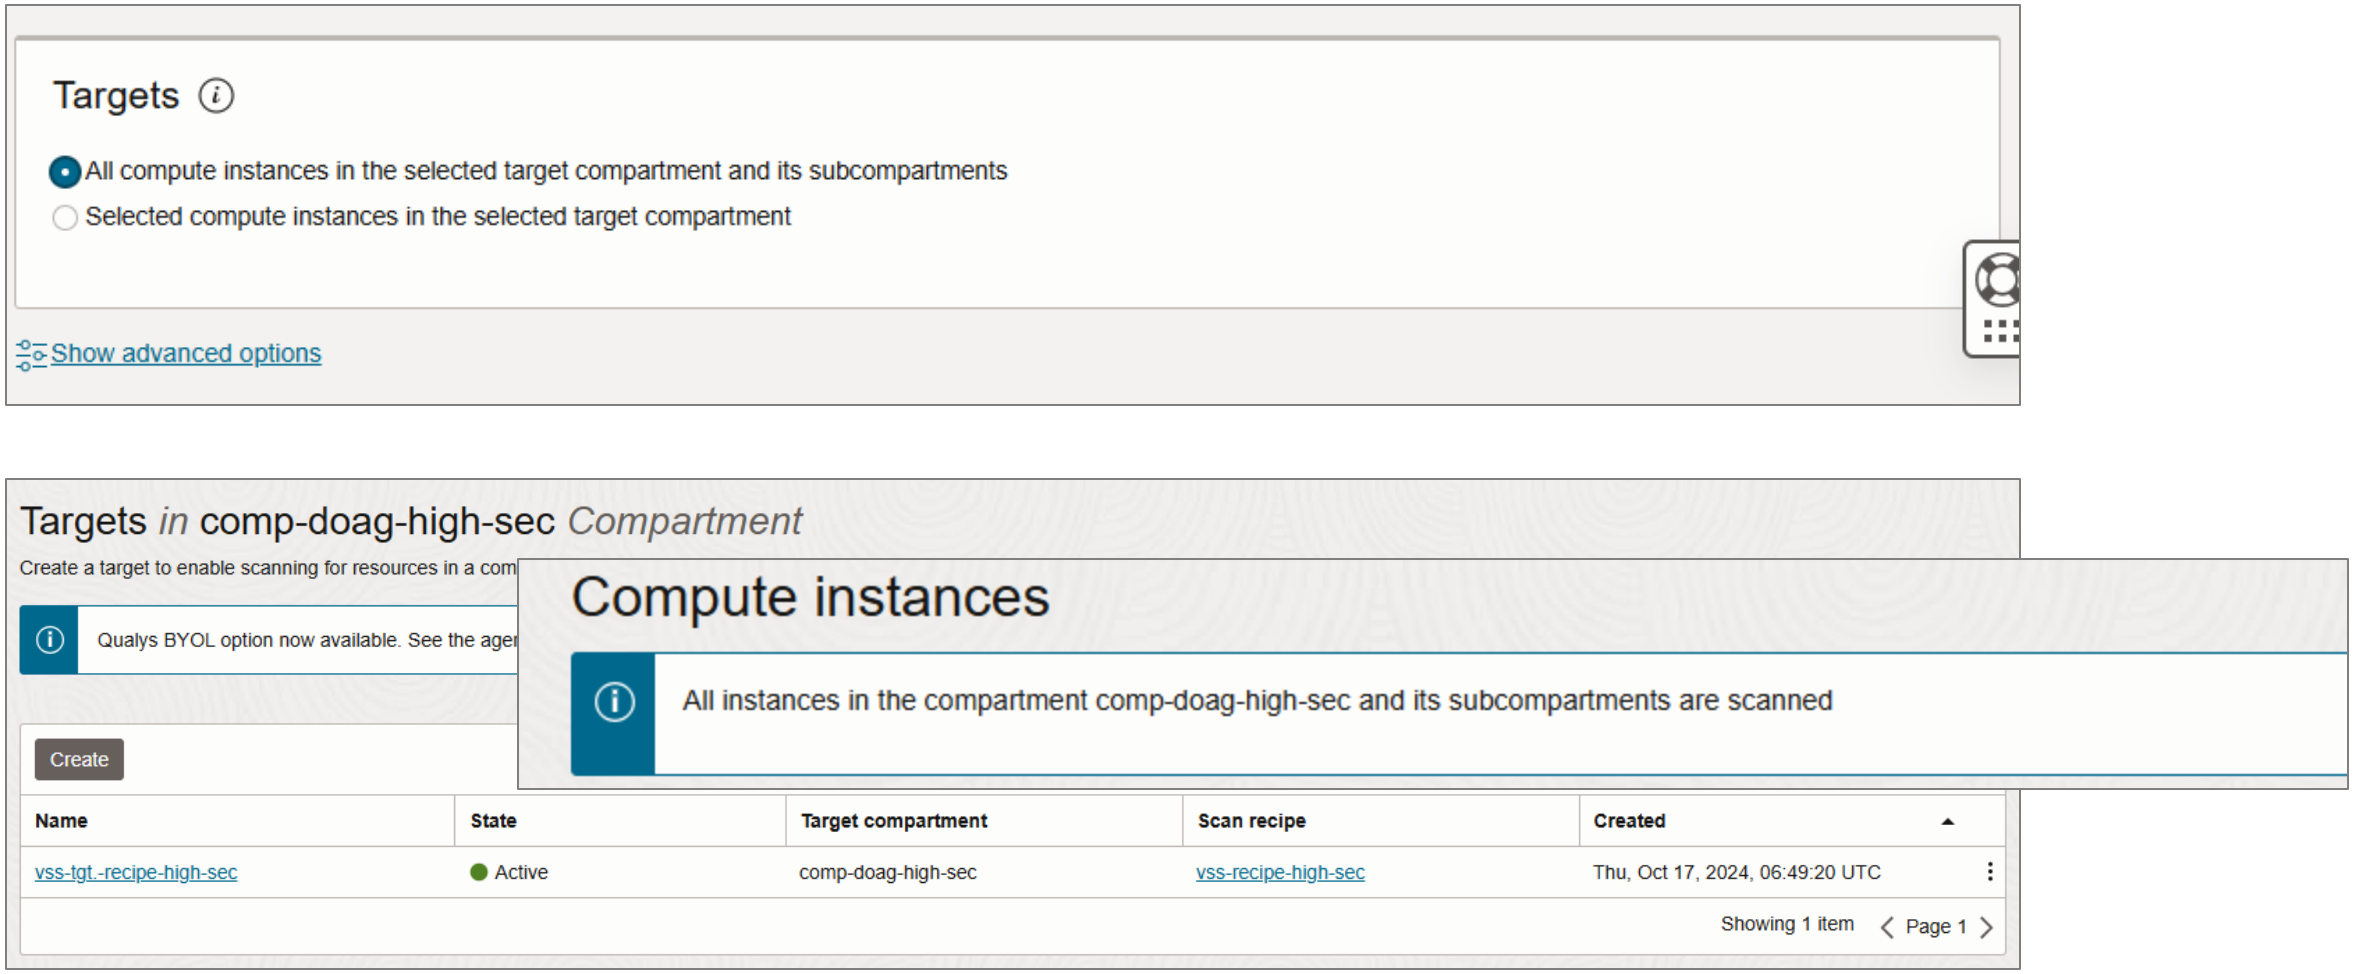

4

Shielded Instances

Security for free.

A combination of **Secure Boot, Measured Boot,** and the **Trusted Platform Module**

Functions:

- Enhanced Security: Oracle Shielded Instances provide an additional layer of security by protecting against unauthorized access and tampering.
- Secure Boot: Ensures that only trusted software is loaded during the boot process, preventing malicious code from running.
- Trusted Platform Module (TPM): Utilizes TPM to securely store cryptographic keys and other sensitive data, enhancing overall security.

Benefits:

- Protection Against Ransomware: Shielded Instances help protect against ransomware attacks by preventing unauthorized changes to system configurations and data.
- Compliance and Regulatory Standards: Meet compliance requirements and regulatory standards with enhanced security features provided by Shielded Instances.
- Integration with OCI Services: Seamlessly integrate with other OCI services such as monitoring, notifications, and identity management for comprehensive security managemen

Components

Shielded Instances	Harden the firmware security on bare metal hosts and virtual machines (VMs) to defend against malicious boot level software.
Secure Boot	Unified Extensible Firmware Interface (UEFI) feature that prevents unauthorized boot loaders and operating systems from booting.
Measured Boot	Measured Boot enhances boot security by taking and storing measurements of boot components, such as bootloaders, drivers, and operating systems. Bare metal instances do not support Measured Boot.
Trusted Platform Module	The Trusted Platform Module (TPM) is used to securely store boot measurements.

Uncompromised Boot Security

- No live migration or reboot migration support.
- If you enable the hardware TPM on a bare metal instance, the instance cannot be migrated, because the hardware TPM is not migratable.
- Custom images are not supported.
- Confidential computing is not supported.

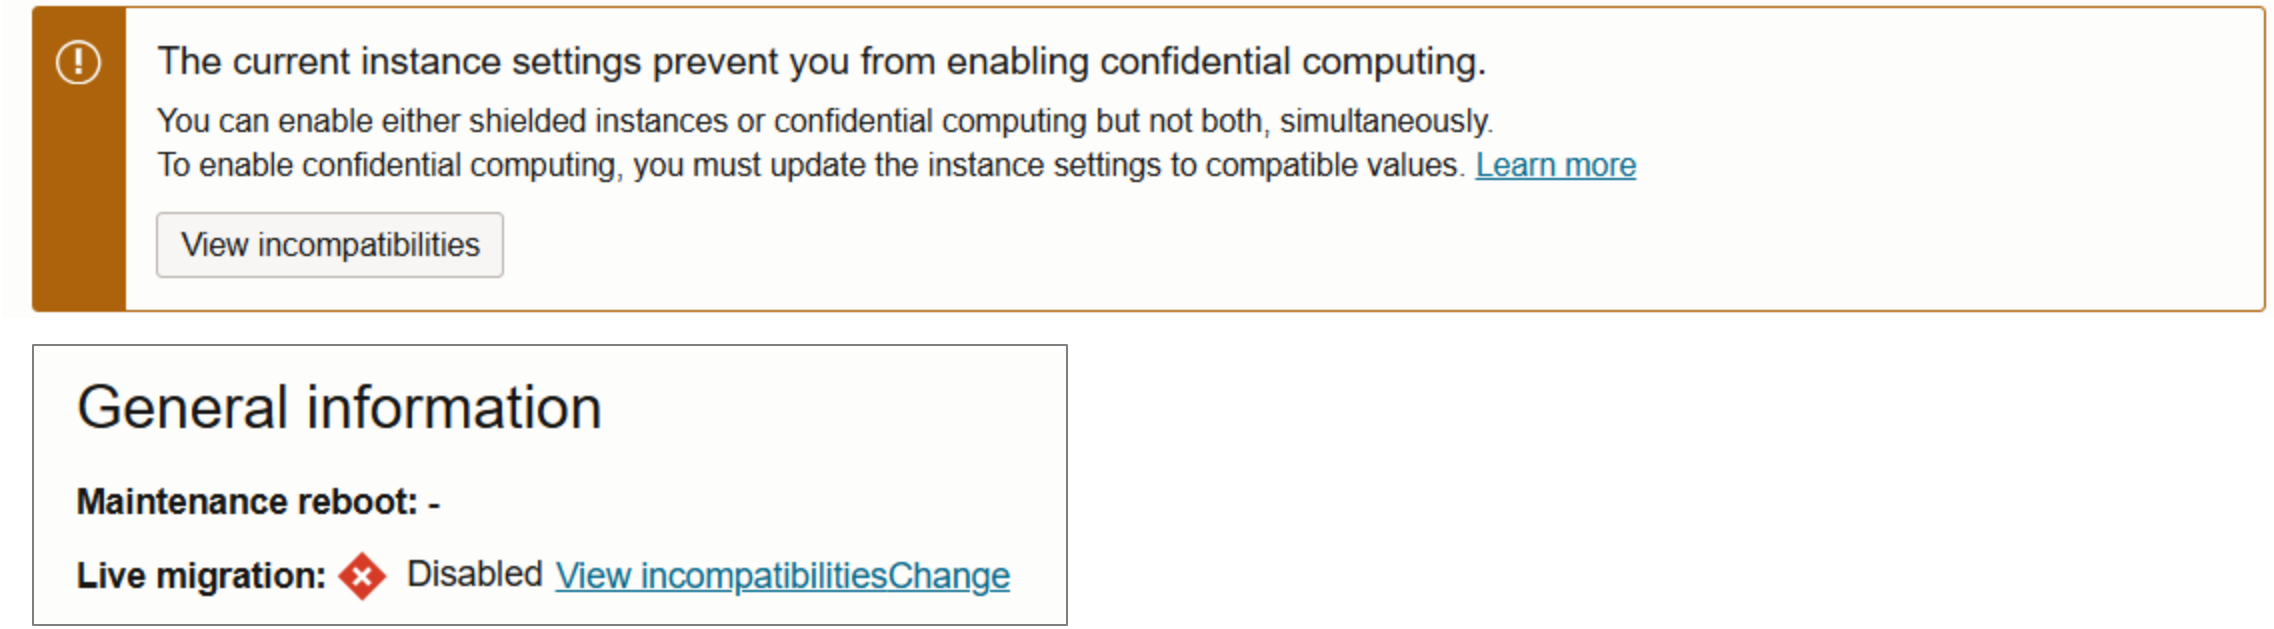

Platform Configuration Register

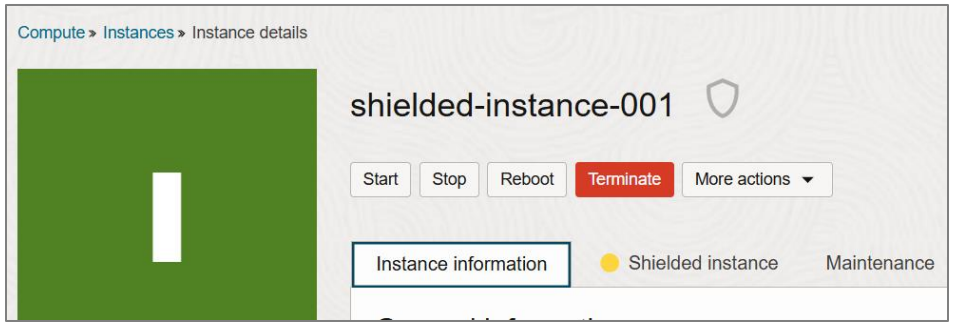

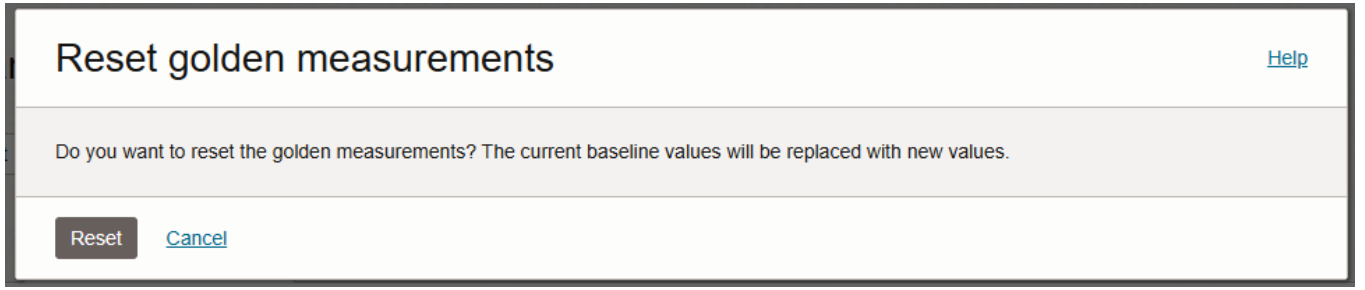

Platform Configuration Register (PCR)

Reset golden measurements | Copy PCR values

Index	Hash algorithm	Value	Status
PCR0	sha1	9ec8a29fbd32ff0aff53771ad5c9d7c1cb02474e	✓
PCR0	sha256	21a59414a7eccc85bf0932c9eb4b51c5b73721605c36934d03f8d53c281e781f	✓
PCR0	sha384	af176b0ab08ba51218ee9653222041d3ebc2ab58bbaa955566fb6e608ce2222918b4c6aa2ecec09ef00b22fb84b997ea	✓

A PCR is an internal memory slot within a Trusted Platform Module (TPM) that attests to the current system configuration or any alterations thereof. PCRs are used to store measurements of software and hardware states, which can then be used to verify the integrity of the system. This ensures that the system has not been tampered with and is running trusted software.

Confidential Computing - AMD EPYC™ Processors

- **Real-time Encryption**: OCI Confidential Computing provides real-time CPU encryption, ensuring that data is encrypted while it is being processed. This helps protect sensitive information from unauthorized access1.
- **Secure Key Management:** The encryption keys used in OCI Confidential Computing reside only in the AMD Secure Processor, making them inaccessible to any user, including administrators1.
- **Enhanced Data Privacy:** By encrypting data in use, OCI Confidential Computing enhances data privacy and security, making it suitable for workloads that require high levels of confidentiality

Virtual Machine Compute Shapes (on Oracle Linux 7.x or 8.x platform images)	Bare Metal Instance Compute Shapes (on any platform image)
VM.Standard.E4.Flex	BM.DenseIO.E4.128
VM.Standard.E3.Flex	BM.Standard.E4.128
	BM.Standard.E3.128

Create Instance

- Create instance, enable confidential computing without addional costs.

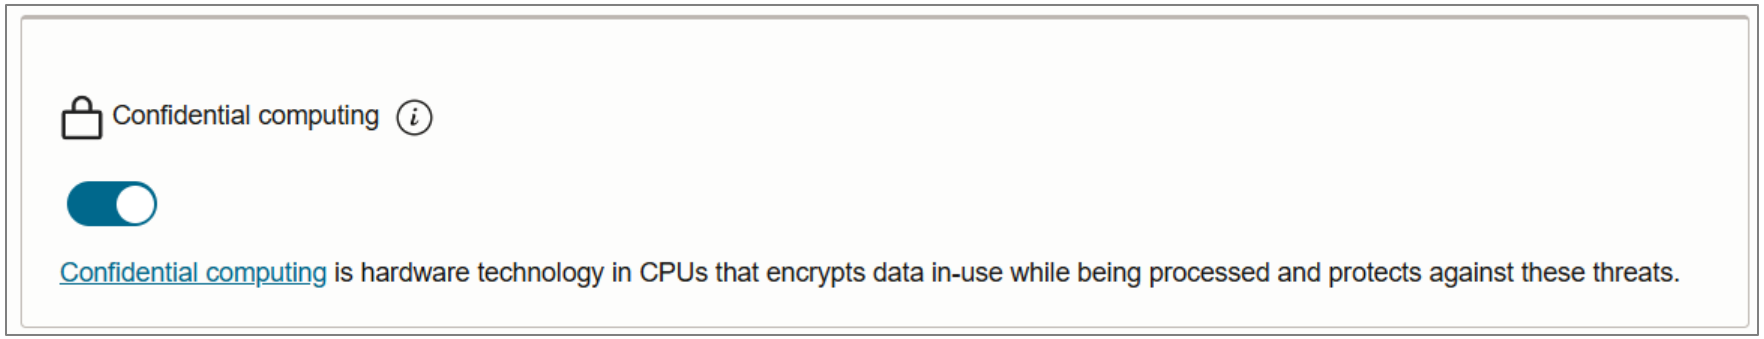

- Select shape – watch for shield icon

Instance with confidential computing enabled can be live migrated. No code change and a minimal performance impact.

Launch options

NIC attachment type: PARAVIRTUALIZED

Remote data volume: PARAVIRTUALIZED

Firmware: UEFI_64

Boot volume type: PARAVIRTUALIZED

In-transit encryption: Disabled

Secure Boot: Disabled

Measured Boot: Disabled

Trusted Platform Module: Disabled

Confidential computing: Enabled

5

Hands-On Labs

What about the Database Security?

Hands-On Labs

List of Hands-On Labs

- **Key Management** – Create and manage encryption keys within a customer-managed key vault.

Oracle Cloud Infrastructure Security

Cloud Guard

Martin Berger
Stefan Oehrli

Cloud Guard

Security for free.

1 Understanding Cloud Guard

2 Hands-On Labs

1

Understanding Cloud Guard

Security for free.

What is Cloud Guard?

Oracle Cloud Guard is a cloud-native service designed to help customers monitor, identify, and maintain a strong security posture on Oracle Cloud Infrastructure. It continuously examines OCI resources for security weaknesses related to configuration and user activities

Real-Time Threat Detection

Cloud Guard provides real-time monitoring and threat detection by continuously scanning OCI resources for misconfigurations, vulnerabilities, and anomalous activities. This ensures that potential security issues are identified and addressed promptly

Automated Risk Assessment

Cloud Guard automatically assesses security risks based on predefined security policies and industry standards, such as the CIS Benchmark. It offers actionable recommendations for addressing detected risks, allowing for quick remediation and enhancing overall security

Customizable Security Policies

Users can configure custom policies and rules to align monitoring with their specific security requirements. Cloud Guard integrates seamlessly with other OCI tools, providing a comprehensive and flexible security solution.

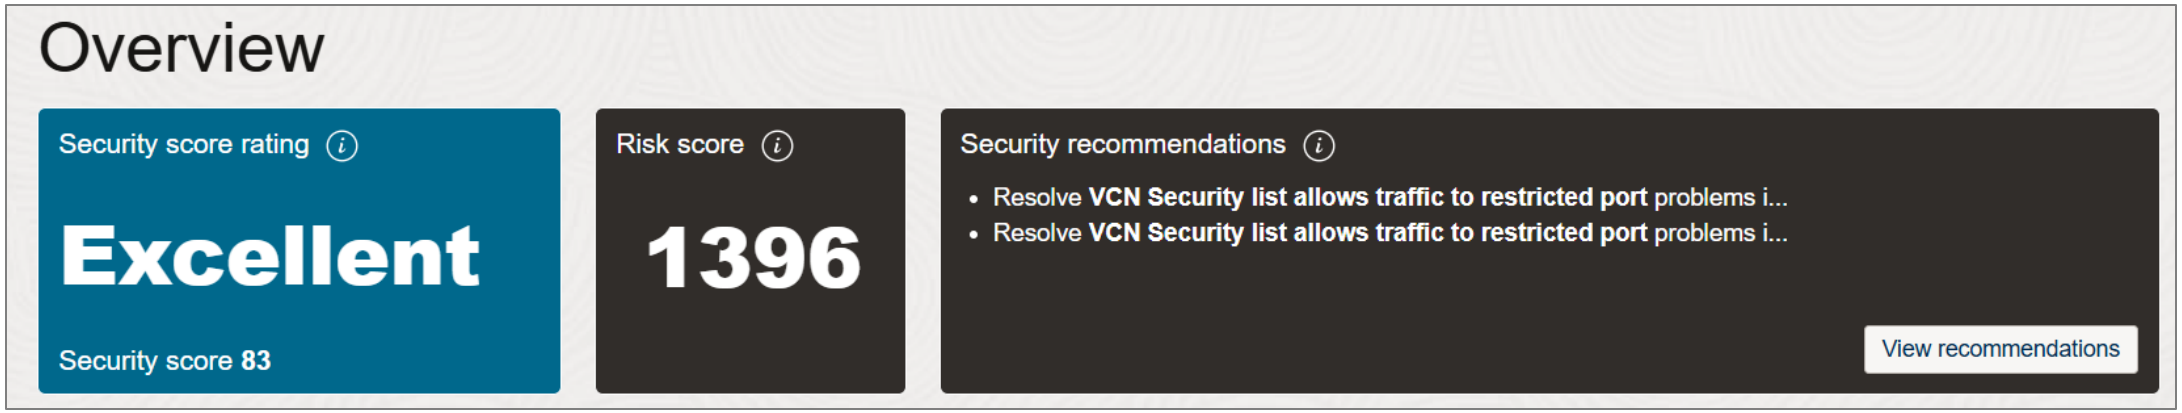

The Concept

- **Detector Rules** provide a class of resources with specific actions or cinfiugurations that can cause a detector to report a problem.
- **Detector Receipes** provide the baselines for examining the resources and activities in the target.

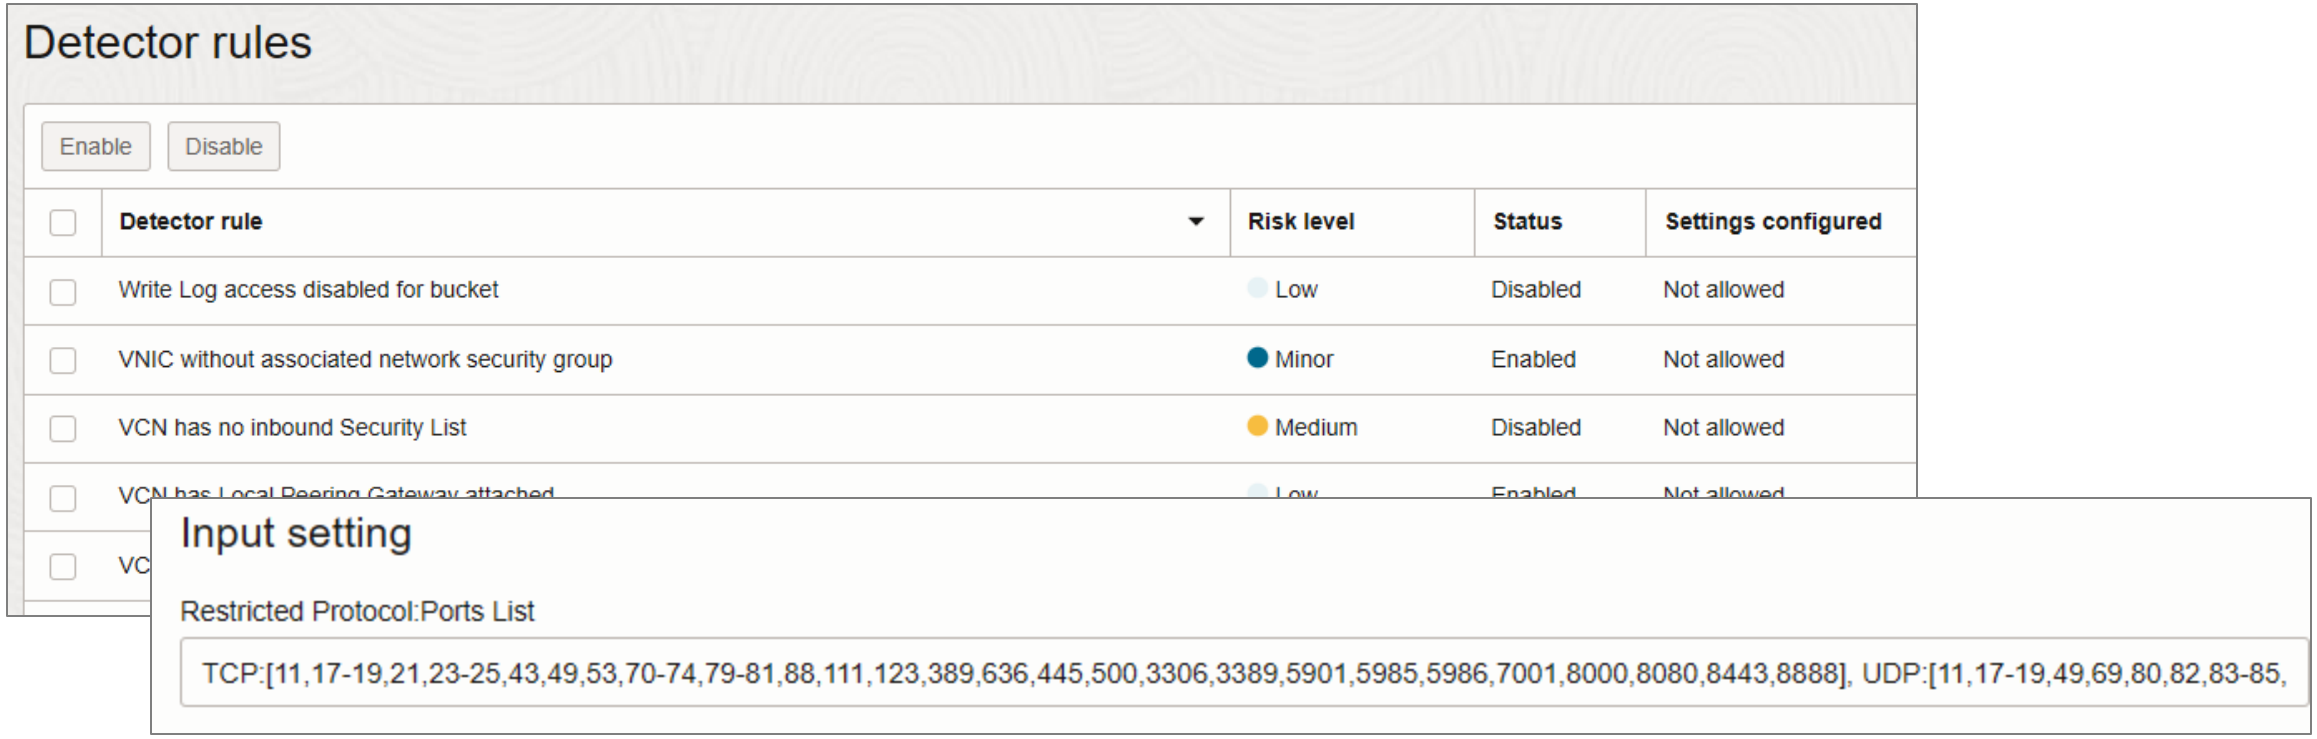

Key Components of Cloud Guard

Targets

Targets define the scope of what Cloud Guard is to check. This scope is tied to the compartment where the target is defined and all the child compartments from that point until another target is encountered. Compute Instances, Object Storage etc.

Detectors

Detectors perform checks and identify potential security problems based on their type and configuration.

Responder Recipes

These provide the baselines for examining the resources and activities in the target. Oracle-managed detector recipes allow setting only the scope of resources for which a rule triggers a problem, while user-managed detector recipes allow more customization.

Detector Recipes

Responder recipes define the actions to be taken when a problem is detected. They can be configured to automatically remediate security issues or to notify administrators for manual intervention.

How Cloud Guard works

Target	Detectors	Problems	Responder
Object Storage Bucket	Visibility Public	Mail Notification	Change Visibility
Security List	0.0.0.0/0	Stream Information	Remove Rule
Compute Instance	Public IP	Slack Channel	Shutdown

Data Collection and Analysis

Alert

Remediation

Enable Cloud Guard

The prerequisites for enabling Oracle Cloud Guard:

- Having a paid Oracle Cloud Infrastructure tenancy, as Cloud Guard is not available for free tenancies.
- Users must create a user group with administrator privileges to work with Cloud Guard, and this group should be restricted to a limited audience.
- Additionally, specific policy statements must be added to enable the Cloud Guard users group to manage Cloud Guard resources.

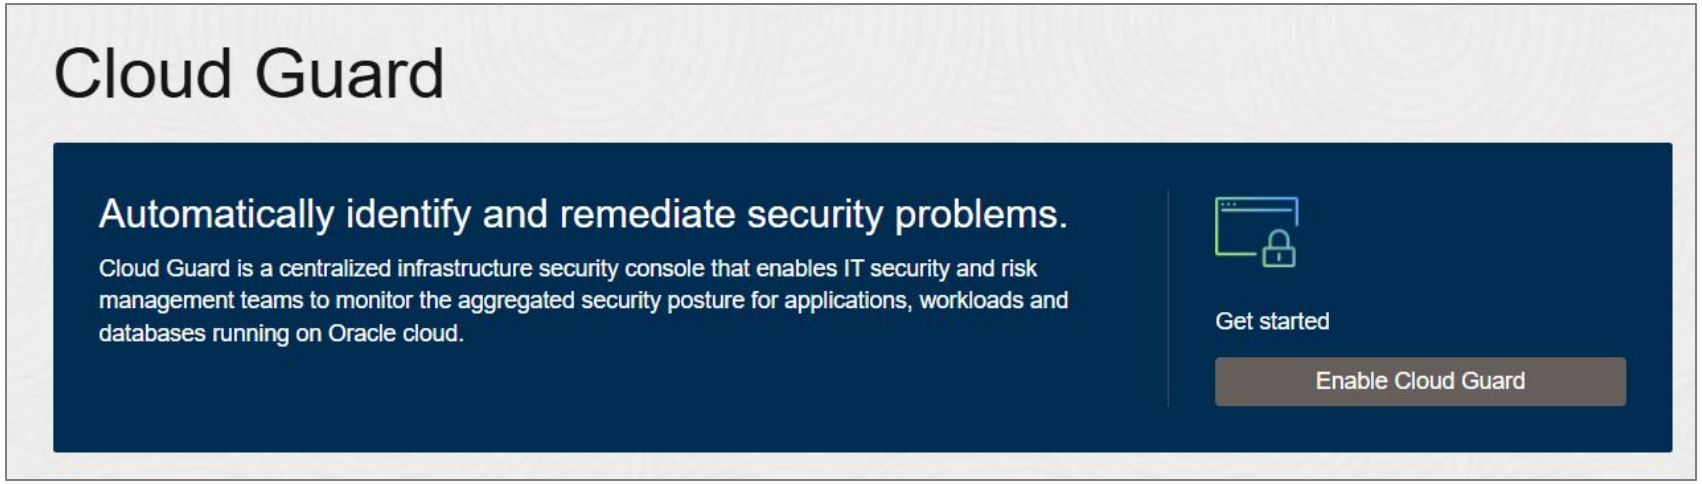

Scores

Security Score: A higher security score is better. A security score of 100 would mean that no problems were detected for any resources. Reflects monitoring for past 30 days.

Risk Score: Related to the number and severity of problems. In general, organizations with many more resources are likely to have more problems, and thus a higher risk score. The risk score is closely related to the "potential surface area" of risk. Updated every 15 Minutes.

About Compartments – Important Notes

All compartments of a target inherit that target's configuration.

- Detector and responder rule settings for a target apply to:
  - The top-level compartment assigned to that target.
  - Any subordinate compartments below it in the hierarchy.

Target defined within an existing target overrides inherited configuration.

- If you want to exclude some compartments from monitoring, create targets below the root level and do not include the root compartment in any target.

Setting Up Cloud Guard Targets

- Defines the resources where Cloud Guard is enabled.
- Setting up more targets allows more specific monitoring
- OCI Cloud Guard service targets are logical objects inspected by Cloud like Compute Instance, Object Storage Bucket, Security Lists etc.
- Rules to detect and resolve issues are applied to targets.

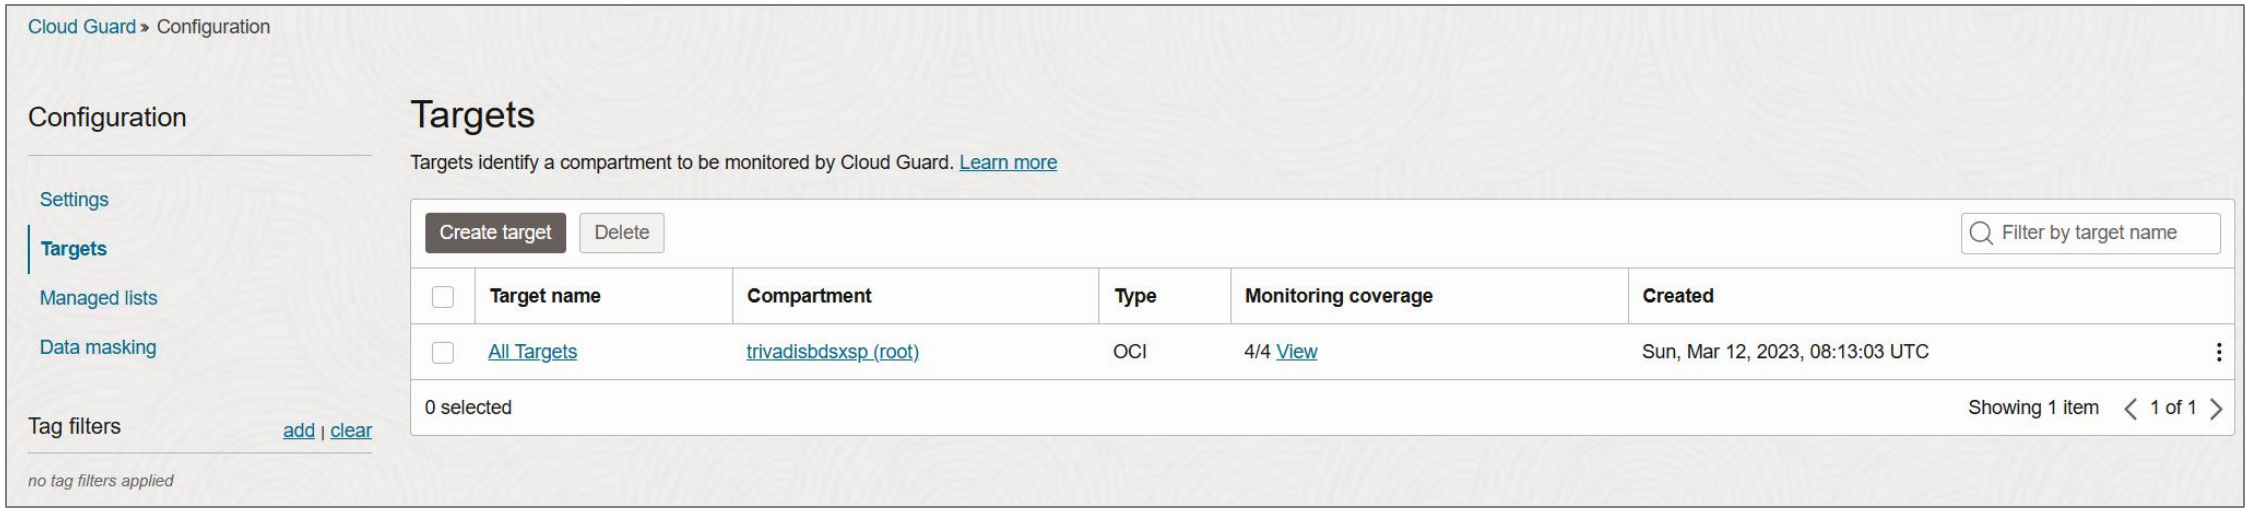

Targets Inside

You don’t see the new created target for a compartment in the configuration list? Search it via **Tenacy Explorer** or **Search Bar**.

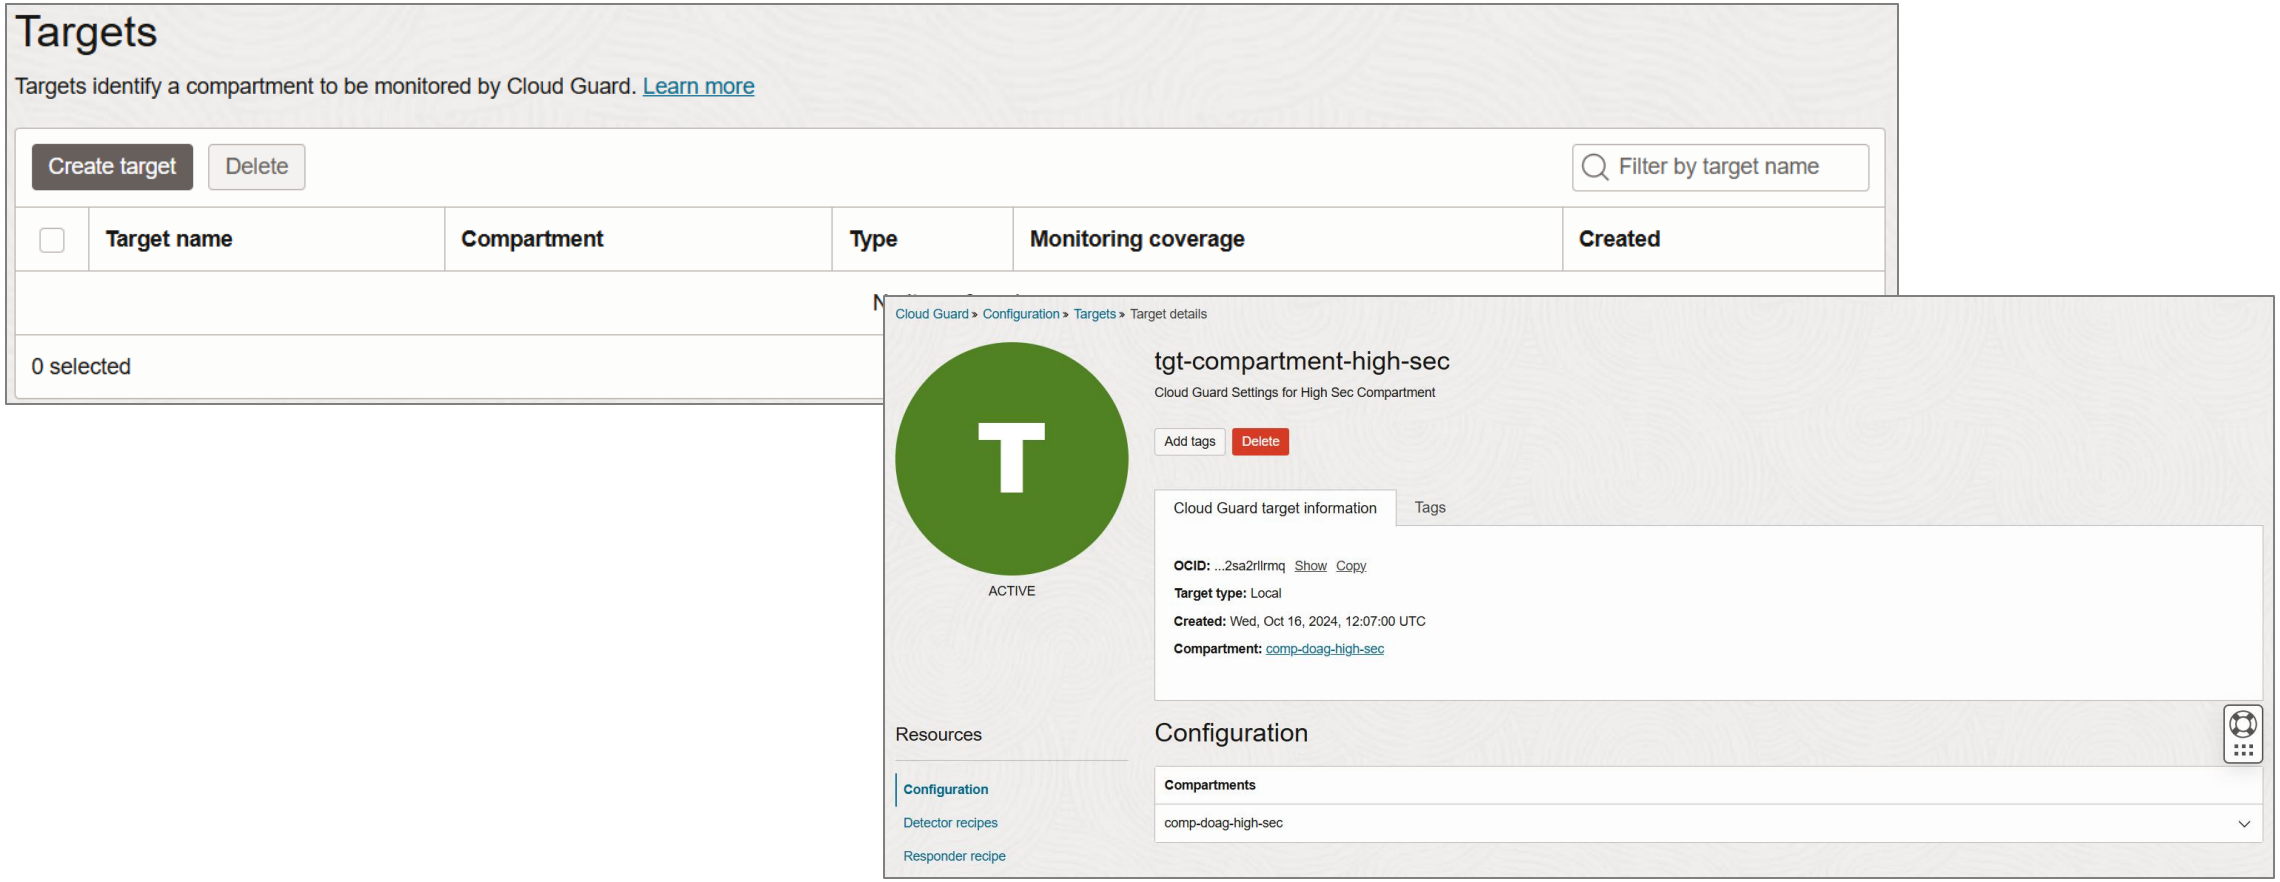

Managed Lists

A managed list is a reusable set of target parameters that simplifies setting the scope for detector and responder rules. It can be used to apply specific configurations to detectors.

- For example, a predefined "**Trusted Oracle IP address space**" list includes all Oracle IP addresses you want to consider trusted when defining rules for detectors and responders.
- Cloud Guard also allows you to create your own managed lists as needed, such as lists of states or provinces, zip or postal codes, OCIDs, or any other criteria you require.

Here are some specific use cases for custom managed lists:

- **Trusted IP addresses**: Exempt listed IP addresses from triggering alerts that should only be triggered by untrusted IP addresses.
- **Resources that should be public**: Exempt listed resources from all detectors related to identifying public configurations.

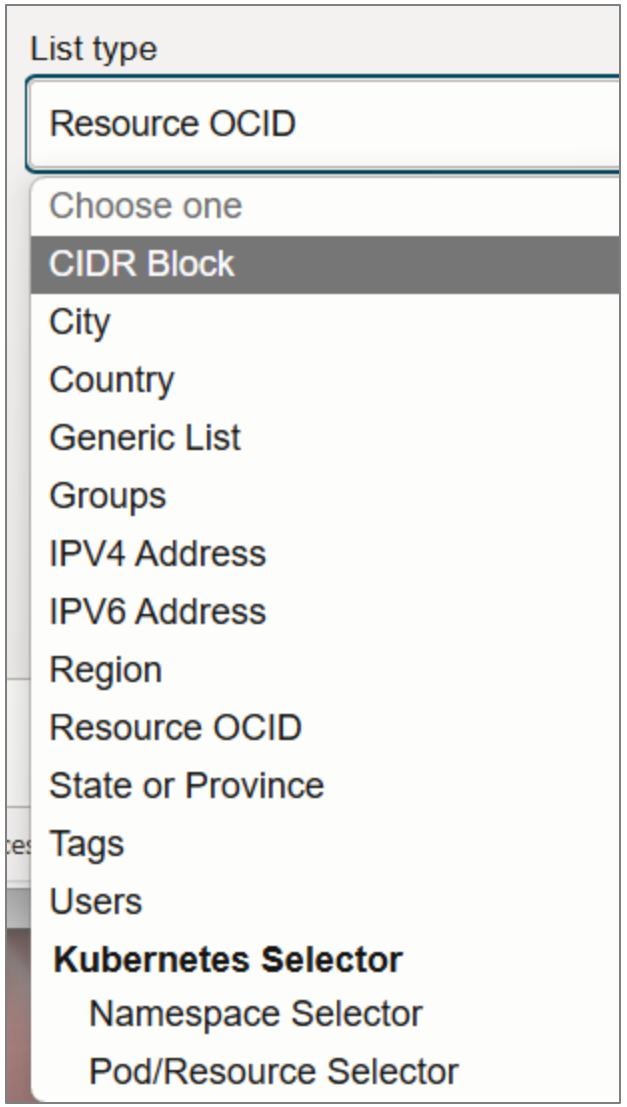

Detection or Responding

Detector Recipes

- These recipes are used to identify potential security issues by continuously monitoring resources for misconfigurations, vulnerabilities, and anomalous activities.
- Detector recipes contain rules that define what constitutes a problem. When these rules are triggered, Cloud Guard creates a problem.

Responder Recipes

- These recipes define the actions to be taken when a problem is detected by a detector recipe.
- Responder recipes contain rules that specify the actions to remediate, resolve, or dismiss the identified problems.

Oracle provides a set of recipes (3+1) – you can clone then for adaption according your security needs.

Detector Recipes

OCI Configuration Detector Recipe

- Designed to detect resource configuration settings that could pose a security problem.
- Helps identify misconfigured resources to ensure compliance with security policies.

OCI Instance Security Recipe

- Monitors compute hosts for suspicious activity, providing runtime security for workloads in Compute virtual and bare metal hosts.
- Collects important security information such as security alerts, vulnerabilities, and open ports to provide actionable guidance for detection and prevention.

OCI Activity Detector Recipe

- Detects actions on resources that could pose a security problem.
- Monitors activities to identify potential security issues based on resource actions.

OCI Threat Detector Recipe

- Designed to detect subtle patterns of activity that could be building up to pose a security problem.
- Continuously monitors OCI audit events for malicious activity using machine learning-based behavioral attack models.

OCI Configuration Detector Recipe – Center of Internet Security

The **Configuration Detector** recipe follows the CIS – *Center of Internet Security* – guidelines in version 1.0 ,1.1 or 2.0.

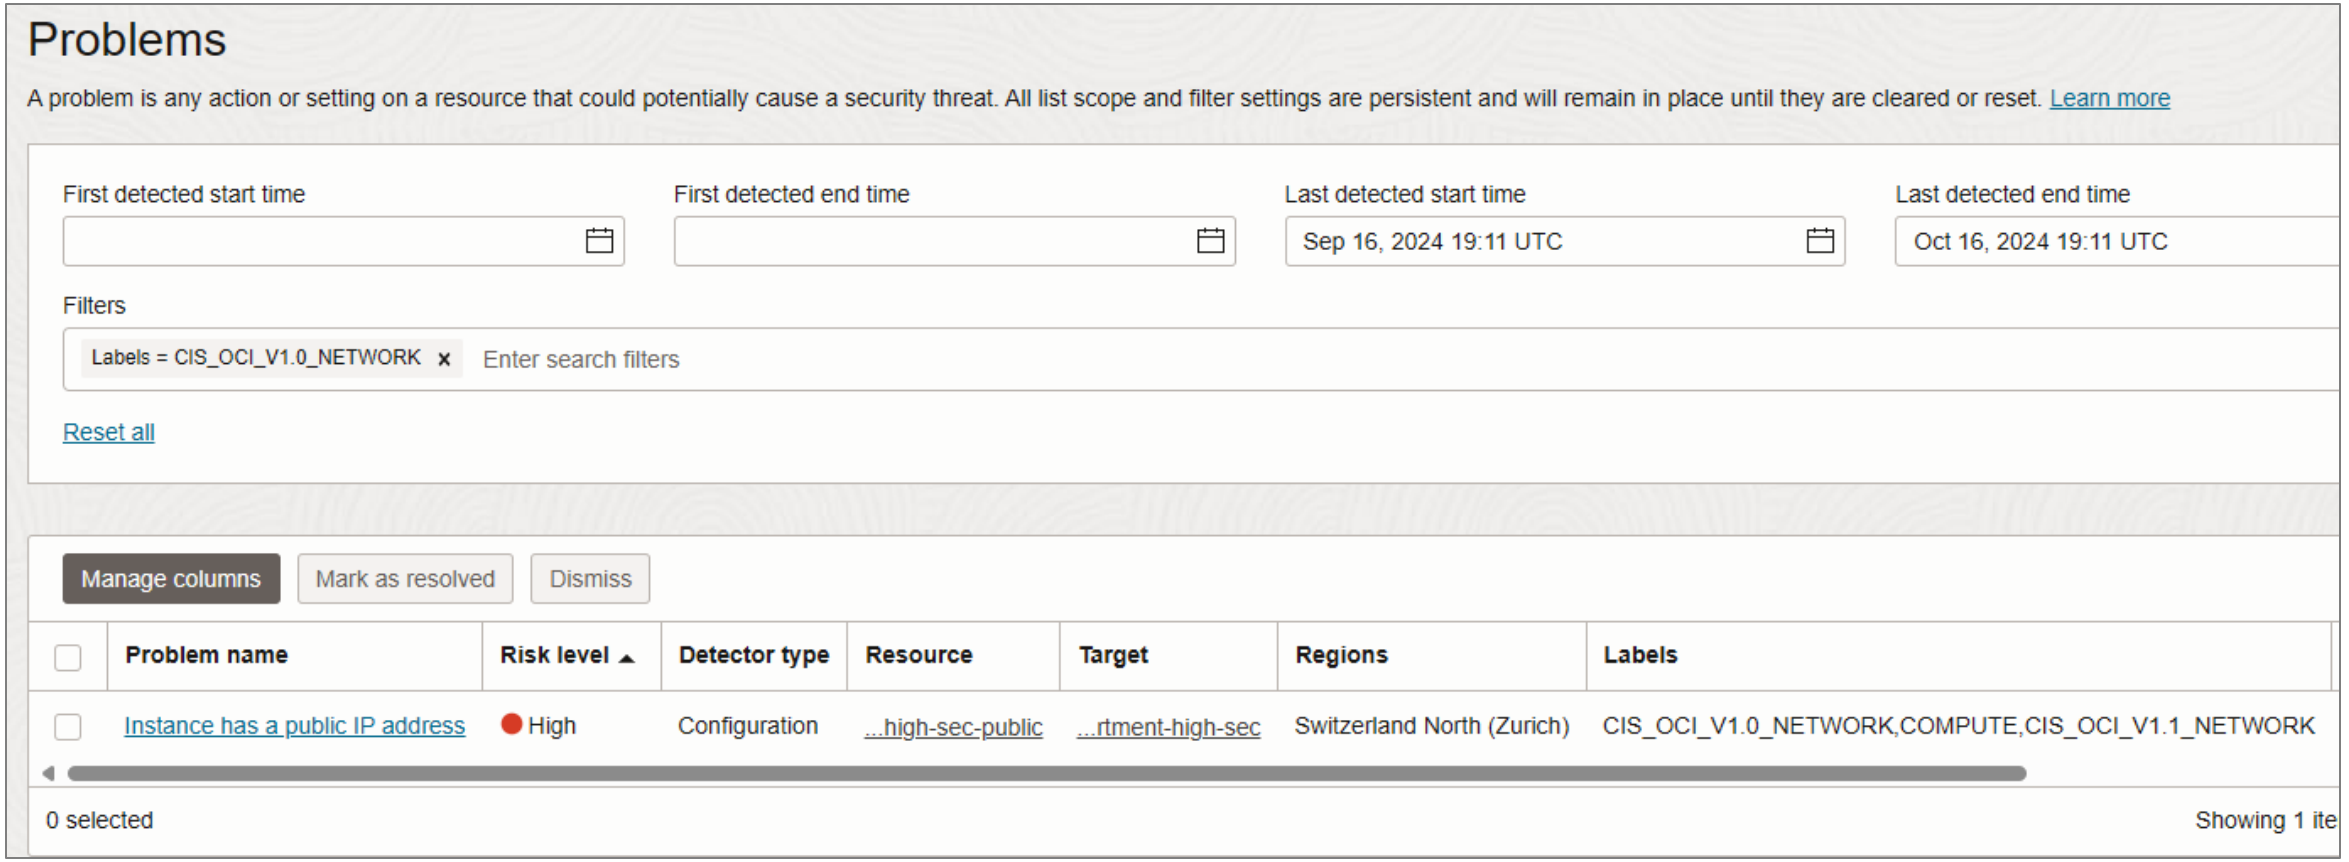

Responder Recipes

- Define actions to be taken when a problem is detected by a detector recipe.
- Can be configured for automatic remediation or to notify administrators for manual intervention.
- Contain multiple responder rules specifying actions to address identified problems.
- Can use Oracle-managed recipes with default rules or customized user-managed recipes.

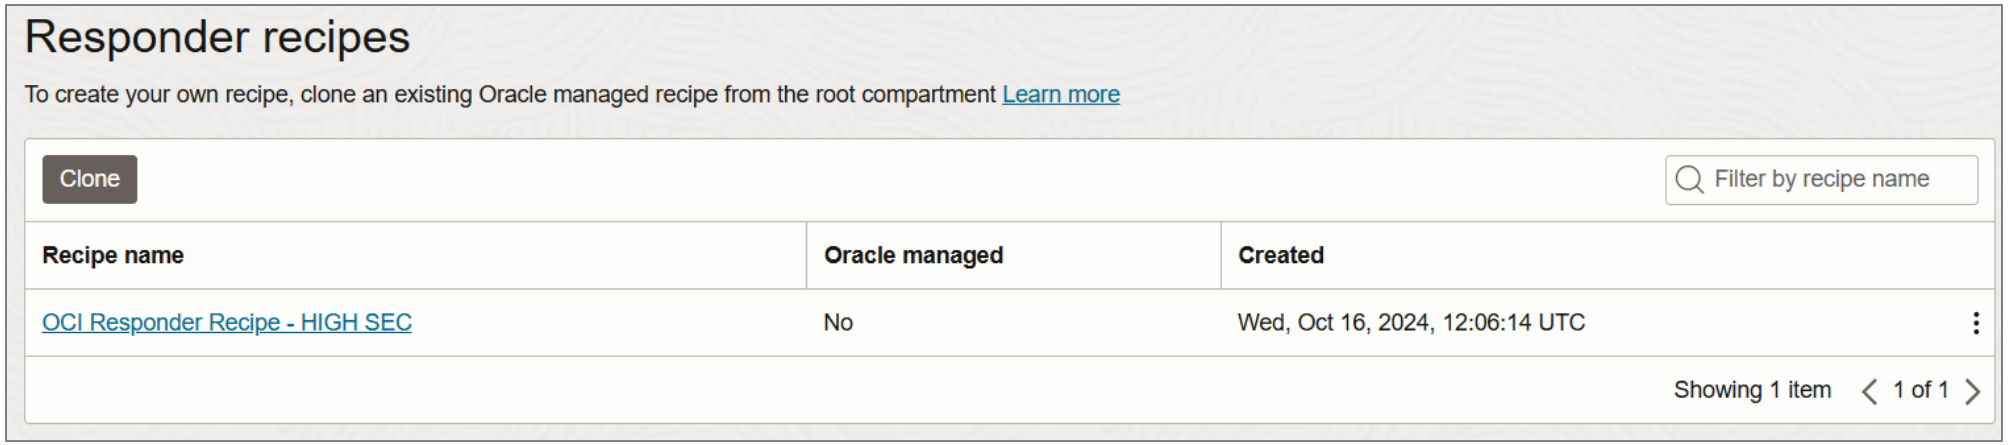

Auto-Resolve

- This is configured on level targets according the **Responder** rules.
- Requires an IAM policy to allow Cloud Guard the action.

```
Allow service cloudguard to manage buckets in compartment
comp-doag-high-sec
```

- The Detector and Responder rule must be enabled.

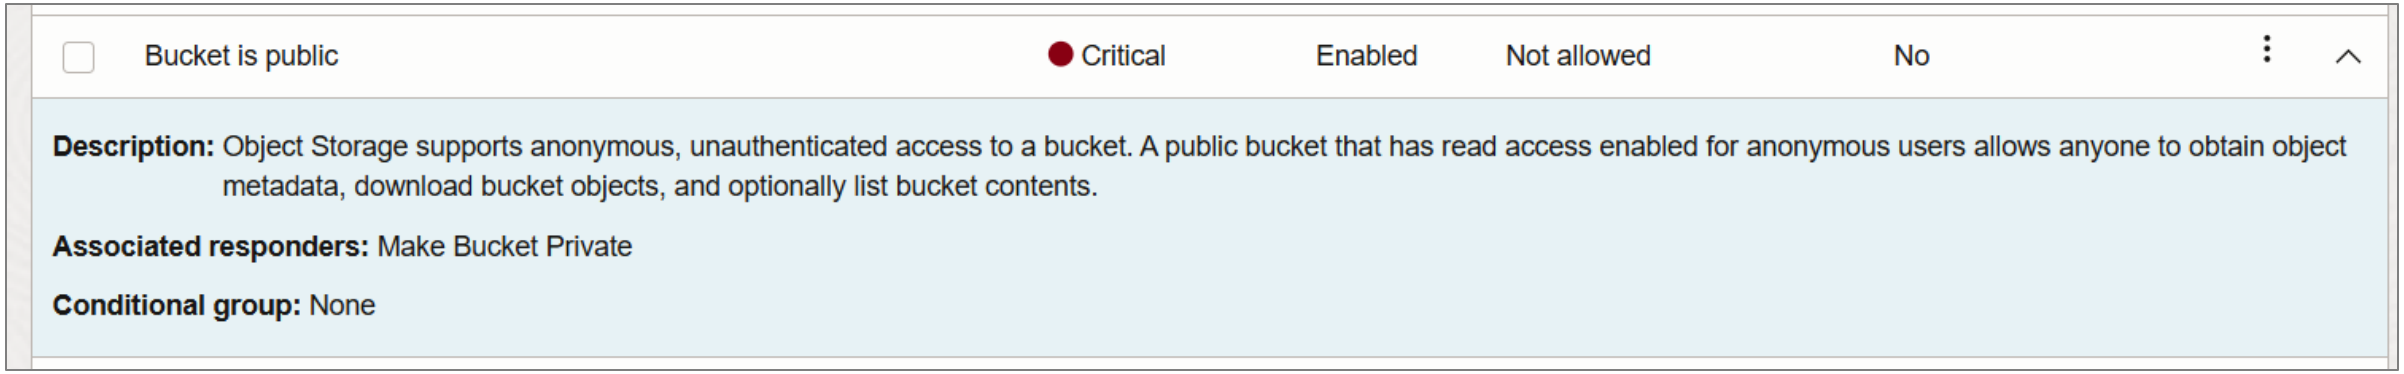

How to enable Auto-Resolve

1. Search **Cloud Guard Target Group**
2. Select **Responder Recipe - Responder Rule** – click on three dots / edit
3. Change **Setting** and enable **checkbox** for confirmation

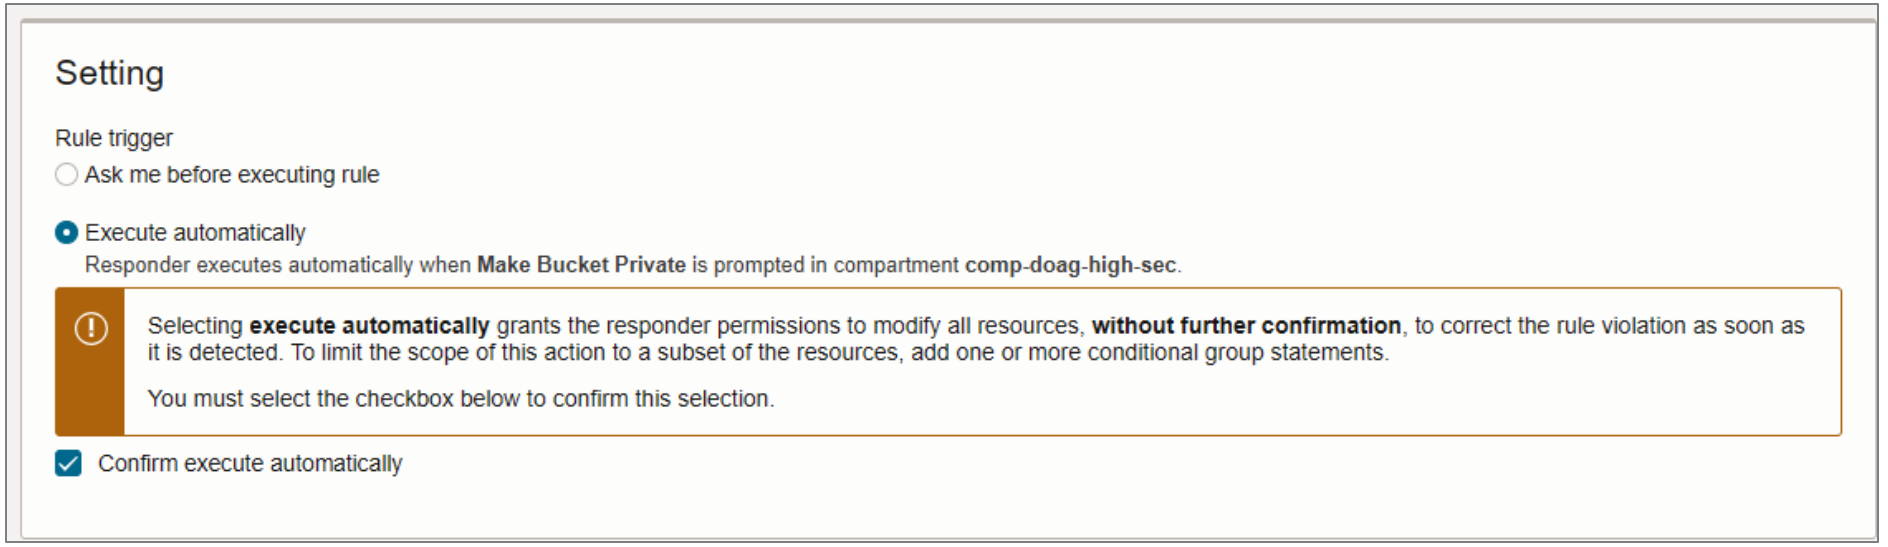

Conditional Execution

Key points about the **conditional execution** of a responder recipe:

- Specify conditions under which a responder rule will be executed.
- Define conditional groups with parameters like region, resource type, or tags.
- The rule is enforced only if all conditions in the group are met.
- Helps create precise, context-aware security responses, reducing unnecessary actions.

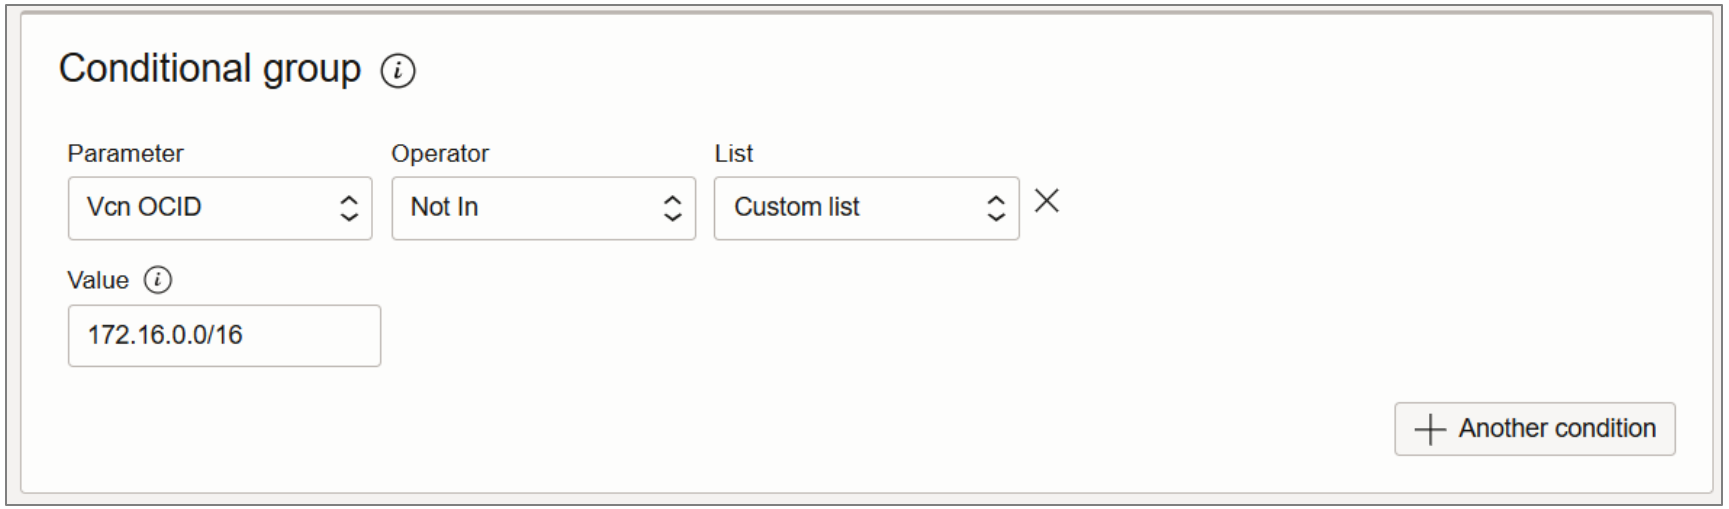

Problem solving

- **Remediate**: When you click the "Remediate" button, you are instructing Cloud Guard to take action to fix the identified problem. This typically involves executing a responder rule that addresses the issue, such as changing a configuration setting or applying a security patch. The goal is to resolve the problem so that it does not occur again.
- **Mark As Resolved:** Used to indicate that a problem has been addressed and resolved. When you mark a problem as resolved, you are telling Cloud Guard that the issue was indeed a problem, but you have taken the necessary actions to handle it.
- **Dismiss**: Clicking the "Dismiss" button indicates that you have reviewed the alert and decided that no action is necessary. This could be because the alert is a false positive, or the issue is not relevant or critical. Dismissing an alert will close it without taking any corrective action.

Processing Problems

- If not auto-resolved, see what you have to do for remediation.
- Remediated: Fixed by Cloud Guard Responder
- Resolved: Fixed by other processes
- Dismissed: Ignored / closed

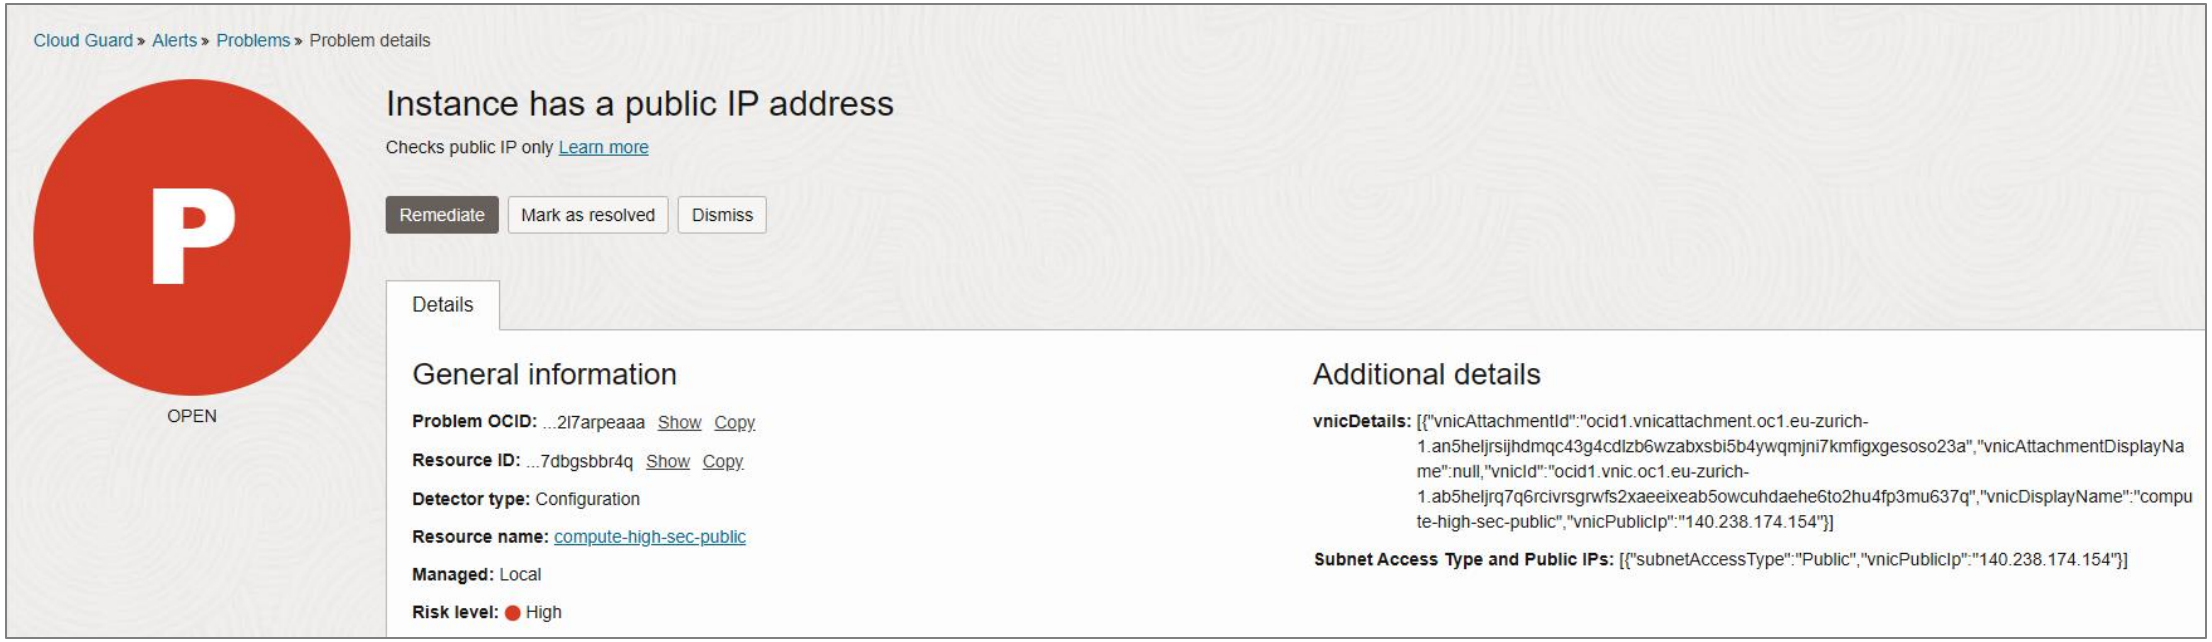

Notification

- Be informed whenever Cloud Guard has detected or auto-resolved the problem by definig Rules.

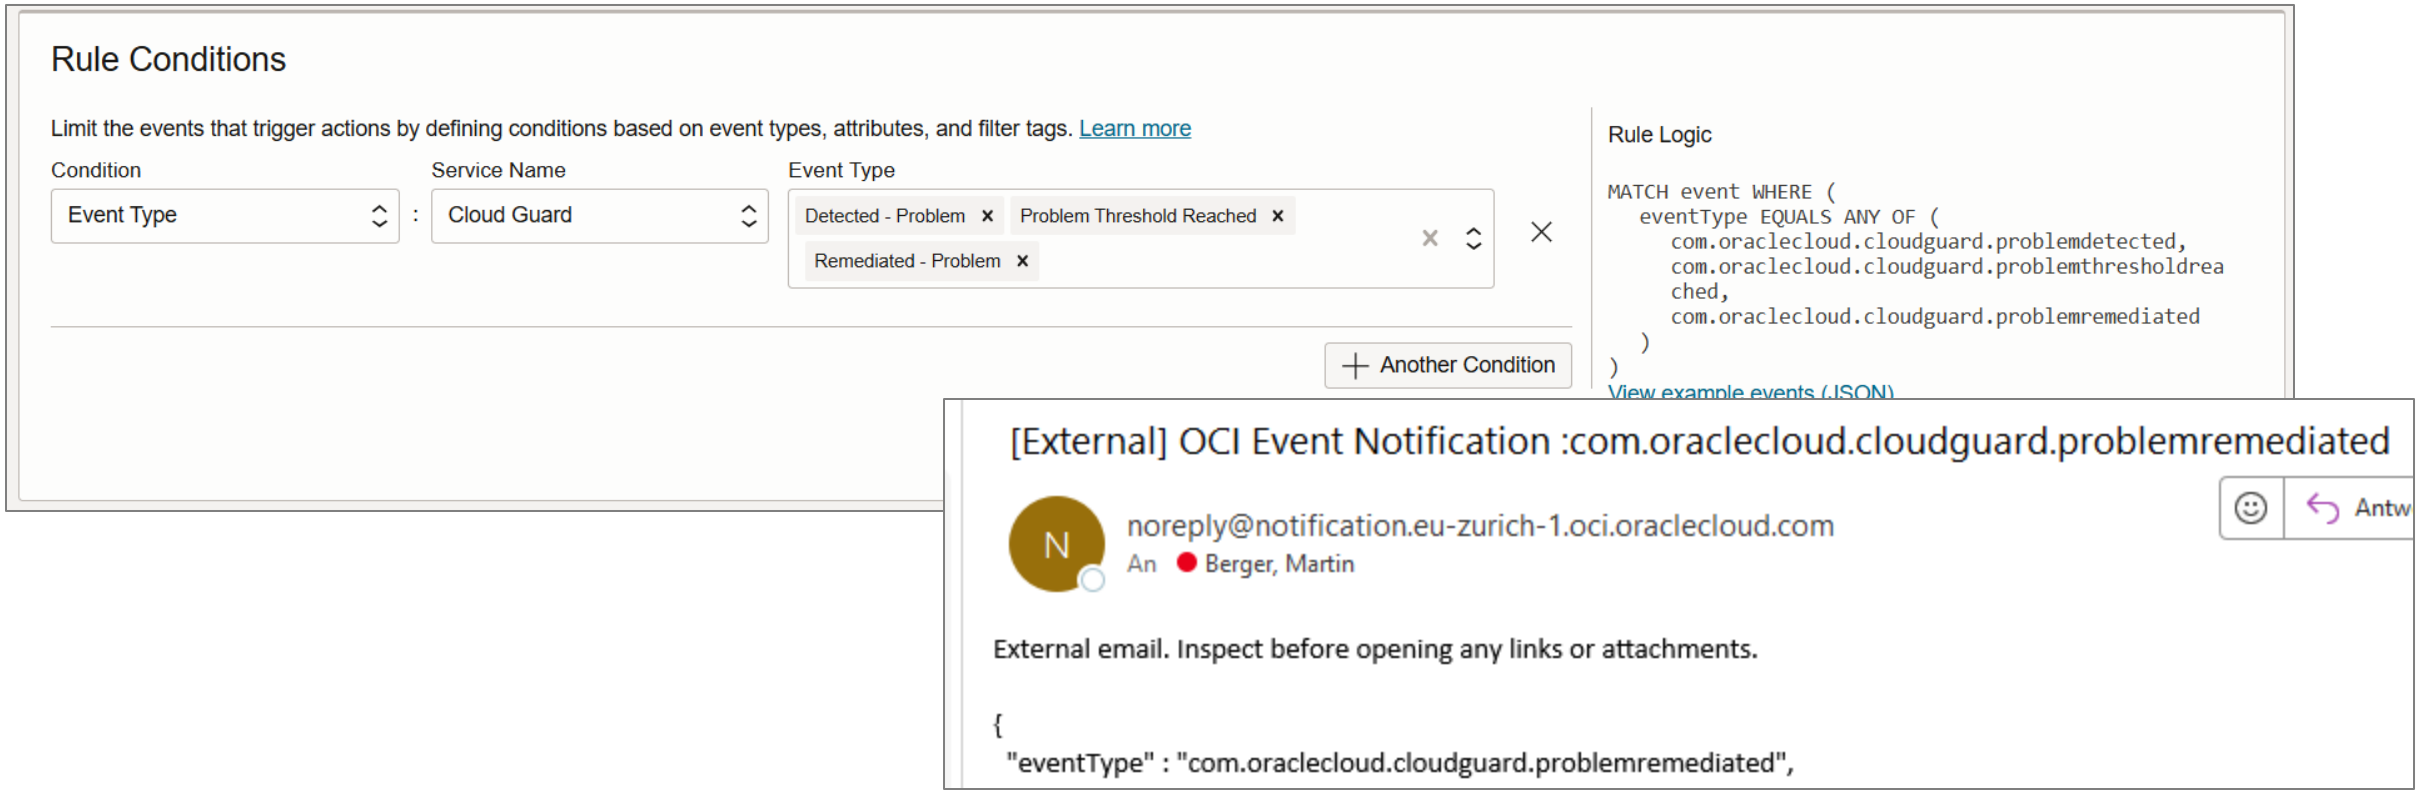

Notification

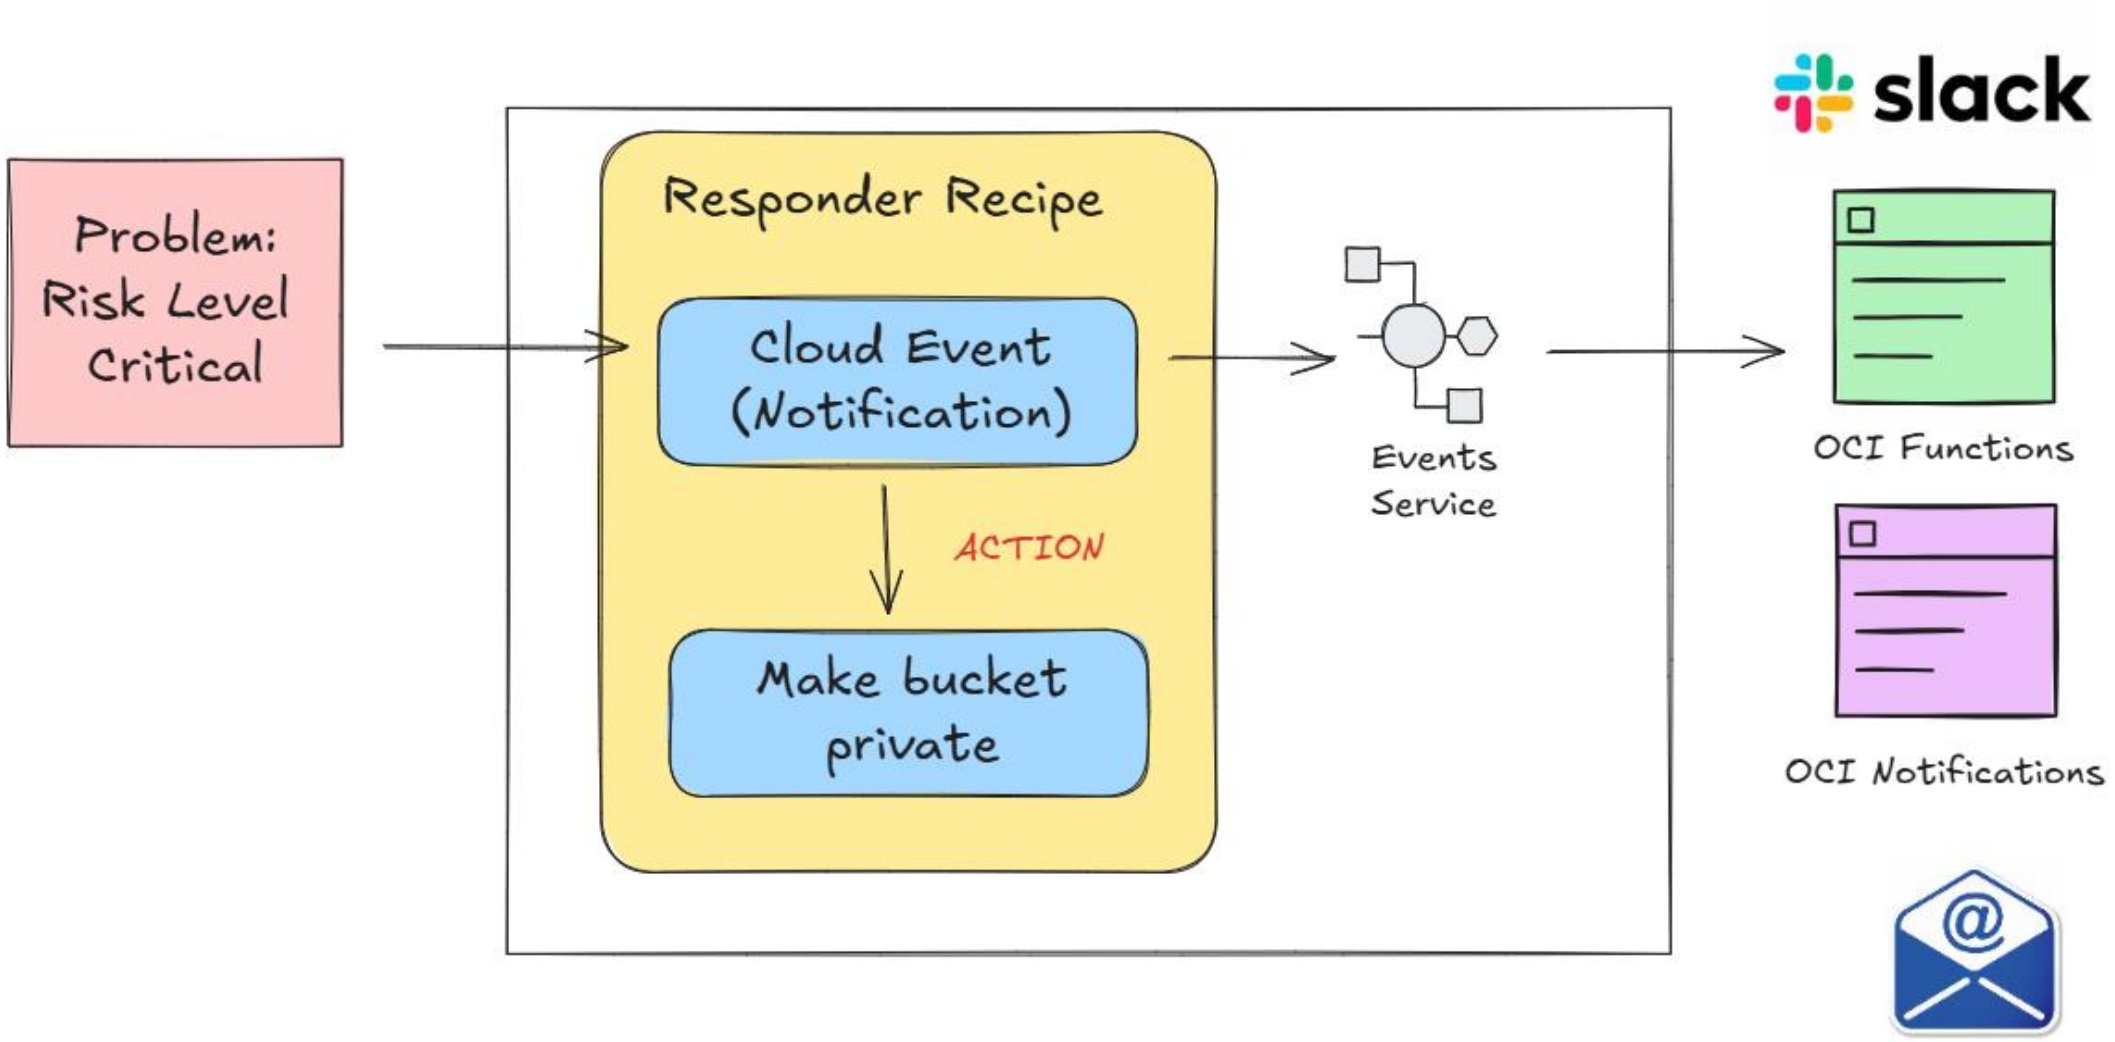

Queries – Verify Instances

- Use queries to get critical information about the current state of your compute instances via **OS Agent** plugin.
- Instance Security uses **Osquery**, which leverages a relational data model to describe an instance.
- A little bit tricky to configure (IAM & Firewall)

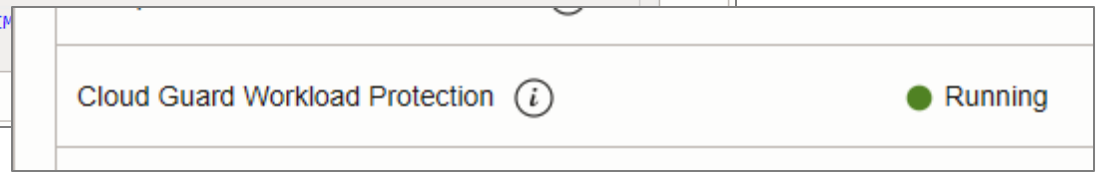

Queries – Results

Results

Running query

1 of 1 instances completed

Showing **1** result.

Instance OCID	Status	Region	Time submitted
...7dbgsbbr4q	Completed	Switzerland North (Zurich)	Wed, Oct 16, 2024, 19:04:06 UTC

memory_used	name	pid
1768.8900000000001	oci-wlp	2378
1764.0	runcommand	5043
1763.53	oci-vulnerabili	2333
1763.28	gomon	4901
1762.5599999999999	agent	2196

Sh

Best Practices for Using OCI Cloud Guard

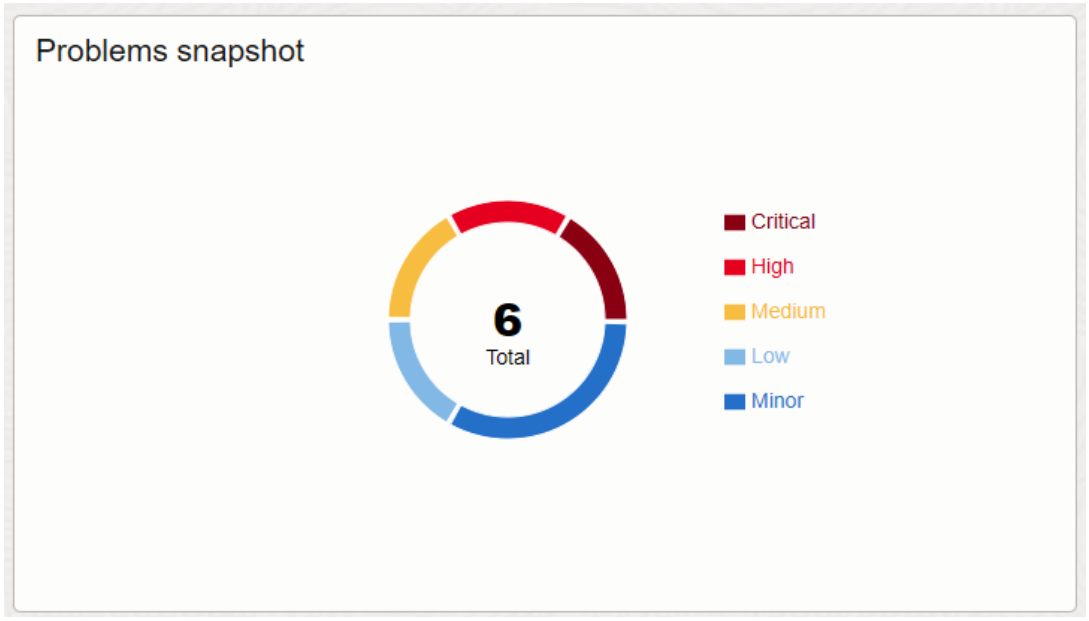

- **Enable** — Enable it, it is for free
- **Define** — Define own recipes and rules
- **Stay** — Stay informed
- **Force** — Force users to remediate issues
- **Create** — Create guidelines for creating OCI ressources

2

Hands-On Labs

What about the Database Security?

Hands-On Labs

List of Hands-On Labs

- **Manual Remediation** – Configure manual remediation for detected security risks.
- **Auto Remediation** – Automate responses to detected security vulnerabilities
- **Notification Setup** – Set up notifications to monitor and alert on security-related events.

Oracle Cloud Infrastructure Security

Data Safe

Martin Berger
Stefan Oehrli

Data Safe

A short Introduction

1. Overview
2. Security Assessment
3. User Assessment
4. Database Audit
5. SQL Firewall
6. Sensitive Data Discovery
7. Data Masking
8. Hands-On Labs

1

Overview

Data Safe in a Nutshell

Data: The Crown Jewel of Your Organization

Protecting Data to Prevent Liability

- **Data is a Key Asset**: While data holds immense value, it can quickly turn into a major liability if not adequately protected.
- **Rising Cybercrime:** Cybercrime is expected to cause $8 trillion in global damages in 2023, with databases being prime targets due to their concentration of valuable information.
- **Increasing Regulations:** New and expanding data protection laws demand stricter security and accountability.

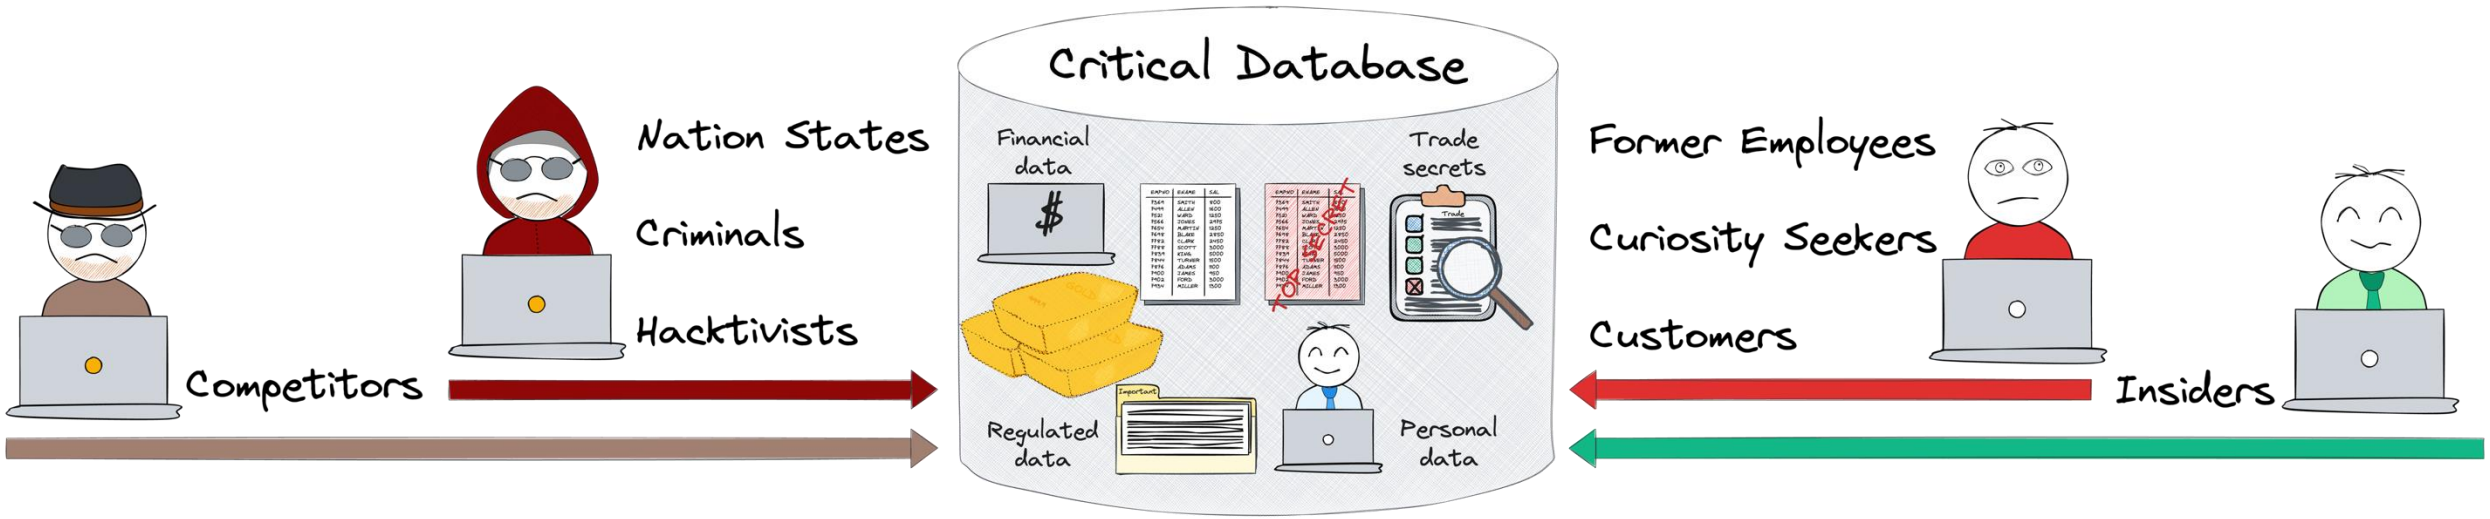

Why Oracle Data Safe?

Beyond Encryption and Patching – Comprehensive Database Security

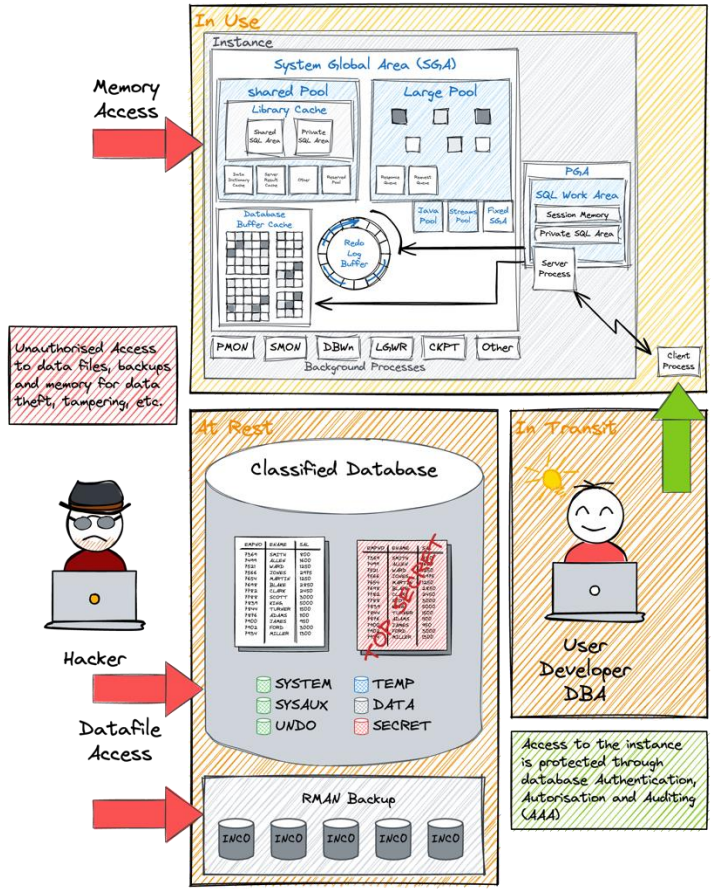

More Than Basic Security – Encryption and patching are crucial but incomplete.

Key Security Focus Areas:

- **Configuration Compliance** – Is the database securely configured?
- **User Risk & Monitoring** – Who are the highest-risk users, and what actions are they performing?
- **Audit & Compliance** – Which activities should be audited, and how do we manage and protect the audit logs effectively?
- **Sensitive Data Control** – What sensitive data is stored, and can exposure be minimized?

Holistic Security Management *– Oracle Data Safe offers centralized tools to address these areas, enhancing control, visibility, and compliance.*

Oracle Data Safe Overview

Centralized Security for Cloud and On-Premises DBs

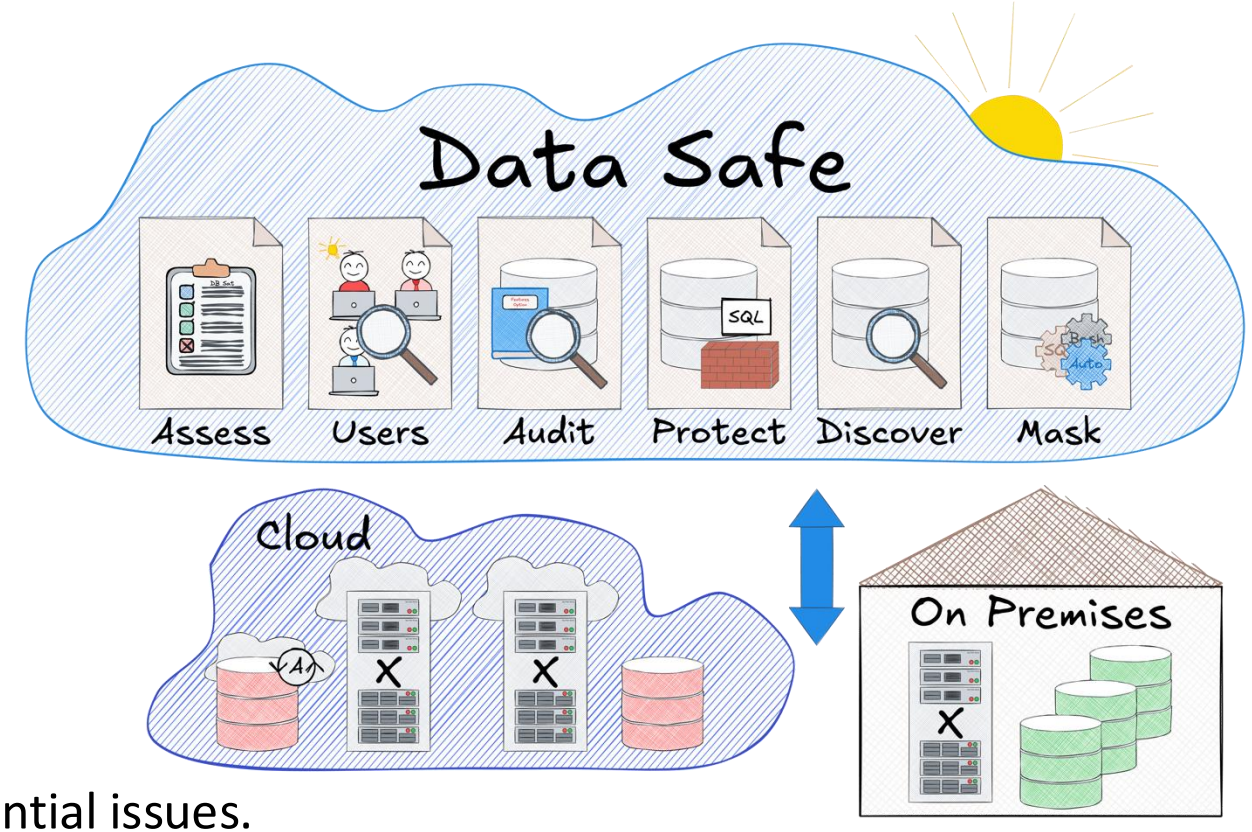

Unified Security Center – Manage and monitor security for cloud and on-premises databases.

Core Capabilities

- **Risk Dashboard:** Visualizes risk across users, data and configurations.
- **User Monitoring:** Audits user activities to detect potential issues.
- **Data Masking:** Protects sensitive data in test environments.
- **Future-Ready:** Regular updates add new security features.

Key Benefits

- **Easy & Efficient:** No special expertise required—quick setup and automated security.
- **Risk Mitigation:** Saves time and reduces exposure to threats.
- **Comprehensive Protection:** Defense-in-depth security for all Oracle databases.

Oracle Data Safe Overview

Key Features for Comprehensive Database Security

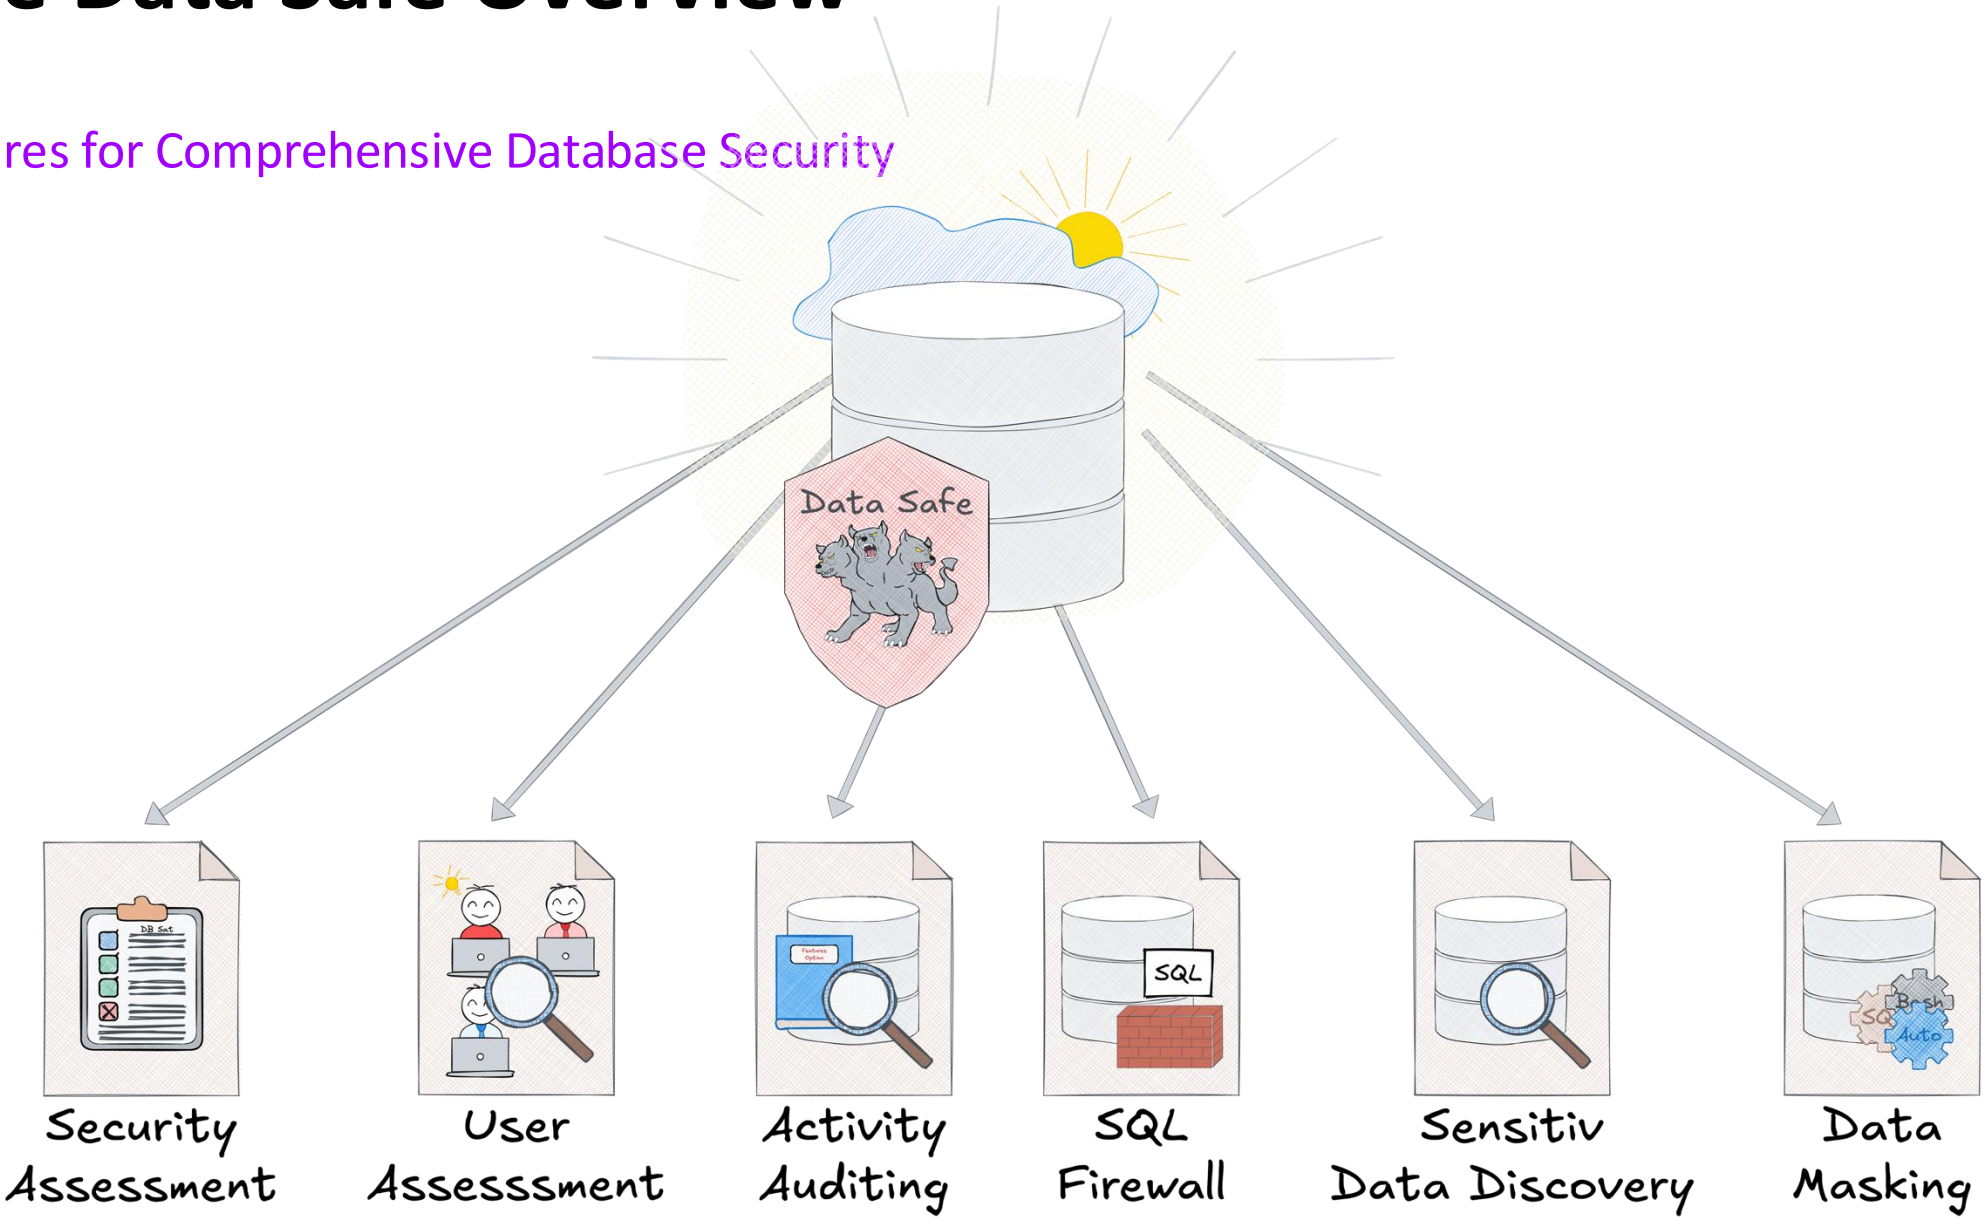

Data Safe Dashboard

Key Security Indicators at a Glance

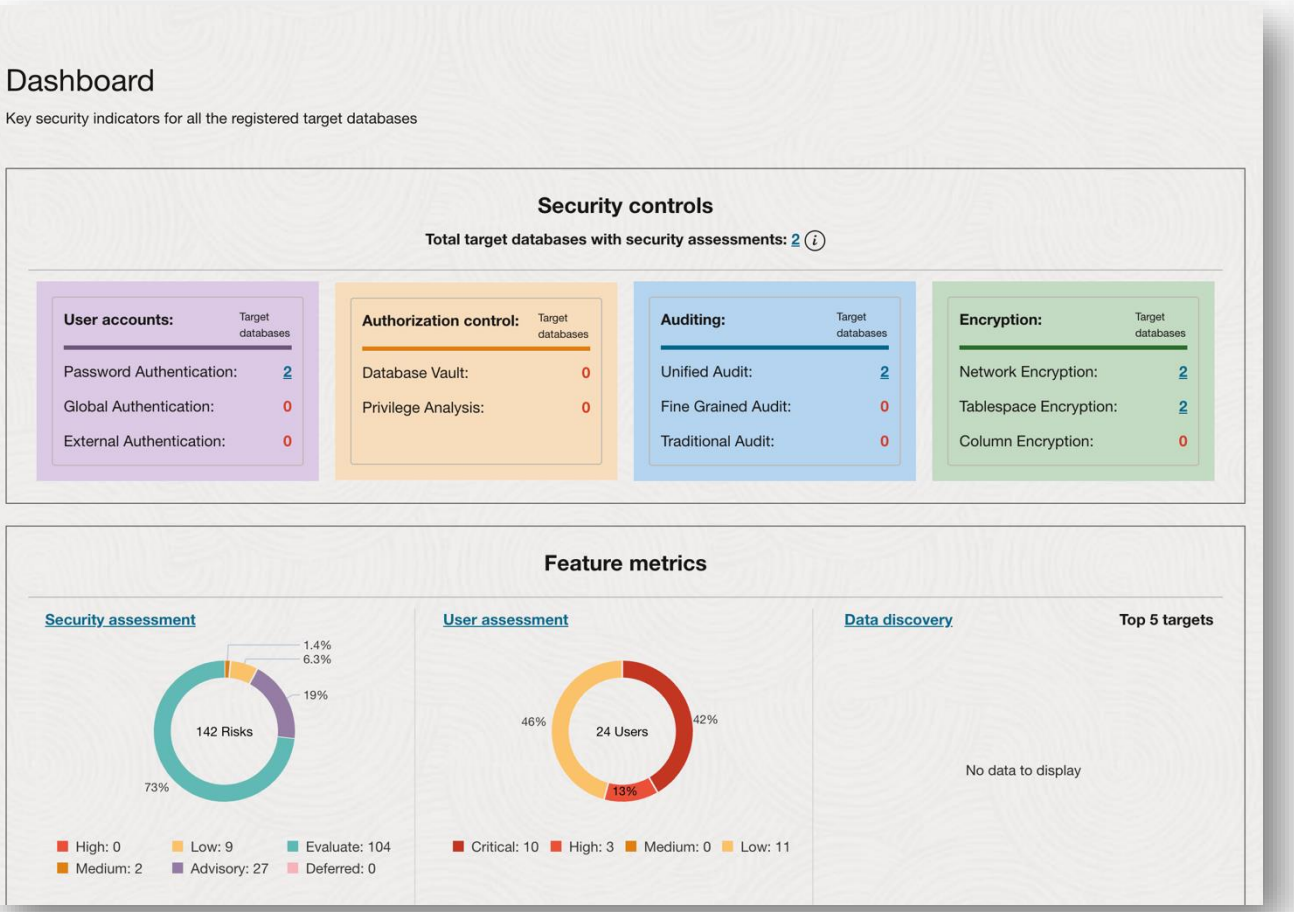

2

Security Assessment

Database Security Assessment

Database Security Assessment

Comprehensive Assessment

- Security parameters
- Security controls in use
- User roles and privileges

Landscape-Wide Risk Overview

- Identified risks with actionable recommendati

Compliance Mappings

- Aligns with GDPR, STIG, CIS standards

Database Security Assessment

Detecting Configuration Drifts

Establish a Security Baseline

- Define a baseline configuration for secure operations.

Automatic Comparisons

- New assessments are automatically checked against the baseline.

Notifications for Drift

- Receive alerts and review any deviations from the baseline to address potential risks.

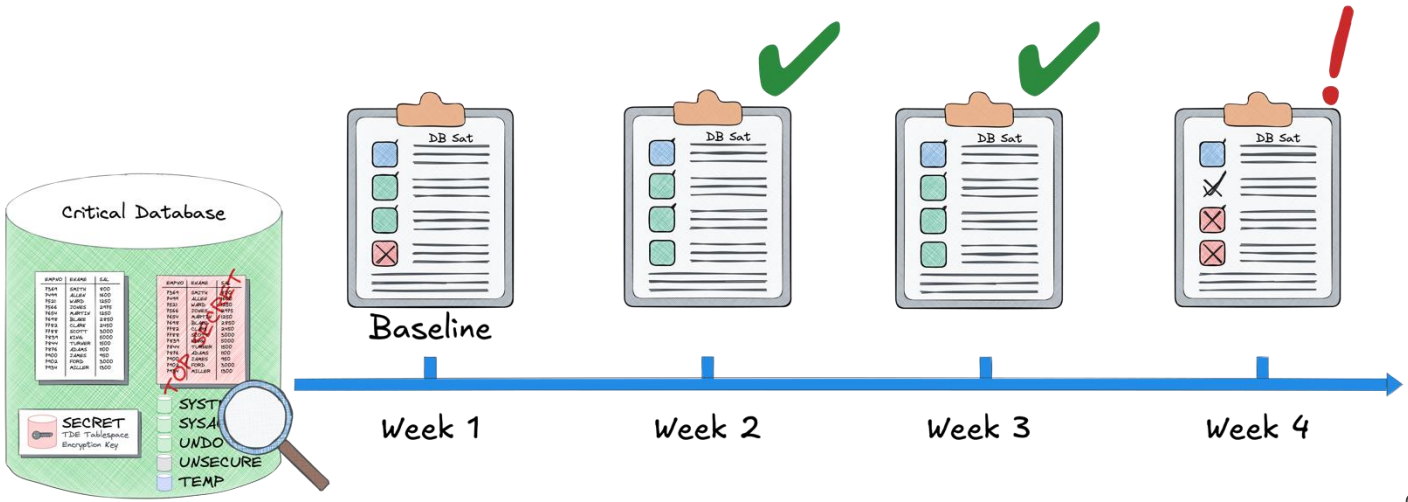

3

User Assessment

User Security Assessment

User Risk Assessment

Reducing Risk Through Role and Privilege Management

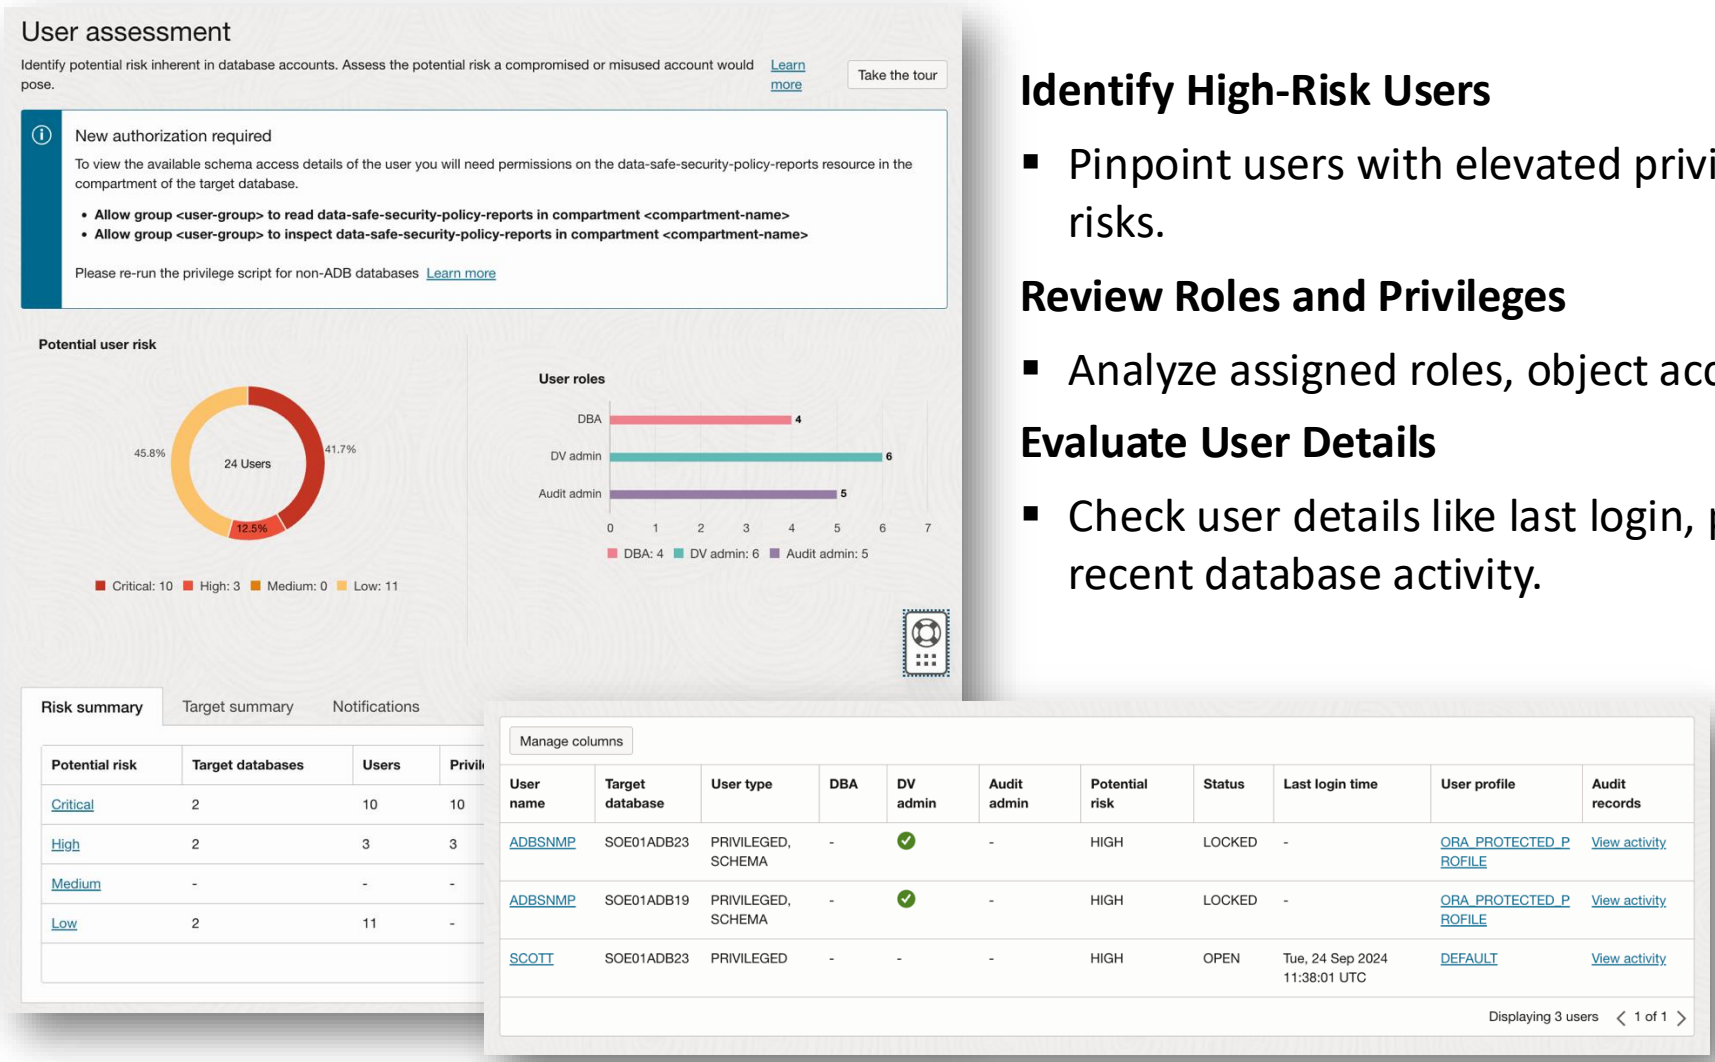

Identify High-Risk Users

- Pinpoint users with elevated privileges who may pose security risks.

Review Roles and Privileges

- Analyze assigned roles, object access, and system privileges.

Evaluate User Details

- Check user details like last login, password change history, and recent database activity.

User Risk Assessment

Detecting User and Entitlement Changes

Periodic User Assessments

- Regularly evaluate user accounts and privileges.

Compare with Previous Results

- Identify changes by comparing new assessments with prior ones.

Alerts for Changes

- Receive notifications for newly added users or modifications in user entitlements.

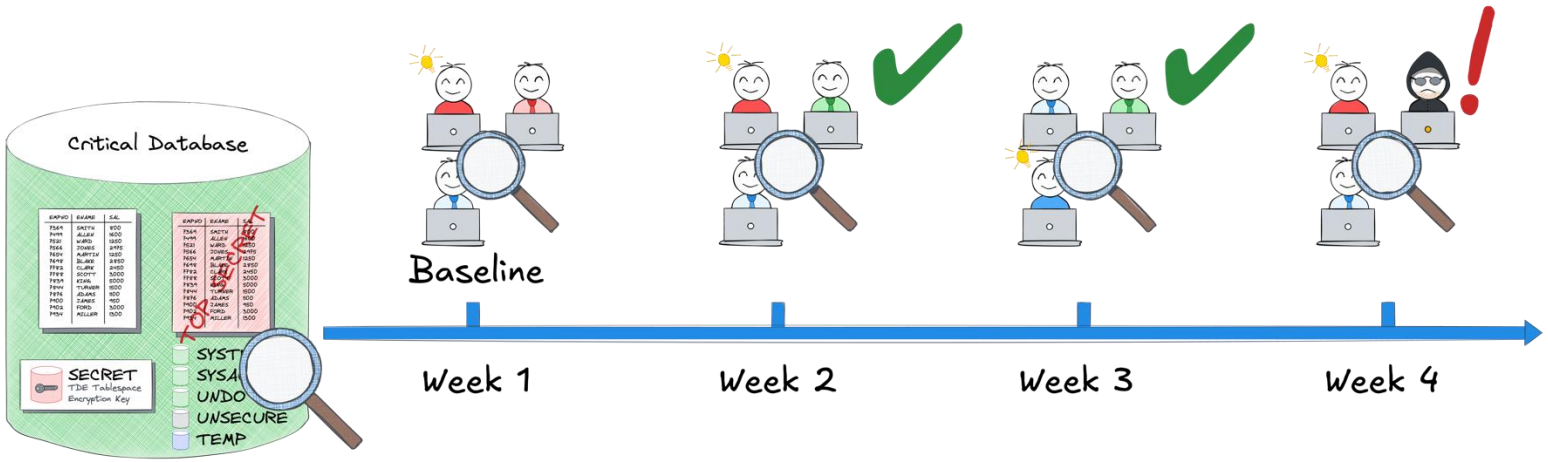

User Profile Insight

Evaluating Password-Related Attributes in User Profiles

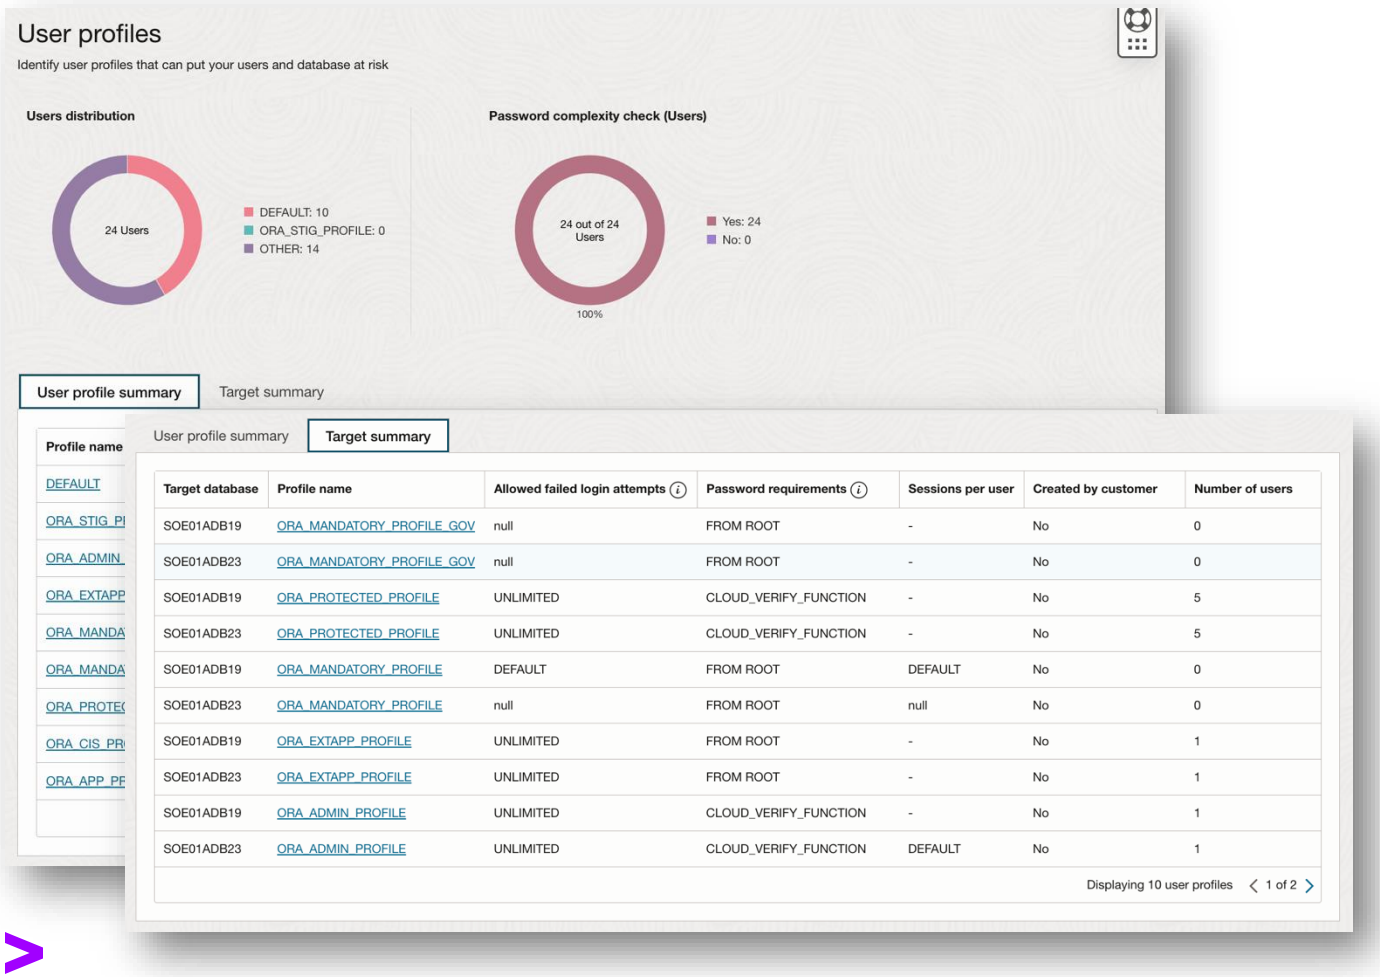

Review User Profiles

- Examine existing profiles and their security parameters.

Map Profiles to Users

- Identify which users are assigned to specific profiles.

Detect Weak Profiles

- Easily spot profiles and users lacking password complexity or other security measures.

4

Database Audit

Central Audit Management and Reporting

User Activity Auditing

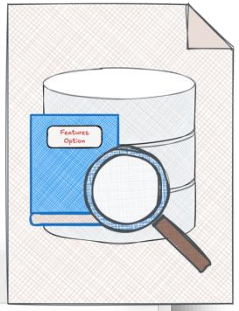

Track User Actions with Comprehensive Auditing and Reporting

Audit and Compliance Policies

- Define policies for auditing, compliance, and alerts.

Data Collection

- Gather audit data from databases, focusing on sensitive operations.

Audit Reports:

- **Interactive Reports –** For in-depth forensic analysis.
- **Summary & Detailed Reports –**Tailored insights for different auditing needs.
- **PDF Compliance Reports –** Ready-made reports for regulatory compliance.

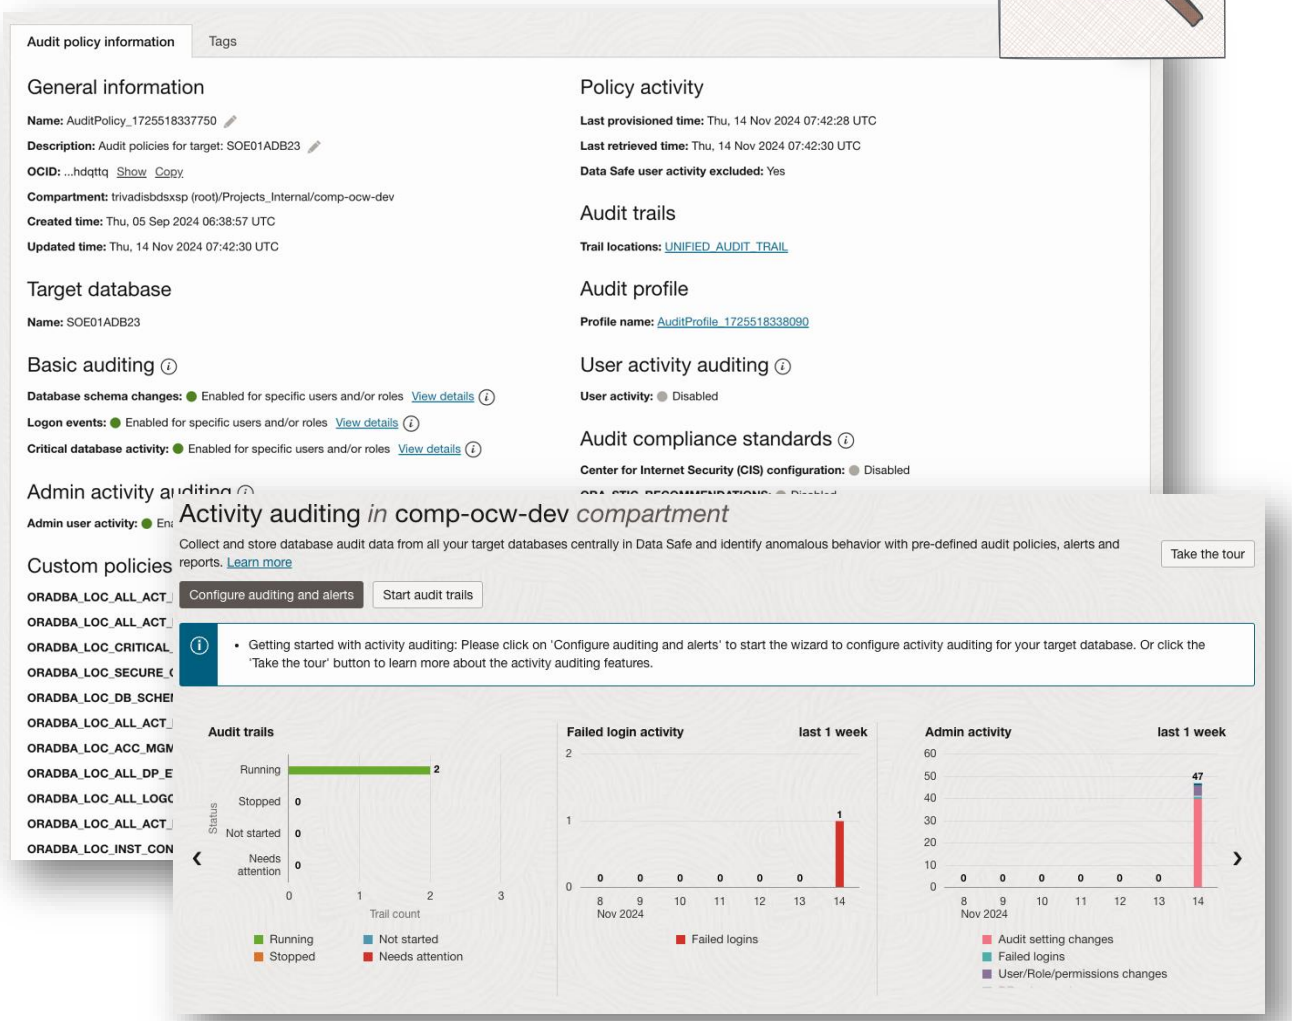

Audit Insights Dashboard

Fine-Tune Your Audit Policies

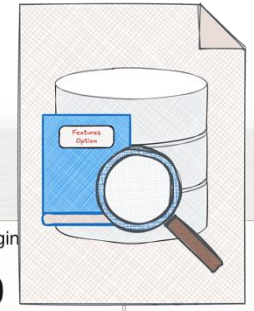

Gain Deeper Insights

- Analyze audit data to optimize policies.

High-Volume Policies

- Identify which audit policies generate the most records.

Top Targets

- See which databases produce the highest audit volumes.

Frequent Access Patterns

- Discover the most-accessed objects and schemas.

Audit Data Retention Management

Ensure Compliance with Retention Policies

Retention Periods

- Store audit data for up to 7 years
- **Online Retention** – Up to 1 year
- **Archive Retention** – Up to an additional 6 years

Configurable Retention

- Set retention periods globally or for specific targets.

Easy Access to Archived Data

- Quickly retrieve audit records as needed.

Fully Managed by Data Safe

- No additional fees for archiving audit data….
  … as long as it is not too much data per month

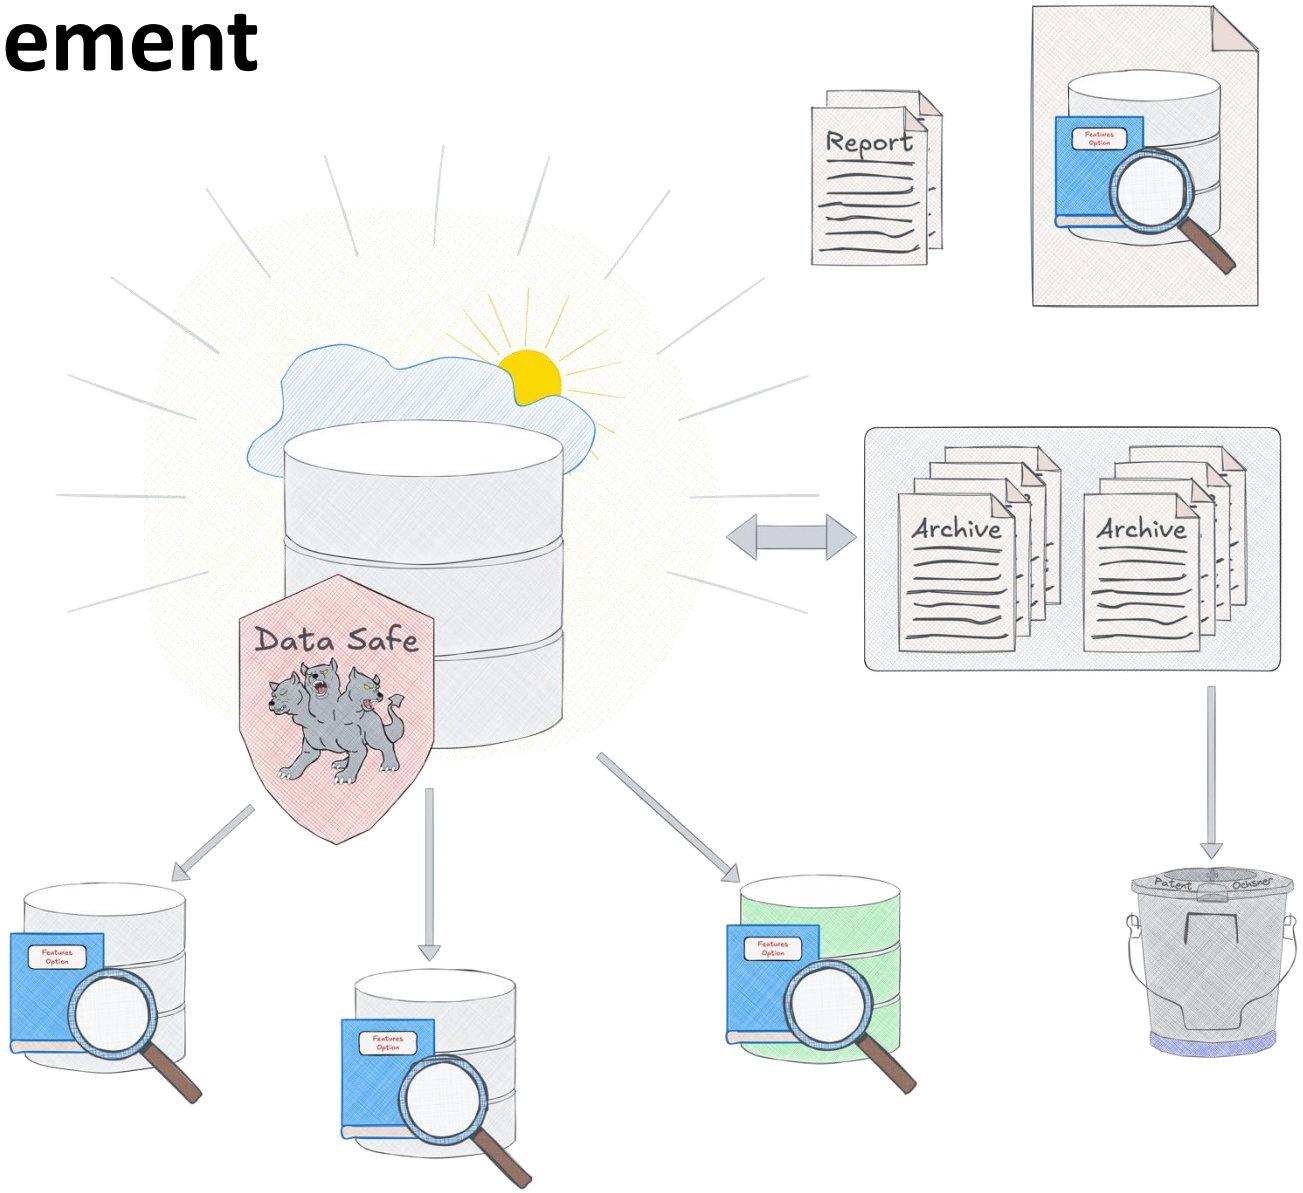

5

SQL Firewall

23ai latest Security Enhancement

SQL Firewall

Prevent SQL Injection and Unauthorized Access

Real-Time Protection

- Restricts database access to only authorized connections and SQL statements.

Monitor or Block Violations

- Choose to either monitor or block any access violations.

Risk Mitigation

- Defends against SQL injection, unauthorized access, and credential abuse.

Currently only available in Oracle 23ai

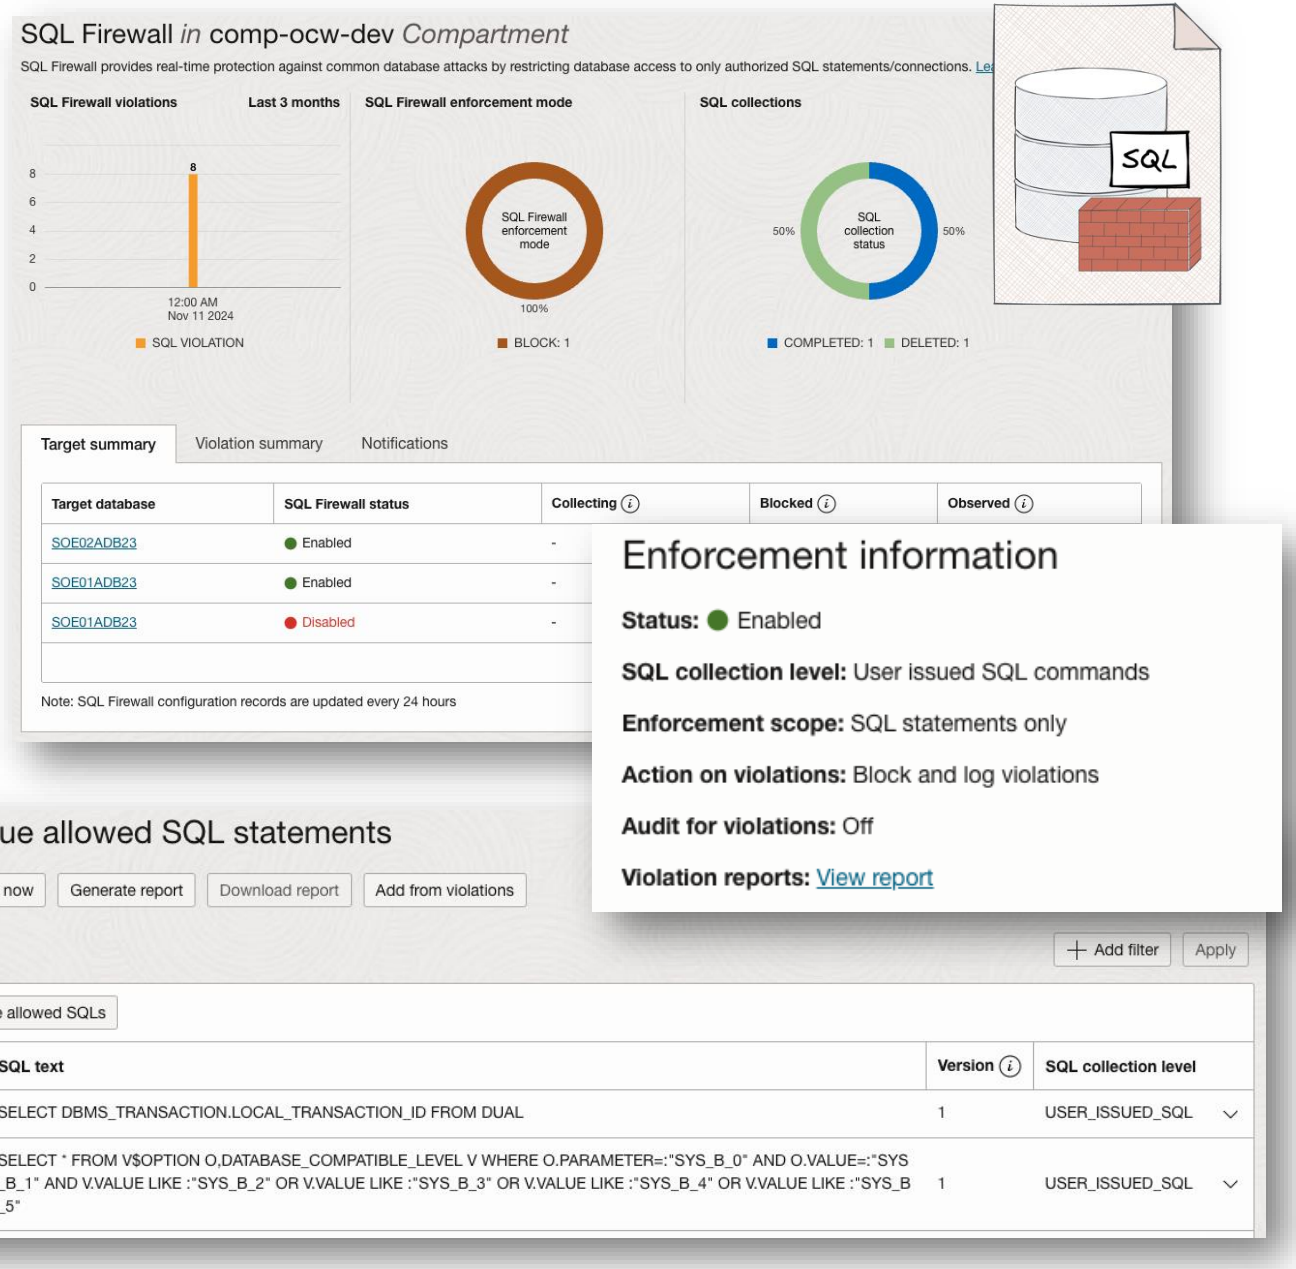

Navigating SQL Firewall – Processes

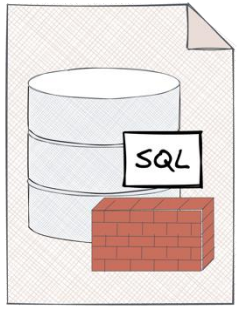

Understanding the Mechanics and Strategies for Optimal Deployment

Learning Stage

- **Collect** Enable collection of SQL statements and user connections.
- **Review & Configure** Review collected SQL statements and connections, then define allowed connections as needed.
- **Set Permissions:** Adjust allowed SQL statements and user connections based on security requirements.

Protecting Stage

- **Enable** the allow-list
- **Monitor violations** SQL Firewall raises violation for any unexpected access patterns to ensure compliance

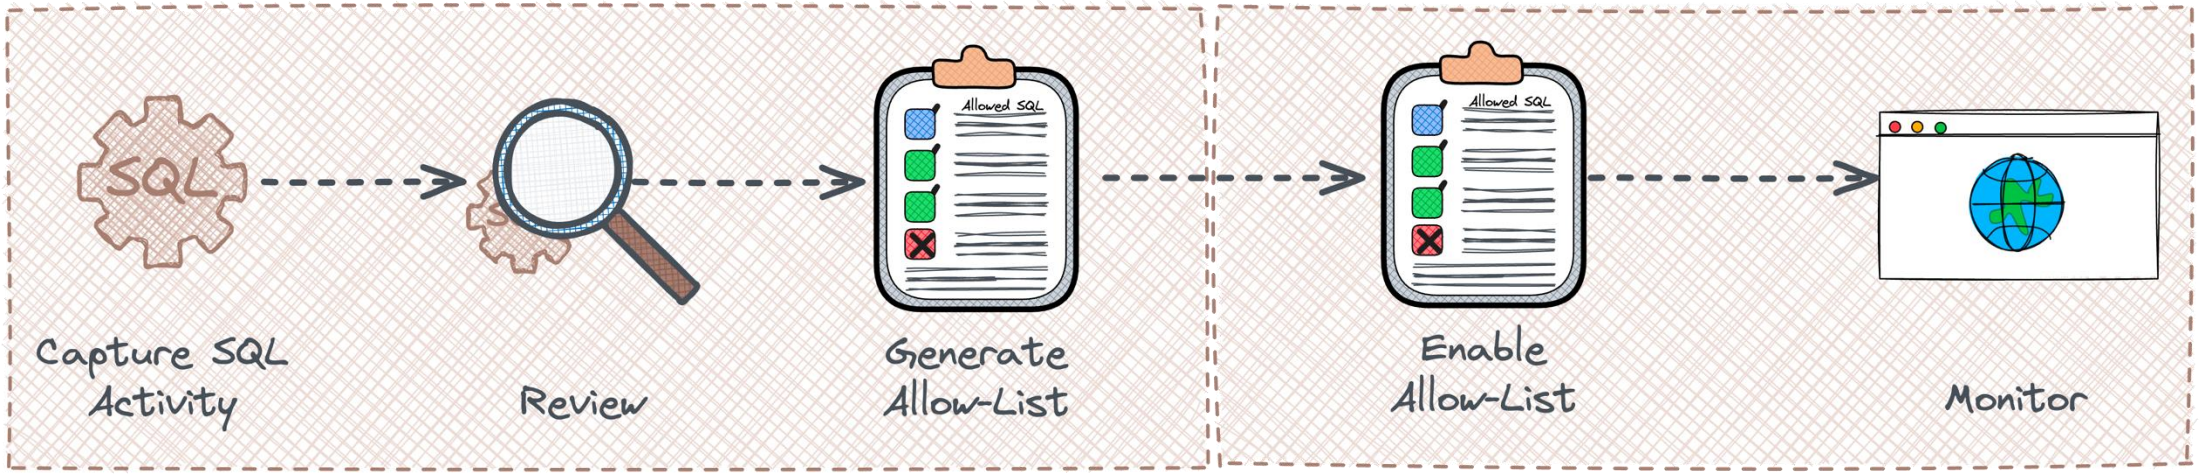

6

Sensitive Data Discovery

Where is my sensitive Data?

Sensitive Data Discovery

Identify and Classify Sensitive Data to Prioritize Security Efforts

Predefined Sensitive Types

- Detects and classifies over 150 sensitive data types.

Custom Sensitivity Definitions

- Allows user-defined sensitive data types for tailored discovery.

Incremental Discovery

- Supports ongoing scans to identify newly added sensitive data.

Detailed Reporting

- Provides reports on the amount and types of sensitive data.

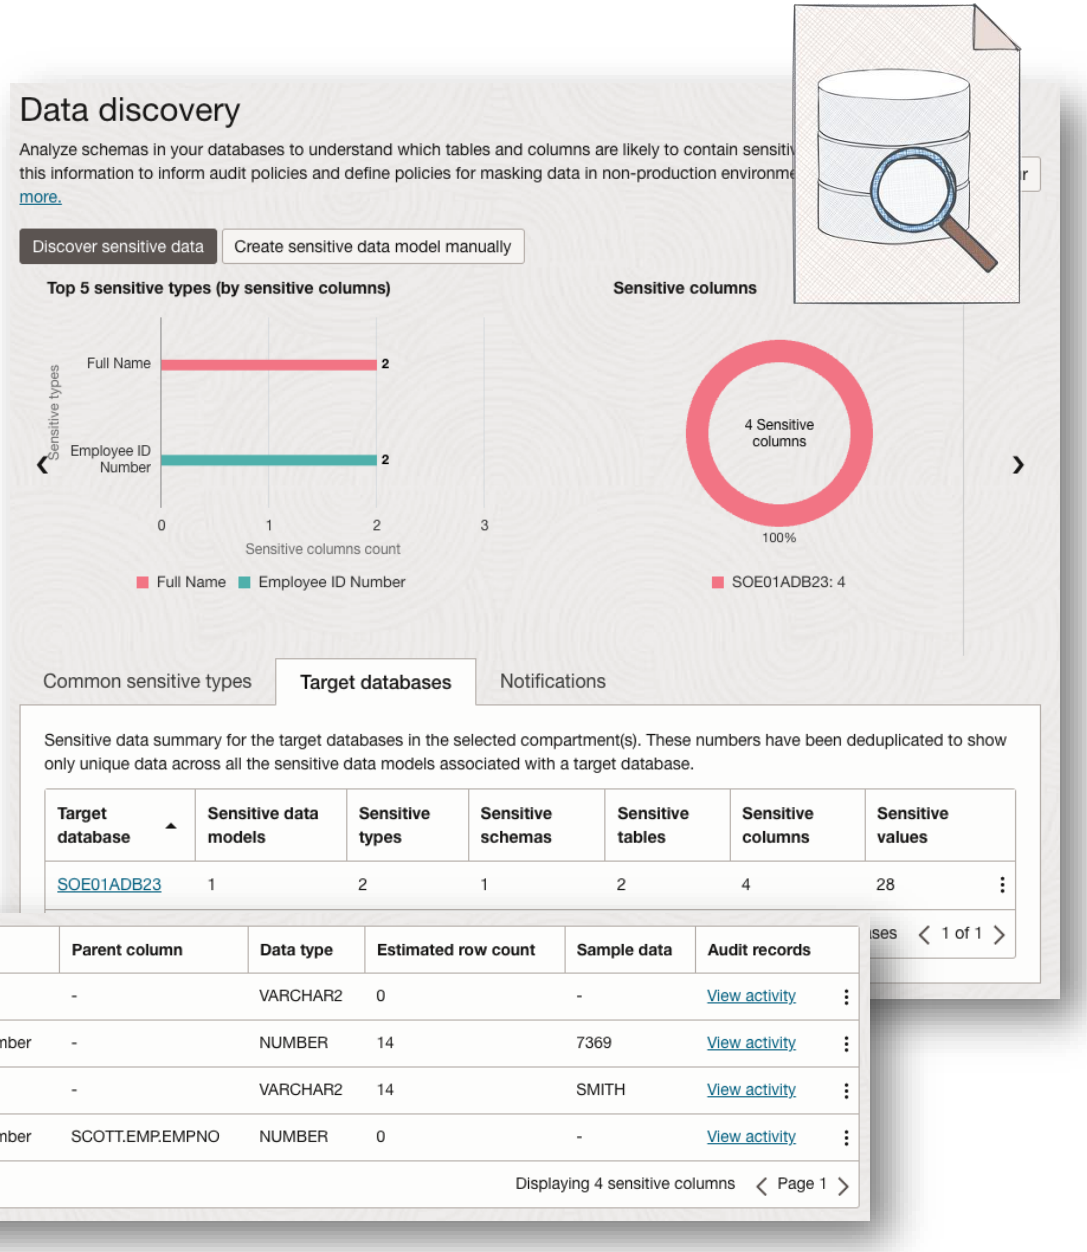

Sensitive Data Discovery

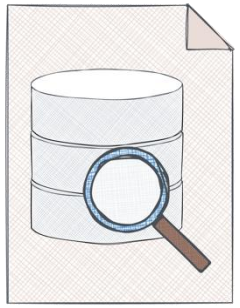

150+ Pre-Defined Sensitive Data Types

Identification	Biographic	IT	Financial	Healthcare	Employment	Academic
SSN Name Email Phone Passport DL Tax ID …	Age Gender Race Citizenship Address Family Data Date of Birth Place of Birth …	IP Address User ID Password Hostname GPS location …	Credit Card CC Security PIN Bank Name Bank Account IBAN Swift Code …	Provider Insurance Height Blood Type Disability Pregnancy Test Results ICD Code …	Employee ID Job Title Department Hire Date Salary Stock …	College Name Grade Student ID Financial Aid Admission Date Graduation Date Attendance …
		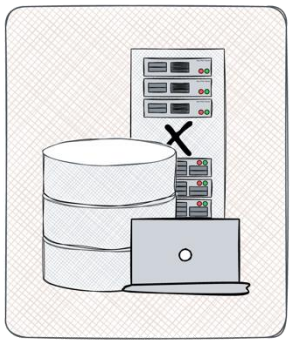		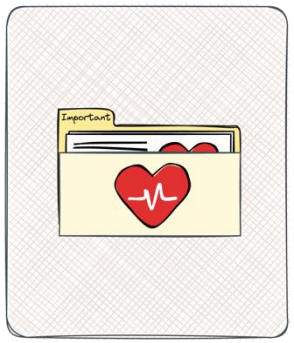	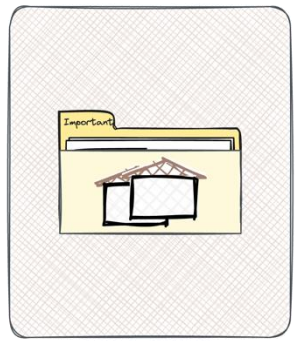	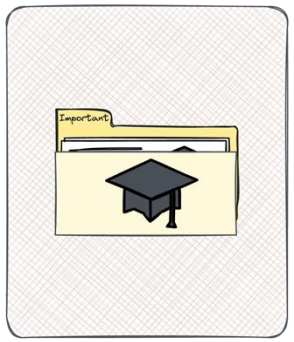

7

Data Masking

Mask test and dev data using sensitive Data Masking

Sensitive Data Masking

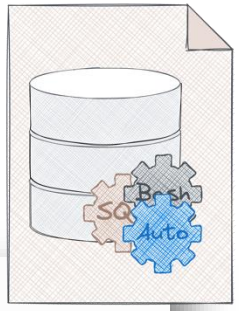

Minimize Sensitive Data Exposure in Development and Testing

Sensitive Data Masking

- Mask data identified as sensitive to protect privacy.

Predefined Masking Formats

- Over 50 predefined formats for common data types.

Automated Format Selection

- Automatically selects appropriate formats based on data type.

Custom Masking Options

- Supports user-defined formats for specialized requirements.

Advanced Transformations

- Offers rich masking options for complex data scenarios.

Comprehensive Masking Report

- Detailed reports on masked data for compliance and review.

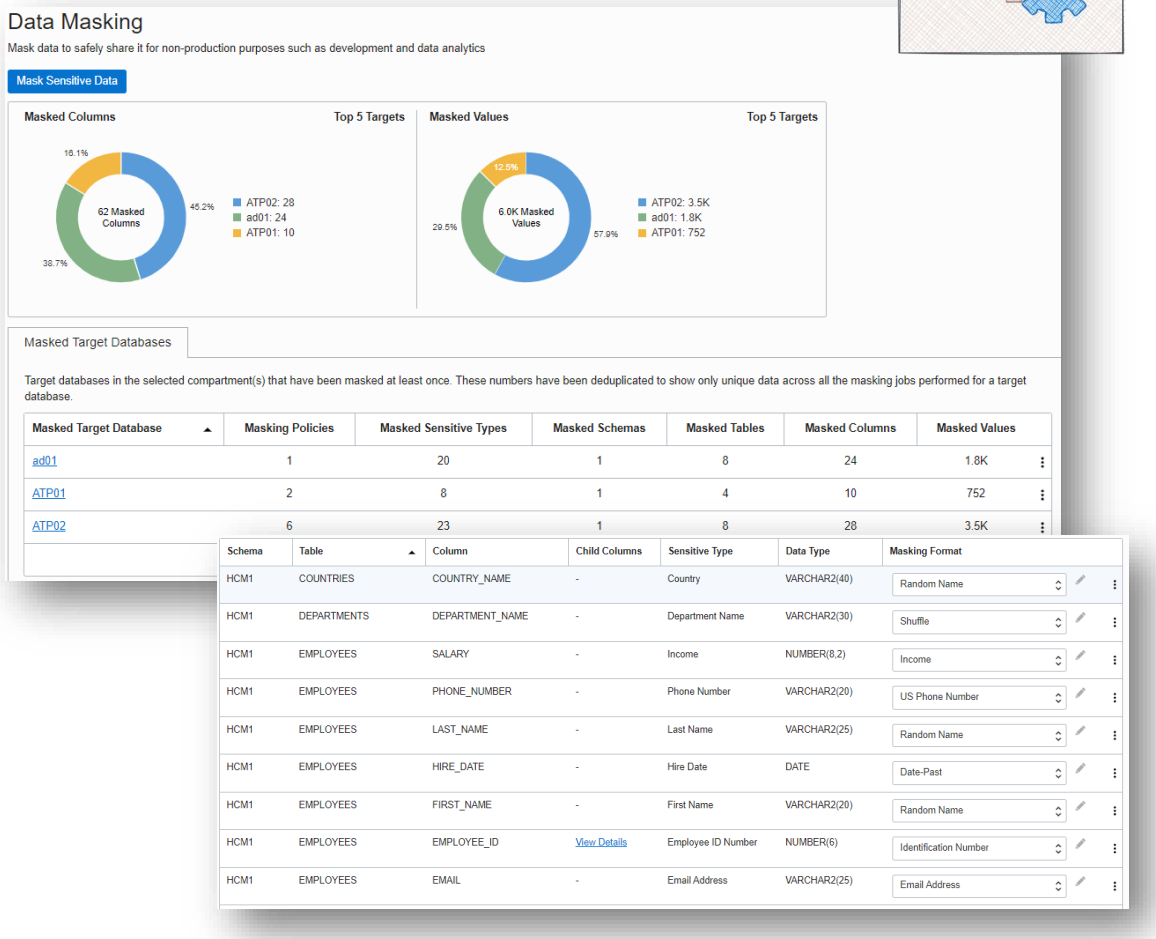

Sensitive Data Masking

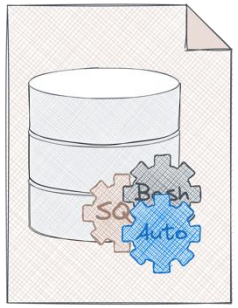

Pre-Defined Masking Formats for Comprehensive Data Protection

50+ Pre-Defined Formats

- Data Safe offers a wide range of masking options, including
  - **Specific Formats –** Social security number, credit card number, email address, and more.
  - **Generic Formats –** Random date, random number, random name, fixed number, fixed string.
  - **Advanced Techniques –** Format-preserving randomization, regular expression, data truncation, group masking.

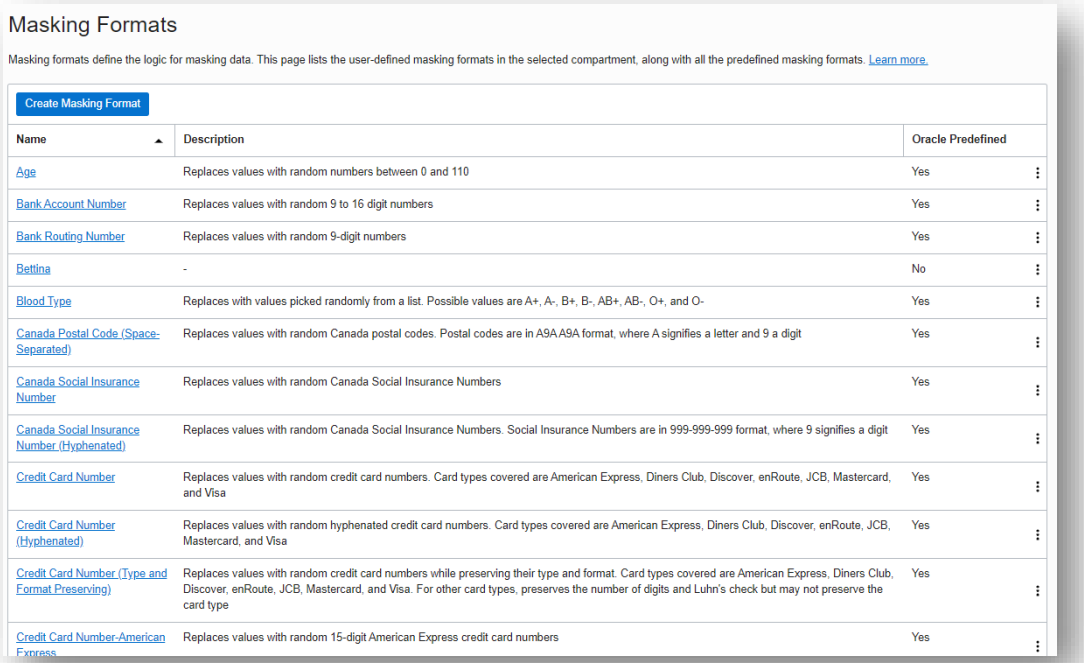
Masking Formats

Masking formats define the logic for masking data. This page lists the user-defined masking formats in the selected compartment, along with all the predefined masking formats. Learn more.

Create Masking Format

Name	Description	Oracle Predefined
Age	Replaces values with random numbers between 0 and 110	Yes
Bank Account Number	Replaces values with random 9 to 16 digit numbers	Yes
Bank Routing Number	Replaces values with random 9-digit numbers	Yes
Bettina	-	No
Blood Type	Replaces with values picked randomly from a list. Possible values are A+, A-, B+, B-, AB+, AB-, O+, and O-	Yes
Canada Postal Code (Space-Separated)	Replaces values with random Canada postal codes. Postal codes are in A9A A9A format, where A signifies a letter and 9 a digit	Yes
Canada Social Insurance Number	Replaces values with random Canada Social Insurance Numbers	Yes
Canada Social Insurance Number (Hyphenated)	Replaces values with random Canada Social Insurance Numbers. Social Insurance Numbers are in 999-999-999 format, where 9 signifies a digit	Yes
Credit Card Number	Replaces values with random credit card numbers. Card types covered are American Express, Diners Club, Discover, enRoute, JCB, Mastercard, and Visa	Yes
Credit Card Number (Hyphenated)	Replaces values with random hyphenated credit card numbers. Card types covered are American Express, Diners Club, Discover, enRoute, JCB, Mastercard, and Visa	Yes
Credit Card Number (Type and Format Preserving)	Replaces values with random credit card numbers while preserving their type and format. Card types covered are American Express, Diners Club, Discover, enRoute, JCB, Mastercard, and Visa. For other card types, preserves the number of digits and Luhn's check but may not preserve the card type	Yes
Credit Card Number-American Express	Replaces values with random 15-digit American Express credit card numbers	Yes

8

Hands-On Labs

What about the Database Security?

Hands-On Labs

List of Hands-On Labs

- **Data Safe Configuration and Register ADB –** Initial setup of Oracle Data Safe, including registration of an Autonomous Database.
- **Assess Database Configurations –** Use Oracle Data Safe to assess database configurations for compliance.
- **Assess Database Users –** Review and analyze database user accounts with Oracle Data Safe.
- **Audit Database Activity** – Monitor and audit database activity for enhanced security visibility.
- **Generate Alerts** – Configure and generate alerts based on Oracle Data Safe findings.
- **Discover Sensitive Data** – Identify and classify sensitive data within the database using Oracle Data Safe.
- **SQL Firewall** – Implement SQL Firewall to manage and restrict SQL execution within the database.

Oracle Cloud Infrastructure Security

Security Zones

Martin Berger
Stefan Oehrli

Security Zones

Security for free.

1. Security Zones
2. Web Application Firewall
3. Hands-On Labs

1

Security Zones

Security for free.

Concept

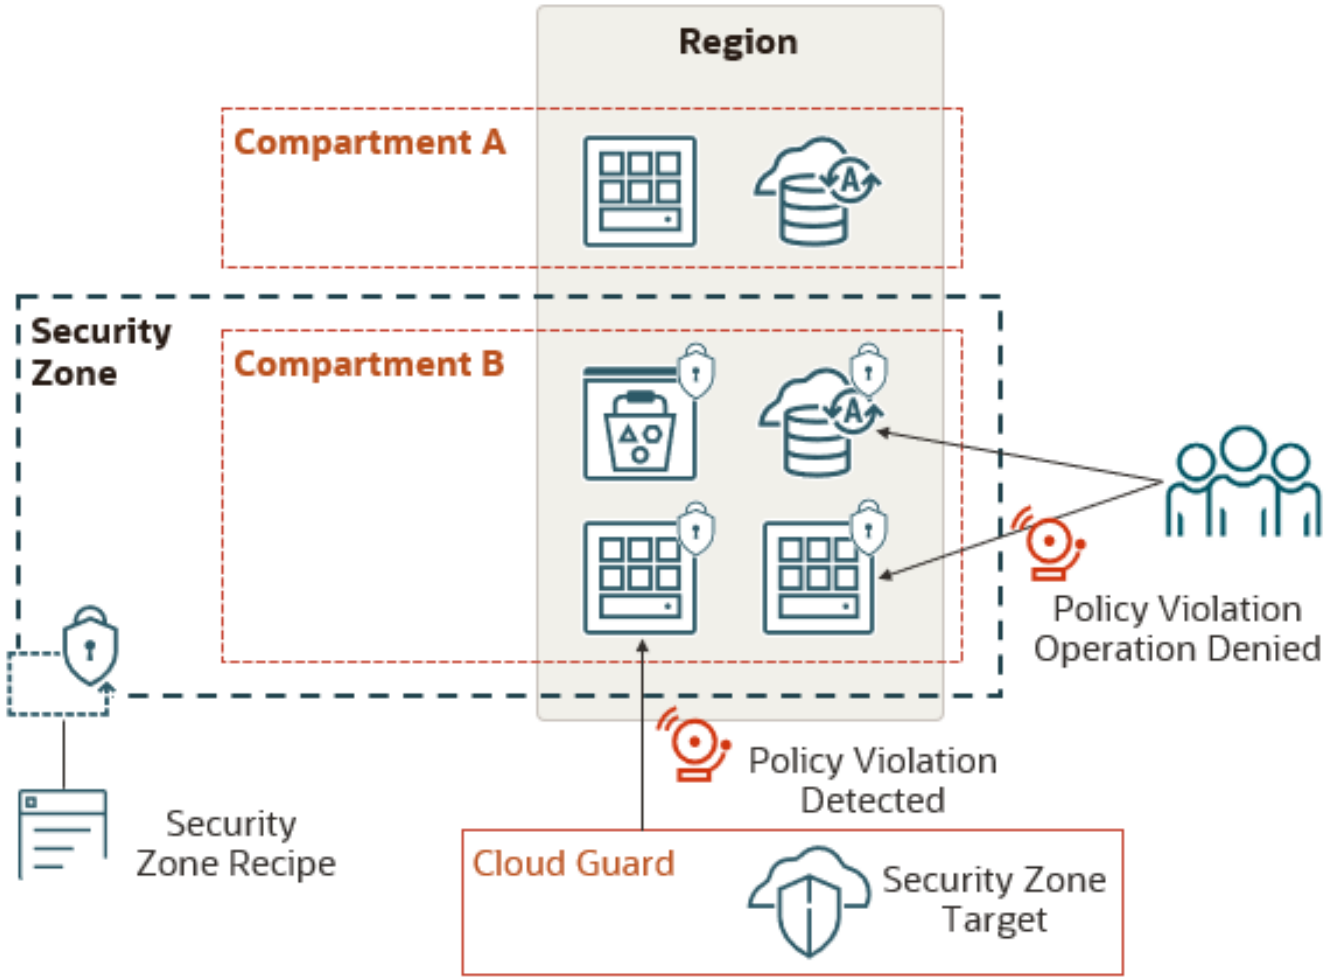

- Deny Public Access
- Ensure Data Security
- Avoid Resource Movement
- No unattached Volumes
- ...

Source: oracle.com

Maximum Security Recipe - 20200914

Policies

Each policy in a Security Zone Recipe prohibits some action that violates a security best practice. Policies are categorized by security principle. Learn more

Policy statement	Type	Resource types
deny instance_without_sanctioned_image	Use Only Configurations Approved by Oracle	COMPUTE, COMPUTEMANAGEMENT
deny instance_in_security_zone_in_subnet_not_in_security_zone	Restrict Resource Association	COMPUTE, COMPUTEMANAGEMENT
deny block_volume_in_security_zone_attach_to_instance_not_in_security_zone	Restrict Resource Association	COMPUTE
deny block_volume_not_in_security_zone_attach_to_instance_in_security_zone	Restrict Resource Association	COMPUTE
deny boot_volume_in_security_zone_attach_to_instance_not_in_security_zone	Restrict Resource Association	COMPUTE
deny boot_volume_not_in_security_zone_attach_to_instance_in_security_zone	Restrict Resource Association	COMPUTE
deny instance_in_security_zone_launch_from_boot_volume_not_in_security_zone	Restrict Resource Association	COMPUTE, COMPUTEMANAGEMENT
deny instance_not_in_security_zone_launch_from_boot_volume_in_security_zone	Restrict Resource Association	COMPUTE, COMPUTEMANAGEMENT
deny instance_in_security_zone_move_to_compartment_not_in_security_zone	Restrict Resource Movement	COMPUTE

Components

- **Proactive** enforcement of security policies in a compartment.
- **Policy Compliance:**
  - Security Zones ensure that resources like Compute, Networking, Object Storage, Block Volume, and Database comply with security policies.
- **Validation and Enforcement:**
  - OCI validates and enforces security policies when creating or updating resources in a Security Zone, denying operations that violate policies.
- **Predefined and Custom Recipes:**
  - Use the **Maximum Security Recipe** provided by Oracle, or create custom recipes to meet specific security needs.
- **Oracle Cloud Guard Integration:**
  - Enable Cloud Guard to detect policy violations in existing resources before creating Security Zones.

Principles

- Resources in a security zone **can't be moved to a compartment outside** of the security zone because it might be less secure.
- All the required components for a resource in a security zone **must also be located in the same security zone**.
  - Example, an instance (Compute) in a security zone **can't use a boot volume that is not in the same security zone**.
- Resources in a security zone must **not be accessible from the public internet**.
- Resources in a security zone must be **encrypted using customer-managed keys**.
- Resources in a security zone must be **regularly and automatically backed up**.
- Data in a security zone is considered privileged and **can't be copied outside of the security zone** because it might be less secure.
- Resources in a security zone must use only **configurations and templates approved by Oracle**.

Create a new Security Zone

- You can create your own recipes for a zone based on a template,
- Any pre-defined Cloud Goard target settings are replaced by Oracle recipes if there is an existing detector configuration.

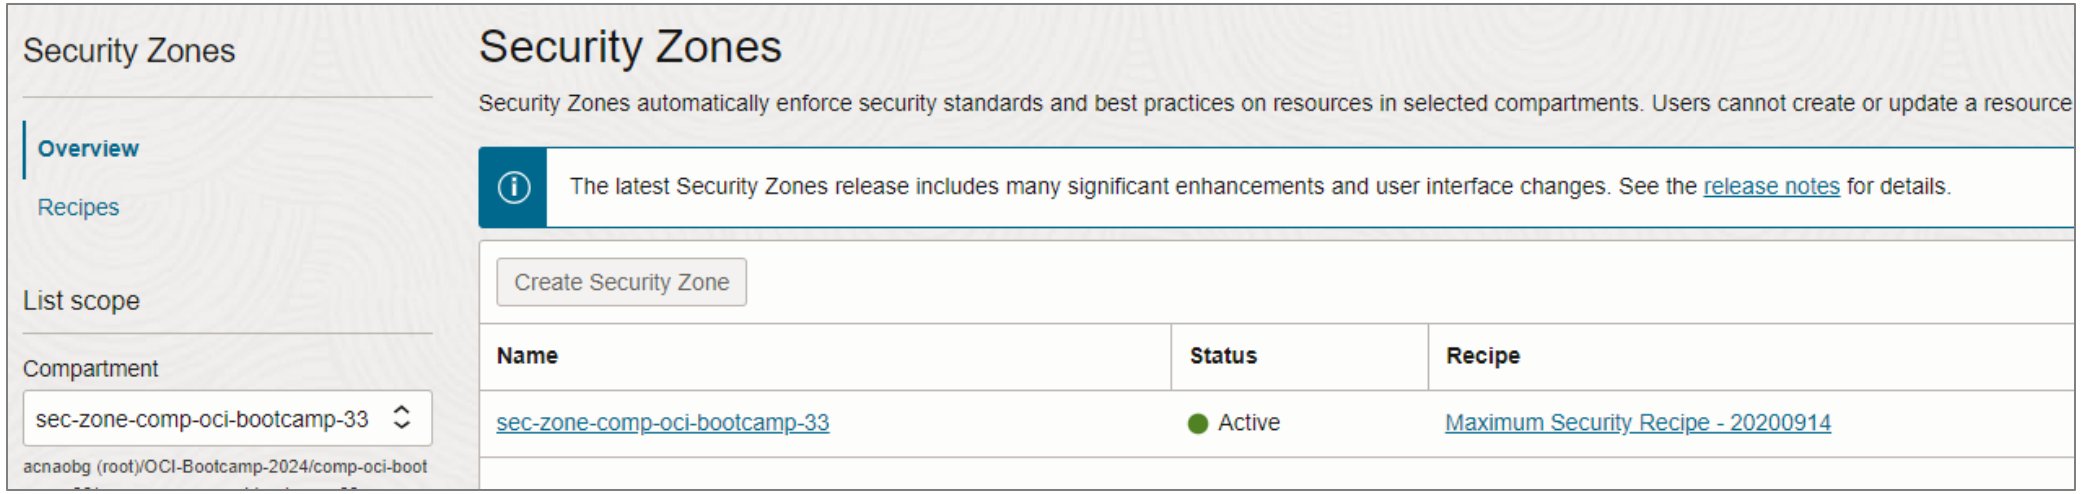

- Rules are immediately active.

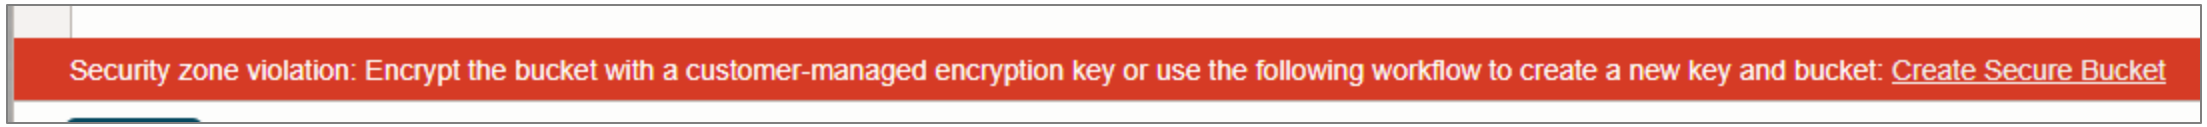

Security Zones and Cloud Guard

- **Compartment and Security Zone Integration**:
  - Security Zones enforce policies on resource operations within compartments, including existing resources created before the zone.
- **Cloud Guard Integration**:
  - Security Zones work with Cloud Guard to identify and address policy violations in existing resources.
- **Centralized Monitoring**:
  - Cloud Guard provides a dashboard to monitor and manage security weaknesses across cloud resources, offering suggestions and corrective actions.

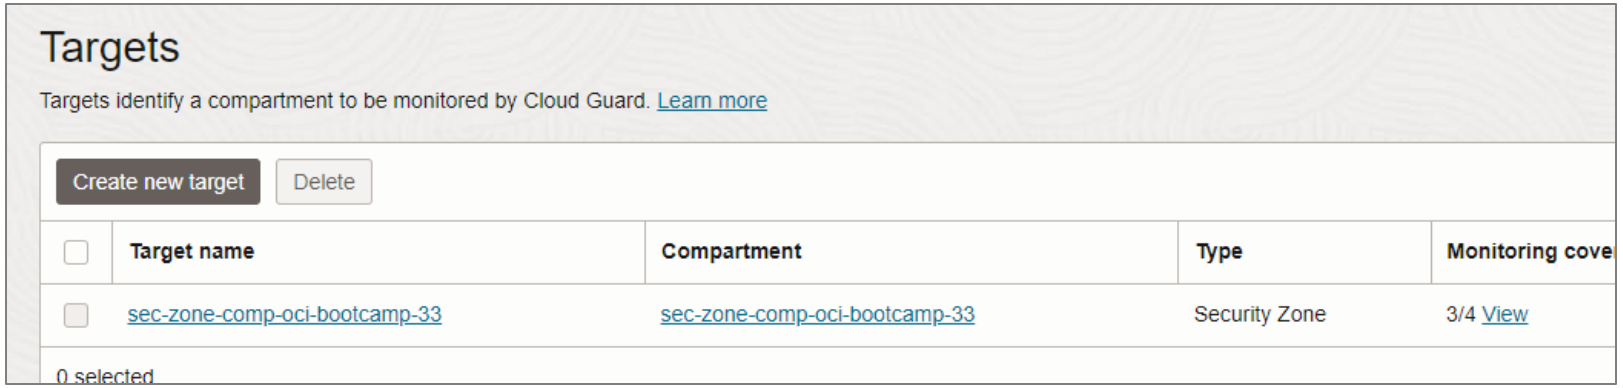

2

Web Application Firewall

Security for free.

Web Application Firewall WAF

- **Comprehensive Protection**:
  - WAF shields applications from malicious internet traffic and enforces consistent security rules across all applications.
- **Advanced Threat Management**:
  - Create and manage rules to defend against threats like XSS, SQL Injection, and other OWASP-defined vulnerabilities; access rules can limit traffic by geography or request signature.
- **Regional and Edge Solutions**:
  - WAF policies are regional and act as load balancer plug-ins, while edge policies provide global protection; allowlist Oracle nodes for edge enforcement.

Web Application Firewall WAF Architecture

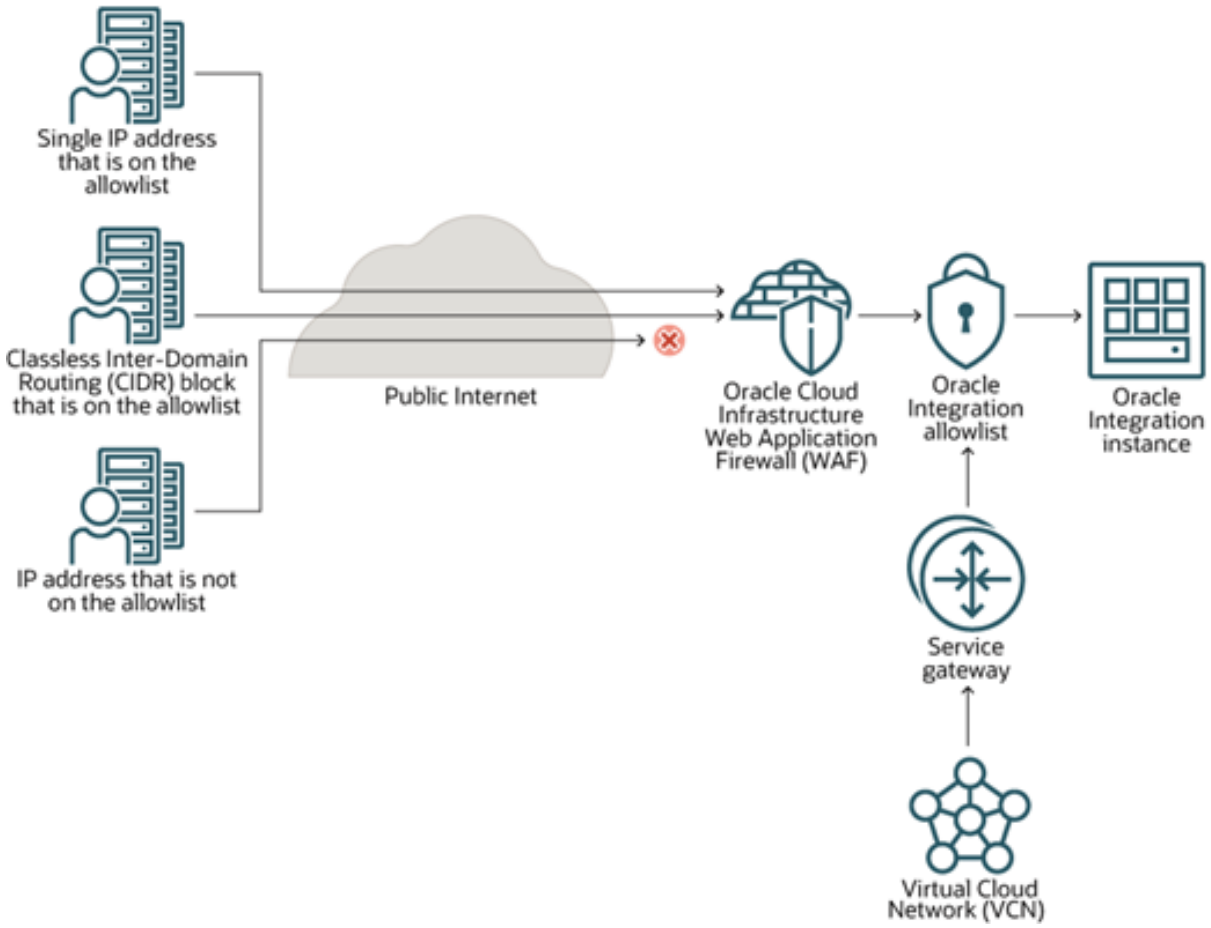

Load Balancer and WAF

- A WAF – Layer 7 - can be configured in front of a load balancer to enforce security
- You can define policies like regions, verify for XSS injections, rate limitations etc.
- Response can be another check or a 401 error

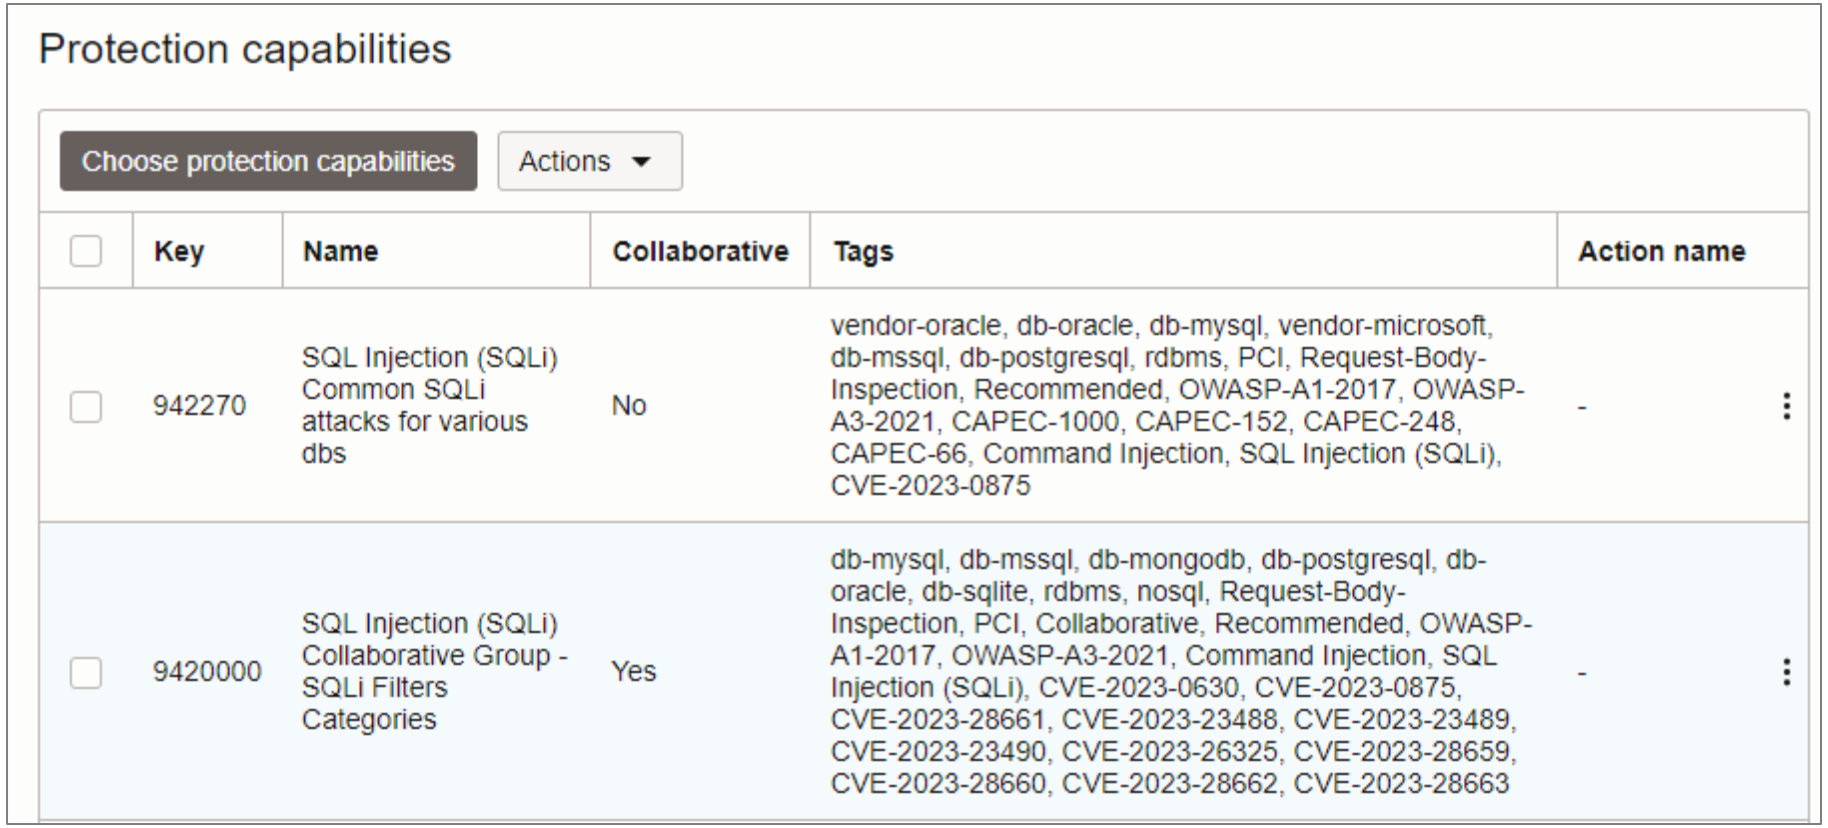

Protection capabilities

Choose protection capabilities | Actions

Key	Name	Collaborative	Tags	Action name
942270	SQL Injection (SQLi) Common SQLi attacks for various dbs	No	vendor-oracle, db-oracle, db-mysql, vendor-microsoft, db-mssql, db-postgresql, rdbms, PCI, Request-Body-Inspection, Recommended, OWASP-A1-2017, OWASP-A3-2021, CAPEC-1000, CAPEC-152, CAPEC-248, CAPEC-66, Command Injection, SQL Injection (SQLi), CVE-2023-0875	-
9420000	SQL Injection (SQLi) Collaborative Group - SQLi Filters Categories	Yes	db-mysql, db-mssql, db-mongodb, db-postgresql, db-oracle, db-sqlite, rdbms, nosql, Request-Body-Inspection, PCI, Collaborative, Recommended, OWASP-A1-2017, OWASP-A3-2021, Command Injection, SQL Injection (SQLi), CVE-2023-0630, CVE-2023-0875, CVE-2023-28661, CVE-2023-23488, CVE-2023-23489, CVE-2023-23490, CVE-2023-26325, CVE-2023-28659, CVE-2023-28660, CVE-2023-28662, CVE-2023-28663	-

3

Hands-On Labs

What about the Database Security?

Hands-On Labs

List of Hands-On Labs

- **Security Zones** – Configure Security Zones to enforce compliance and security policies.
- **Web Application Firewall (WAF)** – Set up and test the Web Application Firewall to protect applications from threats.

5

Wrap-Up

Key Takeaways, Resources, and Next Steps

Basics of OCI Security

Foundational Security Features

Key Management

- Centralized control over encryption keys to protect data at rest and in transit.

OS Management

- Simplifies OS updates and patching for enhanced security.

Vulnerability Scanning

- Automated detection of system vulnerabilities to proactively manage risks.

Shielded Instances

- Ensures uncompromised boot security with Secure Boot, Measured Boot, and Trusted Platform Module (TPM).

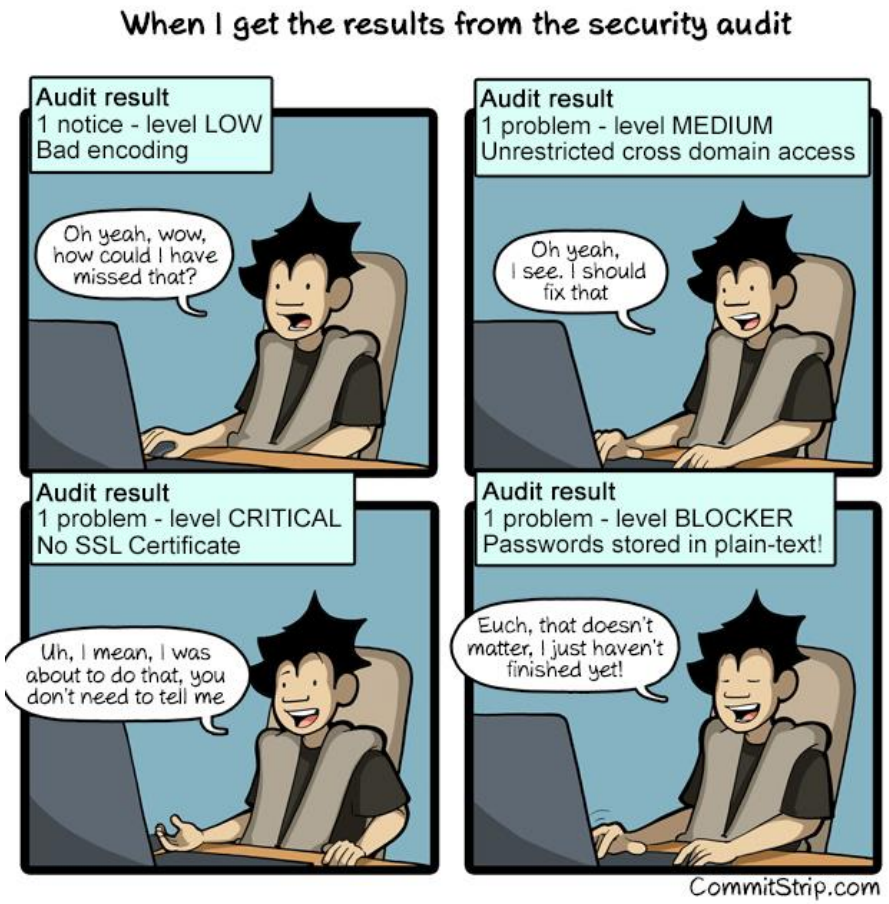

Cloud Guard Overview

Proactive Threat Monitoring

Key Features

- Detect misconfigurations, vulnerabilities, and anomalous activities.
- Responder recipes for automated remediation of identified risks.
- Centralized monitoring and scoring for better risk management.

Use Case

- Monitor resources like Object Storage, Compute Instances, and Networking.

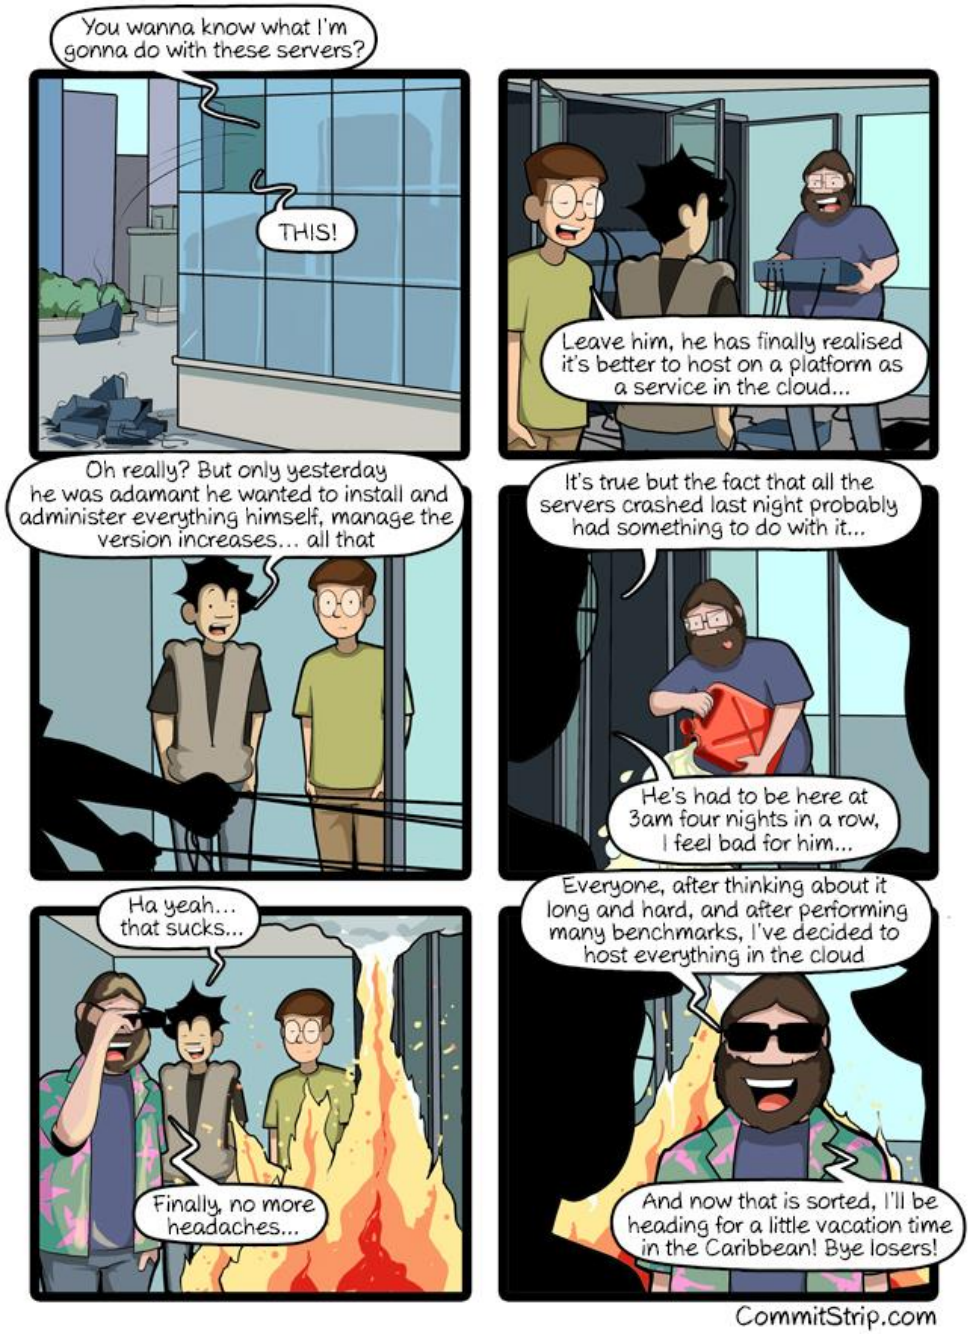

Enhancing Database Security with Data Safe

Centralized Database Security Management

Core Capabilities

- **Security Assessment:** Analyze configurations for compliance.
- **User Assessment:** Monitor high-risk users and their actions.
- **Audit Activity:** Comprehensive logs for user actions and compliance.
- **Sensitive Data Discovery:** Locate and classify sensitive data for protection.

Advanced Feature

- SQL Firewall for real-time SQL execution monitoring and restriction.

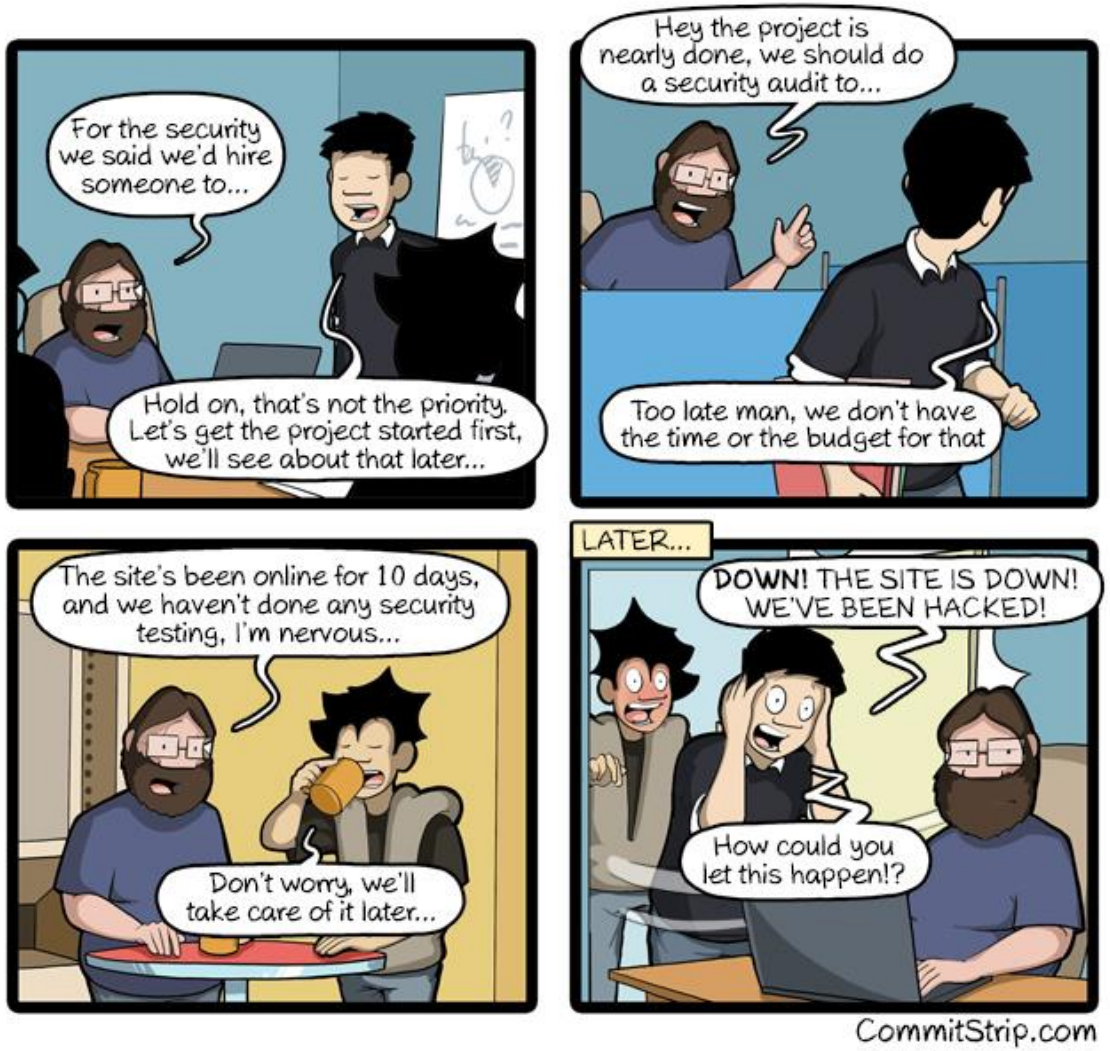

Security Zones and Wrap-Up

Enforcing and Consolidating Security Policies

Security Zones:

- Enforce strict compliance for OCI resources.
- Block public access, mandate encryption, and ensure regular backups.
- Seamlessly integrate with Cloud Guard for enhanced monitoring.

Overall Takeaways:

- Built-in OCI features simplify complex security challenges.
- Proactive tools like Cloud Guard and Data Safe mitigate risks effectively.
- Security Zones ensure consistent policy enforcement across resources.
- **Next Steps:** Explore hands-on labs to apply these concepts.

OCI Access Information

Resources and Cloud Environment Details

The course materials and exercises are available via the GitHub repository/website:

- OCI Walkthrough: https://code.oradba.ch/oci-sec-ws
- PDFs and Course Materials: https://code.oradba.ch/oci-sec-ws/others

Thank You