CLOUDGUARD

02 July 2025

CLOUDGUARD CLOUD NATIVE APPLICATION PROTECTION PLATFORM

Administration Guide

CHECK POINT™

Check Point Copyright Notice

Important Information

Certifications

For third party independent certification of Check Point products, see the Check Point Certifications page.

Latest Version of this Document in English

Open the latest version of this document in a Web browser.
Download the latest version of this document in PDF format.

Patent Notice

CloudGuard is protected by the following patents in the United States and elsewhere.
This page is intended to serve as notice under 35 U.S.C. § 287(a):
US9,531,754, US10,616,235, US10,979,457
US11,431,732, US11,797,685, US11,843,614
US11,481,517, US11,966,466
US11,212,305

Revision History

Date	Description
February 2025	Merging of the Dome9 and Infinity Portal Administration Guides
January 2025	Added: ■ WAF Support to *"Risk Calculation" on page 101* ■ *"Sending Security Events to Microsoft Sentinel" on page 820* Updated: ■ *"Cloud Detection and Response (CDR)" on page 565* ■ *"Onboarding a Google Cloud Platform (GCP) Project and Google Workspace" on page 183*
December 2024	Added: ■ *"Onboarding Docker Hub Container Registry" on page 226* ■ *"Vulnerability Exclusions" on page 442* ■ *"Ignoring Malware from Toxic Combinations" on page 98* ■ *"Base Image" on page 426* and *"Vendor Image" on page 428* to *"Images" on page 425* ■ *"Classifying Assets with Cyera" on page 831* ■ *"Securing Open Source Code" on page 665* ■ *"CI/CD Hardening" on page 709* ■ *"Teams and Asset Mapping" on page 658* ■ *"Code Security Integration with Confluence " on page 747* ■ *"Code Security Integration with Jira" on page 749* ■ *"Code Security Integration with Terraform Cloud" on page 751* ■ *"GitHub Bot" on page 713* ■ *"Gitlab Pre-Receive Hook" on page 734* ■ *"Bitbucket Pre-Receive Hook" on page 736* ■ *"Gitlab Pipeline" on page 731* Updated: ■ *"Integration with Microsoft Teams" on page 817*
October 2024	Added: ■ *"Onboarding Quay.io Container Registry" on page 263* ■ Azure account encryption to *"AWP for Azure Environments" on page 510*

Date	Description
September 2024	Added: ■ *"Pod and Container Requirements" on page 419* ■ *"Toxic Combinations" on page 92* ■ *"Action Hub" on page 93* ■ *"Exclusions for Toxic Combinations" on page 94* ■ *"Security Controls" on page 100* ■ Agent status errors to *"Kubernetes Runtime Protection Troubleshooting" on page 561* ■ The option of ECS scanner to container registries: • *"Onboarding GitHub Container Registry" on page 258* • *"Onboarding Harbor Registry" on page 242* • *"Onboarding JFrog Artifactory" on page 247* • *"Onboarding Sonatype Nexus Registry" on page 253*
July 2024	Merged all types of Exclusions into one topic: *"Configuring CloudGuard Exclusions" on page 80* Renamed *Integrations* to *"Integration Hub" on page 796*
June 2024	Added: ■ CDR and Centralized Account and Sub-Accounts scanning to *"Onboarding of AWS Organizations" on page 146* ■ Exclusions for Vulnerabilities - *"Parameters for Vulnerabilities" on page 82* Updated: ■ *"Integration Hub" on page 796* ■ *"Notifications" on page 852* ■ *"Onboarding an Azure Subscription" on page 61*
May 2024	Added: ■ *"SBOM" on page 276* ■ *"Classifying Assets with Sentra " on page 830* ■ *"Vulnerability Search" on page 447* ■ *"Onboarding Azure Organizations" on page 170* ■ Centralized Account and Sub-Accounts scanning to *"AWP for AWS Environments" on page 497*
March 2024	Added: ■ *"Onboarding GitHub Container Registry" on page 258* ■ *"Scan Engine V2" on page 467*

Date	Description
February 2024	Added: ■ *"Reports" on page 792* ■ Ruleset Versions (*"Viewing a Ruleset Version" on page 313*) Updated: ■ *"Onboarding Azure Subscriptions to Intelligence" on page 601* ■ *"Known Limitations" on page 924* ■ *"Integration with Microsoft Teams" on page 817*
January 2024	Added: ■ *"Agentless Workload Posture" on page 489* support for AWS China accounts
December 2023	Added: ■ *"Environment Risk (Risk Level)" on page 103* Updated: ■ *"ERM Protected Assets" on page 90*
November 2023	Added: ■ Function Apps scanning in *"AWP for Azure Environments" on page 510* (Early Availability) ■ Centralized Account and Sub-Accounts scanning in *"AWP for Azure Environments" on page 510* (Early Availability) ■ *"Onboarding Sonatype Nexus Registry" on page 253* Updated: ■ *"Kubernetes Runtime Protection" on page 549*
October 2023	Added: ■ *"Toxic Combinations" on page 92* Updated: ■ Explanation and lists of supported entities in *" Risk Management" on page 89* ■ Authentication methods in *"Onboarding AWS Elastic Container Registry" on page 214* ■ *"Cloud Infrastructure Entitlement Management (CIEM)" on page 385* ■ *"Findings" on page 392*

Date	Description
September 2023	Changed the configuration and script in *"Configure Policies for Azure Key Vault Entities" on page 178*
August 2023	Added: ■ *"Updating Onboarded AWS Organizations" on page 155* ■ *"Container Environments" on page 413*
July 2023	Added: ■ *"Errors Troubleshooting in ImageScan Agent" on page 458* ■ *"Image Scan Findings" on page 449* ■ ECS scanner to *"Image Assurance" on page 434* ■ ECS scanner to *"Onboarding Container Registries" on page 204*
June 2023	Added: ■ *"Asset Details" on page 273* ■ *"Entitlement Map" on page 387*
May 2023	Added: ■ *"CIEM Policies, Exclusions, and Remediation" on page 397* ■ *"Onboarding of AWS Organizations" on page 146*
April 2023	Added: ■ *"Kubernetes Posture Management" on page 307* ■ *"AWP for Azure Environments" on page 510* (Early Availability)
March 2023	Added: ■ *"Onboarding Oracle Cloud Infrastructure Environments" on page 186* ■ *"Integrating Amazon GuardDuty Findings with CloudGuard" on page 823*
February 2023	Added: ■ *"ERM Protected Assets" on page 90* ■ *"ERM Rulesets" on page 117*

Date	Description
January 2023	Added: ■ *"Serverless Functions" on page 536* ■ In-Account Mode for *"Agentless Workload Posture" on page 489* ■ *"Image Admission" on page 486* ■ *"Onboarding Google Artifact Registry" on page 237* ■ *"Onboarding JFrog Artifactory" on page 247* ■ *"Onboarding Harbor Registry" on page 242* ■ *"Container Registry Scanning" on page 464*

Click here to see the earlier revision history

Date	Description
December 2022	*Identity* is now Cloud Infrastructure Entitlement Management (*"Cloud Infrastructure Entitlement Management (CIEM)" on page 385*) Added: ■ *"Agentless Workload Posture" on page 489* ■ *" Risk Management" on page 89*
November 2022	Added: ■ New *"Code Security" on page 647* (Spectral) functionality The ShiftLeft functionality moved under CI-CD Tool ■ The Google Container Registry type in *"Onboarding Container Registries" on page 204* ■ The function `containSecrets` in *"Governance Specification Language (GSL)" on page 326*
September 2022	Added: ■ *"Which CloudGuard endpoints do I have to allow on my network?" on page 917* ■ *"API Audit Logs" on page 132* ■ *"System Audit Logs" on page 134*
August 2022	Added *"Troubleshooting Kubernetes Onboarding" on page 198*

Date	Description
July 2022	Added: ■ *"Sending Findings to Eventarc" on page 810* ■ *"Sending Findings to Azure Defender for Cloud" on page 826* ■ *"Image Admission" on page 486* (Early Availability) ■ *"Email Notifications" on page 860* ■ *"Configuring CloudGuard as an AWS Security Hub Provider" on page 821*
June 2022	Added: ■ *"Sending System Notifications to AWS SNS" on page 806* ■ *"Removing Intelligence from AWS Environments with API" on page 593*
May 2022	Added: ■ *"Images" on page 425* ■ *"Workloads Settings" on page 861* ■ *"Onboarding Container Registries" on page 204*
April 2022	Added: ■ *"Frequently Asked Questions" on page 917* ■ New wizard - *"Onboarding AWS Environments to Intelligence" on page 572* ■ *"Onboarding AWS Environments to Intelligence with API" on page 590* ■ *"Configuring CloudGuard Policies" on page 78*
March 2022	Added: ■ *"Getting Started with CloudGuard" on page 50* ■ New AWS onboarding experience: *"Unified Onboarding of AWS Environments" on page 54* Added: ■ *"GSL Builder" on page 350* ■ *"Billing Reports" on page 834* ■ *"Intelligence for Kubernetes Containers" on page 645* ■ *"Onboarding Kubernetes Clusters to Intelligence" on page 618* ■ *"Removing Intelligence from Kubernetes Clusters" on page 620*

Date	Description
December 2021	Added: ■ *"Removing Intelligence from AWS Environments" on page 588* ■ *"Removing Intelligence from Azure Subscriptions" on page 608*
November 2021	Added *"CloudGuard Permissions for Intelligence" on page 621*
September 2021	Added: ■ ShiftLeft Integration to CI/CD Pipeline ■ ShiftLeft Integration to CircleCI, ShiftLeft Integration to GitLab, ShiftLeft Integration to Jenkins
July 2021	Added: ■ Assessment Results View ■ ShiftLeft Configuration in CloudGuard ■ *"Kubernetes Containers" on page 415* ■ Integration to the Infinity Portal
June 2021	Updated: *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287*
May 2021	Added: ■ *"Admission Control" on page 476* ■ *"Image Assurance" on page 434* ■ *"Vulnerability Policies (Image Assurance)" on page 444* ■ *"Image Assurance Troubleshooting" on page 452*
March 2021	Added: ■ *"Integration with Microsoft Teams" on page 817* ■ *"Sending Reports to Jira" on page 814*
February 2021	Added: ■ *"Kubernetes Runtime Protection" on page 549* ■ *"Data Handling" on page 547* ■ *"AWS Runtime Protection Implementation" on page 542* ■ *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287* Updated the menu, changed the name from CloudGuard Dome9 to CloudGuard General updates

Date	Description
25 November 2020	Added ShiftLeft
10 August 2020	Updated *"Onboarding Cloud Environments" on page 53*
24 May 2018	First release of this document

Table of Contents

Introduction to CloudGuard CNAPP

Check Point CloudGuard Cloud Native Application Protection Platform (CNAPP) is a SaaS platform that provides unified cloud-native security across your applications, workloads, and network. You can use it to automate security, prevent threats, get compliance and manage posture for all of your cloud environments: from Amazon AWS and Microsoft Azure to GCP, Kubernetes, and more.

Protect your Private and Public Clouds

CloudGuard CNAPP ensures network security and enforces security policy, prevents changes not approved, and enforces the previously defined configuration. Regardless if you use public or private clouds, CloudGuard CNAPP facilitates server configuration management. Its flexible security management tools ensure compliance and decreases configuration errors and possible breaches.

Its powerful layer of Threat Intelligence transforms cloud big data into high-definition, actionable security logic. Customize alerts and built-in queries, quarantine threats, and stop attacks in progress.

Secure Kubernetes Containers

CloudGuard CNAPP offers a depth of coverage for all container types, rich visualization of cloud assets, and an assessment of security posture to quickly identify misconfiguration issues and threats. Understand at a glance what is running in your container environment and how it is configured. Visualize Kubernetes data flows and get visibility of container misconfigurations and anomalies.

Create and Manage Custom Compliance Rules

Create custom compliance rules with intuitive GSL language, and align with NIS and CIS security benchmarks, with the largest number of rulesets and compliance frameworks across cloud environments.

Overview of the Main Menu

The main menu along the left side of the main screen provides navigation to the CloudGuard pages and features. You can search through the menu items with the search bar located above the menu (**Search Navigation**). Start to enter the name of the page, and CloudGuard offers you a list of menu items with this name.

The menu options appear as sections in this Administration Guide.

Note - The search of the Code Security menu items is not supported.

Menu Icon	Section in this Guide	Description
	■ Introduction to CloudGuard ■ *"Getting Started with CloudGuard" on page 50*	Get to know CloudGuard: ■ Learn the user interface ■ Learn how to onboard new environments ■ Go over basic configurations ■ See Dashboards
	*" Risk Management" on page 89*	■ Maximize the productivity of security teams ■ Increase visibility of each asset's risk score
	*"Events" on page 119*	■ See the Posture Findings, Threat and Security Events that occur in your cloud environments, based on configured policies ■ See the finding details in the CloudGuard portal or receive messages to different notification targets, such as email and SNS
	*"Assets" on page 138*	■ Manage protected assets in onboarded environments in CloudGuard ■ Configure and manage Organizational Units
	*"Cloud Security Posture Management (CSPM)" on page 302*	■ Assess the compliance of your cloud environments ■ Select built-in rulesets or build your compliance test rulesets
	*"Network Security" on page 353*	■ Visualize the security policies in your environments ■ Protect and manage your cloud assets, such as Security Groups ■ Control access to cloud assets
	*"Cloud Infrastructure Entitlement Management (CIEM)" on page 385*	■ Reduce your attack surface by ensuring that cloud entitlements or permissions respect the principle of *least privilege* ■ See Cloud Users and Cloud Roles

Menu Icon	Section in this Guide	Description
	*"Workload Protection" on page 410*	■ Manage serverless functions in your AWS environments and Kubernetes clusters ■ Use Image Assurance to analyze Kubernetes and Container Registry images at each stage of their life cycle ■ Use Admission Control to help you configure and control operations for Kubernetes clusters
	*"Cloud Detection and Response (CDR)" on page 565*	Use Intelligence to hunt for and visualize threats and anomalous behavior in your environments through cloud log files
	*"Code Security" on page 647*	Build the Code Security functionality into your CI/CD pipeline to detect and prevent risk in cloud deployments
	*"Reports" on page 792*	Generate reports that show summaries, trends, and insights based on data that CloudGuard collects.
	*"Integration Hub" on page 796*	Enable seamless integration with internal and external third-party applications, APIs, and services.
	*"Settings" on page 832*	■ Manage your CloudGuard account ■ Configure users and roles, organizational units ■ Create and manage notifications for security policies ■ Learn how to install and use the CloudGuard mobile app to gain access to your protected environment (IAM Safety) (for Dome9 accounts)

System Architecture

The diagram below shows the architecture for the CloudGuard portal.

CloudGuard is connected to cloud platforms with the correct platform APIs and platform notification services, such as SNS for AWS. In addition, CloudGuard can connect to logging, ticketing, and email systems, such as ServiceNow and PagerDuty, to forward CloudGuard alerts.

Upstream, corporate systems can connect to CloudGuard with its REST API, to implement automation processes to manage activities on CloudGuard. Moreover, users can use Infrastructure as Code systems, such as Terraform or AWS CloudFormation, to connect to CloudGuard.

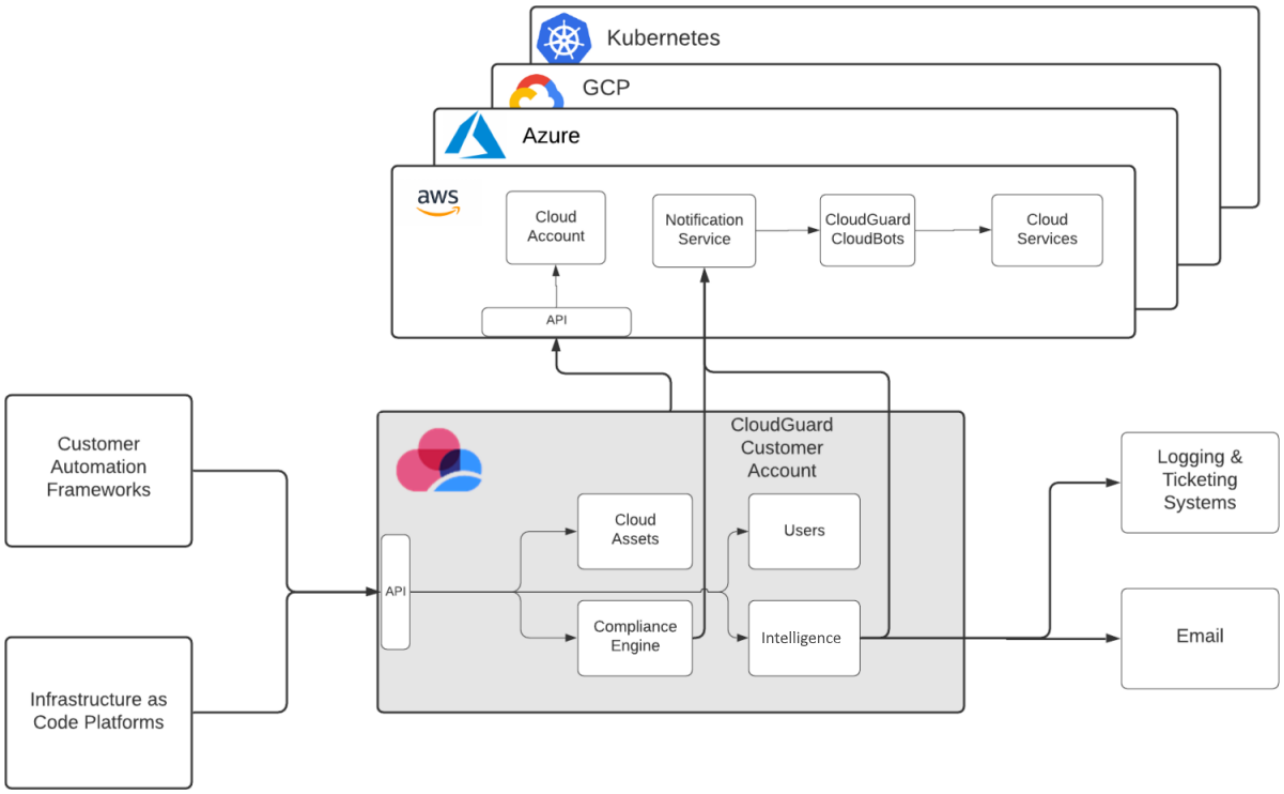

Getting Started with CloudGuard

Before you can start to onboard cloud environments, configure policies, and monitor your environments, you must do these steps:

Step 1: Creating an Account in the Infinity Portal

Check Point Infinity Portal is a web-based interface that hosts the Check Point security SaaS services.

With Infinity Portal, you can manage and secure your IT infrastructures: networks, cloud, IoT, endpoints, and mobile devices.

To create an Infinity Portal account, see the Infinity Portal Administration Guide.

Important - The creation of new accounts from secure.dome9.com and its subdomains is not available.

Step 2: Logging in to your account

For accounts registered on Dome9

Log in to the account in your region:

- United States: https://secure.dome9.com
- Europe: https://secure.eu1.dome9.com
- Australia: https://secure.ap2.dome9.com
- Canada: https://secure.cace1.dome9.com
- India: https://secure.ap3.dome9.com
- Singapore: https://secure.ap1.dome9.com

For accounts registered on the Infinity Portal

To access the Infinity CloudGuard portal:

1. Log in to the Infinity Portal.
2. Click the **Menu** icon in the top left corner of the Infinity Portal window.
3. From the CloudGuard group, select one of the CloudGuard modules.
4. If you access the CloudGuard portal for the first time, select one of the options below.

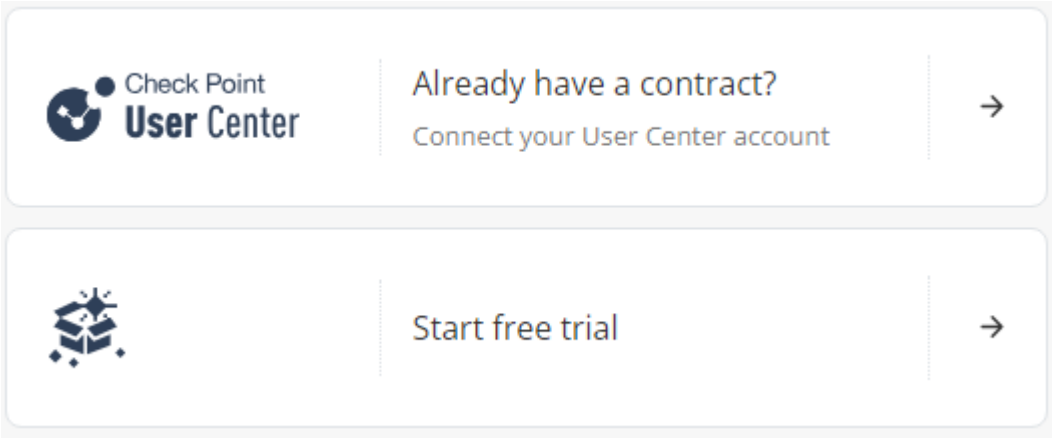

- Connect your User Center account if you already have a Check Point contract. When you select this option, the **Attach Account** window opens. For more information, see **Associated Accounts** in the Infinity Portal Administration Guide.
- Start a free trial if you do not associate CloudGuard with a user account. When you select this option, you can use CloudGuard for a 30-day period.

  CloudGuard starts to load, and the main screen opens.

For **existing accounts**, the main screen shows a dashboard (by default, the Risk Management dashboard) of your environments and resources onboarded to CloudGuard.

Note - You can create a new dashboard from scratch or customize one of the default dashboards to display the most applicable aspects of your cloud environment security. To learn more about CloudGuard dashboards, see *"Dashboards" on page 86*.

Step 3: Onboarding your cloud environment

For onboarding steps, see *"Onboarding Cloud Environments" on page 53*.

Getting Started with CloudGuard Policies

- *"Configuring CloudGuard Policies" on page 78*
- *"Getting Started with Posture Management Policy" on page 306*
- *"Getting Started with Intelligence Policy" on page 569*
- *"Getting Started with Image Assurance Policy" on page 437*
- *"Getting Started with Admission Control Policy" on page 481*

External Resources

- CloudGuard Knowledge Base - The Knowledge Base articles tell how to configure and use different CloudGuard features, how to use third-party services and systems with CloudGuard, and how to use the CloudGuard REST API.
- CloudGuard REST API - You can access CloudGuard programmatically with the CloudGuard REST API. The API has resources to onboard accounts, manage security groups, retrieve findings, run compliance assessments, and more.
- CloudGuard GSL Knowledge Base - The GSL Knowledge Base is a comprehensive repository of CloudGuard GSL rules and compliance rulesets.
- CloudGuard Release Notes - The Release Notes show the latest features and fixes in the CloudGuard portal.

Onboarding Cloud Environments

Onboarding connects your cloud environments or container clusters to Check Point CloudGuard. Onboarding allows CloudGuard monitor and, optionally, manage the security posture of your environment.

The topics in this section explain how to onboard accounts for the different cloud platforms in the CloudGuard portal.

In addition, you can onboard accounts programmatically with the CloudGuard REST API, with Terraform, and with open-source scripts that you run on your account.

More Links

- *"Onboarding AWS Environments" on page 144*
- *"Onboarding Azure Subscriptions" on page 169*
- *"Onboarding Google Cloud Platform Projects" on page 182*
- *"Onboarding Oracle Cloud Infrastructure Environments" on page 186*
- *"Manual Onboarding of Kubernetes Clusters" on page 68*
- *"Onboarding Container Registries" on page 204*
- *"Onboarding Alibaba Cloud Accounts" on page 75*

Unified Onboarding of AWS Environments

This topic describes how to onboard an AWS environment automatically. For other onboarding methods, see *"Onboarding AWS Environments" on page 144*.

Prerequisites

Before onboarding your AWS account, make sure:

- You have Administrator permissions to create and manage resources in this account.

Two Paths: One Click or Advanced Onboarding

Select an onboarding path for your AWS environment:

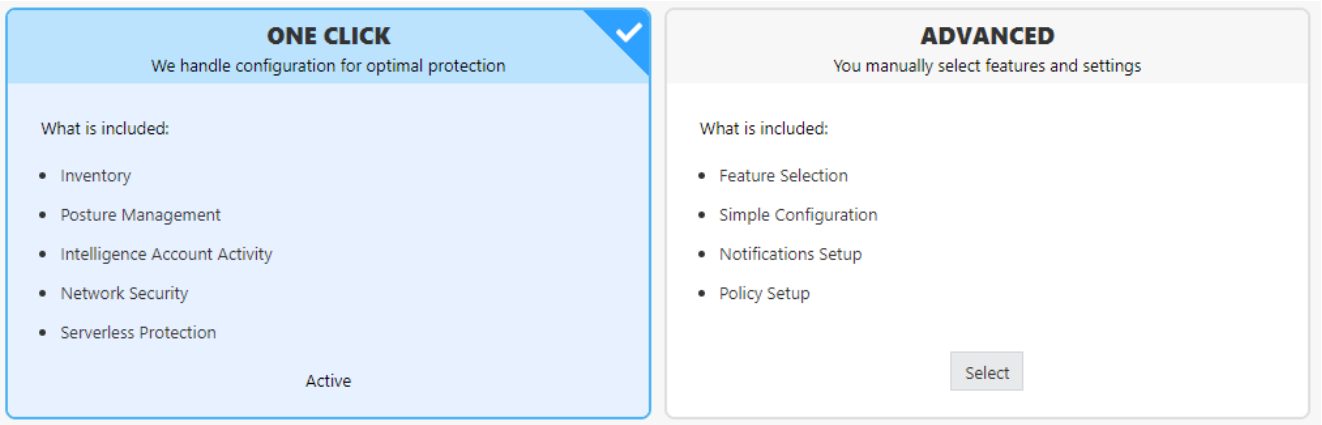

- **ONE CLICK** Onboarding - Automatically onboard your AWS account to CloudGuard. The welcome screen includes the features enabled for your environment. The CloudGuard algorithm decides which resources to onboard and how with minimal involvement from your side. Your initial configuration includes:
  - **Posture Management** - CloudGuard creates these policies from the rulesets recommended by Check Point security experts:
    - AWS CIS Foundation ruleset - Latest version
    - AWS CloudGuard Best Practices
    - AWS CloudGuard CheckUp

- **Intelligence Account Activity** (for Standard AWS Accounts only. GovCloud and China Cloud Accounts are not supported) - Intelligence Account Activity is enabled on a selected S3 bucket that has a CloudTrail.

  Note - In this path, CloudGuard activates Intelligence automatically. Before you start the onboarding process, make sure that your AWS account has an active CloudTrail with an S3 bucket assigned.

  CloudGuard creates your Intelligence policy based on the *AWS CloudGuard Best Practices* Intelligence ruleset.

- **Permissions** - CloudGuard applies the **Monitor** mode to all assets related to this account, for example to the Security Groups.

  CloudGuard can manage your AWS accounts in Monitor or Full Protection modes that determine the type of permissions that CloudGuard receives from AWS.

- **Serverless Protection** (for Standard AWS Accounts only) - CloudGuard enables Serverless Protection on your account by default.

- **ADVANCED** Onboarding - Multiple options for non-standard and customizable environments. Select the functionalities, resources, and settings to configure. During the onboarding process, CloudGuard uses a combination of Lambda functions and CFT deployed on your account to create an optimal configuration.

ONE CLICK Onboarding

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. For first-time onboarding, click **AWS**. Or, if you already onboarded environment(s), then from the top bar, select **Add** > **AWS Environment**.
3. On the **Welcome** page, select an account type:
   - **Standard** AWS Account
   - **GovCloud** Account
   - Amazon Web Services (China) Account - **China Cloud**
4. For **CFT Permissions Management**, if you agree to start the process of permissions update automatically from CloudGuard select **Allow CloudGuard to update and delete its CloudFormation stack resources**. If you select this option, then when the permissions update is required, you have to do it manually in the AWS portal. For more information, see Updating AWS Permissions.

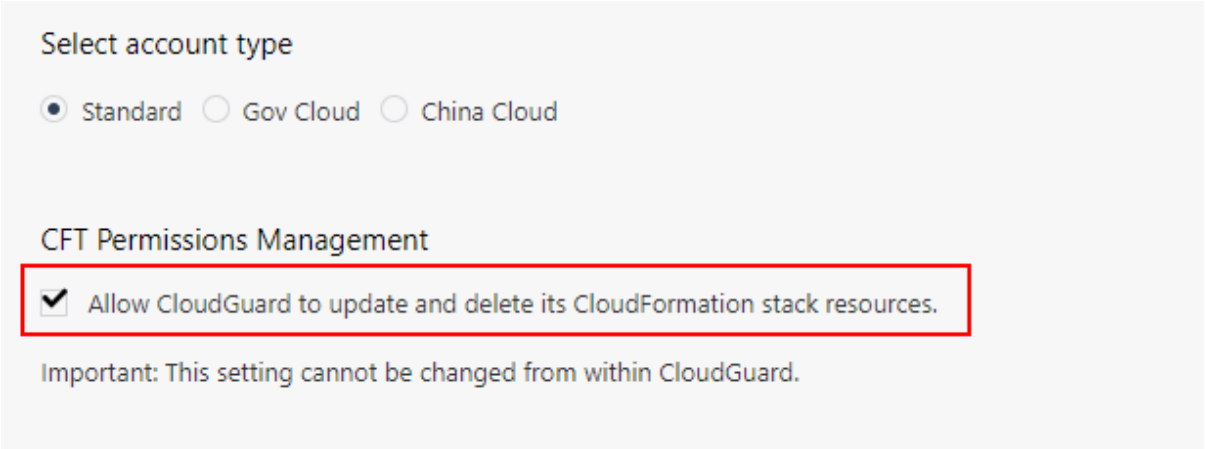

Important - You cannot revoke this consent from the CloudGuard portal after the account onboarding is done.

5. If **ONE CLICK** is not selected, click the **Select** button on the relevant card and click **Next**.
6. On the **Set up your cloud account** page, follow the on-screen instructions:
   a. Click **CFT Template** to review all resources for CloudGuard to deploy on your environment. These resources are organized by components on different tabs for Onboarding, Permissions, Serverless, and Intelligence. Optionally, you can click:
      - **Download CFT** - Save the resource file for each component in YAML on your local drive.
      - **Review Source** - Browse the resources code in the GitHub repository.
   b. Click **Close**.
   c. Open a new browser tab, go to the AWS portal and sign in to your AWS account.
   d. In the CloudGuard onboarding wizard, click **Launch Stack**.

      A new browser tab opens with the CloudFormation stack. CloudGuard automatically enters all required parameters.

      Note - By default, using the **ReadOnlyAccess** policy is enabled, which allows you to receive permissions update requests less frequently. You can manually disable the policy at this stage in the **UseAwsReadOnlyPolicy** field if you prefer not to grant redundant permissions. For more information about policies, see *"Policies" on page 279*.

e. Below **Capabilities**, read the explanation and select the **I acknowledge...** option to accept. Click **Create stack**.

AWS begins to create the stack. This stack creates new roles and resources for initial work and an AWS Lambda to arrange the environment and set up all modules. After deployment, the stack deletes the Lambda function from your environment.

CloudGuard waits for the stack deployment to complete before the last step. This procedure can take several minutes.

7. After the deployment process is complete, in the CloudGuard portal click **Next**.
8. The last screen in the wizard is the Onboarding Summary. Until the process is done, the page shows the current status with the number of active, pending, and inactive features, the number of errors, error details, and suggested remediation. When the onboarding is done, click **Finish**.

Your environment is onboarded to CloudGuard with a default CloudGuard-Managed policy. When the process is done, CloudGuard redirects you to the Environments page that lists your new onboarded environment.

ADVANCED Onboarding

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. For first-time onboarding, click **AWS**. Or, if you already onboarded environment(s), then from the top bar, select **Add** > **AWS Environment**.
3. On the **Welcome** page, select an account type:
   - **Standard** AWS Account
   - **GovCloud** Account
   - Amazon Web Services (China) Account - **China Cloud**
4. For **CFT Permissions Management**, if you agree to start the process of permissions update automatically from CloudGuard select **Allow CloudGuard to update and delete its CloudFormation stack resources**. If you select this option, then when the permissions update is required, you have to do it manually in the AWS portal. For more information, see Updating AWS Permissions.

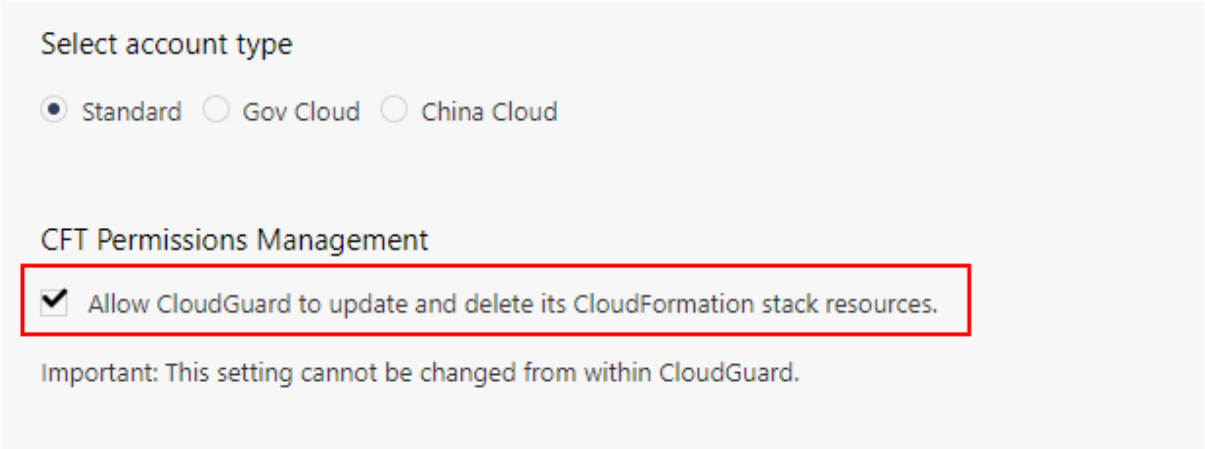

Important - You cannot revoke this consent from the CloudGuard portal after the account onboarding is done.

5. If the **Advanced** option is not selected, click the **Select** button on the relevant card and click **Next**.
6. On the **Permissions** page, click **Select** for the desired operation mode:
   - **Monitor** - Monitor and visualize your environments in CloudGuard, run compliance tests, and receive alerts, notifications and reports of activities and changes to cloud entities. You cannot manage the entities from CloudGuard.
   - **Full Protection** - Contains all the capabilities of the **Monitor** mode. In addition, you can use CloudGuard to enforce access protection and tamper protection on your assets, manage your Security Groups, and control direct access to your cloud assets.

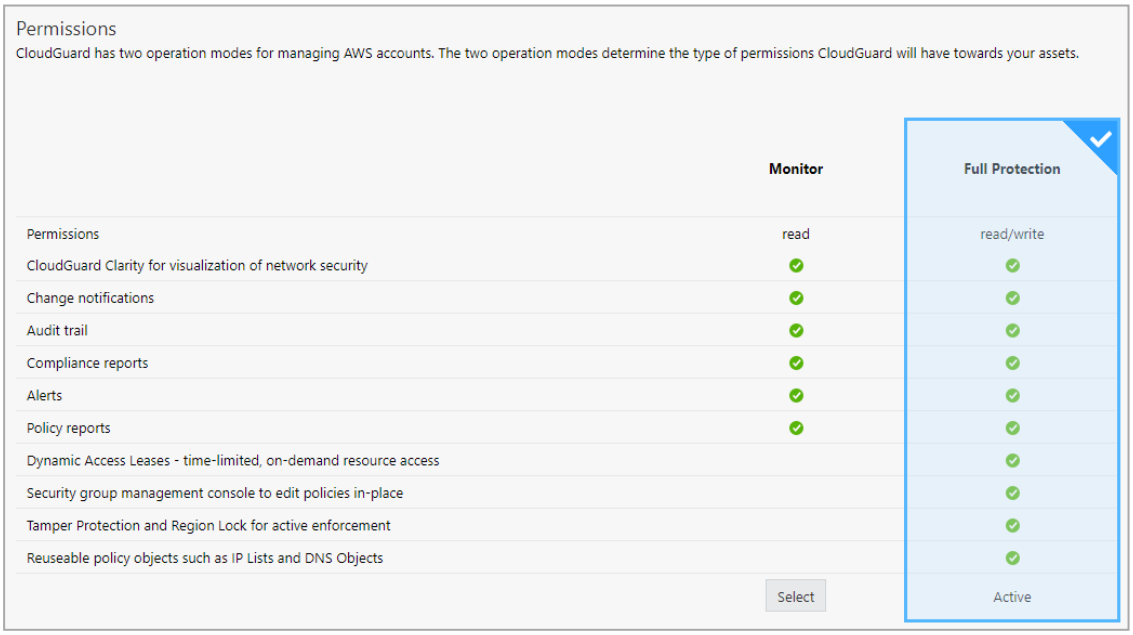

7. Click **Next**.

8. On the **Posture Management** page, select the rulesets to enable on your environment and click **Next**.

   - The **Common** rulesets (click to expand) include CSPM rules to check your environment's compliance with industry standards and best practices.
   - The **Additional** rulesets (click to expand) include all other rulesets customized for an organization's security policy.

     To open and review a ruleset, click the arrow after the ruleset's description.

   When the onboarding process completes, CloudGuard creates a policy for each selected ruleset.

9. On the **Intelligence** page, turn the Intelligence Account Activity on or off. If you enable the Intelligence Account Activity, you can select rulesets from the table. When the onboarding process completes, CloudGuard creates a policy for each selected ruleset. The *AWS CloudGuard Best Practices* ruleset for Intelligence is selected by default.

10. Click **Next**.

    Notes:
    - This capability is not available for GovCloud and China Cloud accounts.
    - To activate Intelligence, before you start the onboarding process, make sure that your account has an active CloudTrail with an S3 bucket assigned.

11. On the **Serverless Protection** page, turn the protection on or off.

12. Click **Next**.

    Note - This function is not available for GovCloud and China Cloud accounts.

13. On the **Set up your cloud account** page, follow the on-screen instructions:

    a. Click **CFT Template** to review all resources for CloudGuard to deploy on your environment. These resources are organized by components on different tabs for Onboarding, Permissions, Serverless, and Intelligence. Optionally, you can click:

       - **Download CFT** - Save the resource file for each component in YAML on your local drive.
       - **Review Source** - Browse the resources code in the GitHub repository.

    b. Click **Close**.

    c. Open a new browser tab, go to the AWS portal and sign in to your AWS account.

d. In the CloudGuard onboarding wizard, click **Launch Stack**.

   A new browser tab opens with the CloudFormation stack. CloudGuard automatically enters all required parameters.

   Note - By default, using the **ReadOnlyAccess** policy is enabled, which allows you to receive permissions update requests less frequently. You can manually disable the policy at this stage in the **UseAwsReadOnlyPolicy** field if you prefer not to grant redundant permissions. For more information about policies, see *"Policies" on page 279*.

e. Below **Capabilities**, read the explanation and select the **I acknowledge...** option to accept. Click **Create stack**.

AWS begins to create the stack. This stack creates new roles and resources for initial work and an AWS Lambda to arrange the environment and set up all modules. After deployment, the stack deletes the Lambda function from your environment.

CloudGuard waits for the stack deployment to complete before the last step. This procedure can take several minutes.

14. The last screen in the wizard is the Onboarding Summary. Until the process is done, the page shows the current status with the number of active, pending, and inactive features, the number of errors, error details, and suggested remediation. When the onboarding is done, click **Finish**.

Your environment is onboarded to CloudGuard with the advanced settings. When the process is complete, CloudGuard redirects you to the Environments page that lists your new onboarded environment.

Onboarding an Azure Subscription

This topic describes onboarding one Azure subscription. For onboarding an Azure Organization, see *"Onboarding Azure Organizations" on page 170*.

Prerequisites

Before onboarding your Azure subscription, make sure that you have required permissions assigned through the **Application Administrator** or **Owner** role. For more information, see *"Permissions" on page 293*.

Onboarding in the Portal

STEP 1 - Welcome

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **Azure** and follow the setup instructions.

   Or, if you already onboarded environment(s), from the top menu, select **Add** > **Azure Subscription**.
3. Select to onboard a **Single account**.
4. Select the cloud platform that hosts the account: **Standard**, **Azure Government**, or **Azure China**.
5. Enter your **Subscription ID**.
6. Select operation mode: **Read-Only** or **Manage**.
7. Click **Next**.

STEP 2 - Configurations

1. On the **Configurations** page, enter a name for the environment (optional) to identify it in the CloudGuard portal.
2. Select an Organizational Unit to which to associate the environment.

   Note - This is an optional step. Later, you can use the **Organizational Units** page in **Assets** to create and change associations.
3. Click **Next**.

STEP 3 - Connect

1. On the **Connect** page, click **Onboarding script** to review it.

   Note - To run the script in a quiet mode that skips all questions to the user, run it with the `--quiet` flag .

2. Log in to the Azure Cloud Shell.
3. Copy and run the command provided in the CloudGuard wizard. As the program runs it prints out its actions and shows resources created in your Azure account.

   Wait until the Cloud Shell shows the outputs for your environment.

   Example

   ```
   -----Outputs-----

   Tenant ID:                    XXXXXXXX-XXXX-XXXX-XXXXXXXXXXXX
   Application (client) ID:      XXXXXXXXXXXXXXXXXXXXXXX
   Secret Key:                   XXXXXXXX-XXXX-XXXX-XXXXXXXXXXXX
   Subscription ID:              XXXXXXXX-XXXX-XXXX-XXXXXXXXXXXX
   ```

4. Copy and paste to the CloudGuard wizard these values:
   - **Tenant ID**
   - **Application ID**
   - **Secret Key**
5. Click **Onboard**. The onboarding starts.

   When CloudGuard finishes onboarding, you get a notification, and the onboarded environment appears on the **Environments** page.

This procedure creates a default CSPM policy with two rulesets: **Azure CloudGuard Checkup** and **Azure CIS Foundation v. 1.5.0**.

Troubleshooting

Admin Consent

For some subscriptions, you need to add the admin consent permission to the application. Follow the instructions in the Microsoft Azure Cloud Shell.

To grant admin consent to the CloudGuard application:

1. In the navigation tree, go to **App registrations** > **All applications** and find the application you created with the script (by default, *CloudGuard-Connect*).
2. Click your CloudGuard app.

3. In the navigation tree, select **Manage** > **API permissions**.
4. The page shows in the permissions status that they need the admin consent.
5. Above the permissions list, click **Grant admin consent for Check Point Software Technologies** and confirm the option.

Unsuccessful onboarding

To remove the CloudGuard resources created in your Azure subscription:

- Run the onboarding script with the `--clean` flag.

  You can use the option for troubleshooting an unsuccessful onboarding. Make sure you use the same parameters as in the initial command.

More Links

- *"Onboarding Azure Organizations" on page 170*
- *"Azure Roles and Permissions" on page 293*
- *"Troubleshooting Azure Onboarding" on page 176*

Onboarding a Google Cloud Platform (GCP) Project and Google Workspace

Prerequisites

- You must have Owner permissions for the GCP project.
- To connect Google Workspace to CloudGuard, you must have Owner permissions for the Google Workspace.

To onboard a GCP project to CloudGuard:

1. In GCP, open the project that you want to onboard to CloudGuard. Keep the GCP project open throughout this procedure.
2. In the CloudGuard UI, from the left menu, expand **Assets** and click **Environments**.
3. In the toolbar above the table, in the top left, click **Add** and then click **GCP Project**.

   The **GCP Onboarding** wizard opens.
4. In the **Welcome** step of the wizard:
   a. Copy the **Project ID** from GCP.
   b. Paste the **Project ID** into CloudGuard.
   c. In CloudGuard, click **Next**.
5. In the **Configurations** step of the wizard:
   a. **Optional -** For **Environment Display Name**, enter a name for the integration to appear in the CloudGuard UI. By default, the **Display Name** is the **Project ID**.
   b. Select an **Organizational Unit** to associate with the integration.

      Note - An integration of GCP with CloudGuard CNAPP always includes the CSPM feature in CloudGuard CNAPP. The **CSPM** slider is on by default, and cannot be turned off.
   c. **Optional -** To onboard GCP to CloudGuard without onboarding Google Workspace, move the **Workspace** slider to "off".

      Note - It is also possible to connect Google Workspace to CloudGuard after you finish the GCP onboarding.
   d. Click **Next**.
6. In the **Connect** step of the wizard:

a. Click **Onboarding Script** to review the GCP onboarding script that CloudGuard generates automatically.

b. Copy the command from CloudGuard.

c. Paste the command into the CLI of the GCP project.

   The CLI of the GCP project generates a JSON that contains GCP credentials.

d. In the CLI of the GCP project, copy the JSON (including the opening and closing brackets).

e. In CloudGuard, paste the JSON into the **Credentials JSON** field.

f. Click **Next**.

7. In the **Workspace** step of the wizard, do **one** of these:

   - If you are onboarding a Google Workspace:

     a. Follow the instructions shown in the CloudGuard UI to configure Google Workspace.

     b. Click **Next**.

   - If you are **not** onboarding a Google Workspace, click **Skip**.

Onboarding Oracle Cloud Infrastructure Environments

You can onboard an OCI account to CloudGuard.

To onboard an OCI account to CloudGuard:

STEP 1 - Welcome

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **OCI** and follow the setup steps.

   Or, if you already onboarded environment(s), then from the top menu, select **Add** > **OCI Environment**.
3. In the onboarding wizard, enter a new name for your OCI environment. This name lets you identify the environment in CloudGuard.
4. Follow the on-screen instructions to enter mandatory (*) and optional information in the fields. Use tooltips to see more explanations.

   Note - After onboarding, you can possibly receive a validation email message to the provided **OCI tenant administrator email address**. CloudGuard creates a new user on your tenant. This does not require action, and you can safely ignore the message.
5. Click **Next**.

STEP 2 - Organizational Unit

1. Select one of the available organizational units to associate with the environment.
2. Click **Next**.

STEP 3 - Set up a New CloudGuard Environment

CloudGuard uses a Terraform template for policy definition and for permissions to read data in your tenancy. Download the template and follow the wizard instructions to configure your OCI account.

At this stage, you create and run a stack in your OCI account.

Click **Done** to complete the onboarding procedure.

The new environment appears in the list of all environments. It takes several minutes to see all assets onboarded to CloudGuard.

To remove the OCI environment data from CloudGuard:

1. On the environment page, click **Remove** to disconnect the selected environment from CloudGuard.

   A verification window opens.

2. Click **Remove** in the verification window.

   If you have active policies on the environment, CloudGuard notifies you that it deletes the policies in line with the environment.

   This process deletes:

   - Policies
   - Notifications
   - Compliance alerts (findings)

Manual Onboarding of Kubernetes Clusters

You can onboard a Kubernetes cluster to CloudGuard. On the process completion, you can see clusters, nodes, pods, and other resources on the CloudGuard Assets page. Then you can run compliance assessments on them and use the data for more security functionality, such as Runtime Protection, Image Assurance, etc.

For more onboarding options, see *"Onboarding Kubernetes Clusters" on page 188*.

For information on Kubernetes versions and container requirements, see *"Kubernetes Containers" on page 415*.

Follow the steps below to manually onboard a Kubernetes cluster to CloudGuard:

STEP 1 - Configuration

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **Kubernetes**. The first window of the wizard to onboard a Kubernetes cluster opens.

   Or, from the top menu, select **Add** > **Kubernetes / OpenShift / Tanzu**.
3. Enter a name for the cluster. This is the name that appears in CloudGuard.
4. Follow the on-screen instructions to complete these steps:
   - Configure a Service Account by one of these methods:
     - Select an existing Service Account with its corresponding API Key.
     - Enter a Service Account manually.
     - Click **Add Service Account** to create a new account.
   - Enter a name for the Kubernetes namespace in which the agent is to be deployed or keep the default name - *checkpoint*.

- Select what type of monitoring and security checks are necessary for your Kubernetes cluster by default. You can add each of these features later. Read more about each feature on a dedicated page:
  - Posture Management - for details, see *"Cloud Security Posture Management (CSPM)" on page 302* (mandatory feature)
  - Image Assurance - for details, see *"Image Assurance" on page 434*
  - Admission Control - for details, see *"Admission Control" on page 476*
  - Runtime Protection - for details, see *"Kubernetes Runtime Protection" on page 549*
  - Threat Intelligence - for details, see *"Intelligence for Kubernetes Containers" on page 645*

5. Click **Next** to continue to the next step.

STEP 2 - Select Organizational Unit

1. Select the *"Organizational Units" on page 297* with which the onboarded cluster will be associated. If no Org Unit is selected, the root (top-level) unit is used.
2. Click **Next**.

STEP 3 - Deploy the agent on the cluster

1. Follow the on-screen instructions and apply Helm. As an alternative, you can follow the Non-Helm instructions to deploy the agents. This generates a YAML file for deployment with `kubectl` commands.

   For more installation options, see *"Installing the Agent" on page 192*.
2. Click **Next**.

STEP 4 - Onboarding Summary

1. Verify the deployment status. The status is dynamically updated as the agents come online.

   CloudGuard informs you that:

   - Your Kubernetes cluster has been successfully created.
   - It is waiting for the agent to start communication.
   - You can skip the validation if you click the **Finish** button.
2. Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

After the agent is deployed, CloudGuard accesses the cluster through the agent to get information about the assets and synchronize with it. This takes several minutes based on the time needed to download the images to the cluster and the number of assets in the cluster.

The Onboarding Summary page is updated automatically with the change of the cluster status.

Cluster Status

Available options of the cluster status:

- **Pending** - CloudGuard has not received communication from the agents.
- **Initializing** - CloudGuard is receiving communication from some of the agents. The progress bar shows how many agents are up and ready.

  Note - During this state, if the number of running pods does not change for 10 minutes, the indicator pauses and the status changes to **TIME OUT**. In this case, verify the agents status on the cluster to make sure they do not have issues. For example, agents can be stuck because of missing resources (memory or CPU). After you resolve the issue, you can continue the validation or skip the validation process entirely.
- **Error** - There are agents in the Error state. Click **Finish** to complete the process. You can go to the cluster page to see which agents have the Error state and browse their Kubernetes logs for issues.

When all the agents are running, the cluster status changes to **SUCCESS**, and the onboarding process finishes successfully.

Agent Status

On the cluster page, for each feature, you can see the status of its agents:

- **Pending** - The agent has never communicated with CloudGuard.

  Note - There is a limitation for DaemonSet agents. During the cluster status calculation, tolerations settings are not considered. Agents from excluded nodes are considered **Pending** which can cause a false error state for the cluster.
- **Initializing** - Status of an agent that comes online and initiates communication with the CloudGuard portal. The agent has a small period to report a successful self-test. If the agent does not report it back on time, the status is changed to Error because of a timeout.
- **Warning** - Status of an agent that successfully finished its initialization, while it is based on an old image. See *"Upgrading the Agent" on page 196* for how to resolve this issue.

- **Error** – Status of agents that
  - failed their self-test
  - sent an error message
  - suffered a loss of connectivity for a minimum of one hour
  - have the version below the minimal version
- **Pending cleanup** - Disabled features that still have an agent that sends data to appear with the Pending cleanup status.

Onboarding Container Registries

A Container Registry is a repository that stores container images. To scan your Container Registry environment with the Image Assurance capability, onboard the Container Registry to CloudGuard.

These are two options to scan your Container Registry in CloudGuard:

- Link it to a Kubernetes cluster that has the ImageScan agents scanning your registry
- Deploy ImageScan with an AWS ECS scanner (available for selected types of registry)

CloudGuard can scan these types of container registries:

- With an **AWS ECS scanner** or a **Kubernetes scanner**:
  - Azure Container Registry (ACR) - See *"Onboarding Azure Container Registry" on page 207*
  - AWS Elastic Container Registry (ECR) - See *"Onboarding AWS Elastic Container Registry" on page 214*
  - Docker Hub Container Registry - See *"Onboarding Docker Hub Container Registry" on page 226*
  - Google Cloud Container Registry (GCR) - See *"Onboarding Google Container Registry" on page 231*
  - Harbor - See *"Onboarding Harbor Registry" on page 242*
  - JFrog Artifactory - See *"Onboarding JFrog Artifactory" on page 247*
  - Nexus - See *"Onboarding Sonatype Nexus Registry" on page 253*
  - GitHub Container Registry - See *"Onboarding GitHub Container Registry" on page 258*
  - Quay.io Container Registry - See *"Onboarding Quay.io Container Registry" on page 263*
- With a **Kubernetes scanner** only:
  - Google Artifact Registry (GAR) - See *"Onboarding Google Artifact Registry" on page 237*

General Workflow

To onboard a Container Registry to CloudGuard, follow these steps on the onboarding wizard:

1. **Registry Configurations** - Configure the registry.
2. **Cluster Configurations** - In this step, it is necessary to provide the CloudGuard Service Account credentials.
3. **Environment Configurations** - In the hosting environment, select to associate the registry with a new or existing cluster. Follow the instructions to configure the environment.
4. **Onboarding Summary** - For onboarding with a Kubernetes cluster only, CloudGuard shows the full details of your newly onboarded registry and its related cluster. If the process includes updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

CloudGuard opens the onboarded registry. For onboarding validation, in the **Scanners** tab, see the status of the registry and the cluster that scans it.

The related Kubernetes cluster page shows information on the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

Inactive Container Registries

CloudGuard deletes inactive environments when a year (365 days) passed since any of the environment's agents has communicated with CloudGuard. An agent is required to communicate with CloudGuard at least once in the past.

Note - Environments with agents that communicated with errors are not removed.

Troubleshooting

Error	Corrective Actions
Failed to create registry worker	1. Make sure you created the pull secret in the same namespace where the CloudGuard agents are located. 2. Make sure you gave the same secret name on the onboarding wizard page. **Note**: If you create or update the pull secret after the agents startup, you must restart the *imagescan-engine* and *imagescan-list* pods.

Error	Corrective Actions
Failed to authenticate	1. Make sure the pull secret key name is correct and is created in the correct namespace. 2. Make sure you entered correctly the username, password, and server URL in the secret definition.

Known Limitations

- By default, CloudGuard adds to Protected Assets and scans only 10 recent images of each repository. You can change the default value with the API call (maximal number is 1000 for a JFrog Artifactory and Sonatype Nexus). For more information, see the API Reference Guide.
- Scanning Windows container images is not supported.
- For JFrog Artifactory, it can take about 20 minutes that the images start to show for the first time.
- For JFrog Artifactory and Sonatype Nexus, the maximal number of tags per repository is 1000. Container images from the repositories with more than 1000 tags are neither shown as protected assets, nor scanned. The number is limited due to extensive API calls and performance considerations.

More Links

- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*

Onboarding Alibaba Cloud Accounts

You can onboard an Alibaba Cloud account to CloudGuard.

To onboard an Alibaba Cloud account to CloudGuard:

STEP 1 User Creation - Credentials

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **Alibaba Cloud** and follow the setup steps.

   Or, if you already onboarded environment(s), then from the top menu, select **Add** > **Alibaba**.
3. Enter a new name for your Alibaba environment. This name allows you to identify the environment in CloudGuard.
4. Follow the on-screen instructions to complete these steps:
   - In your Alibaba Cloud account, navigate to **Products and Services** > **Resource Access Management** (RAM) service > **Users** (or, search for RAM in the search field) and click **Create User**.

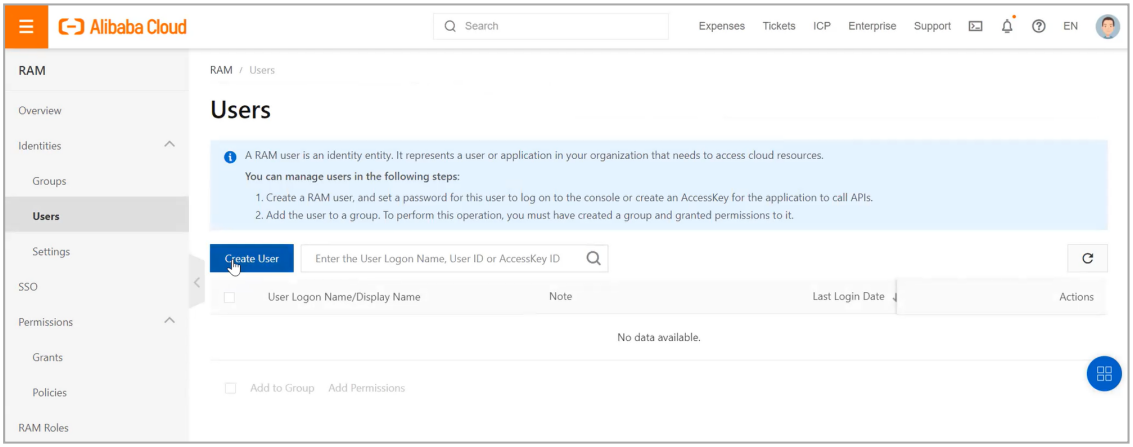

   - Below **Logon Name**, enter the name, for example, **CloudGuard-Connection**. Below **Display Name**, enter the name, for example, **CloudGuard-Connection User**.

- Below **Access Mode**, select **Programmatic Access** and click **OK**. Alibaba Cloud generates an AccessKey ID and an AccessKey Secret.

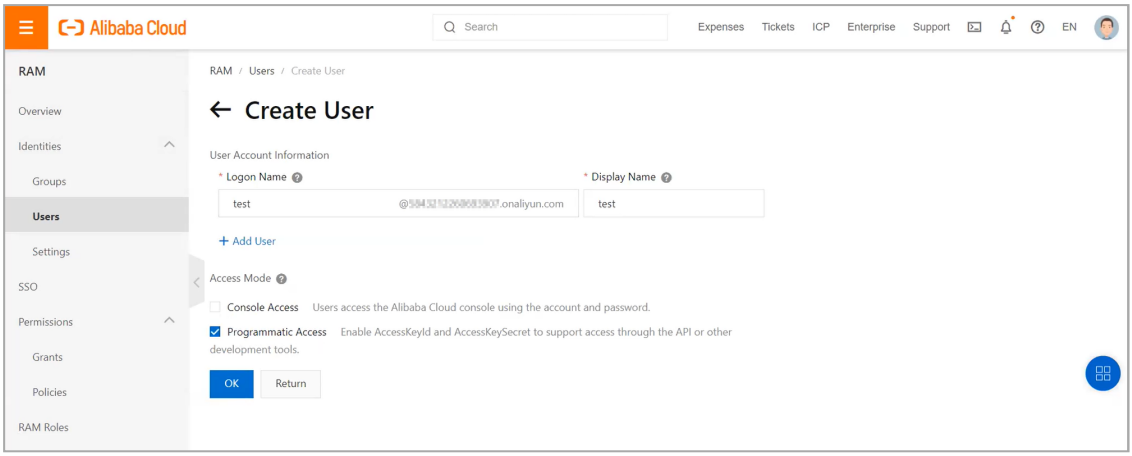

- Copy the AccessKey ID and the AccessKey Secret and enter them in CloudGuard.

5. Click **Next**.

STEP 2 User Creation - Permission

1. In your Alibaba Cloud account, select the newly created user and click **Add Permissions**.

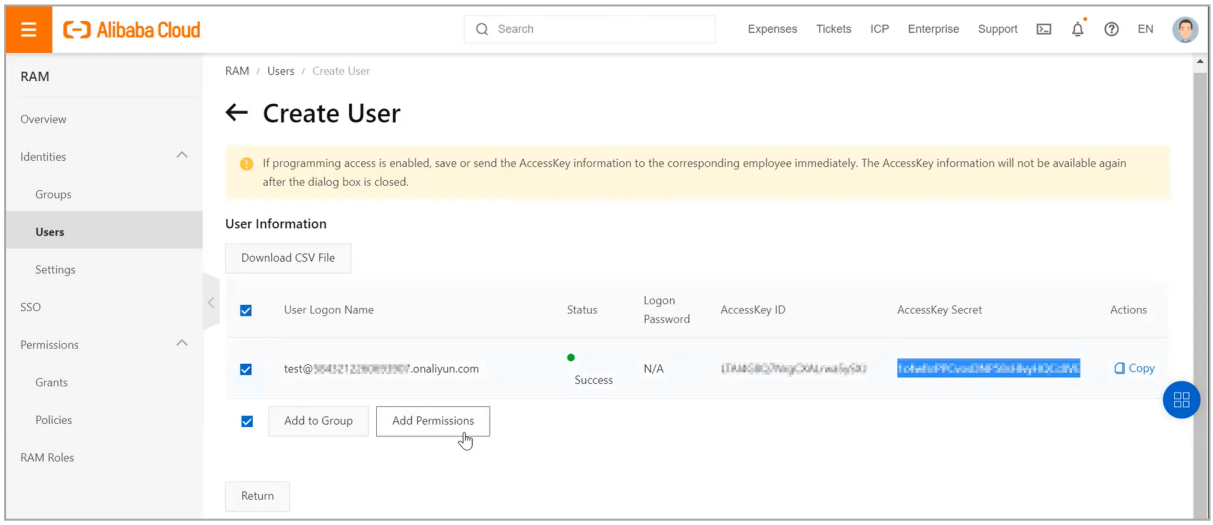

2. Make sure that below the **Authorized Scope**, the **Alibaba Cloud Account** is selected, so all resources are permitted.
3. Below **System Policy**, search for **ReadOnlyAccess** and select this option. Click **OK**.
4. Make sure the **ReadOnlyAccess** is granted and click **Complete**.
5. In the CloudGuard portal, click **Next**.

STEP 3 Summary

After CloudGuard synchronizes with the new environment, a summary of the assets shows the environment name and the Access Key ID.

Click **Finish** to complete the onboarding procedure.

The new environment appears from the list of all environments.

Alibaba's environment is an asset with several known limitations. To learn more about the limitations, see *"Alibaba Cloud Accounts" on page 924*.

Configuring CloudGuard Policies

In CloudGuard, you configure a Policy with three components:

- **Environment** - After you onboard your environment to CloudGuard, it appears on the Environments page of the Assets menu.
- **Ruleset** - CloudGuard provides a wide range of rules for each security feature (Intelligence, Compliance, Admission Control, etc.) and each cloud platform. CloudGuard applies these rules as a bundle called a ruleset, for the selected environments. You can customize the predefined rulesets and create your own rules.
- **Notification** - All your means and channels to receive information about violated rules:
  - All events and findings listed in the CloudGuard portal (Alerts console)
  - Email messages or reports
  - HTTP endpoints (ServiceNow, QRadar, SumoLogic, etc.)
  - Slack and Teams channels
  - Security management systems
  - Issue management systems

General Workflow

To configure a policy in CloudGuard:

1. Select a platform, environment, or both.
2. Select rulesets. For more details, see *"Rules and Rulesets" on page 309*.
3. Select notifications. For more details, see *"Notifications" on page 852*.

After you have onboarded your environments and clusters to CloudGuard, you can configure these policies:

- *"Getting Started with Posture Management Policy" on page 306*
- *"Getting Started with Intelligence Policy" on page 569*
- *"Getting Started with Image Assurance Policy" on page 437*
- *"Getting Started with Admission Control Policy" on page 481*

Policy Deletion

Each policy in CloudGuard is a combination of an environment, a ruleset, and a notification. When you delete a policy, you break the association between these three components. After it, none of the components is deleted, that is:

- Your environment remains onboarded to CloudGuard.
- The ruleset exists in the list of the available rulesets based on the applicable feature, for example, Intelligence or Image Assurance.
- The notification exists in the list of available notifications.

However, the events generated because of the policy application are no longer valid. The findings created before the policy deletion become *resolved* or *passed* because no rule is violated. The *passed* indication is sent to all targets (except for the email) enabled in the associated notification, such as:

- SNS
- HTTP endpoints, such as:
  - JSON-based third-party applications
  - Splunk
  - ServiceNow
  - QRadar
  - Sumo Logic
  - Jira
- Slack channel
- Teams channel
- Security Management systems
- Issue Management systems

To delete a policy:

1. Navigate to the **Policies** page in one of the CloudGuard features, for example, **CSPM** > **Continuous Posture**, **CIEM** > **Policies**, or **Workload Protection** > **Admission Control** > **Policies**.
2. Select one or more policies to delete.
3. On the top menu, click **Unassociate**.
4. Click **Yes** to confirm the operation.

Configuring CloudGuard Exclusions

You can select to exclude specific findings that appear in the results of assessments or vulnerability scanning, manually triggered compliance assessments, and Continuous Posture assessments, including these CloudGuard solutions:

- CSPM
- CIEM
- CDR
- Image Assurance (Vulnerabilities)
- Admission Control

With exclusions, you can control the findings and show only those applicable to you. After you create an exclusion, the findings that match the exclusion parameters do not appear in the calculation of the assessment result statistics. Excluded findings are not sent as notification messages (by email, SNS, etc.) to external systems.

Some typical cases to make exclusions are:

- Exclude findings from unrelated rules, for specific or for all environments. For example, when you use preconfigured CloudGuard rulesets, possibly some rules do not apply to your environments, and you can create exclusions to adjust them.
- Provide temporary correction for rules that require adjustments.
- Stop generation of findings for specific entities.

Best Practice - Do not overuse exclusions. If it is necessary to have a large number of exclusions to control your assessment results, then perhaps make adjustments to your rulesets. As a result, the rulesets fit better the current state of your cloud environments.

In the **Exclusions** page, use the *"Filter and Search" on page 863* toolbar to select parameters to filter out from the exclusion table. Only exclusions that match the parameters show up in the exclusion table.

You can use these preconfigured filters:

- Platform - Select an environment platform.
- Environment/OU - Select one or more environments or organizational units.
- Rulesets - Select from the available rulesets.
- Rules - Select from the available rules.

- Status - Select currently **Active** exclusions (in the Date Range) or **Inactive** exclusions (out of the Date Range).
- Severity - Select from the available alert severity objects.

Creating an Exclusion

There are two methods to create an exclusion:

- *Full* - Create a new empty exclusion and enter all the required parameters manually.
- *Based on assessment results or findings* - Some of the parameters exist; you can edit them and complete the missing parameters.

To create a new exclusion with the full procedure:

1. Navigate to one of the relevant menu items (**CSPM**, **CIEM**, **Workload Protection** > **Vulnerabilities or Admission Control**, or **CDR** > **Threat Monitoring**) and click **Exclusions**.
2. Click **Create New Exclusion** in the top right.
3. Select the **Ruleset** to which you apply the exclusion. CloudGuard shows only applicable rulesets. This parameter is mandatory.
4. Enter your comment to distinguish between different exclusions. The comment is a mandatory parameter.
5. Select at least one characteristic whose finding should be excluded from the results. The list below contains all available parameters that your type of exclusion may have or may have not:
    - **Environment or Organization unit** - Exclude findings that correspond to an asset from a specific environment or organization unit. The field shows only environments that match the platform of the selected ruleset.
    - **Regions** - Select one or more regions where you want the exclusion to apply.
    - **Date range** - Select during which time frame the exclusion takes effect. If you do not select the date range, the exclusion applies permanently.
    - **Rule** - Exclude findings that correspond to a specific rule. Select the rule from the list based on the selected ruleset. If you do not select a rule, the exclusion applies to all rules. The rule severity applies to the exclusion automatically, so you cannot configure it separately.

- **Entity** - Exclude findings that correspond to specific entities. Enter the entity name or ID. You can enter one or more entity names. Start to type the entity name to see and select a matching option. You can include the wildcard '%' in the entity name, to include a group of entities. For example, `%s3%` matches all entities with 's3' in their name.
- **Account number** - Exclude findings that correspond to an AWS account with a specific number.
- **Tags** - Exclude findings that contain specific tags (key + value).
- **Alerts severity** - Exclude findings that have specific severity. You cannot select a rule when you select the alert's severity, because each rule has its severity level.

Note - The exclusion characteristics apply to the finding with the AND logic. For example, if you set the date range, rule, and account number, the finding is excluded from the assessment if it matches all the parameters at the same time. That is, it matches the configured ruleset **and** the specified date range **and** the specified rule **and** it has the specified account number. To apply the characteristics with the OR logic, create more exclusions.

6. Click **Save**.

To create an exclusion based on existing parameters:

A simpler procedure to create an exclusion is to start the procedure directly from an assessment (*"Creating Exclusions from Assessment" on page 305*), finding (*"Creating exclusion for findings" on page 126*), or GSL rules. In the **Create New Exclusion** window, some parameters appear configured as they are in your assessment or finding.

Parameters for CDR

CDR exclusions have these additional parameters:

- **Source or Destination IP** - Exclude findings that have a specific Source or Destination IP address, IP range, or use saved IP lists. For more information about IP lists, see *"Custom Resources" on page 300*.
- **Source or Destination Port Range** - Exclude findings that have a specific Source or Destination Port or Port Range.

Parameters for Vulnerabilities

Vulnerability exclusions have these additional parameters:

- **Finding Type** - Select one of these types:
  - **Package** - Enter the package name, version, or path.
  - **Malware** - Enter the name or path.
  - **Insecure Content** - Enter the path or payload-sha256.

Parameters for Admission Control

Admission Control exclusions have these additional parameters (filters):

- Annotation
- Namespace - The namespace where the agents are installed can be referred to with the variable `CHECKPOINT_NAMESPACE`.
- Role
- Service Account
- Label

Important - To optimize the ruleset, the workload object model considers only the containers section of the YAML. This ensures the enforcement of a rule on all workload types: pods, deployments, daemonsets, etc. This means that the labels on the top-level metadata are not considered, but only labels on the container metadata are inspected.
For example, if the label *"app: test"* is added as an exception, the workload and its pods are created.
With this, if the label *"app: test-pod"* is added, the workload is created, but its pod is not.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: deployment-test
  labels:
    app: test
spec:
  replicas: 2
  selector:
    matchLabels:
      app: test-pod
  template:
    metadata:
      labels:
        app: test-pod
    spec:
      containers:
        - name: front-end
          image: nginx
          ports:
            - containerPort: 80
```

Editing an Exclusion

1. Navigate to one of the relevant menu items (**CSPM**, **CIEM**, **Workload Protection** > **Vulnerabilities or Admission Control**, or **CDR** > **Threat Monitoring**) and click **Exclusions**.
2. Select an exclusion to edit and click **Edit** on top bar.

Deleting an Exclusion

1. Navigate to one of the relevant menu items (**CSPM**, **CIEM**, **Workload Protection** > **Vulnerabilities or Admission Control**, or **CDR** > **Threat Monitoring**) and click **Exclusions**.
2. Select an exclusion to delete and click **Delete** on the top bar.

Using API

You can use API to configure a new exclusion. For more information, see the CloudGuard API Reference Guide - https://docs.cgn.portal.checkpoint.com/reference/complianceexclusion_post_post_v2complianceexclusion.

Vulnerability Exclusions

For `logicExpressions` strings, use a combination of these expressions:

- `name like 'PackageName' and category = 'Package'`
- `version like 'PackageVersion' and category = 'Package'`
- `package-manager.path like 'PackageManagerPath' and category = 'Package'`
- `(scannedAsset.entityName like 'entityName1' or scannedAsset.entityName like 'entityName2') and category = 'Package'` (up to 10 entities)
- `files contain [file-path like 'FilePath'] and category = 'InsecureContent'`
- `files contain [contents contain [payload-sha256 like 'InsecureContentPayloadSha256']] and category = 'InsecureContent'`
- `name like 'MalwareName' and category = 'Malware'`
- `files contain [file-path like 'FilePath'] and category = 'Malware'`

More Links

- *"Cloud Security Posture Management (CSPM)" on page 302*
- *"CIEM Policies, Exclusions, and Remediation" on page 397*
- *"Image Assurance" on page 434*
- *"Admission Control" on page 476*
- *"Cloud Detection and Response (CDR)" on page 565*
- *"Kubernetes Runtime Protection Rules and Exclusions" on page 556*
- *"Creating Serverless Functions Exclusions" on page 539*

Dashboards

The CloudGuard Dashboard provides summary and aggregates data from different CloudGuard data sources, such as Events, Protected Assets, Intelligence, and external integrations such as Tenable.io, Serverless, and Kubernetes.

The dashboard shows widgets, which show specific CloudGuard statistics, such as alerts or assets. They can be graphical, showing, for example, the issuing of alerts by severity or account, or lists, showing, for example, the top protected asset types in an account.

You can create custom dashboards and set up the layout of widgets in them. You can customize the type of graphical display for widgets, such as pie charts, histograms, and gauges.

The Dashboards feature:

- At-a-glance view of your cloud posture, assets, and alerts
- Customizable layout and graphical widgets
- Creating and sharing multiple, customized dashboards
- Information display at the Organizational level
- Click-thru interface to show more detail in applicable CloudGuard pages

Home Dashboard

A preconfigured default Dashboard is included in CloudGuard. This is the Home Dashboard shown as the CloudGuard home page. It shows a broad summary of your CloudGuard statistics. You cannot change the preconfigured dashboards.

Viewing and Configuring Dashboards

Viewing Dashboards

1. Navigate to the **Overview** and select **Home Dashboard** to open the default Dashboard.
2. Select one of the default Dashboards from the list in the top right corner, which includes all the Dashboards you create.
3. Click an item in a widget to open the related page in CloudGuard and show more details.
4. Widgets with no data may not be shown based on the settings for the widget (see below *"Adding Widgets to a Dashboard" on the next page*). You can override this and show all widgets. Select Show All, in the options bar at the top of the dashboard, to show all widgets.

Creating a Dashboard

You can create custom dashboards and include widgets of your choice to show custom information. You can share these dashboards with other users.

1. On the **Dashboard** page, click **Add Dashboard** in the options menu.
2. Enter a name for the dashboard. Click **Public** if it is necessary to share the dashboard with other users on the same CloudGuard account. A new dashboard page appears, without any widgets.
3. Follow the steps below to add widgets and customize the layout of the dashboard.
4. As an alternative, you can clone an existing dashboard, such as the Home dashboard, and make changes to the clone. Open the source dashboard and select **Clone Dashboard** from the options menu to clone the dashboard.

Adding Widgets to a Dashboard

You can add widgets to custom dashboards, but not to configured dashboards.

1. In the top right corner of a Dashboard, click **Add new widget**.
2. Enter these details for the widget (the details can change based on the widget source and type):
   - **Title** - A title for the widget that appears at the top of the widget.
   - **Source** - The type of information to show, such as alerts, assets, events, or compliance statistics.
   - **Type** - The type of display (pie chart, histogram, etc). This varies based on the source type.
   - **Aggregation** - The specific information that is shown in the widget. For example, for a widget showing assets, the aggregation is the *Asset Type*.
   - **Size** - The size of the widget as it appears on the page, in units. The dashboard page is 8 units wide.
   - **Visibility** - If the widget is displayed when there is no data for it.
   - **Filters** - Filter the information shown in the widget to specific accounts, regions, platforms, etc.
   - **Description** - More information that appears next to the widget title (optional field).
   - **Widget Preview** - Small example of the widget appearance after you fill in all mandatory fields.
3. Click **Save**. The widget appears on the dashboard page.

Changing a Dashboard Layout

Change the location and size of widgets on a custom dashboard.

1. Click the options menu of a widget and then select **Settings**.
2. Change settings for the widget such as the size, type, and data source.
3. Click **Save**.

To change the location of a widget, move it to a new location.

Saving Dashboards to a Preference List

Save custom dashboards to a list of preferred dashboards. This list appears in the Dashboard menu, for fast selection.

To set a dashboard as preferred, select **Pinned** from the options menu at the top of the dashboard page.

Setting the Default Dashboard

To set a dashboard as the Default dashboard, select **Default** from the options menu at the top of the dashboard page. The dashboard you set is saved as default only on the user level. Only one dashboard, built-in or custom, can be the default dashboard. Initially, the Home Dashboard is the default one.

Risk Management

CloudGuard provides many alerts from sources such as the Compliance engine, Intelligence, CIEM, and AWP. Security teams do not always know which alerts are most important.

Effective Risk Management (ERM) helps you prioritize remediation and ensures you make the highest impact to decrease the risk for your cloud environments. CloudGuard calculates a risk score for cloud assets based on many inputs.

Benefits

- Maximizing the productivity of security teams.
- Visibility of each asset's risk score based on all risk vectors, context, and business priority.

ERM Protected Assets

ERM applies to these assets:

AWS	Microsoft Azure	GCP	Kubernetes
■ API Gateway ■ API Gateway V2 ■ AppSync ■ Auto Scaling Group ■ DynamoDB Table ■ EC2 Instance ■ ECR Repository ■ ECS Service ■ EKS Cluster ■ IAM Role ■ IAM User ■ KMS ■ Lambda ■ RDS ■ Redshift ■ S3 Bucket ■ Secrets Manager ■ SNS Topic ■ SQS	■ Cache for Redis ■ Cosmos DB ■ Function App ■ Group ■ SQL Managed Instance ■ SQL Server ■ Storage Account ■ Storage Blob Container ■ User ■ User Assigned Identity ■ App Registration ■ Virtual Machine ■ Virtual Machine Scale Set ■ Web App ■ MySQLDBSingleServer ■ MySQLDBFlexibleServer ■ PostgreSQL ■ PostgreSQLFlexibleServer	■ Cloud Functions ■ Cloud Storage Bucket ■ GCP IAM User ■ VM Instance	■ CronJob ■ DaemonSet ■ Deployment ■ Pod ■ ReplicaSet ■ StatefulSet

The **Protected Assets** page shows the Risk Score of your assets. The score ranges from 0 to 10. The assets with the highest risk score are necessary to be addressed first. CloudGuard recalculates the risk score every few hours or in case of a business priority update.

Click an asset to see its details.

The **Overview** tab summarizes information about the security posture of the asset. This includes the most important (top five) remediation actions that are necessary to reduce the asset risk.

More Links

- *" Risk Management" on page 89*
- *"Agentless Workload Posture" on page 489*

Toxic Combinations

Toxic Combinations show potential attack paths. CloudGuard shows Toxic Combinations automatically. You do not need to configure a policy to show Toxic Combinations.

CloudGuard combines insights from different features to identity Toxic Combinations. If you enable more features, then CloudGuard can identify more Toxic Combinations. Enable one or more of these CloudGuard features to generate insights for Toxic Combinations:

- *"IAM Exposure" on page 112*
- *"Data Sensitivity" on page 113*
- Data Sensitivity and Posture Management (DSPM) with Sentra
- *"Workload Protection" on page 410*
- *"Cloud Detection and Response (CDR)" on page 565*

Toxic Combinations appear in a table. To open the table, in the left menu expand **Toxic Combinations** > click **Issues**. You can sort the table by adding filters in the **Group By** menu. The **Overview** button sorts the table by **Severity** and **Title**.

Click a Toxic Combination in the table to see details and remediation suggestions.

Known Limitations

- CloudGuard shows Toxic Combinations for AWS and Azure environments only.

Action Hub

On this page, you can create notifications based on Toxic Combinations in your environment.

Prerequisite

You must configure an Integration in the **Integration Hub** to receive the notification from CloudGuard. For more information, see *"Integration Hub" on page 796*. These Integrations are supported for Toxic Combinations:

- Generic Webhook
- Microsoft Teams
- Slack
- Email

To create a notification based on Toxic Combinations:

1. From the left menu, expand **Risk Management** > click **Action Hub**.
2. Click **Add**.

   The **New Automated Action** sliding window opens.
3. Enter a **Name** for the notification.
4. Finish creating the notification.
5. Click **Save**.

Exclusions for Toxic Combinations

You can exclude Toxic Combinations from the table on the **Toxic Combinations** page. This can help you do a more focused investigation of Toxic Combinations in your environment. Exclusions apply to all users of the CloudGuard tenant.

Exclusions apply only to the table on the **Toxic Combinations** page. If you exclude a Toxic Combination from the table, CloudGuard continues to scan your environment for that Toxic Combination and enters findings into logs.

To exclude a specific Toxic Combination from the table

1. From the left menu, expand **Risk Management** > click **Toxic Combinations**.
2. If **Overview** is selected, remove all filters from the table.
3. Select a checkbox to the left of a Toxic Combination.
4. Above the table, click **Exclude**.

   The **New Exclusion** sliding window opens. The relevant options are selected automatically to exclude the Toxic Combination.
5. Click **Save**.

To create a general Exclusion

1. From the left menu, expand **Risk Management** > click **Exclusions**.
2. Click **Add**.

   The **New Exclusion** sliding window opens.
3. Enter a **Name** for the exclusion.
4. Finish creating the exclusion.
5. Click **Save**.

To edit a general Exclusion

1. From the left menu, expand **Risk Management** > click **Exclusions**.
2. Select an exclusion.
3. Click **Edit**.
4. Edit the exclusion.
5. Click **Save**.

To delete an Exclusion

1. From the left menu, expand **Risk Management** > click **Exclusions**.
2. Select an exclusion.
3. Click **Delete**.

Ignoring CVEs from Toxic Combinations

When CloudGuard calculates the Risk Score for a Toxic Combination, CloudGuard considers known vulnerabilities (CVEs) it identifies in your cloud assets. You can configure CloudGuard to ignore a specific CVE from the Toxic Combinations calculation for an Entity, Organizational Unit, or Environment. Ignore a CVE if you do not want to focus on it in your investigation. After you ignore a CVE, CloudGuard may assign a lower risk score to a Toxic Combination or remove it from the Toxic Combinations table.

To ignore a CVE from a Toxic Combination for a specific entity

1. From the left menu, expand **Risk Management** > click **Toxic Combinations**.
2. Select a Toxic Combination.

   A sliding window opens.
3. In the sliding window, expand the **Vulnerabilities** section.
4. To the right of the relevant CVE, click **Ignore**.

   The **New CVE Ignore Item** window opens.
5. **Optional -** Enter or change one or more of these attributes of the CVE Ignore Item:
   - **Name**
   - **Description**
   - **Expiration Date** - By default, the exclusion is permanent.

   Note - The **CVE Details** section and the **Vulnerable Entity** section are filled automatically to ignore the CVE for the entity. To add more CVEs or entities to the CVE Ignore Item, see *"To ignore one or more CVEs from Toxic Combinations for multiple entities" below*.
6. Click **Save**.

To ignore one or more CVEs from Toxic Combinations for multiple entities

1. From the left menu, expand **Risk Management** > expand **Toxic Combinations**.
2. Click **CVEs Ignore List**.
3. Click **Add**.

   The **New CVE Ignore Item** window opens.
4. Enter a name for the CVE Ignore Item.
5. **Optional -** Enter a Description.
6. **Optional -** Enter an Expiration date. By default, the exclusion is permanent.

7. In the **CVE Details** section, fill **one** of these fields to identity one or more CVEs to ignore:
   - **CVE IDs** - Enter one or more CVE IDs to ignore.
   - **Package name** - Enter one or more package names. CloudGuard ignores all CVEs that it finds in these packages.
   - **Package path** - Enter one or more package paths. CloudGuard ignores all CVEs that it finds in these package paths.

   Note - You can use % as a wildcard. For example, `arn:aws:lambda:us-east-1:123456789012:function:%` applies to all Lambda function names in the us-east-1 region for the AWS account number `123456789012`.
8. In the **Vulnerable Entity** section, select where to ignore the CVE(s):
   - To ignore the CVE from all entities in one or more Organizational Units, select **Organizational Unit** and enter the names of the Organizational Unit(s).
   - To ignore the CVE from all entities in one or more Environments, select **Environment** and enter the names of the Environment(s).
   - To ignore the CVE from one or more specific entities, select **Entity Name** or **Entity ID** and enter the names or IDs of one or more Entities.
9. Click **Save**.

To edit a CVE Ignore Item

1. From the left menu, expand **Risk Management** > expand **Toxic Combinations**.
2. Click **CVEs Ignore List**.
3. Click the name of the CVE Ignore Item.

   A sliding window opens.
4. Edit the CVE Ignore Item.
5. Click **Save**.

To delete a CVE Ignore Item

1. From the left menu, expand **Risk Management** > expand **Toxic Combinations**.
2. Click **CVEs Ignore List**.
3. Select the checkbox to the left of the name of the CVE Ignore Item.
4. Click **Delete**.
5. In the confirmation window, click **Delete**.

Ignoring Malware from Toxic Combinations

When CloudGuard calculates the Risk Score for a Toxic Combination, CloudGuard considers malware families it identifies in your cloud assets. You can configure CloudGuard to ignore a specific malware family from the Toxic Combinations calculation for an Entity, Organizational Unit, or Environment. Ignore a malware family if you do not want to focus on it in your investigation. After you ignore a malware family, CloudGuard may assign a lower risk score to a Toxic Combination or remove it from the Toxic Combinations table.

To ignore a malware family from a Toxic Combination for a specific entity

1. From the left menu, go to **Risk Management** > **Toxic Combinations**.
2. Select a Toxic Combination.

   A sliding window opens.
3. In the sliding window, expand the **Vulnerabilities** section.
4. To the right of the relevant malware family, click **Ignore**.

   The **New Malware Ignore Item** window opens.
5. **Optional -** Enter or edit one or more of these attributes of the Malware Ignore Item:
   - **Name**
   - **Description**
   - **Expiration Date** - By default, the Malware Ignore Item is permanent.

   Note - The **Malware IDs** section and the **Vulnerable Entity** section are filled automatically. To add more malware families or entities to the Malware Ignore Item, see *"To ignore one or more malware families from Toxic Combinations for multiple entities" below*.
6. Click **Save**.

To ignore one or more malware families from Toxic Combinations for multiple entities

1. From the left menu, go to **Risk Management** > **Toxic Combinations**.
2. Click **Malware Ignore List**.
3. Click **Add**.

   The **New Malware Ignore Item** window opens.
4. Enter a name for the Malware Ignore Item.
5. **Optional -** Enter or edit one or more of these attributes of the Malware Ignore Item:

- **Description**
- **Expiration Date** - By default, the Malware Ignore Item is permanent.

6. In the **Malware Details** section, click the **+** (plus sign) button.
7. Enter the relevant Malware ID for the malware family.
8. **Optional** - Add more malware IDs to the Malware Ignore Item.
9. In the **Vulnerable Entity** section, select where to ignore the malware:
   - To ignore the malware from all entities in one or more Organizational Units, select **Organizational Unit** and enter the names of the Organizational Unit(s).
   - To ignore the malware from all entities in one or more Environments, select **Environment** and enter the names of the Environment(s).
   - To ignore the malware from one or more specific entities, select **Entity Name** or **Entity ID** and enter the names or IDs of one or more Entities.
10. Click **Save**.

To edit a Malware Ignore Item

1. From the left menu, expand **Risk Management** > expand **Toxic Combinations**.
2. Click **Malware Ignore List**.
3. Click the name of the Malware Ignore Item.

   A sliding window opens.
4. Edit the Malware Ignore Item.
5. Click **Save**.

To delete a Malware Ignore Item

1. From the left menu, go to **Risk Management** > **Toxic Combinations**.
2. Click **Malware Ignore List**.
3. Select the checkbox to the left of the name of the Malware Ignore Item.
4. Click **Delete**.
5. In the confirmation window, click **Delete**.

Security Controls

The **Security Controls** page shows the rules that CloudGuard uses to look for Toxic Combinations in your environment. Check Point creates and manages these rules automatically. The table is for informational purposes only. You can filter the table to learn more about specific types of Security Controls.

Risk Calculation

CloudGuard assesses cloud risk based on findings, exposure, privilege levels, and other factors.

CloudGuard gives a **Risk Score** to cloud assets. The **Risk Score** of a cloud asset is a number between **0.1** and **10.0**.

CloudGuard gives a **Risk Level** to cloud environments. The **Risk Level** of a cloud environment is **Low**, **Medium**, **High**, or **Critical**.

This table shows the correspondence between Risk Levels and Risk Scores (and their background colors):

Risk Score	Risk Level	Background Color
0.1 - 3.9	Low	
4.0 - 6.9	Medium	
7.0 - 8.9	High	
9.0 - 10.0	Critical	

Asset Risk (Risk Score)

CloudGuard analyzes your cloud assets and gives a risk score to each supported asset. CloudGuard considers these factors:

- The attack surface of each asset (for example, Common Vulnerabilities and Exposures (CVEs), misconfigurations, or Toxic Combinations)
- The likelihood that the asset is a target for attacks (for example, publicly exposed assets or WAF availability)
- The possible impact if the asset is compromised (for example, business priority)

CloudGuard recalculates the risk score after you change the rules for risk calculation and after you change the business priority of an asset. If you do not do one of these, CloudGuard recalculates the risk score once every several hours. For more information about rulesets and configuration instructions, see *"ERM Rulesets" on page 117*.

To see the risk score of your assets, navigate to **Risk Management** > **Protected Assets**. For more information, see *"ERM Protected Assets" on page 90*.

How CloudGuard Calculates a Risk Score

1. CloudGuard uses these findings to calculate a risk score:
   - toxic combinations
   - CVEs
   - misconfigurations
   - threats
   - secrets

   CloudGuard supports several vulnerability scanners. CloudGuard's AWP solution finds CVEs, threats, and secrets. In AWS environments, Amazon Inspector v2 also scans assets for CVEs. The risk score calculation does **not** include vulnerabilities with **Informational** and **Unknown** severity.
2. CloudGuard classifies each finding by severity: **Low**, **Medium**, **High**, or **Critical**.
3. CloudGuard uses a formula to determine the risk score for an asset. These are general principles of how the formula compares assets:
   a. An asset with a higher severity finding gets a higher risk score than an asset with lower severity findings.

      Example - Asset A has 1 **High** severity finding, and **Asset B** has 20 **Medium** security findings. **Asset A** gets a higher risk score.

   b. If two assets have a highest-level security finding at the same level, the one with a larger number of findings at the highest level gets priority.

      Example - Asset C has 2 **High** severity findings, and **Asset D** has 4 **High** severity findings. **Asset D** gets a higher risk score.

   c. If two assets have the same number of findings at the highest common level, the asset with more findings at the next highest level gets a higher risk score.

      Example - Asset E has 2 **High** severity findings and 2 **Medium** severity findings. **Asset F** has 2 **High** severity findings and 1 **Low** severity finding. **Asset E** gets a higher risk score.

How CloudGuard Modifies the Risk Score based on Context Modifiers

After CloudGuard takes findings into account, CloudGuard adjusts the risk score based on the business priority of the asset. For more information about Business Priority, see *"Business Priority" on page 118*.

In addition, CloudGuard modifies the Risk Score based on contextual data, including:

- **Network Exposure** - The level of network accessibility from the public domain. If a network is partially public or private, CloudGuard reduces the risk score by a constant magnitude.
- **IAM Exposure** - The level of asset accessibility from the public domain. If an asset is partially public or private, CloudGuard reduces the risk score by a constant magnitude.
- **IAM Sensitivity** - The possible damage caused to the cloud environment because of IAM permissions. The less sensitive the asset, the more CloudGuard reduces the risk score. For more information, see IAM_Sensitivity.
- **Data Sensitivity** - Indicates if the data in the asset is sensitive or not. If the asset does not hold sensitive data, CloudGuard reduces the risk score by a constant magnitude.
- **WAF Protection** - Adds a layer of security and reduces the risk score by a constant magnitude. For more information, see *"WAF Protection" on page 116*.

Environment Risk (Risk Level)

CloudGuard calculates a risk level for an environment. The risk level is based on the risk scores of the assets in the environment. The risk level calculation does not consider assets that are stopped and assets that have a risk score of zero. Environment Risk is supported for these platforms:

- AWS
- Azure
- GCP
- Kubernetes

How CloudGuard Calculates Environment Risk

CloudGuard uses a formula to determine the environment risk. These are general principles of how the formula compares environments:

1. An environment with higher-risk assets gets priority over an environment with lower-risk assets.

   Example - The highest-scoring asset in **Environment A** is in the **Critical** range. The highest-scoring asset in **Environment B** is in the **High** range. **Environment A** gets priority.

2. If two environments have assets at the same highest risk level, the environment with a larger number of assets at the highest level gets priority.

   Example - The highest-scoring assets in **Environment C** and in **Environment D** are in the **High** range. **Environment C** has 10 assets in the **High** range **Environment D** has 4 assets in the **High** range. **Environment C** gets priority.

3. If two environments have the same number of assets at the highest common risk level, the environment with more assets at the next-highest level gets priority.

   Example - The highest-scoring assets in **Environment E** and **Environment F** are in the High range. **Environment E** and **Environment F** each have 10 assets in the **High** range. **Environment E** has 25 assets in the **Medium** range. **Environment F** has 15 assets in the **Medium** range. **Environment E** gets a higher environment risk.

More Links

- *"ERM Rulesets" on page 117*
- *"ERM Protected Assets" on page 90*
- *"Agentless Workload Posture" on page 489*
- *"Business Priority" on page 118*
- *"Network Exposure" on page 105*
- *"IAM Exposure" on page 112*
- *"Data Sensitivity" on page 113*
- *"WAF Protection" on page 116*
- Amazon Inspector documentation

Network Exposure

Network exposure is the level of accessibility of the asset from the public domain. CloudGuard considers the network exposure of your assets and defines each of them as:

- **Public** - Accessible from the Internet.
- **Partially public** - Accessible from specific addresses on the Internet.
- **Private** - Confirmed as having no exposure.
- **Unknown** - CloudGuard cannot determine the asset exposure based on the available data.

Supported Asset Types

CloudGuard analyzes these asset types to calculate Network Exposure:

AWS Asset Types

Asset Type	Supported Statuses	Assets Used for Calculation
EC2 instance	■ Public ■ Partially public ■ Private ■ Unknown	■ Subnet ■ VPC ■ Route Table ■ NACL ■ Security Group ■ Load Balancers (Network / Application) ■ Elastic Load Balancer ■ Target Group
ECS service	■ Public ■ Partially public ■ Private ■ Unknown	■ Subnet ■ VPC ■ Route Table ■ NACL ■ Security Group ■ Load Balancers (Network / Application) ■ Target Group
Lambda	■ Public ■ Partially public ■ Unknown	■ API Gateway V1/2 ■ Application Load Balancer ■ Security Group

Asset Type	Supported Statuses	Assets Used for Calculation
RDS	■ Public ■ Partially public ■ Private ■ Unknown	■ Subnet ■ VPC ■ Route Table ■ NACL ■ Security Group
EKS Cluster	■ Public ■ Partially public ■ Private	■ Cluster endpoint ■ CIDR

Microsoft Azure Asset Types

Asset Type	Supported Statuses	Assets Used for Calculation
Virtual Machine	■ Public ■ Partially public ■ Unknown	■ Subnet ■ Network Security Group ■ Application Security Group ■ Application Gateway ■ Public IP Address ■ NIC ■ Load Balancers
Storage account*	■ Public ■ Partially public ■ Unknown ■ Private	For partially public access: ■ Selected Virtual Networks (VNet) ■ Selected IP addresses
WebApp	■ Public ■ Partially public ■ Private ■ Unknown	■ WebApp Access Restrictions ■ Application Gateway ■ Subnet ■ Network Security Group
FunctionApp	■ Public ■ Partially public ■ Private ■ Unknown	■ FunctionApp Access Restrictions ■ Application Gateway ■ Subnet ■ Network Security Group ■ Network Triggers ■ Functions Authorization Level

Asset Type	Supported Statuses	Assets Used for Calculation
SQL Server	■ Public ■ Partially public ■ Private ■ Unknown	■ SQL Server Public Access ■ SQL Server Firewall Rules
Storage Blob Container	■ Public ■ Partially public ■ Private ■ Unknown	■ Storage Account Networking ■ Storage Account Blob Anonymous Access ■ Blob Container Anonymous Access Level
Redis Cache	■ Public ■ Partially public ■ Private ■ Unknown	■ Redis Cache Public Endpoint and Firewall Rules
Cosmos DB	■ Public ■ Partially public ■ Private	■ CosmosDB Account Public Access ■ CosmosDB Account Firewall Rules
MySQLDBFlexibleServer	■ Public ■ Partially public ■ Private ■ Unknown	■ Server Public Access ■ Server Firewall Rules
PostgreSQLFlexibleServer	■ Public ■ Partially public ■ Private ■ Unknown	■ Server Public Access ■ Server Firewall Rules
SQL Managed Instance	■ Public ■ Partially public ■ Private ■ Unknown	■ SQL Managed Instance Public Endpoint ■ SQL Managed Instance Subnet NSG

* When you configure Azure storage accounts networking, you select one of these options for public network access:

- **Enable from all networks** - The storage account can be accessed from any network or IP address on the public Internet.
- **Enable from selected virtual networks and IP addresses** - The storage can be accessed from the selected IP addresses and VNet.
- **Disable and use private access** (Private endpoint connections) - The storage cannot be accessed from a public endpoint.

GCP Asset Types

Asset Type	Supported Statuses	Assets Used for Calculation
VMInstance	■ Public ■ Partially public ■ Unknown	■ NIC ■ Effective Firewall Rules ■ Effective Firewall Policies

Kubernetes Asset Types

The calculation of Kubernetes network exposure depends on the cloud platform of the onboarded environment. CloudGuard supports Kubernetes assets on the AWS cloud platform. For other unsupported environments, it shows the network exposure as **Unknown**.

Asset Type	Supported Statuses	Assets Used for Calculation
Kubernetes Pod	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes Deployment ■ Kubernetes ReplicaSet ■ Kubernetes StatefulSet ■ Kubernetes DaemonSet ■ Kubernetes CronJob ■ Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group

Asset Type	Supported Statuses	Assets Used for Calculation
Kubernetes Deployment	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes Deployment ■ Kubernetes ReplicaSet Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group
Kubernetes ReplicaSet	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes Deployment ■ Kubernetes ReplicaSet Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group

Asset Type	Supported Statuses	Assets Used for Calculation
■ Kubernetes StatefulSet	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes StatefulSet ■ Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group
Kubernetes DaemonSet	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes DaemonSet ■ Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group

Asset Type	Supported Statuses	Assets Used for Calculation
Kubernetes CronJob	■ Public ■ Partially public ■ Private ■ Unknown	■ Kubernetes Pod ■ Kubernetes CronJob ■ Kubernetes Service ■ Kubernetes Ingress ■ EC2 instance • Subnet • VPC • Route Table • NACL • Security Group • Load Balancers (Network / Application) • Elastic Load Balancer • Target Group

IAM Exposure

Identity and Access Management (IAM) exposure is the level of accessibility of the asset from the public domain based on its access permissions.

The table below shows the AWS assets that CloudGuard uses for the IAM exposure calculation and their status.

Platform	Asset Type	Possible Status	Assets Used for Calculation
AWS	S3 bucket	■ Public - accessible to a wide range of principals ■ Partially Public – specific objects in the asset can be public ■ Private – the asset is accessible to a limited set of principals ■ Unknown - CloudGuard cannot determine the IAM exposure based on the available data	■ S3 bucket Resource Policy ■ S3 bucket ACL policy ■ Access Point ■ Multi-Region Access Point ■ Block Public Access settings on account and resource levels
	IAM Role	■ Public - Roles can be assumed by a wide range of AWS principals ■ Partially Public - Roles can be assumed by a third-party principal ■ Private - Roles can be assumed only by a limited set of principals	Resource Policy
	Lambda	■ Public - A wide range of AWS principals can do at least one action on the asset ■ Private - Only a limited set of principals can take actions on the resource	
	SQS		
	SNS Topic		
	ECR Repository		

Data Sensitivity

Data sensitivity shows if data in the asset is sensitive or not. Data is considered sensitive when the asset contains, for example:

- **Credentials**, such as private keys or secret access keys
- **Financial information**, such as credit card numbers or bank account numbers
- **Sensitive personal information**, such as health insurance or medical identification numbers

The risk score considers the data sensitivity of your assets and defines each of them as:

- **Sensitive** - The asset contains sensitive data.
- **Not sensitive** - The asset does not contain sensitive data.
- **None** - CloudGuard cannot calculate the data sensitivity of the asset based on the available information.

CloudGuard assigns one of these **Data Classification** categories to each asset:

- **PII** (Personal Identifiable Information)
- **PCI** (Payment Card Industry)
- **PHI** (Protected Health Information)
- **Credentials**
- **Other**

The table below shows the sources that CloudGuard uses for the data sensitivity classification.

Data Security Posture Management (DSPM) Provider	Platform	Asset Type	Sources for Data Classification
AWS Macie	AWS	S3 bucket	CloudGuard uses the sensitivity score calculated by Amazon Macie to find the data sensitivity of the S3 bucket.

Data Security Posture Management (DSPM) Provider	Platform	Asset Type	Sources for Data Classification
Microsoft Purview	Azure	Storage Account	For each Microsoft Purview account that you connect to CloudGuard, you must grant a **Data Reader** role in **Root Collection** to the **App Registration** that you created during CloudGuard onboarding.
		Cosmos DB	
		MySQLDBFlexibleServer	
		SQL Server	
Cyera	Azure	Storage Account	Create a Cyera integration in the CloudGuard Integration Hub. See *"Classifying Assets with Cyera" on page 831*.
		Storage Blob Container	
		SQL Server	
		Cosmos DB	
		MySQLDBFlexible Server	
		PostgreSQLFlexibleServer	
		Virtual Machine	
		SQL Managed Instance	
		Cache for Redis	

Data Security Posture Management (DSPM) Provider	Platform	Asset Type	Sources for Data Classification
Sentra	AWS	S3 Bucket	Create a Sentra integration in the CloudGuard Integration Hub. See *"Classifying Assets with Sentra " on page 830*.
		RDS	
		DynamoDB Table	
		Redshift	
		EC2 Instance	
	Azure	Storage Account	
		SQL Server	
		PostgreSQL	
		PostgreSQL FlexibleServer	
		Cosmos DB	
		Virtual Machine	
	GCP	Cloud Storage Bucket	

WAF Protection

Users with CloudGuard accounts created in the Infinity Portal can see which of their assets are protected by Check Point WAF. On the **Protected Assets** page, **WAF Mode** column, the asset has one of these indications:

- **Prevent** - WAF protects the asset
- **Detect** - WAF only shows the risk
- **Disabled** - WAF is disabled

The **Detect** and **Disabled** indications do not impact the risk score of the assets.

With WAF protection, CloudGuard:

- Identifies assets protected by CloudGuard WAF and the configured WAF mode.
- Reduces the risk score if the asset has the **Prevent** indication.

  Note - CloudGuard reduces the risk score only if all public network paths to the asset pass through WAF.

- Considers the CVEs and *"Toxic Combinations" on page 92* with the security domain of **Vulnerabilities** as less severe.

CloudGuard supports SaaS deployment and Gateway deployment of WAF on AWS EC2 instances. It identifies protected AWS EC2 instances, ECS services, and Auto Scaling Groups. For these assets protected by WAF, CloudGuard can better identify the network topology visualized in the Context Graph (see *"Asset Details" on page 273*).

Note - CloudGuard WAF is not available for users with accounts created in the Dome9 portal.

One of the ERM dashboard widgets shows the impact of WAF on exposed assets.

For more information, see the CloudGuard WAF documentation.

ERM Rulesets

The **Rulesets** page shows the CSPM findings that are considered misconfigurations for risk calculation, for each platform. To calculate the risk score of your protected assets, CloudGuard uses by default all CSPM findings as misconfigurations. You can limit them to a specific ruleset of your interest to focus on selected security tasks. For this, replace the default rulesets with one designated ruleset for each cloud platform (AWS, Microsoft Azure, GCP).

CloudGuard considers CSPM exclusions, which means that it does not use excluded findings for the risk score calculation. For more details on exclusions, see *"Configuring CloudGuard Exclusions" on page 80*.

Note - Make sure all your important environments are part of a policy that includes the selected ruleset. Otherwise, ERM does not take their findings as misconfigurations for the risk score calculation and does not show misconfigurations. You can still see other information, such as CVEs and business priority.

To replace a ruleset:

1. In the CloudGuard menu, navigate to **Risk Management** > **Rulesets**.
2. Select your platform and click **Replace Ruleset** on the platform card. The list of applicable rulesets opens.
3. Select a ruleset. You can filter the rulesets by management type - CloudGuard-managed or user-managed.
4. Optionally, you can automatically create a policy for the ruleset if a different policy does not use it. Click to select this option. It is not available for the default option where all posture findings are considered in the risk score calculation.
5. Click **Save**.

Best Practice - Check Point recommends to add new environments to one of the existing policies with the selected ruleset. The policy must have a notification with enabled **Alerts Console**. If such policy does not exist, you can use the option of automatic creation of a policy when you replace the default ruleset.

Business Priority

Your input is crucial for defining the real significance of a particular asset (or a group of assets) in your business. You can add your input in **Business Priority Rules** to indicate the most important or less significant assets. When you define a number of rules with different business priority that contain the same asset, CloudGuard selects the highest applied priority for the asset.

If you do not configure explicitly a business priority value for an asset, CloudGuard sets it as **Undefined**. CloudGuard recalculates the risk score of an asset with every change of its business priority.

To add a rule for Business Priority:

1. In the CloudGuard menu, navigate to **Risk Management** > **Business Priority Rules** and click **Add Rule**.
2. Enter the rule name.
3. Select one or more parameters:
   - **Organizational Unit** - Select one or more Organizational Units (the root Organizational Unit cannot be selected).
   - **Environment** - Select one or more environments onboarded to CloudGuard.
   - **Asset Tags** - (optional) - Add as many tags with applicable values as you need, or use separately only a key or only a value.
   - **Asset name contains** - Enter one or more strings that the asset name contains. Use commas (,) to separate the strings.

   When you select multiple parameters, CloudGuard searches assets with the combination of these parameters (AND logic). Within each combination, it searches assets with one of the options selected, for example, one of the asset tags entered (OR logic).
4. Select **Business Priority** for the rule based on its importance:
   - **Crown Jewel**
   - **High Importance**
   - **Important**
   - **Minor Importance**
5. Click **Save**.

Events

CloudGuard generates alerts for findings on your cloud environments based on policies. These findings and events can be viewed in the CloudGuard portal and sent as messages to different notification targets, such as email and SNS.

CloudGuard engines show the found Events on pages below the **Events** menu. On the **All Events** page, you can find a summary table (**All**) for all security events. From this page, you can select events of a specific type (**Posture Findings**, **CIEM**, or **Vulnerabilities** tab) and drill down to learn more details about the event, add remarks for the event, or assign it to specific users for remedial actions.

You can search and filter the view for specific events of interest, based on the environment, event type, entity type, ruleset, and other parameters.

Benefits

- Enterprise view across all platforms, environments, and entities
- System messages view on a separate page
- Customizable by search or filter view for Organizational Unit, environment, platform, source, etc.
- Actionable from the table menu (acknowledge, set up a remediation or exclusion)
- Links to referenced entities in CloudGuard

Use Cases

- For enterprise security managers: high-level summary of security posture and key metrics of security findings across the organization - see *"Dashboards" on page 86*
- For security engineers:
  - High-level summary of security posture and key metrics of security findings for specific environments - see *"Dashboards" on page 86*
  - Can review security findings for the applicable environments and apply remediations - see *"Creating a remediation for findings" on page 127*

All Events

The **All Events** page aggregates CloudGuard events generated from these sources:

- Posture Management
- CDR
- Serverless Runtime Protection
- Kubernetes Admission Control
- Image Assurance of containers and registries
- Containers Runtime Protection

This page does not show usual system or account events, such as account sign-ins or configuration issues, which appear below the **Operational** section of the **Events** menu.

Configuring Events

You have to configure alerts from Compliance Engine, CDR, and other sources to appear on the Events page. For this, configure *"Notifications" on page 852* for these events. You have to do this for each policy separately, so you can control which rulesets and which environments generate events. To receive notifications from all rulesets and environments, configure this in each Policy.

In the Notification configuration window, make sure to select the **Include in the CloudGuard Events page** option.

Aggregated Events

The main (**All**) page shows a summary table for all security events in your environments. As the list can contain a large number of results, you can move through the list and view new findings that align with the search and filter criteria.

You have many options to set up the page so that you can conveniently view the applicable findings. Use the tabs to show separate pages with **Posture Findings**, **CIEM**, **Threat & Security** events, and more. CloudGuard automatically saves the changes you make to each table.

Filter and Search Area

The filters bar is at the top of the page and includes:

- Preconfigured filters
- Free text filter

- Time frame filter
- Saved filters (**Favorites**)

With the preconfigured filters, you can select to view events based on various parameters. To add a preconfigured filter, click the icon and select as many filters as necessary from the list. The selected filters appear in the Search field, and the table is updated automatically based on them.

The free text filter allows you to enter text and use it as a filter. The entered text applies immediately to the current table.

With the time frame filter, you can view the findings according to their creation time. You can select one of the preconfigured periods or click **Custom** and select a custom date range.

Action Menu

When you select a finding, the actions applicable to this finding become available. In addition, you can access the action menu from the finding sliding window. Not all actions are available for all findings.

Use the action menu for these actions:

- Create an exclusion for a finding
- Close the finding
- Acknowledge or unacknowledge a finding
- Add comments to the event
- Archive the finding
- Change the severity of the event
- Assign an event to a CloudGuard user
- Report issues related to the finding
- Immediately remediate an issue with a CloudBot
- Create remediation for a finding

For more actions and detailed information about them, see the steps below in *"Actions" on page 125*.

Findings Table

You can select one or more findings when you click the check box in their row.

Organize the table columns as necessary and adjust these parameters:

- **Visibility** - To select which columns to see in the table, click **Columns** on the right.
- **Position** - To change the column's location, click the column header and drag it to the desired location.
- **Width** - To change the column width, drag the right separator line of its header in the desired direction. To adjust the width by the longest column value, double-click the right separator.
- **Sorting** - To change between the default, ascending, or descending order of the entries, click the column header.

To restore the default settings of the table, click **Reset** in the **Columns** menu.

Group Arrangement

With the **Group By** menu, you can set up findings with the same parameters together, so they appear in the table below the same group title.

Arrange the findings by:

- Action
- Title
- Severity
- Environment, and so on

To group the findings, select a category from the **Group By** list. All findings are arranged by applicable groups.

Click the arrow on the left of each group name to expand the group and see its contents. Click the arrow again to close the group.

Finding Details

Click a finding in the table to open a sliding window with the finding details. The details can be different based on the finding type and include a different number of tabs.

The finding details contain basic information such as the finding source, environment, applied ruleset, etc. On the **Overview** tab, see where the finding was found and which rule was triggered. On the *"Entity Viewer" on page 125* tab, see the underlying finding structure.

Posture Findings

The information can include:

- **Severity** as defined in the applicable rule or use case
- **Date** of creation
- **Assignee** - a user assigned to manage the finding, for example, set a remediation

- **Title** of the applicable rule or use case
- **Ruleset** that contains the applicable rule or use case
- **Remediation** - See the actions that CloudGuard recommends
- **GSL** expression - for more information on GSL, see *"Governance Specification Language (GSL)" on page 326*

Note - The actual rule is sometimes more complex than a GSL-code representation, so CloudGuard does not show the GSL code in Rulesets > Rule.

CIEM

On the **Overview** tab, see the finding description, remediation (if available), and the analysis period.

On the **Permissions** tab, see the *"Entitlement Map" on page 387*.

For more information on CIEM findings, see *"Findings" on page 392*.

CDR

The Event Graph is a visual representation of account activity and network activity. It is supported for AWS and Azure. The location icon appears above the entity for which you opened the Event Graph. To see details about an asset, hover over its icon. To see a table that shows the occurrences of an event, hover over the **Occurrences** icon.

For an activity event, the graph shows the relationship between the **Identity**, **Issuer**, and **Target**.

For a network event, the graph shows the network topology between **Internal** and **External** networks.

Based on the asset type, permissions, and connected services, asset details can have these symbols:

Icon	Meaning	Explanation
9.6	Risk score	Read more in *"Risk Calculation" on page 101*
56	IAM sensitivity	Read more in *"Entitlement Map" on page 387*
Important	Business priority	Read more in *"Business Priority" on page 118*
Public	Network exposure	Read more in *"Network Exposure" on page 105* and *"IAM Exposure" on page 112*

Icon	Meaning	Explanation
Base Image	Base Image	Read more in *"Base Image" on page 426*
Vendor Image	Vendor Image	Read more in *"Vendor Image" on page 428*

To see a table of events for an asset, click its icon.

The table shows other events that occurred for the asset within the past 30 days. The table row of the selected event is highlighted in blue.

The logs investigation is not available if the related logs passed the retention period.

For more information, see *"Intelligence Security Events" on page 634*.

Threat & Security

For Threat and Security events, on the **Occurrences** tab, you see the details of the event occurrences.

Each time the rule discovers a finding, CloudGuard registers this finding as a separate occurrence. CloudGuard aggregates the findings if they have the same environment, entity, and event (same GSL code). The time interval to group all occurrences in the same security event is 30 minutes. Every five minutes CloudGuard checks the traffic of the previous 30 minutes and alerts if necessary. As CloudGuard does not include the occurrences that were displayed before, some occurrences of the same event can overlap after the others in time intervals.

Click **Investigate** to open the event log in the *"Traffic Explorer" on page 359* and examine the actual log information for the selected entity at the event's time frame. To have a clearer view, drag the Source and Destination headers to the grouping bar.

The logs investigation is not available if the related logs passed the retention period.

For more information, see *"Intelligence for Kubernetes Containers" on page 645*.

Vulnerabilities

On the **CVE** tab, see the list of CVEs found in the container registry images sorted by severity.

Fields for Kubernetes images:

- **Title** - The specific ID or type for which the finding is created based on the finding category.

- ImageScan findings have the title with the name of the image
- Common Vulnerabilities and Exposures (CVE) findings have the title with the CVE ID

- **Description** - The issue description, for example, the CVE description as it appears in the National Vulnerability Database (NVD).
- **Environment** - The Kubernetes cluster that contains the image with the finding.

For more information, see *"Vulnerability Findings (Image Assurance)" on page 439*.

Entity Viewer

The Entity Viewer tab contains information about the configuration of the protected asset. Use the menu buttons to customize this view.

The entity can have the N/A (not available) status when:

- The resource creation event was blocked, and it is not possible to create the entity on the environment.
- The resource creation violation was detected. It is not possible to create the resource on the environment, but the Event can appear before the protected asset update in the CloudGuard backend. It can take up to five minutes for the entity link to appear in the finding.

Events Deletion

Posture findings and security events are deleted from the Events page when they are considered resolved, that is, the rule is not violated anymore (passed). This happens when:

- Users correct or remediate the issue that triggered the event.
- Users voluntarily close the finding - see *"Closing findings" on page 128*.
- Users delete the policy (break association between the environment, rules, and notification) - see *"Policy Deletion" on page 78*.
- Users delete (offboard) the environment for which the finding is created.
- Users delete the rule that generated the finding or the ruleset that contains the rule.
- Users delete the notification to be sent when the finding is generated.

  Note - The *passed* notification is sent only to the valid (not misconfigured) integrations available at the time of the notification deletion.

Actions

You can also access the action menu from the finding sliding window.

Creating exclusion for findings

You can create an exclusion from a finding, to exclude more findings equivalent to it. You can exclude the specific ruleset, rule, and entity, or widen the exclusion to include all entities in the rule, all accounts, or all rules in the ruleset.

You cannot apply this action to more than one finding at a time.

1. Select the finding in the table.
2. From the menu, select **Exclude**.
3. In the **Create New Exclusion** window, the specific ruleset, rule, environment, and entity are selected, which excludes only this specific combination.

   Clear the choices to widen the exclusion and cover all rules, environments, or entities.
4. Add a comment to distinguish the exclusion. The comment field is mandatory.
5. Click **Save** to create the exclusion.

   The Exclusion icon appears on the **Overview** header of the Entity Card.
6. You can manage the exclusion on the **Exclusions** page in the CSPM/CIEM/CDR/Workload Protection > Admission Control or Vulnerabilities menu. See *"Configuring CloudGuard Exclusions" on page 80*.

Best Practice - If you clone a ruleset and run an assessment with it, you can see duplicate findings for the initial and cloned rulesets. Create an exclusion to hide the duplicate findings.

Acknowledging findings

You can acknowledge a finding to show it as read. This does not close the finding or indicate that it is resolved.

1. Select one or more findings in the table.
2. From the menu, select **Acknowledge**.
3. In the **Acknowledge Finding** box, optionally, add a comment with the reason for acknowledgment.

   These comments are seen by all users who can see the finding.
4. Click **Acknowledge**. The Acknowledgment icon appears on the **Overview** header of the Entity Card.

   In addition, your comment appears in the Comments section of the Entity Card, with information of the date, time, and user that made it.

Creating a remediation for findings

It is possible to create a remediation and associate it with the rules underlying findings. These remediations are applied to cloud resources to correct the issues that caused the finding. CloudGuard Cloudbots are an example of remedies.

You cannot apply this action to more than one finding at a time.

1. Select the finding in the table.
2. From the menu, select **Remediate**.
3. Complete the details for the remediation (see *"Adding Remediation" on page 319*) and then click **SAVE**.
4. The Remediation icon appears on the **Overview** header of the Entity Card.

Applying a CloudBot immediately (Fix it)

Apply a CloudBot solution to a found issue immediately, directly from the findings and events. This is applicable to AWS and Azure environments, where you already deployed one or more CloudBots.

You cannot do this action for more than one finding at a time.

1. Select a finding or event in the table.
2. From the menu, select **Fix it**.
3. In the **Remediate Now** window, select a CloudBot and enter required parameters (if applicable).
4. Optionally, to add one or more CloudBots, click **Add**.
5. Click **Execute**. The **Remediation Execution Status** window opens. This window shows the results of the bot application in the **Bot Trigger** and **Execution Status** columns.

   See the table below:

Bot Trigger	Execution Status	Description
Failed	-	The bot was not triggered because of an endpoint problem.

Bot Trigger	Execution Status	Description
Success	Failed	You triggered the bot successfully but its execution failed. Click to open the *"System Audit Logs" on page 134* and see e the audit log entry of the bot event.
Success	Success	You triggered the bot successfully, and its execution was successful. Click to open the *"System Audit Logs" on page 134* and see the audit log entry of the bot event.

Adding comments

You can add a text comment to the finding. Users who can view the finding on the Events page, based on their permissions, can see the comment.

1. Select one or more findings in the table.
2. From the menu, select **Comment**.
3. Enter text in the Comment field.
4. Click **Add** to save the comment.

You can repeat the steps to add more than one comment to a finding.

Closing findings

You can close a finding found by a source, which is not Compliance, external findings, or Qualys source. This action deletes the alert from the elastic search. You cannot recover the deleted alert in the future.

1. Select one or more findings in the table.
2. From the menu, select **Close Alert**.
3. Click **Close Alert** to confirm the action.

Archiving findings

You can change the view and see only applicable and important findings if you keep less applicable findings in an archive. Archived findings are not seen in the findings page, and you cannot apply an action to them. Resolved findings are deleted even if they are archived.

To archive findings:

1. Select one or more findings in the table.
2. From the menu, select **Archive**. CloudGuard moves the selected findings to the Archive View.
3. Toggle the **Archive View** slider from OFF to ON to see the archived findings.

To change archived findings:

1. Toggle the **Archive View** slider from OFF to ON to see the archived findings.
2. Select one or more findings that are necessary to change.
3. Click **Unarchive** to restore the findings in the primary findings view.
4. Click **Close** to permanently close the findings.

Reporting issues

If you do not agree with one of the alert parameters or think it is erroneous, you can report the alert and provide a reason.

1. Select one or more findings in the table.
2. Click the three-dots menu and select **Report an issue** to report an alert.
3. Select a reason from the list:
    a. False positive
    b. Wrong severity
    c. Remediation issue
    d. Incorrect information
    e. Other
4. Optionally, you can add a comment to provide more details.
5. Click **Report**.

Changing severity

You can set the severity level for the finding from the list:

- High
- Medium
- Low

- Critical
- Informational

For more information on severity levels, see *"Severity Levels" on page 309*.

Users who can view the finding on the Events page, based on their permissions, can see this attribute. It is useful for filtering the list of findings.

1. Select one or more findings in the table.
2. Click the three-dots menu and select **Change Severity**.
3. Select the new Severity of the finding from the list. Initially, it is the severity of the rule that found it.
4. Click **Save**.

Assigning findings

You can assign the finding to a CloudGuard user to take more steps, such as remedial actions.

Users who can view the finding on the Events page, based on their permissions, can see this attribute to filter the list of findings.

To assign findings:

1. Select one or more findings in the table.
2. Click the three-dots menu and select **Assign**.
3. Select a user email address from the **Assign user** list. Possible assignees are all users of the CloudGuard account.
4. Click **Save**.

To remove findings assignment:

1. Select one or more findings in the table.
2. From the menu, select **Assign**.
3. From the **Assign user** list, select **Unassigned**. This removes the assignment from all users it was assigned to.
4. Click **Save**.

Exporting findings

You can export selected findings to a CSV file. You can export all findings or those shown in a filtered view on the Events page.

In the findings table, select a time interval and filter the view to show the necessary findings to export (or skip this step to show all findings for the time interval).

1. Click **Export** in the top right.
2. Select how to receive the report file - download directly from the CloudGuard portal or from an email message.
   a. Click **CSV Report - Download** to save the file on your computer.

      Note - You cannot download the file if the table contains more than 10,000 entries. Apply a filter to decrease the number of entries.

   b. Click **CSV Report - Email** to receive the download link by email. Enter your email address and follow the instructions in the message received.

Use the filter action buttons on the filter bar to save or clear your search criteria.

Saving filters

You can save all your currently applied filters. The saved filters are stored in two groups: public or private.

1. In the filters area, click **Saved Filters**.
2. Enter the filter name.
3. Select **Public** for other users to see the filter. If not, the filter is private and only seen by you.
4. Click **Save**.
5. To use the filter, click **Saved Filters** and select it from the list.

Clearing filters

To clear all applied filters, click **Clear All** in the filters area.

Known Limitations

The **Fix it** option is not applicable to GCP environments.

More Links

- *"API Audit Logs" on page 132*
- *"System Audit Logs" on page 134*

API Audit Logs

API Audit Logs record all API actions done by users through the user interface or API on a specific CloudGuard account. To illustrate, an action is when a user creates a new CloudGuard user, updates the Security Groups, creates a new rule, renames a ruleset, and more. The logs show who did the action, the action's status, and full details of the action.

CloudGuard keeps your audit logs for three years.

To see the API audit logs:

1. Navigate to **Events** > **Operational** > **API Audit Logs**. The API Audit Log table opens and shows these columns:
   - **Username** - The user that did the action.
   - **Event Name** - For example, "Fetch a specific AWS cloud account".
   - **Time** - The date and time the event occurred.
   - **HTTP Method** - An API request sent to a resource, such as `Get` or `Post`.
   - **HTTP Response** - The server's response to the API request, as in "200" or "204".
   - **Request URL** - The URL to which the API request was sent.
2. In the table, below the **Username** column, select a log to see its details.

   In addition to the log's username and event time, the log's details show this important information:
   - **Body** - The response body (which includes the key and value) to the API request.

     Note - Not all requests include a response body.
   - **Parameters** - The parameters used in the API request.
   - **Client IP** - The API's client ID.

To filter the API audit logs:

1. Navigate to **Events** > **Operational** > **API Audit Logs**.
2. Click the down arrow on the time filter and select a time or a **Custom** time.
3. To use GSL, select one of these options:

- Enter a GSL query or click **GSL** and then click  to open the GSL editor and select **Builder** or **Free Text**. For more information about GSL, see *"Governance Specification Language (GSL)" on page 326*.
- Click **Add Filter** and select a filter option such as **HTTP Method** or **Event Name**.

Sample GSL search:

```
auditapi where http.method='POST'
```

Returns: All API requests that used the 'POST' HTTP method.

4. Click **OK**.
5. To run the query, click **Run**.

To export the logs:

1. Navigate to **Events** > **Operational** > **API Audit Logs**.
2. Below the logs table, click  .

System Audit Logs

CloudGuard keeps a full audit log of all accesses to your environments and of each action done on the account. This page shows a record of all actions taken by the system. For example, you opened a lease for two hours. After two hours, CloudGuard closed the lease.

CloudGuard keeps your system audit logs for three years.

To see the system audit logs:

1. Navigate to **Events** > **Operational** > **System Audit Logs**.
2. In the table, below the **Cloud Account ID** column, select an account ID to see its details. The Details window opens.

   In addition to the event name and time, the log's details show this important information:

   - **Cloud Account ID** - The cloud account ID on which the action was done.
   - **Event Name** - The type of event on the cloud account, such as "API key created" or "Security group change detected".
   - **Time** - The date and time the event occurred.
   - **Description** - Details of the system audit such as changed tags, owner, and compliance.

To export the system audit logs:

1. Navigate to **Events** > **Operational** > **System Audit Logs**.
2. Below the logs table, click  .

Note - The export table is limited to 10,000 events.

System Events

Events

You can configure CloudGuard to send audit log messages to an AWS SNS topic. The table below lists the messages.

Message-EventType	Audit parent type	Audit child type	Description
AccountLicenseUpdatedEvent	CloudGuard account	Account license updated	The licensing plan was updated.
CrossAccountIdentifierCreatedEvent	CloudGuard account	Cross account identifier was generated	A cross-account identifier was generated for the account (for MSP)
AlertTriggeredEvent	Alerts events	Alert triggered	Alert was triggered on a security group
AlertClosedEvent	Alerts events	Alert resolved	Alert was resolved on a security group
AlertUpdatedEvent	Alerts events	Alert updated	Alert content was updated on a security group
InvalidAwsCredentialsEvent	Cloud account	Invalid cloud credentials	The cloud account has invalid credentials
AwsCredentialsValidatedEvent	Cloud account	Cloud credentials validated	The cloud account that had invalid credentials is now valid

Message-EventType	Audit parent type	Audit child type	Description
CloudSecGroupTamperDetectedEvent	Cloud security groups	Security group tamper detected and handled	A change was detected on a fully protected security group and it was reverted
CloudSecGroupChangesDetectedEvent	Cloud security groups	Security group change detected	A change was detected on a read-only security group
CloudSecGroupImportedEvent	Cloud security groups	Security group imported	Security group was imported from your cloud account
AwsLeaseEndedEvent	Cloud access leases	Access lease ended	An access lease was ended when the period finished
LeaseTerminatedEvent	Cloud access leases	Access lease terminated	An access lease was terminated manually by the user
ServerStateChangedEvent	CloudGuard Agents	Agent state changed	Agent state changed from state to state
AwsLeaseEndedEvent	CloudGuard access leases	Access lease ended	An access lease was ended when the period finished

Message-EventType	Audit parent type	Audit child type	Description
LeaseTerminatedEvent	CloudGuard access leases	Access lease terminated	An access lease was terminated manually by the user
ApiKeyCreatedEvent	Users management	API Key created	API key was created for a user
SSOUserLogOnFailureEvent	Users	SSO login failed	SSO login failed by a user
UserRoleCreatedEvent	User role event	User role created	The new role was created
UserRoleUpdatedEvent	User role event	User role updated	Role permissions were updated
AzureCloudAccountAddEvent	Azure Cloud Account	Azure Cloud Account created	New Azure cloud account was added to CloudGuard Console
AzureSecurityGroupImportedEvent	D9 Azure security group event	Azure network security group imported	New Azure security group imported
AzureSecurityGroupUpdatedEvent	D9 Azure security group event	Azure network security group change detected	Change detected on network security group

Assets

This section describes how to manage cloud assets in CloudGuard. It explains how to:

- See assets that CloudGuard protects
- Manage other custom resources, such as text and IP lists

To add (onboard) environments to CloudGuard, see *"Onboarding Cloud Environments" on page 53* in the *"Getting Started with CloudGuard" on page 50* section.

Environments

The Environments page shows your CloudGuard-managed cloud accounts and Kubernetes clusters.

If CloudGuard fully manages your environments, you can set the protection of your Security Groups from here.

In the Environments section, you can see all of your environments, on all platforms, in a single pane of glass. For managed accounts, you can configure and apply changes centrally to all these environments in one area.

Use Cases

Here are some typical use cases to illustrate the control of Environments from one central location.

- Search for environments - To quickly search for specific environments across all your cloud presence, see *"Filter and Search" on page 863*.
- Review security posture - To assess your security posture effectively and review all your security groups protection state in one view, see *"Security Groups" on page 366*.
- Apply equal changes - To expand your cloud presence, you can change the security policies for all regions from one portal, see *"Cloud Security Posture Management (CSPM)" on page 302*.
- Respond to environment permissions behavior - To receive a notification about changes to one of your environments and then take corrective steps, see *"Notifications" on page 852*.

Actions

Viewing Your Environments

The primary page shows a list of all your environments, on all cloud providers.

Filtering the List of Environments

To filter the list of environments, use the Filter and Search bar at the top of the page. As filter criteria, use Platform, IAM Safety status, Intelligence status, the number of assets, or other available parameters.

See *"Filter and Search" on page 863*.

Adding an Environment

In the Environments page, add (onboard) environments to CloudGuard, for cloud platforms. This adds the accounts to the CloudGuard Console. You do not create accounts on the cloud provider here (as an alternative, use the cloud provider site). When you add an environment to CloudGuard, you can select to manage it from CloudGuard (Full Protection) or monitor it (Read-Only).

1. Navigate to **Assets** > **Environments**. This shows a list of the environments added to CloudGuard.

   Note - For first-time onboarding, select a cloud platform from the Environment's primary page and follow the onboarding steps. For more information, see *"Onboarding Cloud Environments" on page 53*.

2. Click **Add** and select the cloud platform.

3. Follow the instructions to onboard an environment to CloudGuard, for the selected cloud platform. For more details, see *"Onboarding Cloud Environments" on page 53*.

Viewing Environment Details

You can view details for an environment.

From the primary page, click an environment link to show more details. The details are organized by region (based on the cloud provider regions).

They show general information for the account, with the environment number, the date of adding to CloudGuard, the number of instances, and security groups. The information varies depending on the cloud platform.

Viewing AWP Details

The **AWP** tab shows AWP details if the environment is onboarded to AWP. For more information, see *"Viewing AWP Details" on page 491*.

Updating Environment Permissions

If the environment policies have missing permissions to allow CloudGuard to see or manage your environment, the warning message appears: **Missing 10 permissions for CloudGuard-Connect**.

These permissions relate to the CloudGuard-Connect policy (an AWS policy, which enables CloudGuard to connect and manage your AWS accounts).

1. Click **Show more** to see the missing permissions. The list shows the cloud resources that are missing each permission (CustomDomainName, for example), the permission type (tags), and the action for the resource that you must add (ListTags). In addition, it shows the number of resources missing this permission (# Affected Entities). Click **Show Entities** in the last column to see the specific resources.
2. Click **Validate Permissions** to add the missing permissions to your account.
3. To verify that the policies are updated for your AWS accounts, see *"Updating AWS Permissions" on page 281*.

   Note - CloudGuard cannot fetch updated data for entities that have missing permissions.

Renaming an Environment

You can change the name of an environment. This changes the name as it appears on the CloudGuard portal, but not on the cloud provider.

1. Enter the environment.
2. Select **Rename** from the top right menu.
3. Make your changes.
4. Click **Save** to save the changes (or close to cancel the changes).

Editing Role Credentials

It is possible to change the AWS IAM Role for an environment. The role must exist in your AWS account.

1. Click on an account from the list of accounts on the primary **Environments** page.
2. Select **Edit Credentials** from the top right menu.
3. In the AWS console, open your AWS account and navigate to the IAM page. Select **Roles** and copy the ARN for the role to be applied to the account in CloudGuard. See AWS IAM Roles.
4. Enter (or paste) the ARN value in the **Role ARN** field.
5. Click **Confirm**.

Removing an Environment

Click **Remove** to delete the selected environment from CloudGuard. This does not delete the environment or its resources on the cloud provider.

Selecting the default Protection Mode

You can select the Protection Mode that CloudGuard applies to new security groups detected in AWS environments. CloudGuard defines and applies Security Groups in AWS for each region separately.

You can select from these options:

- **Read-Only** - CloudGuard includes new Security Groups in Read-Only mode, without changes to the rules
- **Full Protection** - CloudGuard includes new Security Groups in Full Protection mode, without changes to the rules
- **Region Lock** - CloudGuard includes new Security Groups in Full Protection mode and clears all inbound and outbound rules

You can set or change the Protection Mode for existing Security Groups, in all regions, for all of your AWS accounts.

To set or modify the Protection Mode:

1. Navigate to **Assets** > **Environments** and select an environment from the list. The **Network** tab shows the regions for the environment and the number of Security Groups defined for each region.
2. Click one of the regions. This shows a list of the Security Groups defined for the region.
3. Select the Protection mode to apply by default to new Security Groups in the region.
4. Select the Protection mode for each of the existing Security Groups in the region. Click **select entire region** to apply the mode to all Security Groups in the region.

   Note - The account must have a *CloudGuard-write-policy* to apply Full Protection to a Security Group (see *"Setting an AWS Security Group to Full Protection" on page 369*).

5. Click **Save**.

Identifying Identical Entity Roles

An organization with a large cloud environment may come close the cloud platform's maximum number of roles. The Identical Entity Roles feature identifies roles in your cloud environment with duplicate permissions. This feature is supported for AWS and Azure.

Use Case

Use the **Identical Identities** window to identify roles in your cloud environment that have duplicate permissions. Use this information to combine roles in your cloud environment safely and efficiently.

1. Click on an account in the list of accounts on the main **Environments** page.
2. In the upper right, above the table, click the three dots menu and select **Show Identical Identities**.

   The **Identical Identities** window opens. Each **Identities Group** is a group of identities that have the same permissions.
3. Expand an Identities Group to see the identities in the group.
4. In the **Permissions** column, to the right of the Identities Group, click the button to see the permissions shared by all members of the group.
5. In your AWS or Azure environment, combine identities that have the same permissions.

Onboarding AWS Environments

This topic describes all available methods of onboarding AWS environments to CloudGuard.

You can select one of the methods below depending on the type and number of your environments:

- To onboard the entire Organization - *"Onboarding of AWS Organizations" on page 146*
- To onboard one AWS account
  - Unified (all available features) - *"Unified Onboarding of AWS Environments" on page 54*
  - Manually
    - Standard account - *"Manual Onboarding of AWS Environments" on page 162*
    - GovCloud or AWS China account - *"Manual Onboarding of AWS GovCloud or AWS China Environments" on page 165*
  - Using Terraform - *"Onboarding with Terraform" below*

Onboarding with Terraform

You can use the Check Point CloudGuard (Dome9) Terraform provider to onboard and update AWS environments in CloudGuard. First, you need to prepare Terraform files for your AWS environments. For more information, see the Terraform documentation at https://www.terraform.io/docs/providers/dome9/.

The dedicated resource at https://registry.terraform.io/providers/dome9/dome9/latest/docs/resources/aws_unified_onbording includes Intelligence and Serverless configurations, and the rulesets for Posture Management (Compliance) and Intelligence.

The full source code is at https://github.com/dome9/terraform-provider-dome9.

CloudGuard Features

Learn more about each functionality that CloudGuard provides:

- For Posture Management, see *"Cloud Security Posture Management (CSPM)" on page 302*
- For Intelligence, see *"Cloud Detection and Response (CDR)" on page 565*
- For operational modes, see *"AWS Security Group Management Modes: Full Protection or Read-Only" on page 371*
- For Serverless Protection, see *"Enabling Serverless Protection" on page 533*

- For Agentless Workload Posture, see *"Agentless Workload Posture" on page 489*
- For Permissions updates, see *"Updating AWS Permissions" on page 281*

Troubleshooting

Intelligence

Issue: CloudGuard cannot onboard your AWS account to Intelligence during the environment onboarding. The corresponding status and error message appear on the Onboarding Summary page.

Possible causes:

- CloudGuard cannot find CloudTrail logs on your account
- CloudGuard cannot find an applicable log destination, because your S3 bucket already has a configured Event Subscription

  Note - The preferred type of CloudTrail is a trail that applies to all Regions. If CloudGuard finds that your AWS account contains multiple globally applied trails, it selects one on a random basis. A warning message on the Onboarding Summary page notifies you that other buckets were found but not onboarded.

Solution: Onboard your AWS account to Intelligence separately, with Custom Onboarding. See *"Custom Onboarding" on page 595* for more information.

See also *"Troubleshooting" on page 585*.

Onboarding of AWS Organizations

This topic describes how to onboard an AWS organization automatically. For other onboarding methods, see *"Onboarding AWS Environments" on page 144*.

Prerequisites

Before onboarding your AWS Organization, make sure:

- You have Administrator permissions to create and manage resources in this Organization.

How it Works

After you onboard an AWS Organization to CloudGuard, every new AWS account added to the Organization is automatically onboarded to CloudGuard.

In general, after onboarding with the Unified procedure (*"Unified Onboarding of AWS Environments" on page 54*), the environment has a set of CloudGuard features (configuration) defined by the CloudFormation Template. Similarly, when you onboard an Organization, its configuration is defined by a CloudFormation Template, which is used as a blueprint. Nonetheless, an AWS account added (onboarded) to the Organization acquires the configuration defined in the CFT and **not** the configuration currently existing in the Organization. To learn how you can onboard organizations with different configurations, see examples in *"Updating Onboarded AWS Organizations" on page 155*.

Onboarding

STEP 1 - Welcome

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. For first-time onboarding, click **Amazon Web Services**.

   Or, if you already onboarded environment(s), from the top menu, select **Add** > **AWS Environment**.
3. On the **Welcome** page, select to onboard an **Organization**.
4. In the **Management Account ID** field, enter the ID of your AWS Management Account.
5. Click **Next**.

STEP 2 - Management Account Stack

On the second wizard page, you create a management account stack. Follow the on-screen instructions.

1. Click **CloudFormation Template** to review all resources for CloudGuard to deploy on your management account. Optionally, you can click **Download CFT** to save the resource file on your local drive.
2. Open a new browser tab, go to the AWS portal, and sign in to your AWS account.
3. In the CloudGuard onboarding wizard, click **Launch Stack**.

   A new browser tab opens with the CloudFormation stack. CloudGuard automatically enters all required default parameters. Change the default AWP, Serverless, and CDR parameters as you need.

   ⚠ **Caution** - If the management account is already onboarded to CloudGuard through unified onboarding, do not change the existing AWP, Serverless, and CDR parameters in the stack, otherwise the Organization onboarding fails. To change the parameters, see *"Updating Onboarded AWS Organizations" on page 155*.
4. Below **Capabilities**, read the explanation and select the **I acknowledge...** option to accept. Click **Create stack**.
5. AWS begins to create the stack. Wait until it creates the IAM role (*CrossAccountRole*) and the stack status becomes **Create_Complete**.
6. Enter the stack details in the CloudGuard wizard:
   - Management account external ID (entered automatically)
   - ARN of *CrossAccountRole*
   - Organization name (optional)
7. Click **Next**.

STEP 3 - Create Stackset

On the third wizard page, you create a stackset for member accounts.

In AWS console:

1. Open **CloudFormation**, navigate to **StackSets** and click **Create StackSet**. Follow these five steps below:

a. **Choose a template**

   i. In the **Permissions** section, select **Service-managed permissions**.

   ii. In the **Prerequisite – Prepare template** section, select **Template is ready**.

   iii. In the **Specify template** section:

   - Select **Upload a template file** and upload the attached CloudFormation Template.
   - Or select **Amazon S3 URL** and paste the S3 address from the CloudGuard wizard to Amazon S3 URL.

   iv. Click **Next**.

b. **Specify StackSet details**

   i. Enter *CloudGuardOnboarding* for the StackSet name.

   ii. Optionally, enter a **StackSet description**.

iii. In the **Parameters** section, select these values:

- In **AwpMode**:
  - Select **Disabled** to disable AWP for all accounts in the organization.
  - Select **InAccount** (default) to enable AWP scanning within your account.
  - Select **Saas** to enable AWP scanning of your snapshots on CloudGuard's account.

  For more details, see *"Agentless Workload Posture" on page 489*.

- In **CDR**:
  - Select **Enabled** to onboard the organization to CDR. You can select to onboard up to three S3 buckets. For this, enter the values:
    - `CloudAccountId1`, `CloudAccountId2`, `CloudAccountId3` (mandatory)
    - `KmsDecryptArn1`, `KmsDecryptArn2`, `KmsDecryptArn3` (optional)
    - `S3BucketArn1`, `S3BucketArn2`, `S3BucketArn3` (mandatory). If the management account was previously onboarded to CDR, provide the same bucket that it used before.
    - `SnsTopicArn1`, `SnsTopicArn2`, `SnsTopicArn3` (optional)

    Note - For the management account previously onboarded to CSPM and to CDR, you need to enable CDR manually to have it for the entire Organization. Follow the link at the end of the onboarding process.
  - Select **Disabled** to skip onboarding the organization to CDR.

  For more details, see *"Onboarding AWS Environments to Intelligence" on page 572*.

  For troubleshooting, see *"Troubleshooting" on page 585*.

- In **Serverless**:
  - Select **Enabled** (default) to enable Serverless Runtime Protection.
  - Select **Disabled** to disable Serverless Runtime Protection.

  For more details, see *"AWS Serverless Function Runtime Protection" on page 530*.
- In **ExternalId** - Copy the value from the CloudGuard wizard.
- In **UseAwsReadOnlyPolicy** - Set **Disable** if you prefer not to grant redundant permissions.

  Note - By default, using the **ReadOnlyAccess** policy is enabled, which allows you to receive permissions update requests less frequently. You can manually disable the policy at this stage in the **UseAwsReadOnlyPolicy** field if you prefer not to grant redundant permissions. For more information about policies, see *"Policies" on page 279*.

iv. Click **Next**.

c. **Configure StackSet options**

Click **Next**.

d. **Set deployment options**

i. In the **Add stacks to stack set** section, select **Deploy new stacks**.

ii. In the **Deployment targets** section, select one of two options.

If you select **Organizational units (OU)**, enter the OU ID in the **AWS OU ID**. To see this value, go to the organization page in **AWS Organizations** > **AWS accounts** and copy the **ID** value from the **Organizational unit details** section.

iii. In the **Auto-deployment options** section:

- For **Automatic deployment**, select **Enable**.
- For **Account removal behavior**, select **Delete stacks**.

iv. In the **Specify regions** section, select one region that matches your CloudGuard Data Center. For more information, see *"Region Selection" on the next page*.

Important - Do not select **Add all regions** for this option.

v. In the **Deployment options** section:

- For **Maximum concurrent accounts - optional**, select **Percentage** and enter **100**.
- For **Failure tolerance - optional**, select **Percentage** and enter **100**.
- For **Region concurrency**, select **Sequential**.

vi. Click **Next**.

e. **Review**

i. Review all the details.

ii. Select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.

iii. Click **Submit**.

2. AWS redirects you to the stackset page. On the **StackSet info** tab, find the **StackSet ARN**.

In the CloudGuard wizard:

1. Paste the StackSet ARN in the wizard's field.
2. Click **Next**.

STEP 4 - Onboard Organization

On this page, you see the onboarding summary.

To onboard the management account, start the unified onboarding process. For more information, see *"Unified Onboarding of AWS Environments" on page 54*.

Your organization is onboarded to CloudGuard. When the process is done, CloudGuard redirects you to the **Environments** page that lists your onboarded environments. The new organization appears on the **Assets** > **Organizational Units** page under the root OU as its child, like the manually created CloudGuard OUs. All actions available for regular OUs (creating sub-OU, renaming, moving, and deletion) are available for the onboarded AWS organization.

Note - For Organization onboarding configured with CloudGuard Organizational Unit where the CloudGuard OU was deleted, the newly discovered cloud account is onboarded under the root CloudGuard Organizational Unit.

You can change some of the configured parameters after the Organization onboarding is completed. For more information, see *"Updating Onboarded AWS Organizations" on page 155*.

Region Selection

Region selection is relevant for organizations onboarded with AWP or Serverless Runtime Protection and not with CSPM only.

When you onboard an organization with enabled AWP or Serverless Runtime Protection, make sure to specify the AWS region that matches the Data Center of your CloudGuard account (appears in **Settings** > **Account** > **Account Info** > **Data Center**).

See available CloudGuard Data Centers and their corresponding AWS regions in the table below.

Data Center	Region
United States	us-east-1
Ireland	eu-west-1
India	ap-south-1
Singapore (for Dome9 accounts only)	ap-southeast-1
Australia	ap-southeast-2
Canada	ca-central-1

Onboarding with API

To onboard AWS organizations with API, make these API calls and changes in your AWS account:

1. Make the **first call**: GET - https://api.dome9.com/v2/aws-organization-management-onboarding/management-stack (Link).
2. In AWS, create a management stack with the `managementCftUrl` field obtained from the response.
3. Make the **second call**: GET - https://api.dome9.com/v2/aws-organization-management-onboarding/member-account-configuration (Link).
4. In AWS, create a stackset with the `content` from the second API call.
5. Make the **third call**: POST - https://api.dome9.com/v2/aws-organization-management (Link).

   Data:

   - secret - Use the `externalId` field from the first API call.
   - roleArn – Take from the management stack outputs.
6. Make the **forth call**: PUT - https://api.dome9.com/v2/aws-organization-management/{id}/stackset-arn (Link)
   - Id - Use the ID that returns from the second call.
   - stackSetArn - Use the created stackset ARN from AWS.

Wait about one hour until all your AWS accounts are onboarded to CloudGuard.

More Links

- *"Onboarding AWS Environments" on page 144*
- *"Updating Onboarded AWS Organizations" on page 155*

Updating Onboarded AWS Organizations

Sometimes, it is necessary to change some of the configured parameters of the onboarded Organizations. These parameters are stored in the stackset of your AWS account. You can change the parameters only from the AWS portal.

Note - The procedures below show the minimal set of parameters that is necessary to configure on the AWS portal to change your CloudGuard Organization. You can configure and change more AWS parameters based on your needs.

Use Cases

- Managing the deployment scope (adding or removing OU)
- Managing the deployment parameters (for example: enabling or disabling AWP)

Managing the Organization Onboarding Scope

CloudGuard organization onboarding uses AWS CloudFormation stackset to manage the scope that defines which account to onboard to CloudGuard.

Each AWS account with a successful stack instance starts an automatic onboarding attempt. Hence, to add (or remove) accounts to CloudGuard, it is necessary to add (or remove) the accounts to the current stackset scope.

Adding Organizational Units to StackSet

The changes made with this procedure are applicable to all accounts that you add to the stackset in the future.

1. In a new browser tab, go to the AWS portal and sign in to your AWS account.
2. In **CloudFormation**, navigate to **StackSets** and select the stackset used for onboarding the Organization.
3. Click **Actions** and select **Add stacks to StackSet**.
4. On the **Set deployment options** page, select these options:
    a. **Deploy new stacks** (default).
    b. For **Deployment targets**, select **Deploy to organizational units (OUs)**.
    c. Optional - to add one or more new OUs, enter the OU ID below **AWS OU ID**. For this:

i. Open a new browser tab and open **AWS Organizations**.

ii. In the **AWS accounts** page, below **Organizations**, in the Organizational structure, find the organization that is necessary to add and copy its ID.

iii. Click **Add another OU** if needed.

d. In the **Specify regions** section, select one region that matches your CloudGuard Data Center. For more information, see *"Region Selection" on page 161*.

Important - Do **not** select **Add all regions** for this option.

e. In the **Deployment options** section:

- For **Maximum concurrent accounts – optional**, select **Percentage** and enter **100**.
- For **Failure tolerance - optional**, select **Percentage** and enter **100**.
- For **Region concurrency**, select **Sequential**.

f. Click **Next**.

5. Optionally, on the **Specify overrides** page, in the **Parameters** section, change these values:

- In **AwpMode**:
  - Select **Disabled** to disable AWP for all accounts in the organization.
  - Select **InAccount** (default) to enable AWP scanning within your account.
  - Select **Saas** to enable AWP scanning of your snapshots on CloudGuard's account.

  For more details, see *"Agentless Workload Posture" on page 489*.

- In **CDR**:
  - Select **Enabled** to onboard the organization to CDR. You can select to onboard up to three S3 buckets. For this, enter the values:
    - CloudAccountId1, CloudAccountId2, CloudAccountId3 (mandatory)
    - KmsDecryptArn1, KmsDecryptArn2, KmsDecryptArn3 (optional)
    - S3BucketArn1, S3BucketArn2, S3BucketArn3 (mandatory)
    - SnsTopicArn1, SnsTopicArn2, SnsTopicArn3 (optional)
  - Select **Disabled** to skip onboarding the organization to CDR

  For more details, see *"Onboarding AWS Environments to Intelligence" on page 572*.
- In **Serverless**:
  - Select **Enabled** (default) to enable Serverless Runtime Protection.
  - Select **Disabled** to disable Serverless Runtime Protection.

  For more details, see *"AWS Serverless Function Runtime Protection" on page 530*.
- In **ExternalId** - Copy the value from the CloudGuard wizard.
- In **UseAwsReadOnlyPolicy** - Set **Disable** if you prefer not to grant redundant permissions.

The updated values of AWP and Serverless Runtime Protection apply only to the newly onboarded accounts in the organization. Already onboarded accounts stay with their initial settings.

Caution - Do not change the values of the **RoleName** and **ExternalId** parameters.

6. Click **Next**.
7. On the **Review** page, examine all the parameters and click **Submit**.
8. Go to the **Stack instances** tab to examine the added stack.
9. In the CloudGuard portal, navigate to **Assets** > **Environments**. After approximately 15 minutes, this page shows the updated scope with new accounts.

Removing Organizational Units from StackSet

CloudGuard does not try to onboard again an account that was removed from the scope.

1. In the AWS portal, sign in to your AWS account.
2. In **CloudFormation**, navigate to **StackSets** and select the stackset used for onboarding the Organization.
3. Click **Actions** and select **Delete stacks from StackSet**.
4. On the **Set deployment options** page, select these options:
    a. Enter the OU ID below **AWS OU ID**. For this:
        i. Open a new browser tab and open **AWS Organizations**.
        ii. In the **AWS accounts** page, below **Organizations**, in the Organizational structure, find the organization that is necessary to remove and copy its ID.
        iii. Click **Add another OU** if needed.
    b. In the **Specify regions** section, select the suggested region.
    c. In the **Deployment options** section:
        - For **Maximum concurrent accounts - optional**, select **Percentage** and enter **100**.
        - For **Failure tolerance - optional**, select **Percentage** and enter **100**.
        - For **Region concurrency**, select **Sequential**.
    d. Click **Next**.
5. On the **Review** page, examine all the parameters and click **Submit**.
6. Go to the **Stack instances** tab to make sure that the stack is deleted.
7. In the CloudGuard portal, navigate to **Assets** > **Environments**.
8. To remove the accounts from CloudGuard, see *"Removing an Environment" on page 141*.

Changing Onboarded AWS Organizations

Managing Automatic Deployment

With automatic deployment, when an account is added to an OU, the stackset automatically deploys more stack instances to this account. When you remove an account from the OU, the stackset automatically deletes stack instances in this account.

1. In the AWS portal, sign in to your AWS account.
2. In **CloudFormation**, navigate to **StackSets** and select the stackset used for onboarding the Organization.
3. Click **Actions** and select **Edit automatic deployment**.

4. In the Edit automatic deployment window, select **Activated** to enable it or **Deactivated** to disable.
5. Select one of the options for account removal behavior (**Delete stacks** or **Retain stacks**).
6. Click **Save**.

Changing Deployment Parameters for the Entire Organization

When you update an onboarded AWS Organization, it is necessary to change parameters in the Organization onboarding stackset. This update usually means that you enable, disable, or change mode for one of the CloudGuard features.

1. In the AWS portal, sign in to your AWS account.
2. In **CloudFormation**, navigate to **StackSets** and select the stackset used for onboarding the Organization.
3. Click **Actions** and select **Edit StackSet details**.
4. On the **Choose a template** page, click **Next**.
5. On the **Specify Stackset details** page, in the **Parameters** section, you can change the deployment parameters.

   Caution - Do not change the values of the **RoleName** and **ExternalId** parameters.

6. Click **Next**.
7. On the **Configure StackSet options** page, click **Next**.
8. On the **Set deployment options** page, select these options:
   a. For **Deployment targets**, select **Deploy to organizational units (OUs)**.
   b. Enter the onboarded OU ID below **AWS OU ID**.
   c. In the **Specify regions** section, select the suggested region.
   d. In the **Deployment options** section:
      - For **Maximum concurrent accounts - optional**, select **Percentage** and enter **100**.
      - For **Failure tolerance - optional**, select **Percentage** and enter **100**.
      - For **Region concurrency**, select **Sequential**.
   e. Click **Next**.
9. On the **Review** page, examine all the parameters and click **Submit**.

The changes apply in CloudGuard after 15 minutes.

Changing Deployment Parameters for Specific OUs

When you update an onboarded AWS Organization, it is necessary to change parameters in the Organization onboarding stackset. This change overrides parameters on specific OUs only, while other OUs stay with initially configured parameters. This does not affect OUs to be onboarded in the future.

1. In the AWS portal, sign in to your AWS account.
2. In **CloudFormation**, navigate to **StackSets** and select the stackset used for onboarding the Organization.
3. Click **Actions** and select **Override StackSet details**.
4. On the **Set deployment options** page, select these options:
    a. **Deploy new stacks** (default).
    b. For **Deployment targets**, select **Deploy to organizational units (OUs)**.
    c. Optional - to add one or more new OUs, enter the OU ID below **AWS OU ID**. For this:
        i. Open a new browser tab and open **AWS Organizations**.
        ii. In the **AWS accounts** page, below **Organizations**, in the Organizational structure, find the organization that is necessary to add and copy its ID.
        iii. Click **Add another OU** if needed.
    d. In the **Specify regions** section, select the suggested region.
    e. In the **Deployment options** section:
        - For **Maximum concurrent accounts – optional**, select **Percentage** and enter **100**.
        - For **Failure tolerance - optional**, select **Percentage** and enter **100**.
        - For **Region concurrency**, select **Sequential**.
    f. Click **Next**.
5. On the **Specify overrides** page, in the **Parameters** section, change the deployment parameters.

    Caution - Do not change the values of the **RoleName** and **ExternalId** parameters.

6. On the **Review** page, examine all the parameters and click **Submit**.

The changes apply in CloudGuard in 15 minutes.

Region Selection

Region selection is relevant for organizations onboarded with AWP or Serverless Runtime Protection and not with CSPM only.

When you onboard an organization with enabled AWP or Serverless Runtime Protection, make sure to specify the AWS region that matches the Data Center of your CloudGuard account (appears in **Settings** > **Account** > **Account Info** > **Data Center**).

See available CloudGuard Data Centers and their corresponding AWS regions in the table below.

Data Center	Region
United States	us-east-1
Ireland	eu-west-1
India	ap-south-1
Singapore (for Dome9 accounts only)	ap-southeast-1
Australia	ap-southeast-2
Canada	ca-central-1

Manual Onboarding of AWS Environments

This topic describes how to onboard an AWS environment with the legacy onboarding experience. For the new experience, with the Unified onboarding procedure, see *"Unified Onboarding of AWS Environments" on page 54*.

The onboarding procedure adds all regions, Security Groups, and assets in the AWS account to the CloudGuard portal. It enables you to manage the AWS Security Groups from CloudGuard.

This is a must and prerequisite step to managing CloudGuard regions, security groups, and assets.

CloudGuard Operation Modes for AWS Accounts

CloudGuard has two operation modes to manage AWS accounts. The procedure of onboarding your environment to CloudGuard varies based on the operation mode you select.

- **Monitor** - Monitor and visualize your accounts in CloudGuard, run compliance tests on them, and receive alerts, notifications, and reports of activities and changes to cloud entities, but you cannot actively manage them from CloudGuard.
- **Full-Protection** - Contains all the capabilities of the **Monitor** mode. Use CloudGuard to enforce access and tamper protection on your assets, manage your Security Groups, and control access to your cloud assets.

See *"AWS Security Group Management Considerations" on page 371* for more details on operation mode considerations.

You can change the operation mode for an environment after it is onboarded to CloudGuard.

Notes Before Starting

Select the operation mode for the account. See *"AWS Security Group Management Considerations" on page 371*.

- You can select an operation mode for each account separately, so some can be Read-Only, while others are Full-Protection.
- If you use the Read-Only mode for an account, then all Security Groups in the account become Read-Only in CloudGuard (you can actively manage them in the AWS console or some other application). But if you use the Full-Protection mode for the account, you can choose to manage each Security Group separately as Read-Only or Full-Protection.
- At the end of the onboarding procedure, all Security groups are set to Read-Only mode in CloudGuard, regardless of the operation mode for the account. You can then change individual Security Groups to Full-Protection (for accounts in Full-Protection); see *"Full Protection" on page 373* in CloudGuard for details.

- The CloudGuard operation mode (Monitor or Full Protection) can be changed after your account has been onboarded.

For details about policies, see CloudGuard *"AWS Policies and Permissions" on page 279*.

For onboarding an AWS GovCloud account, see *"Manual Onboarding of AWS GovCloud or AWS China Environments" on page 165*.

Onboarding Options

You can onboard AWS accounts to CloudGuard in these ways:

- **Using the CloudGuard web portal and AWS console** - Onboard one AWS account following on-screen instructions, in CloudGuard portal and the AWS console.
- **Using automation batch scripts, from your AWS account** - Onboard an AWS account and, optionally, all child accounts, using scripts run from the AWS command line.
- **Using Terraform and the Terraform CloudGuard Dome9 provider** - Onboard one or more AWS accounts with Terraform files (one for each account) and the CloudGuard Dome9 Provider.
- **Using the CloudGuard REST API** - Onboard one or more AWS accounts with the CloudGuard REST API. You must first create a CloudGuard account and get an API Key and Secret in the CloudGuard web portal.

Onboarding from the CloudGuard Portal

The onboarding procedure is done on the CloudGuard portal, with step-by-step instructions presented on-screen for two modes: Monitor and Full-Protection. In the course of this procedure, you have instructions to perform some actions on the AWS Console and some on the CloudGuard portal.

CloudGuard does not make changes to the permissions or roles definitions in your AWS account. The actions you perform when you follow the on-screen instructions.

To onboard an AWS account:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**, click **Add** and select **AWS Environment**. The onboarding wizard with the new method of onboarding opens.
2. On the top of the screen, read the note and click **Switch to the manual onboarding**.

   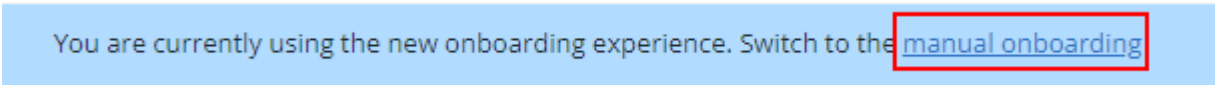

   The onboarding wizard with the legacy onboarding procedure opens.
3. Select the mode, **Monitor** or **Full-Protection**.

4. In your AWS account, prepare the IAM Policy that grants appropriate permissions to CloudGuard to access your AWS account for information about resources. The policy details are different for Read-Only and Full-Protection onboarding.
5. In your AWS account, create an IAM Role that CloudGuard has to use to access your environment (together with the IAM permissions defined in Step 4). It is necessary to provide details for this role of the CloudGuard AWS account, which uses the role.
6. Optionally, select the *"Organizational Units" on page 297* in CloudGuard with which the onboarded environment is associated. These associations can always be changed, from the Organizational Units page in the **Assets** menu.
7. Click **Finish**. The onboarding procedure starts. It can take a few minutes, based on the number of entities in your environment.

Using the CloudGuard Latest Policies

CloudGuard uses the *readonly-policy* to access information from your AWS account, for two operation modes. All CloudGuard functions, such as Posture Management, Network Security, and others, use this information.

Best Practice - Check Point recommends to use the latest version of the *readonly-policy*, which you can download from GitHub.

Additional Onboarding Methods

Onboard Using Automation Scripts

Use this open-source set of scripts to onboard accounts to CloudGuard from your AWS CLI https://github.com/dome9/onboarding-scripts/tree/master/AWS/full_automation.

These scripts create a CFT stack that creates the IAM policies required by CloudGuard, and then onboard the AWS accounts to CloudGuard. If your AWS accounts are organized as an AWS Organization, you can onboard the organization. The script automatically discovers the individual organization member accounts.

Onboard Using the CloudGuard REST API

You can onboard one or more AWS accounts to CloudGuard with the CloudGuard REST API. This requires an API Key and Secret for a CloudGuard account.

See the CloudGuard REST API reference guide and Onboard an AWS account to CloudGuard using the REST API for more details and examples.

More Links

- New onboarding experience - *"Unified Onboarding of AWS Environments" on page 54*
- API Reference

Manual Onboarding of AWS GovCloud or AWS China Environments

This topic explains how to add an AWS GovCloud or AWS China environment to CloudGuard with the legacy procedure. For the new experience with the Unified onboarding process, see *"Unified Onboarding of AWS Environments" on page 54*.

This onboarding process adds all regions and Security Groups in the AWS environment to the CloudGuard console and enables you to manage the AWS Security Groups from CloudGuard.

The onboarding process for these environments is equivalent to that for regular accounts (see *"Manual Onboarding of AWS Environments" on page 162* for details), only that permissions to CloudGuard to access the account are user-based, while for regular accounts they are role-based. An IAM user is created in the AWS GovCloud or China account, which CloudGuard uses to access the account.

You can manage AWS GovCloud or China accounts in CloudGuard as *Monitor* or *Full-Protection*, as for regular AWS accounts.

To onboard AWS GovCloud or China accounts manually:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**, click **Add** and select **AWS Environment**.
2. Select platform and mode. Select **GovCloud** or **AWS China** as the platform, and select the operation mode, Read-Only or Full-Protection.

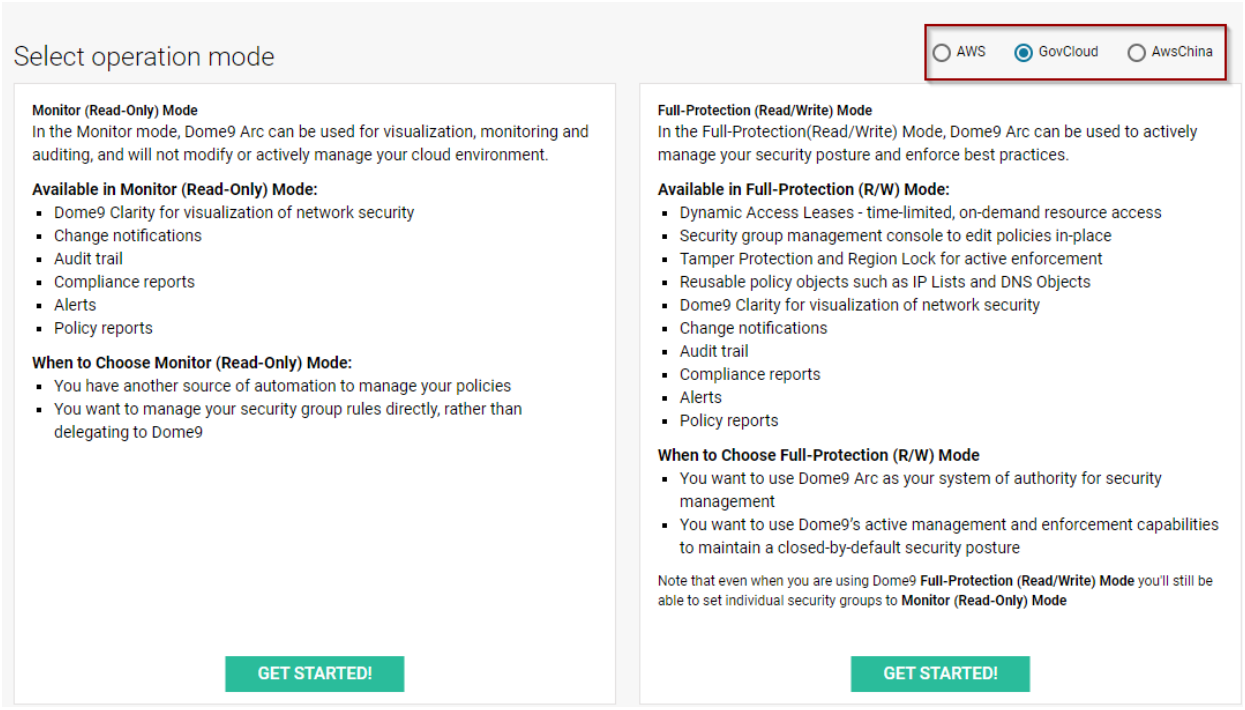

3. Follow these steps to prepare an IAM policy for CloudGuard.

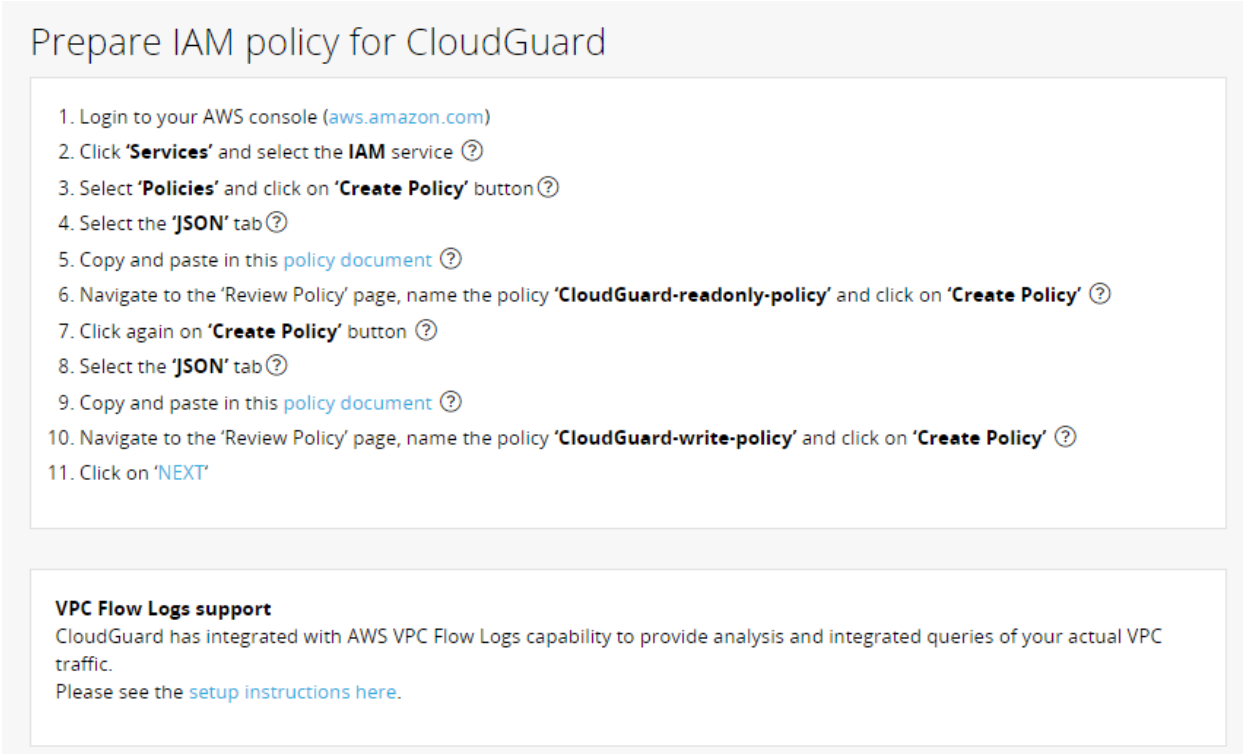

4. Follow these steps to create an IAM user for CloudGuard - GovCloud or AWS China.

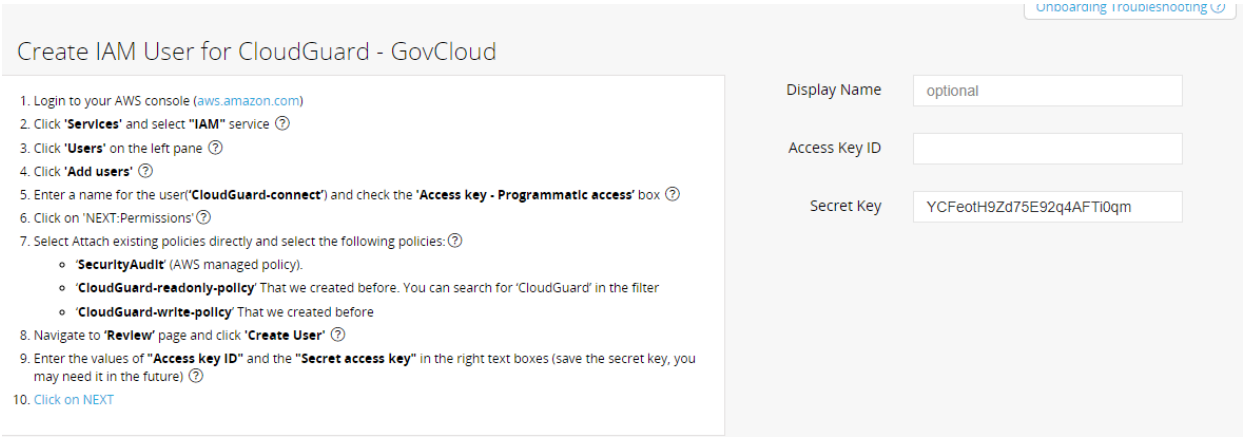

5. Optionally, select the *"Organizational Units" on page 297* in CloudGuard with which the onboarded environment is related. These associations can always be changed from the Organizational Units page.

6. Click **Finish**. The onboarding process starts. It can take a few minutes, based on the number of entities in the account.

Troubleshooting AWS Onboarding

This topic explains error messages and scenarios related to onboarding AWS accounts.

"Unable to add cloud account" Error

This error indicates that there may be a permissions problem.

It can indicate that the AWS IAM Role is missing a mandatory policy, or that the "External ID" is different from the "External ID" given to the AWS IAM Role.

Solution

1. Log in to the AWS console (aws.amazon.com).
2. Select **Services** and select the IAM service.
3. Click **Roles** and search for the Role created for CloudGuard (Usually, *CloudGuard-Connect*).
4. On the Role **Permissions** tab, make sure that you have all the **required polices**:
    a. **AmazonInspectorReadOnlyAccess** (AWS managed policy) - mandatory policy required for AWS Inspector information
    b. **CloudGuard-readonly-policy** (Created for CloudGuard) - mandatory policy
    c. **CloudGuard-write-policy** (Created for CloudGuard) - optional, required only for Full Protection mode
5. If one of the required policies is not attached, click **Attach Policy** to attach the missing policies.
6. To verify the External ID on the Role, click the **Trust relationships** tab.
7. Verify that the External ID is the same as given on the CloudGuard portal.

    Note - The External ID must not be empty.

8. If the **External ID** is empty or needs to be modified, click **Edit trust relationship** and correct it as required.
9. Copy the **Role ARN** and External ID and paste them to the CloudGuard portal.
10. Click **Finish**.

"Account is already protected by CloudGuard" Error

This error indicates that the AWS environment is already protected by CloudGuard.

It can be on the CloudGuard account you are currently trying to add this environment to or on another CloudGuard account.

Solution

First, make sure that you can find this environment on the **Environments** page.

If you cannot, contact your system administrator to verify if there is another CloudGuard account for the company.

"You are not subscribed to this service" Error

This error indicates that the AWS environment you are trying to connect is not in a valid state.

In most cases, it means that the registration process to AWS was not finished or that there is no verified defined payment method in the AWS environment.

When the AWS environment is not in a valid state, its functionality is limited.

Solution

First, make sure the AWS environment registration is completed.

Then, if the registration is correct, make sure that the payment method is valid.

Onboard the Account Again

If an exception persists, delete all the created policies and start onboarding from the beginning. See *"Onboarding AWS Environments" on page 144*.

Contact Check Point Support

If all these steps do not resolve the issue, contact *Check Point Support Center*.

Onboarding Azure Subscriptions

Azure offers infrastructure, platform, and software as services for cloud-based data management and applications (virtual machines, storage accounts, virtual networks, web apps, databases, or database servers).

To identify misconfiguration and compliance risks in Azure resources, you can onboard your Azure subscriptions to CloudGuard.

- For onboarding one subscription, see *"Onboarding an Azure Subscription" on page 61*.
- For onboarding an Azure Organization, see *"Onboarding Azure Organizations" on page 170*.

Azure Account Management Modes

There are two ways to manage your Azure account in CloudGuard.

- **Read-Only** - In this mode, you can view details for your Azure subscription in CloudGuard, run compliance tests on them, and receive alerts, notifications, and reports of activities and changes to cloud entities, but you cannot actively manage them from CloudGuard.
- **Manage** - In this mode, you have all the capabilities of Read-Only mode but, in addition, you can use CloudGuard to actively manage your Network Security Groups.

Onboarding Azure Organizations

Prerequisites

Before onboarding your Azure subscription, make sure:

- You have the **Owner** permissions for the onboarded scope and the **Global Administrator** role. For more information, see *"Permissions" on page 293*.
- You allow sufficient time (about one hour) for Azure to synchronize all created resources. If you create the resources during the onboarding, the synchronization takes some time and impacts the onboarding process. For example, if you create a storage account immediately before the onboarding, Azure cannot export diagnostics logs to the storage because the storage or the subscription does not yet exist.

How it Works

When you onboard Azure organizations to CloudGuard, you select onboarding scope according to the Azure hierarchy:

- **Tenant** - the highest node in the organization hierarchy, that is, a management group that contains all other management groups
- **Management Group** - a unit to organize your subscriptions created under one tenant

In the onboarding scope, CloudGuard provides automatic onboarding for all subscriptions and nested management groups added to the selected management group.

CloudGuard allows onboarding for nested management groups with custom configurations. A subscription configuration in CloudGuard is defined by the configuration (Step 2 of the onboarding wizard) of its nearest Management Group:

- CDR analyzes the selected types of logs for all subscriptions under the selected Management Group.
- AWP scans all subscriptions under the selected Management Group with the selected scan mode.

However, after you onboard a nested management group, you cannot onboard management groups that precede and succeed it in the hierarchy. In other words, if you onboarded a management group, you cannot onboard its parent or its children.

Best Practice - Check Point does not recommend that you move nested management groups between management groups with different onboarding configurations. The resulting configuration is unpredictable and is not supported by CloudGuard.

All management groups under the same tenant must be onboarded to the same CloudGuard account. In other words, you cannot onboard a management group to a CloudGuard account, if this group's tenant contains a management groups which is already onboarded to another CloudGuard account.

CloudGuard automatically adds Key Vaults of each of the types:

- Azure RBAC
- Access Policy (legacy)

The organizational credentials are automatically managed by the CloudGuard application.

After successful onboarding, CloudGuard gradually applies the selected configuration to all the subscriptions in the Management Group, and the process can take approximately one hour.

Onboarding in the Portal

STEP 1 - Welcome

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **Azure** and follow the setup instructions.

   Or, if you already onboarded environment(s), from the top menu, select **Add** > **Azure Subscription**.
3. Select to onboard an **Organization**.
4. Select the onboarding scope: **Tenant** or **Management Group**.
5. Log in to the Azure management portal.
6. Enter your **Tenant ID**.
7. For onboarding management groups, enter the **Management Group ID**. For more information about this parameter, see the Azure documentation.
8. Click **Next**.

STEP 2 - Configurations

1. On the **Configurations** page, enter a name for the organization (optional) to identify it in the CloudGuard portal.

   Note - The **CSPM** option is always enabled.

2. **CDR - Account activity** - enabled by default.
   a. Select the applicable log types. See the on-screen instructions to enable each type of log.

b. Enter the **Storage Account ID**:

- In the Azure portal, go to **Storage Accounts** and select the storage account, which subscription belongs to the onboarding scope (Management Group or Tenant).
- On the right, click **JSON View**.
- From the **Resource JSON** page, copy the **ID** value.

You can use three storage accounts at most.

3. **AWP - Agentless Workload Posture** - disabled by default.
   a. Click the toggle button to enable AWP.
   b. Select a scan mode. By default, the scan mode is **In-Account Centralized**.
   c. For **In-Account Centralized** mode, enter **Centralized Subscription ID**. Make sure the centralized subscription belongs to the onboarded management group.

   For more information about AWP, see *"AWP for Azure Environments" on page 510*.
4. Click **Next**.

Caution - After successful onboarding, moving subscriptions or nested groups between management groups onboarded with different configurations can have unpredictable effects, so make sure to avoid these cases:

- Moving a subscription from the onboarded organization to another organization which is *not* onboarded.
- Moving a subscription from the onboarded organization to another organization onboarded *without* AWP/CDR.
- Moving a subscription from the onboarded organization to another organization onboarded to AWP/CDR but *with a different scan mode* or *a different log type*.

Possible effects:

- AWP/CDR can lose its permission and stop scanning the subscription or analyzing the logs.
- Some resource groups that contain AWP/CDR resources can remain unmanaged.
- If AWP/CDR still has the required permissions, it can continue to scan a subscription with the scan mode of the original management group or analyze the original type of logs.

STEP 3 - Azure Network

If you select to onboard to CDR, and your Azure Storage account is private, CloudGuard requires access to the Storage account.

To allow connectivity to the Storage Account:

1. In your Azure Storage account, navigate to **Security + networking** > **Networking**.
2. On the **Firewall and virtual networks** page, Public network access is **Disabled**. Change it to the **Enabled from selected virtual networks and IP address** option.
3. In the **Firewall** section, add an IP address to allow access to based on your Data Center location. For the list of allowed IP addresses, see FAQ.
4. In the CloudGuard onboarding wizard, click the **Check Now** button to verify connectivity.

STEP 4 - Connect Management Group / Subscription / Tenant

1. On the **Connect Management Group / Subscription** page, click **Onboarding script** to review it. For additional script options, see *"Special Modes for Onboarding Script" on the next page*.
2. Log in to the Azure Cloud Shell.
3. Copy and run the command provided in the CloudGuard wizard. As the program runs it prints out its actions and shows resources created in your Azure account.
4. In the CloudGuard wizard, select:
   - **For Azure China and Azure Gov subscriptions**

     Provide additional information.

     From the command output, copy and paste two values into the CloudGuard wizard.

     a. **Application ID**

     b. **Secret Key**

   - **For Standard Azure subscriptions**

     Select the check box to approve that the script successfully completed its job.
5. Click **Onboard**. The onboarding starts.

In a couple of minutes, you receive a notification that the management groups is successfully onboarded. Wait about 15 minutes (up to one hour in case of AWP In-Account Centralized mode) until the onboarded subscriptions appear on the **Environments** page.

The new organization appears on the **Assets** > **Organizational Units** page under the root OU as its child, like the manually created CloudGuard OUs. All actions available for regular OUs (creating sub-OU, renaming, moving, and deletion) are available for the onboarded Azure organization.

Note - For Organization onboarding configured with CloudGuard Organizational Unit where the CloudGuard OU was deleted, the newly discovered subscription is onboarded under the root CloudGuard Organizational Unit.

This procedure creates a default CSPM policy with two rulesets: **Azure CloudGuard Checkup** and **Azure CIS Foundation v. 1.5.0**.

Onboarding with API

To onboard Azure organizations with API:

1. Make the **first call**: POST - https://api.dome9.com/v2/AzureCloudAccount/OnboardingExecutionCommand (Link). The output is the command string for execution from the Azure Cloud Shell.
2. In the Azure Cloud Shell, run the command obtained from the response.
3. Make the **second call**: POST - https://api.dome9.com/v2/azure-organization-management. Set the ActiveBlades parameter according to the enabled active blades.

Special Modes for Onboarding Script

- To run the script in a quiet mode that skips all questions to the user, run it with the `--quiet` flag.
- To remove the CloudGuard resources created in your Azure organization, run the onboarding script with the `--clean` flag. You can use the option for troubleshooting an unsuccessful onboarding. Make sure you use the same parameters as in the initial command.

Known Limitations

CloudGuard supports scanning of Function Apps in Azure Organizations only with AWP In-Account Centralized mode.

For more limitations related to AWP, see *"Known Limitations" on page 492*.

More Links

- *"Onboarding Azure Subscriptions" on page 169*
- *"Onboarding an Azure Subscription" on page 61*
- *"Azure Roles and Permissions" on page 293*
- *"Troubleshooting Azure Onboarding" on page 176*

Troubleshooting Azure Onboarding

After successfully onboarding your Azure subscription, make sure there are no errors that prevent CloudGuard from evaluating your Azure assets.

Invalid Credentials or Missing Permissions

Most errors are related to invalid credentials (access denied) or missing permissions.

To troubleshoot:

1. Open **Assets** > **Environments**.
2. Click the filter icon.
3. Select:
   - Status: **Error** and **Warning**
   - Platform: Azure
4. To see tool-tip warnings, put the cursor on **Status** column items.
   - *Invalid credentials* - The *Cloudguard-Connect* app does not have sufficient permissions to connect with your Azure environment.
   - *Missing permissions* - The *Cloudguard-Connect* app does not have sufficient permissions to get data from specific Azure resources.

   To correct issues:
5. Select the related environment.
6. Click **Show more** for a list of assets for which *Cloudguard-Connect* does not have permissions.
7. To try to validate the environment permissions, click **Validate Permissions**.
8. To open the CloudGuard troubleshooting wizard, click **permissions wizard**.
   a. Select your operation mode - **Read-Only** or **Manage**.
   b. Complete each step to make sure the *Cloudguard-Connect* app has the correct permissions.
   c. In CloudGuard, click **FINISH**.

Missing Permissions for Azure Web App or Function App

To resolve missing permissions, it is necessary to add a custom role to your *CloudGuard-Connect* application. Create the Azure custom role with action permissions.

1. Log in to the Azure management portal and select your onboarded subscription.
2. In the navigation tree, select **Access control (IAM)**.
3. Click **Add** and select **Add custom role**.
4. Use one of two methods below for the new role:
   - In the web portal, enter these details:
     a. **Role name**: CloudGuard Additional Permissions Role (example)
     b. **Description**: Action permissions for CloudGuard (example)
     c. **Permission**: Microsoft.Web/sites/config/list/Action
   - Paste the JSON with the applicable subscription ID:

```
{
    "properties": {
        "roleName": "CloudGuard Additional Permissions Role",
        "description": "Action permissions for CloudGuard",
        "assignableScopes": [
            "/subscriptions/a1a1a1a1-a1a1-a1a1-a1a1-
a1a1a1a1a1a1"
        ],
        "permissions": [
            {
                "actions": [
                    "Microsoft.Web/sites/config/list/Action"
                ],
                "notActions": [],
                "dataActions": [],
                "notDataActions": []
            }
        ]
    }
}
```

5. In **Access control (IAM)**, click **Add** and select **Add role assignment**.
6. Select the **CloudGuard Additional Permissions Role** created before and add it to your application.
7. In CloudGuard, click **Validate Permissions** and wait approximately 30 min for changes to take effect.

Missing Permissions for Azure Key Vaults

If you use Azure key vaults to store and protect authentication credentials, CloudGuard needs access to your Azure Key Vault metadata to access the vault and its contents for compliance.

To protect the key vault resource and its contents, the Azure role-based permission mechanism does not grant access to this metadata. Key vault permissions must be granted on an individual Key Vault Basis.

In **Assets** > **Environments**, CloudGuard responds to this initial lack of access by showing *missing permission* errors in the **Status** column.

To resolve these missing permission errors and maintain the security of your key vaults, you must implement additional permission-delegation steps (see: *"Configure Policies for Azure Key Vault Entities" below*).

After CloudGuard has the required permissions, the Compliance engine uses a list of pre-defined GSL rules to make sure that:

- Azure Key vaults are used to store and protect authentication credentials (keys, secrets, certificates) in the environment
- An expiration date is set on all keys (such as cryptographic keys)
- An expiration date is set on all Secrets (such as passwords, database, and connection strings)
- The key vault is recoverable to protect against accidental deletion by a user or malicious activity
- The key vault has purge protection enabled to protect against insider attacks
- All interactions with Key Vault instances are available in key vault event logs

To increase the security of your key vaults, create GSL rules modified according to requirements for your Azure Key Vaults and their contents. For example, rules that evaluate data points unique to your key vaults, such as how many persons can issue certificates and who they are.

Configure Policies for Azure Key Vault Entities

Azure Key Vaults have entities that are not accessible with the policy that is set up when the Azure account is onboarded to CloudGuard. This is because, by default, Azure does not grant access rights to vaults, secrets, certificates, and keys. In addition, new entities may be created from time to time. The CloudGuard Compliance Engine, for example, needs to access these entities when evaluating the compliance of your Azure environments.

Note - Azure Storage accounts cannot have expiration or renewal checks of the Access Keys. This is possible only with Azure Key Vaults.

Follow the steps below to set up an Automation account and runbook in your Azure account that, at intervals, grants rights to CloudGuard to access these new entities.

Step 1 - Create a managed identity on Azure

In this step, you create a user-assigned managed identity with access to the subscriptions containing the Key Vaults.

1. Log in to the Azure management portal.
2. Create a new user-assigned managed identity. See Microsoft Entra ID documentation for more details.
3. On the **Azure role assignments** page, click **Add role assignment** and assign the **Key Vault Contributor** role to the managed identity.

Important - Each Azure subscription with a Key Vault must have its own Automation account and runbook. Make sure to repeat the steps below for each subscription. It is enough to create only one managed identity for all configurations.

Step 2 - Create an automation account and runbook on Azure

In this step, you create an Automation account and a runbook, with access to the subscription containing the Key Vaults to the entities to be updated in the CloudGuard policies.

1. Navigate to **Automation Accounts** and create an Azure Automation account. See Azure Automation documentation for details about creating an Automation account.
2. Navigate to **Automation account** > **Identity** > **System assigned** and click **Azure role assignment** to add a role to the applicable automation account.
3. Navigate to **Automation account** > **Identity** > **User assigned** and click **Add** to add the managed identity created in Step 1.
4. Navigate to **Process Automation** > **Runbooks** > **Create a runbook**.
    a. Enter a name for the runbook, for example, *KeyVault*.
    b. For **Runbook type**, select **PowerShell**.
    c. For **Runtime version**, select **7.2**.
    d. Optionally, add a description, then click **Create**.

Step 3 - Prepare the runbook script

In this step, you prepare a script that grants permissions to the CloudGuard application to access the entities in the Key Vaults. This script is then scheduled as a job run at intervals. The script used in this step is an example of a script. You can change it or replace it with your script.

1. Download this script: keyvault_script.ps1
2. In the **Runbook** page, open the runbook created in Step 2 and then click **Edit**.
3. Copy the script from the file and paste into the edit pane.

4. Search for the variable *$excludedKeyVaults*, and set its value to exclude specific Key Vaults (if there are no exclusions, then keep it empty):

   $excludedKeyVaults = "DBKeyvault", "VMKeyvault"

5. Save the changes and close the edit pane.
6. From the Azure dashboard, select **Enterprise applications** and then select *CloudGuard-Connect* (the application created as part of the onboarding process).
7. Copy the *ObjectID* value.
8. Go back to the **Runbook** page, select again the runbook, and open the edit pane.
9. Search for the variable *$objectIds* and paste the value copied above.
   `$objectIds = "12a34bc5-d5e6-78fg-9hi0-1234abcd5678efgh"`
10. Save the runbook script.
11. Optionally, click **Test pane** to open the Test pane and test the script.
12. Click **Start** to start the test. When it completes, the window must show that the access policies are updated for the vaults.

Step 4 - Publish & schedule the runbook as an Azure job

1. In the Runbook page, select the runbook created in Step 3 and open the edit pane.
2. Click **Publish**.
3. In the Runbook, select the runbook, then select **Schedules** in the **Resources** section, and then click **Add a schedule**.
4. Select **Link a schedule to your runbook** and then **Create a new schedule**.
5. Enter a schedule, for example, hourly. The repetition period shows the frequency with which new values are added to the Key Vaults.
6. Click **Create** to save the schedule. The runbook runs as a job based on the schedule that you set up for it.
7. Repeat the above steps for each runbook if there is more than one.
8. In the **Runbook** page, select **Jobs** in the left pane to check the status of the jobs. A new line appears each time the scheduled runbook completes.

Set Vault Access Policy Permission

Sometimes, CloudGuard can show the *Missing permission* errors if your Azure subscriptions host one or more Key Vaults.

An example of the error message appears below.

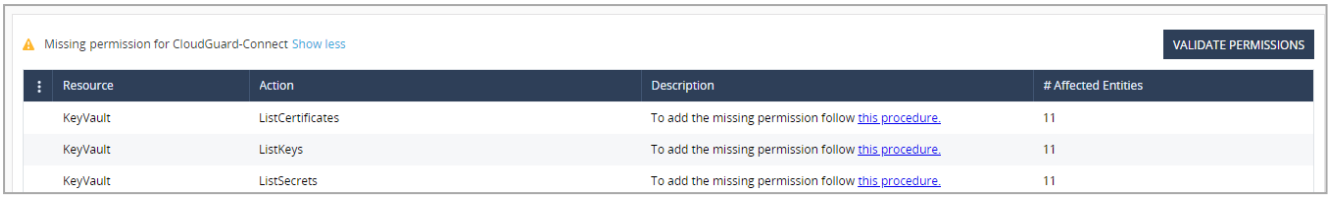

This occurs because Azure supports multiple overlapping access authorization mechanisms that apply to:

- different cloud resource types
- different generations of the same cloud resource-type
- different levels of cloud resource organization

Some of these permission delegation methods are mutually compatible, but many are not.

Most of the permissions required for CloudGuard to collect metadata on Azure-based resources are granted by a Role-based permission mechanism that applies at the Azure subscription level. But access to Azure Key Vault metadata cannot be granted through that process. Key Vault permissions must be granted on an individual Key Vault basis.

CloudGuard collects non-sensitive Key Vault metadata with an Automation Account and Runbook. Recent improvements in CloudGuard extend the scope of required metadata collection to some Azure Graph API endpoints, and the current Runbook + Automation Account framework does not delegate permissions to access Azure Graph.

To enable CloudGuard enhanced support for posture management evaluation of Keys, Secrets, and Certificates stored in Key Vault, complete more permission delegation steps. You have to use Key Vault **Vault Policy** to allow the CloudGuard service to issue read-only commands to the Azure Graph API.

For more instructions about these permissions configurations, see sk173403.

Note - The use of Key Vault **Vault Policies** as a permission delegation mechanism is not compatible with the use of Key Vault's Firewall (network access rule creation) feature. The creation of even one Key Vault Firewall access rule effectively disables each Vault Policy access permission that is not covered by an equivalent Firewall access rule.

Onboarding Google Cloud Platform Projects

To identify misconfiguration and compliance risks in Google Cloud Platform (GCP) projects, you can onboard one project or a group of GCP projects to your CloudGuard account.

Prerequisites

- Google Administrator permissions
- Enabled Cloud Resource Manager API

General Workflow

To successfully onboard CloudGuard, you must:

- Complete the CloudGuard onboarding wizard
    - Open Google Cloud Platform
    - Enable APIs needed by CloudGuard
    - Create a service account in GCP for CloudGuard
    - Select roles for CloudGuard
    - Create keys
    - Upload keys to CloudGuard
    - Select CloudGuard organizational units for the account
- Troubleshoot to remove initial errors

Onboarding a Google Cloud Platform (GCP) Project and Google Workspace

Prerequisites

- You must have Owner permissions for the GCP project.
- To connect Google Workspace to CloudGuard, you must have Owner permissions for the Google Workspace.

To onboard a GCP project to CloudGuard:

1. In GCP, open the project that you want to onboard to CloudGuard. Keep the GCP project open throughout this procedure.
2. In the CloudGuard UI, from the left menu, expand **Assets** and click **Environments**.
3. In the toolbar above the table, in the top left, click **Add** and then click **GCP Project**.

   The **GCP Onboarding** wizard opens.
4. In the **Welcome** step of the wizard:
   a. Copy the **Project ID** from GCP.
   b. Paste the **Project ID** into CloudGuard.
   c. In CloudGuard, click **Next**.
5. In the **Configurations** step of the wizard:
   a. **Optional -** For **Environment Display Name**, enter a name for the integration to appear in the CloudGuard UI. By default, the **Display Name** is the **Project ID**.
   b. Select an **Organizational Unit** to associate with the integration.

      Note - An integration of GCP with CloudGuard CNAPP always includes the CSPM feature in CloudGuard CNAPP. The **CSPM** slider is on by default, and cannot be turned off.
   c. **Optional -** To onboard GCP to CloudGuard without onboarding Google Workspace, move the **Workspace** slider to "off".

      Note - It is also possible to connect Google Workspace to CloudGuard after you finish the GCP onboarding.
   d. Click **Next**.
6. In the **Connect** step of the wizard:

a. Click **Onboarding Script** to review the GCP onboarding script that CloudGuard generates automatically.

b. Copy the command from CloudGuard.

c. Paste the command into the CLI of the GCP project.

   The CLI of the GCP project generates a JSON that contains GCP credentials.

d. In the CLI of the GCP project, copy the JSON (including the opening and closing brackets).

e. In CloudGuard, paste the JSON into the **Credentials JSON** field.

f. Click **Next**.

7. In the **Workspace** step of the wizard, do **one** of these:

   - If you are onboarding a Google Workspace:

     a. Follow the instructions shown in the CloudGuard UI to configure Google Workspace.

     b. Click **Next**.

   - If you are **not** onboarding a Google Workspace, click **Skip**.

Troubleshooting GCP Onboarding

After successfully onboarding your GCP Projects, make sure there are no errors that prevent CloudGuard from evaluating your GCP projects. Most of the errors relate to invalid credentials (access denied) or missing permissions.

To troubleshoot:

1. Open **Assets** > **Environments**.
2. Click the filter.
3. Select:
   - Status: **Error** and **Warning**
   - Platform: GCP
4. Put the cursor on the **Status** column items to see tool-tip warnings.
   - *Invalid credentials* - The *Cloudguard-Connect* service account you created does not have sufficient permissions to communicate with your GCP environment
   - *Missing permissions* - The *Cloudguard-Connect* service account does not have sufficient permissions to get data from specific GCP projects
5. Click on the related GCP asset.
6. Click **Show more** for possible remediation steps.
7. Click **Validate Permissions**.
8. Click **VALIDATE**.

   CloudGuard automatically validates the permissions, dismisses the warning, and updates the environment status in two to three minutes.

Note - If automatic validation fails, for troubleshooting purposes, onboard your GCP project again.

Onboarding Oracle Cloud Infrastructure Environments

You can onboard an OCI account to CloudGuard.

To onboard an OCI account to CloudGuard:

STEP 1 - Welcome

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **OCI** and follow the setup steps.

   Or, if you already onboarded environment(s), then from the top menu, select **Add** > **OCI Environment**.
3. In the onboarding wizard, enter a new name for your OCI environment. This name lets you identify the environment in CloudGuard.
4. Follow the on-screen instructions to enter mandatory (*) and optional information in the fields. Use tooltips to see more explanations.

   Note - After onboarding, you can possibly receive a validation email message to the provided **OCI tenant administrator email address**. CloudGuard creates a new user on your tenant. This does not require action, and you can safely ignore the message.
5. Click **Next**.

STEP 2 - Organizational Unit

1. Select one of the available organizational units to associate with the environment.
2. Click **Next**.

STEP 3 - Set up a New CloudGuard Environment

CloudGuard uses a Terraform template for policy definition and for permissions to read data in your tenancy. Download the template and follow the wizard instructions to configure your OCI account.

At this stage, you create and run a stack in your OCI account.

Click **Done** to complete the onboarding procedure.

The new environment appears in the list of all environments. It takes several minutes to see all assets onboarded to CloudGuard.

To remove the OCI environment data from CloudGuard:

1. On the environment page, click **Remove** to disconnect the selected environment from CloudGuard.

   A verification window opens.

2. Click **Remove** in the verification window.

   If you have active policies on the environment, CloudGuard notifies you that it deletes the policies in line with the environment.

   This process deletes:

   - Policies
   - Notifications
   - Compliance alerts (findings)

Onboarding Kubernetes Clusters

You can onboard a Kubernetes cluster to CloudGuard. On the process completion, you can see clusters, nodes, pods, and other resources on the CloudGuard Assets page. Then you can run compliance assessments on them and use the data for more security functionality, such as Runtime Protection, Image Assurance, etc.

The cluster can be on an on-premises host or in a cloud environment with managed Kubernetes environments such as AKS on Azure, EKS on AWS, and GKE on GCP Cloud.

As part of the onboarding process, CloudGuard agents are deployed on the cluster. The CloudGuard agents send encrypted information back to the CloudGuard server over the Internet.

For information on Kubernetes versions and container requirements, see *"Kubernetes Containers" on page 415*.

Onboarding a Cluster Manually

Follow the steps below to manually onboard a Kubernetes cluster to CloudGuard:

STEP 1 - Configuration

1. In the CloudGuard portal, open **Assets** > **Environments**.
2. For first-time onboarding, click **Kubernetes**. The first window of the wizard to onboard a Kubernetes cluster opens.

   Or, from the top menu, select **Add** > **Kubernetes / OpenShift / Tanzu**.
3. Enter a name for the cluster. This is the name that appears in CloudGuard.
4. Follow the on-screen instructions to complete these steps:
   - Configure a Service Account by one of these methods:
     - Select an existing Service Account with its corresponding API Key.
     - Enter a Service Account manually.
     - Click **Add Service Account** to create a new account.
   - Enter a name for the Kubernetes namespace in which the agent is to be deployed or keep the default name - *checkpoint*.

- Select what type of monitoring and security checks are necessary for your Kubernetes cluster by default. You can add each of these features later. Read more about each feature on a dedicated page:
  - Posture Management - for details, see *"Cloud Security Posture Management (CSPM)" on page 302* (mandatory feature)
  - Image Assurance - for details, see *"Image Assurance" on page 434*
  - Admission Control - for details, see *"Admission Control" on page 476*
  - Runtime Protection - for details, see *"Kubernetes Runtime Protection" on page 549*
  - Threat Intelligence - for details, see *"Intelligence for Kubernetes Containers" on page 645*

5. Click **Next** to continue to the next step.

STEP 2 - Select Organizational Unit

1. Select the *"Organizational Units" on page 297* with which the onboarded cluster will be associated. If no Org Unit is selected, the root (top-level) unit is used.
2. Click **Next**.

STEP 3 - Deploy the agent on the cluster

1. Follow the on-screen instructions and apply Helm. As an alternative, you can follow the Non-Helm instructions to deploy the agents. This generates a YAML file for deployment with `kubectl` commands.

   For more installation options, see *"Installing the Agent" on page 192*.
2. Click **Next**.

STEP 4 - Onboarding Summary

1. Verify the deployment status. The status is dynamically updated as the agents come online.

   CloudGuard informs you that:
   - Your Kubernetes cluster has been successfully created.
   - It is waiting for the agent to start communication.
   - You can skip the validation if you click the **Finish** button.
2. Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

After the agent is deployed, CloudGuard accesses the cluster through the agent to get information about the assets and synchronize with it. This takes several minutes based on the time needed to download the images to the cluster and the number of assets in the cluster.

The Onboarding Summary page is updated automatically with the change of the cluster status.

Cluster Status

Available options of the cluster status:

- **Pending** - CloudGuard has not received communication from the agents.
- **Initializing** - CloudGuard is receiving communication from some of the agents. The progress bar shows how many agents are up and ready.

  Note - During this state, if the number of running pods does not change for 10 minutes, the indicator pauses and the status changes to **TIME OUT**. In this case, verify the agents status on the cluster to make sure they do not have issues. For example, agents can be stuck because of missing resources (memory or CPU). After you resolve the issue, you can continue the validation or skip the validation process entirely.
- **Error** - There are agents in the Error state. Click **Finish** to complete the process. You can go to the cluster page to see which agents have the Error state and browse their Kubernetes logs for issues.

When all the agents are running, the cluster status changes to **SUCCESS**, and the onboarding process finishes successfully.

Agent Status

On the cluster page, for each feature, you can see the status of its agents:

- **Pending** - The agent has never communicated with CloudGuard.

  Note - There is a limitation for DaemonSet agents. During the cluster status calculation, tolerations settings are not considered. Agents from excluded nodes are considered **Pending** which can cause a false error state for the cluster.
- **Initializing** - Status of an agent that comes online and initiates communication with the CloudGuard portal. The agent has a small period to report a successful self-test. If the agent does not report it back on time, the status is changed to Error because of a timeout.
- **Warning** - Status of an agent that successfully finished its initialization, while it is based on an old image. See *"Upgrading the Agent" on page 196* for how to resolve this issue.

- **Error** - Status of agents that
  - failed their self-test
  - sent an error message
  - suffered a loss of connectivity for a minimum of one hour
  - have the version below the minimal version
- **Pending cleanup** - Disabled features that still have an agent that sends data to appear with the Pending cleanup status.

Onboarding a Cluster with Automation

Automation with the CLI

For the onboarding automation, you need a CloudGuard service account with onboarding permissions, so the service account must have a role with the **Manage Resources** *"Direct Permissions" on page 843* (or, at least, the **Onboarding** Permissions).

Follow these steps to automate the onboarding process from the command line:

1. With the above-mentioned service account, create or update these environmental variables: $API_KEY, $API_SECRET, $CLUSTER_NAME, where the API Key and Secret are generated on the CloudGuard portal (For CloudGuard Dome9 accounts, see *"V2 API" on page 840*; for Infinity Portal accounts, see API Keys.).
2. Run this command to create a Kubernetes account on CloudGuard:

```
curl -s -X POST
https://api.us1.cgn.portal.checkpoint.com/v2/kubernetes/accou
nt --header 'Content-Type: application/json' --header 'Accept:
application/json' -d "{\"name\" : \"$CLUSTER_NAME\"}" --user
$API_KEY:$API_SECRET)
```

   Note - This and other commands below use `api.us1.cgn.portal.checkpoint.com` as an API endpoint for Infinity Portal users in the US region. For the full list of the API server endpoints in your region, see *"Which CloudGuard endpoints do I have to allow on my network?" on page 917*.

3. Extract the Cluster ID from the response:

```
CLUSTER_ID=$(echo $CREATION_RESPONSE | jq -r '.id')
```

4. Enable the required features:

```
curl -X POST
https://api.us1.cgn.portal.checkpoint.com/v2/kubernetes/account/$
CLUSTER_ID/imageAssurance/enable --user $API_KEY:$API_SECRET
curl -X POST
https://api.us1.cgn.portal.checkpoint.com/v2/kubernetes/account/$
CLUSTER_ID/admissionControl/enable --user $API_KEY:$API_SECRET
curl -X POST
https://api.us1.cgn.portal.checkpoint.com/v2/kubernetes/account/$
CLUSTER_ID/runtimeProtection/enable --user $API_KEY:$API_SECRET
```

Inactive Kubernetes Clusters

CloudGuard deletes inactive environments when a year (365 days) passed since any of the environment's agents has communicated with CloudGuard. An agent is required to communicate with CloudGuard at least once in the past.

Note - Environments with agents that communicated with errors are not removed.

Installing the Agent

During the onboarding process, you install the CloudGuard agent on the cluster with Helm. For the agent installation, permissions of the preconfigured **Kubernetes Agent** role are sufficient (see *"Roles" on page 843*). The Helm command is shown on the third page of the onboarding wizard; see *" STEP 3 - Deploy the agent on the cluster" on page 189*.

Example:

```
helm install asset-mgmt cloudguard --repo
https://raw.githubusercontent.com/CheckPointSW/charts/master/repositor
y/ --set-string credentials.user=$API_KEY --set-string
credentials.secret=$API_SECRET --set-string clusterID=$CLUSTER_ID --
set addons.imageScan.enabled={true|false} --set
addons.admissionControl.enabled={true|false} --set
addons.runtimeProtection.enabled={true|false} --namespace $NAMESPACE
```

You can set the `*.enabled` flags to `false` or omit them if it is not necessary to enable the corresponding features.

If you do not have Helm installed in your environment, use the command below to generate a YAML file for the agent installation with `kubectl`.

```
kubectl run cloudguard-install --rm --image alpine/helm --tty --stdin
--quiet --restart=Never --command – helm template asset-mgmt
cloudguard --repo
https://raw.githubusercontent.com/CheckPointSW/charts/master/repositor
y/ --set credentials.user=$API_KEY --set credentials.secret=$API_
SECRET --set clusterID=$CLUSTER_ID --set addons.imageScan.enabled=
{true|false} --set addons.admissionControl.enabled={true|false} --set
addons.runtimeProtection.enabled={true|false} --namespace $NAMESPACE -
-set containerRuntime=containerd --kube-version <KUBERNETES-VERSION> >
cloudguard-install.yaml

kubectl apply -f cloudguard-install.yaml
```

- If your cluster uses a Docker or CRI-O runtime environment, change the `containerRuntime` flag to:

  ```
  --set containerRuntime=docker or --set containerRuntime=cri-o
  ```

- If your cluster platform is OpenShift 4+, Amazon EKS, or Tanzu, before output redirection, add:

  ```
  --set platform=openshift or --set platform=eks or --set
  platform=tanzu
  ```

- With the `kube-version` flag, set the Kubernetes version of your cluster, for example, for version 1.25, set:

  ```
  --kube-version 1.25
  ```

Installation with a Values File

You can use a YAML file as an alternative or in addition to the `--set` command line parameters during the Helm chart installation. Use the *--values <file>* or *-f <file>* flags in the Helm installation command. This is helpful when you have many changes to the default installation parameters or when it is necessary to specify complex or nested values.

See the default file format in CloudGuard repo, as well as the description of the configurable values.

Example

This *values* file with `clusterId` and credentials (`userId`, `secretKey`) enables the *imageScan* and *flowLogs* features on the cluster and requests a custom configuration of compute resources for the daemon.

```
## Check Point CloudGuard  cluster ID
```

```
## Example: "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX"
clusterID: clusterId

## CloudGuard datacenter: usea1 [default], euwe1, apso1 etc.
datacenter: usea1

## Check Point CloudGuard Credentials
## Example
## API Secret: "abcdefghijklmnopqrstvuwxyz"
## API ID: "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX"
credentials:
  secret: secretKey
  user: userId

addons:
  imageScan:
    enabled: true
  flowLogs:
    enabled: true
    daemon:
      resources:
        requests:
          cpu: 200m
          memory: 60Mi
        limits:
          cpu: 400m
          memory: 200Mi
  admissionControl:
    enabled: false
  runtimeProtection:
    enabled: false
```

Run this command:

```
helm install asset-mgmt cloudguard --repo
https://raw.githubusercontent.com/CheckPointSW/charts/master/rep
ository/ -f values.yaml --namespace $NAMESPACE
```

Heterogeneous Node Pools

When a cluster contains multiple node pools with different configurations, it is sometimes necessary to configure the CloudGuard agent differently for each node pool. For example, one node pool can have small nodes (for example, four CPUs per node), while another can have very big nodes (32 CPUs per node). In such a cluster, it is practical to adjust the configuration of the CloudGuard agent's DaemonSets for each node pool. Below are some examples of when different DaemonSets configurations in different node pools are beneficial:

- Different resource allocation (for example, allocate more CPU for runtime daemon on nodes with more CPUs)
- Different container runtimes (for example, nodes running Docker against nodes running containerd)
- Different architecture

The CloudGuard agent's Helm chart allows to set up multiple DaemonSet configurations for different node pools with the use of the *daemonConfigurationOverrides* property. It is available under each `addons.<feature>.daemon` section of the YAML file. This property is an array that specifies multiple *override* configurations in addition to the *default* configuration specified under the `daemon` section.

For each section of overrides, Helm creates a new DaemonSet in the cluster, with the specified configuration.

In addition:

- The overrides inherit from the default daemon configuration (`addons.<feature>.daemon` object), and each value set on it applies likewise to the override configurations, unless explicitly changed.
- The name of each configuration override must be unique (case-insensitive). Non-unique names like “configExample” and “ConfigEXAMPLE” overwrite one another.
- Each configuration must have a `nodeSelector` field defined, otherwise, the command fails.
- Make sure that the `nodeSelector` fields do not overlap, and a node fits only one configuration. The node that matches more than one configuration will have additional daemons.

Example

This *values* file enables the *flowLogs* feature and sets two different configurations for daemons based on the values of the `nodeSizeKey` custom label on the nodes.

```
addons:
  flowLogs:
    enabled: true
    daemonConfigurationOverrides:
      exampleConfigOneSmall:
        nodeSelector:
          nodeSizeKey: small
        resources:
          requests:
            cpu: 100m
            memory: 30Mi
          limits:
```

```
        cpu: 200m
        memory: 100Mi
    exampleConfigOneLarge:
      nodeSelector:
        nodeSizeKey: large
      resources:
        requests:
          cpu: 200m
          memory: 60Mi
        limits:
          cpu: 400m
          memory: 200Mi
```

Agent Version Life Cycle

Each CloudGuard agent has its recommended and/or minimal required version, which CloudGuard recommends use. New versions of the agents are released when they accumulate significant content, including new capabilities, fixed vulnerabilities, etc.

To verify the agent's version:

1. Select an environment and click to open its page.
2. In the Blades tab, expand the module's details. In the **Version** column, see the version number.
3. In the Status column, see the agent's status:
   a. **Warning** / Agent is not up to date - The agent version is below recommended
   b. **Error** / Agent version is not supported - The agent version is below minimal

   This status appears on the environment page and in the applicable API (agentSummary APIs).

When an agent accumulates significant content, CloudGuard recommends upgrading it - see *"Upgrading the Agent" below*. The agent status changes from **OK** to **Warning**.

When an agent has many issues or sufficient time passes after the outdated agent status is moved to **Warning**, CloudGuard changes the minimal version. The agent status changes from **Warning** to **Error**.

Upgrading the Agent

Assumptions:

- The environment variables $API_KEY, $API_SECRET, $CLUSTER_ID, $NAMESPACE have the same values as during onboarding
- Image Assurance and Admission Control are enabled

For agents installed with Helm 3, use the command below to upgrade all agents to the latest version:

```
helm upgrade --install asset-mgmt cloudguard --repo
https://raw.githubusercontent.com/CheckPointSW/charts/master/repos
itory/ --set-string credentials.user=$API_KEY --set-string
credentials.secret=$API_SECRET --set-string clusterID=$CLUSTER_ID
--set addons.imageScan.enabled={true|false} --set
addons.admissionControl.enabled={true|false} --set
addons.runtimeProtection.enabled={true|false} --set
datacenter=usea1 --namespace $NAMESPACE --create-namespace
```

Downgrading the Agent

If you want to use a previous version of the agent (not recommended), you can downgrade the agent with standard Helm procedures, specifying the desired Helm chart version. Use the `helm rollback` or `helm upgrade` commands.

Uninstalling the Agent

During the process of onboarding, CloudGuard generates the *cloudguard-install.yaml* file that you use to uninstall the agents.

With Helm:

```
helm uninstall asset-mgmt --namespace $NAMESPACE
```

With kubectl:

```
kubectl delete -f cloudguard-install.yaml --namespace $NAMESPACE
```

Note - To install agents again after you have uninstalled them, follow *" STEP 3 - Deploy the agent on the cluster" on page 189* and not the upgrade procedure.

Troubleshooting Kubernetes Onboarding

Deploying the Agent

1. To see if the pod exists, run:

   ```
   kubectl -n <namespace> get pods
   ```

   If the pod exists, but is not ready, get the pod's details:

   ```
   kubectl describe -n <namespace> pod <pod_name>
   ```

2. Do a review of the Kubernetes logs for error(s). To get the logs, run:

   ```
   kubectl -n <namespace> logs <pod_name> [-c container_name]
   ```

3. Make sure the pods have connectivity to CloudGuard.

   The pod must have HTTPS (port 443) connectivity to `https://api-cpx.dome9.com`.

   Check these entities for possible configuration issues that prevent connectivity to the agent:

   - Proxy
   - Network Policy
   - Security Groups
   - Firewall rules

   You can install a different pod that has curl (such as an Alpine pod, for example) in the same namespace, with the same labels, exec into it, and do a connectivity check to the above URL with curl.

   ```
   apk update ; apk add curl ; curl -k https://api-
   cpx.dome9.com/namespaces -X POST
   ```

   If there is connectivity with the CloudGuard backend, the response is `{"message":"Unauthorized"}`.

4. Make sure the nodes can connect to the Image Registry.

   The node needs HTTPS (port 443) connectivity to the Quay registry. If you see an image pull error, make a connectivity check to:

   ```
   https://quay.io/checkpoint/
   ```

Cluster Behind a Gateway

If the traffic passes from the cluster to the Internet through a Security Gateway with HTTPS inspection, you have to configure a customer CA (Certificate Authority) certificate for the agents.

1. Put the customer Base64 PEM-encoded CA certificate in a configmap in the applicable namespace.

   For example, if the CA certificate is in file `ca.cer`:

   ```
   kubectl -n <namespace> create configmap ca-store --from-file=
   ca.cer=<PATH_TO_CA_CERTIFICATE_FILE>
   ```

2. Install the file on the containers at `/etc/ssl/cert.pem`.

   a. To add a volume, edit the applicable workload:

   ```
   - name: ca-volume
     configMap:
       name: ca-store
   ```

   b. Add `mountPath` below `volumeMount` for the applicable container in the workload:

   ```
   - name: ca-volume
     mountPath: /etc/ssl/cert.pem
     subPath: ca.cer
   ```

Blocked or Unreported Clusters

It is necessary to onboard each cluster with its Environment ID. If the same Environment ID is used for several clusters in parallel, then they are blocked and do not report. An example of this is seen in the agent's logs:

```
[error] api-cpx.dome9.com:443, HTTP status=403
cloudAccount marked as blocked
```

An equivalent error shows on the **Environment** page of the CloudGuard portal.

To correct the issue:

1. Offboard CloudGuard agents that share the same ID from the unnecessary cluster with this command:

   ```
   helm uninstall asset-mgmt --namespace <namespace>
   ```

2. After you offboard the unnecessary clusters, use the API request to correct multiple onboardings and make sure the issue is resolved.

Important - The API request requires special CloudGuard privileges. The credentials (username and password) in this request must belong to a CloudGuard Service Account with *Manage Resources* permission. This is opposed to the username and password used when you configured the agents, which only allows the data to be reported to the backend. To configure the privileges, follow the steps below.

To configure CloudGuard privileges for Manage Resources:

1. In the CloudGuard menu, navigate to **Settings** > **Service Accounts** and select **Add Account**.
2. In **Selected Roles**, select a role with the **Manage Resources** permissions or create a new role with these permissions.

   For more information about roles and service accounts, see *"Users & Roles" on page 842*.

Installation of Agents Fails in Clusters with OPA Gatekeeper

When OPA (Open Policy Agent) Gatekeeper is configured in a cluster with custom *block* policies, the installation of the CloudGuard agents can fail, for example, because of the required permissions. For a successful installation, you must exclude the CloudGuard agents from OPA Gatekeeper enforcement.

For this, add an exclusion of the CloudGuard agents namespace to the Gatekeeper configuration as described in the Gatekeeper documentation.

- Create Gatekeeper config with the statement below:

```
apiVersion: config.gatekeeper.sh/v1alpha1
kind: Config
metadata:
  name: config
  namespace: GATEKEEPER-NAMESPACE
spec:
  match:
    - excludedNamespaces: ["CLOUDGUARD-NAMESPACE"]
      processes: ["*"]
```

- If the Gatekeeper `config` exists, update it to include the statement below:

```
 - excludedNamespaces: ["CLOUDGUARD-NAMESPACE"]
     processes: ["*"]
```

Notes:

- Change GATEKEEPER-NAMESPACE to the Gatekeeper installation namespace.
- Change CLOUDGUARD-NAMESPACE to the CloudGuard installation namespace.

How to Enable Debugging

1. Edit the deployment, and set the debug level:

   ```
   kubectl -n <namespace> set env deployment <deployment> LOG_
   LEVEL=debug
   ```

2. Make sure the logs are correct.

How to Collect CloudGuard Container Release Information

To collect this information for more troubleshooting, download the `cloudguard-container-info-collect-v2.sh` shell script from the Download Center. Before you run the script, use this command:

```
chmod +x cloudguard-container-info-collect-v2.sh
```

Script prerequisites:

- The user must have *kubectl* and *helm* installed on the server that runs this script and `kubeconfig` context set to the related cluster.
- The user must have the correct permissions to run the *helm* and *kubectl* commands for the applicable cluster.
- These common Linux commands must be available: `rm`, `tar`, and `mkdir`.

Default assumptions:

- Helm release name: `asset-mgmt`
- Namespace: `checkpoint`

The script collects CloudGuard CRDs (Custom Resource Definitions) `./cloudguard-container-info-collect.sh`

Syntax:

```
./cloudguard-container-info-collect.sh [-h | -c | -d | -m | - n | -o | -r]"
```

Parameters:

Parameter	Description
`-h/--help`	Shows the built-in help.
`{-c | --crd} yes` `{-c | --crd} no`	Specifies if the script has to collect CloudGuard CRDs. Default value: `yes`.
`{-d | --debug}`	Runs the script in debug mode.
`{-m | --metrics}`	Specifies if the script has to collect metrics for the CloudGuard agent containers. If metrics collection is enabled, the '`kubectl exec`' runs on fluentbit containers. Default: disabled.
`{-n | --namespace} <namespace>`	Specifies the namespace.
`-o <Name of Output File>`	Specifies the custom name for the output TAR archive.

Parameter	Description
{-r | --release} <*Name of Helm Release*>}	Specifies the Helm release name.

Example

In the command below, the script collects debug information for the CloudGuard solution deployed in a custom namespace (`cp`) with custom release name (`cp-asset`). The script does not collect CloudGuard CRDs (Custom Resource Definition) and saves debug archive with the provided name (`cp-debug.tar.gz`):

```
chmod +x cloudguard-container-info-collect-v2.sh
./cloudguard-container-info-collect-v2.sh --release cp-asset --
namespace cp --crd no --o cp-debug.tar.gz
```

More Links

API Reference Guide

Gatekeeper documentation

Onboarding Container Registries

A Container Registry is a repository that stores container images. To scan your Container Registry environment with the Image Assurance capability, onboard the Container Registry to CloudGuard.

These are two options to scan your Container Registry in CloudGuard:

- Link it to a Kubernetes cluster that has the ImageScan agents scanning your registry
- Deploy ImageScan with an AWS ECS scanner (available for selected types of registry)

CloudGuard can scan these types of container registries:

- With an **AWS ECS scanner** or a **Kubernetes scanner**:
  - Azure Container Registry (ACR) - See *"Onboarding Azure Container Registry" on page 207*
  - AWS Elastic Container Registry (ECR) - See *"Onboarding AWS Elastic Container Registry" on page 214*
  - Docker Hub Container Registry - See *"Onboarding Docker Hub Container Registry" on page 226*
  - Google Cloud Container Registry (GCR) - See *"Onboarding Google Container Registry" on page 231*
  - Harbor - See *"Onboarding Harbor Registry" on page 242*
  - JFrog Artifactory - See *"Onboarding JFrog Artifactory" on page 247*
  - Nexus - See *"Onboarding Sonatype Nexus Registry" on page 253*
  - GitHub Container Registry - See *"Onboarding GitHub Container Registry" on page 258*
  - Quay.io Container Registry - See *"Onboarding Quay.io Container Registry" on page 263*
- With a **Kubernetes scanner** only:
  - Google Artifact Registry (GAR) - See *"Onboarding Google Artifact Registry" on page 237*

General Workflow

To onboard a Container Registry to CloudGuard, follow these steps on the onboarding wizard:

1. **Registry Configurations** - Configure the registry.
2. **Cluster Configurations** - In this step, it is necessary to provide the CloudGuard Service Account credentials.
3. **Environment Configurations** - In the hosting environment, select to associate the registry with a new or existing cluster. Follow the instructions to configure the environment.
4. **Onboarding Summary** - For onboarding with a Kubernetes cluster only, CloudGuard shows the full details of your newly onboarded registry and its related cluster. If the process includes updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

CloudGuard opens the onboarded registry. For onboarding validation, in the **Scanners** tab, see the status of the registry and the cluster that scans it.

The related Kubernetes cluster page shows information on the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

Inactive Container Registries

CloudGuard deletes inactive environments when a year (365 days) passed since any of the environment's agents has communicated with CloudGuard. An agent is required to communicate with CloudGuard at least once in the past.

Note - Environments with agents that communicated with errors are not removed.

Troubleshooting

Error	Corrective Actions
Failed to create registry worker	1. Make sure you created the pull secret in the same namespace where the CloudGuard agents are located. 2. Make sure you gave the same secret name on the onboarding wizard page. **Note**: If you create or update the pull secret after the agents startup, you must restart the *imagescan-engine* and *imagescan-list* pods.

Error	Corrective Actions
Failed to authenticate	1. Make sure the pull secret key name is correct and is created in the correct namespace. 2. Make sure you entered correctly the username, password, and server URL in the secret definition.

Known Limitations

- By default, CloudGuard adds to Protected Assets and scans only 10 recent images of each repository. You can change the default value with the API call (maximal number is 1000 for a JFrog Artifactory and Sonatype Nexus). For more information, see the API Reference Guide.
- Scanning Windows container images is not supported.
- For JFrog Artifactory, it can take about 20 minutes that the images start to show for the first time.
- For JFrog Artifactory and Sonatype Nexus, the maximal number of tags per repository is 1000. Container images from the repositories with more than 1000 tags are neither shown as protected assets, nor scanned. The number is limited due to extensive API calls and performance considerations.

More Links

- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*

Onboarding Azure Container Registry

To configure container registry scanning of an Azure Container Registry (ACR), you need to onboard the registry to CloudGuard.

Prerequisites

Before onboarding your Container Registry for scanning with a Kubernetes scanner, select an authentication method:

- **Service Principal** - A user identity for applications, hosted services, and automated tools to access Azure resources. This option lets CloudGuard scan Azure Container Registries from linked clusters not necessarily in Azure.
- **Managed Identity** - An identity for applications to access resources that support Microsoft Entra ID authentication. This option allows CloudGuard to scan Azure Container Registries from Azure clusters in the same tenant.

Note - Only Azure Service Principal authentication is available for onboarding with an ECS scanner.

Onboarding

To onboard a Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.
3. In the Container Registry Onboarding wizard, enter the registry details:
    a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
    b. **Environment Description** - Optionally, enter a description.
    c. Select an **Organizational Unit**.
    d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS, select from the list of all AWS environments onboarded to CloudGuard.

f. **Choose Registry type** - Select **Azure Container Registry (ACR)**.

g. **Registry URI** - Enter the approved endpoint name of your Azure Registry in `<acrName>.azurecr.io` format.

h. **Authentication Method** - For more information, see the Azure documentation in *"More Links" on page 212*.

- **Azure Service Principal**

  i. **Pull Secret Name** - Create an image pull secret in the same namespace where the Image Assurance agents are deployed and enter it in this field.

  Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

  To create the secret, run:

```
kubectl create secret docker-registry <secret-
name> \
    --namespace <namespace> \
    --docker-server=<container-registry-
name>.azurecr.io \
    --docker-username=<service-principal-ID> \
    --docker-password=<service-principal-password>
```

  ii. **Tenant ID** - Enter your Azure AD tenant ID.

  To get the tenant ID, run:

```
az account list --query [].tenantId --output tsv
```

- **Azure Managed User Identity**

  - **Application client ID** - Enter the ID.

  To get the Application client ID, run:

```
az aks show --resource-group <resource-group> -
-name <cluster-name>
--query
identityProfile.kubeletidentity.clientId --
output tsv
```

  Note - Make sure that the role has the permissions to pull from the registry (AcrPull). For this, use the `check-acr` command (see https://learn.microsoft.com/en-us/cli/azure/aks?view=azure-cli-latest#az-aks-check-acr).

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with a Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with an ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <ACR_URI>
   - **Value**: <service-principal-ID>: <service-principal-password>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.

f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:

- **Environment ID**
- **CloudGuard API Key ID**
- **CloudGuard API Key Secret**

Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:

- **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
- **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In **AWS**, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- Azure documentation:

- How to use Service Principal credentials for image pull secret
- How to find the Microsoft Entra tenant ID
- How to get the Application client ID (Configuring ACR Integration)

Onboarding AWS Elastic Container Registry

To configure container registry scanning of an AWS Elastic Container Registry (ECR), you need to onboard the registry to CloudGuard.

Prerequisites

Before onboarding your Container Registry for scanning, select a type of the hosting environment and an applicable authentication method:

For the Kubernetes scanner

- **AWS User Access Key** - Standard AWS user access key with the policy *AmazonEC2ContainerRegistryReadOnly* required for read-only access to the applicable ECR.
- **AWS Node Group Role** - For an EKS cluster created in the AWS environment, CloudGuard needs the IAM role of the EKS node to access the AWS services.

  To use this option, make sure that:

  - Your host cluster is an EKS cluster.
  - Your EKS cluster is on the same AWS account as the registry.
  - The Amazon EKS worker node IAM role (*NodeInstanceRole*) has sufficient permissions to access the ECR.
  - Metadata API is enabled on the EC2 node. Enable it again if it is disabled. For more details, see Amazon EC2 documentation.

  To learn more about these requirements, see *"AWS Node Group Role for Amazon ECR" on page 220*.
- **AWS Custom Role for Amazon ECR (Kubernetes)** - For an EKS cluster created in the AWS environment, CloudGuard needs the custom IAM role to scan ECR registries in the same or across different AWS accounts.

  To use this option, make sure that:

  - Your host cluster is an EKS cluster.
  - You have a custom role with sufficient permissions to access the ECR.
  - The IAM role *NodeInstanceRole* attached to the EKS cluster has the `sts:AssumeRole` permissions.
  - Metadata API is enabled on the EC2 node. If it is disabled, enable it again. For more details, see Amazon EC2 documentation.
  - Minimum required agent version: 2.23.0

To learn more about these requirements, see *"AWS Custom Role for Amazon ECR (Kubernetes)" on page 220*.

For the ECS scanner

- **AWS User Access Key** - Standard AWS user access key with the policy *AmazonEC2ContainerRegistryReadOnly* required for read-only access to the applicable ECR.
- **AWS ECS Task Role** - For an ECS cluster created in the AWS environment, CloudGuard needs the IAM role of the ECS cluster to access the AWS services.

  To use this option, make sure that:

  - Your host cluster is an ECS cluster.
  - Your ECS cluster is on the same AWS account as the registry.
  - The Amazon ECS task IAM role (*TaskRole*) has sufficient permissions to access the ECR.
  - Minimum required agent version: 2.28.0.
- **AWS Custom Role for Amazon ECR (AWS ECS)** - For an ECS cluster created in the AWS environment, CloudGuard needs the custom IAM role to scan ECR registries in the same or across different AWS accounts.

  To use this option, make sure that:

  - Your host cluster is an ECS cluster.
  - You have a custom role with sufficient permissions to access the ECR.
  - The IAM role *TaskRole* attached to the ECS cluster has the `sts:AssumeRole` permissions
  - Minimum required agent version: 2.28.0.

  To learn more about these requirements, see *"Special Roles" on page 220*.

Onboarding

To onboard a Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.
3. In the Container Registry Onboarding wizard, enter the registry details:

a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.

b. **Environment Description** - Optionally, enter a description.

c. Select an **Organizational Unit**.

d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS, select from the list of all AWS environments onboarded to CloudGuard.

f. **Choose Registry type** - Select the **AWS Elastic Container Registry (ECR)**.

g. **Registry URI** - Enter the approved endpoint name of your ECR in this format: `<aws_account_id>.dkr.ecr.<region>.amazonaws.com`.

   Important - Make sure the URI does not contain `/<subfolder>` after `amazonaws.com`.

h. **Authentication Method** - Select one of the methods:

- **AWS User Access Keys**

  Pull Secret Name - Create the Kubernetes secret in the same namespace where the Check Point Image Assurance agents are deployed. The secret must contain the image pull credentials.

  Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

  To create the secret, run:

  ```
  kubectl create secret docker-registry <secret-name>
  \
  --namespace <namespace> \ # must be the same
  namespace as the CloudGuard agent
  --docker-server=<registry-uri> \ # <aws_account_
  id>.dkr.ecr.<region>.amazonaws.com
  --docker-username=<AWS_ACCESS_KEY> \
  --docker-password=<AWS_SECRET_KEY>
  ```

- **AWS Node Group Role** or **AWS ECS Task Role**
- **AWS Custom Role**

  Role ARN - Use the ARN of the custom role created in *"Special Roles" on page 220*.

  To use these methods, make sure the hosting cluster satisfies all the prerequisites in *"Prerequisites" on page 214*.

4. Click **Next** to continue to **Cluster Configurations**.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the AWS ECS scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <ECR_URI>
   - **Value**: <AWS_ACCESS_KEY>:<AWS_SECRET_KEY>

5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
    a. On the **Stacks** page, click **Create stack**.
    b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
    c. For **Template source**, select **Amazon S3 URL**.
    d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
    e. In **Step 2 Specify stack details**, enter a name for the stack.
    f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
        - **Environment ID**
        - **CloudGuard API Key ID**
        - **CloudGuard API Key Secret**

        Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:
        - **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
        - **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy
    g. In AWS, enter these details:
        - **Subnet** - Select a subnet.
        - **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
        - **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.
7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

Special Roles

For the Kubernetes scanner:

AWS Node Group Role for Amazon ECR

The worker node running on your hosting cluster needs the IAM permissions to access the ECR. Kubernetes clusters created with automation like EKS ETL have these permissions by default. Kubernetes clusters created manually may not have the permissions, so you have to add them.

To verify the cluster configuration:

1. Open the Amazon Elastic Kubernetes Service console.
2. Select the cluster to use as a hosting cluster.
3. From the **Compute** tab, add a new node group or select an existing one.
4. In the selected node group, go to the **Details** tab.
5. Below **Node IAM role ARN**, click the ARN link to open the Node Group Role configuration for the attached Role in the EKS cluster.
6. In the **Permissions** tab, make sure that the *AmazonEC2ContainerRegistryReadOnly* managed policy is attached to the role.

AWS Custom Role for Amazon ECR (Kubernetes)

The worker node running on your hosting cluster needs the custom IAM permissions to access an ECR.

1. On your AWS console, open the IAM Service.
2. Create a custom IAM role on the AWS account of your ECR:

   Sample Role Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecr:GetDownloadUrlForLayer",
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchGetImage",
```

```
        "ecr:DescribeImages",
        "ecr:DescribeRepositories",
        "ecr:ListImages",
        "ecr:ListTagsForResource"
        ],
      "Resource": "arn:aws:ecr:*:<AWS-Account-ID where ECR
located>:repository/*"
    },
    {
      "Effect": "Allow",
      "Action": ["ecr:GetAuthorizationToken"],
      "Resource": "*"
    }
  ]
}
```

3. Create a role trust relationship to give the EKS account access to the ECR account:

   Sample Role Trust Relationship

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<AWS-Account-ID where EKS
Located>:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

4. Open the Amazon Elastic Kubernetes Service console.
5. Select the cluster to use as a hosting cluster.
6. On the **Compute** tab, add a new node group or select an existing one.
7. In the **Node group** section, select the created node group or the default node group.
8. In the selected node group, go to the **Details** tab.

9. Below **Node IAM role ARN**, click the ARN link to open the Node Group Role configuration for the attached Role in the EKS cluster.
10. In the **Permissions** tab, click **Add permissions** to create a new IAM policy.
11. Select **Create inline policy**, click **JSON** to edit the policy, and add the sample policy provided below.

    Sample Policy

    This policy uses an ARN of the *Custom-IAM-Role* that you created in Step 2.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["sts:AssumeRole"],
      "Resource": "<ARN-Custom-IAM-Role>"
    }
  ]
}
```

12. Click **Next**.
13. Give a name to the policy and click **Create policy**. The policy is added and attached to the IAM Node Group Role.

For the AWS ECS scanner:

AWS ECS Task Role for Amazon ECR

The ECS Task Role attached to the Service needs the IAM permissions to access the ECR. ECS clusters created manually may not have the permissions, so you have to add them.

To verify the cluster configuration:

1. Open the Amazon Elastic Container Service console.
2. Select the cluster to use as a hosting cluster.
3. From the **Task** tab, select one of the tasks.
4. In the Configuration panel, click **Task definition**.

5. Below **Task role**, click the link to open the Task Role configuration for the attached Role in the ECS cluster.
6. In the **Permissions** tab, make sure that the *AmazonEC2ContainerRegistryReadOnly* managed policy is attached to the role.

AWS Custom Role for Amazon ECR (AWS ECS)

The ECS Task Role attached to the Service needs the custom IAM permissions to access the ECR.

1. On your AWS console, open the IAM Service.
2. Create a custom IAM role on the AWS account of your ECR:

   Sample Role Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecr:GetDownloadUrlForLayer",
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchGetImage",
        "ecr:DescribeImages",
        "ecr:DescribeRepositories",
        "ecr:ListImages",
        "ecr:ListTagsForResource"
        ],
      "Resource": "arn:aws:ecr:*:<AWS-Account-ID where ECR
located>:repository/*"
    },
    {
      "Effect": "Allow",
      "Action": ["ecr:GetAuthorizationToken"],
      "Resource": "*"
    }
  ]
}
```

3. Create a role trust relationship to give the EKS account access to the ECR account:

   Sample Role Trust Relationship

```
{
```

```
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<AWS-Account-ID where EKS
Located>:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

4. Open the Amazon Elastic Container Service console.
5. Select the cluster to use as a hosting cluster.
6. From the **Task** tab, select one of the tasks.
7. In the Configuration panel, click **Task definition**.
8. Below **Task role**, click the link to open the Task Role configuration for the attached Role in the ECS cluster.
9. In the **Permissions** tab, click **Add permissions** to create a new IAM policy.
10. Select **Create inline policy**, click **JSON** to edit the policy, and add the sample policy provided below.

    Sample Policy

    This policy uses an ARN of the *Custom-IAM-Role* that you created in Step 2.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["sts:AssumeRole"],
      "Resource": "<ARN-Custom-IAM-Role>"
    }
  ]
}
```

11. Click **Next**.
12. Give a name to the policy and click **Create policy**. The policy is added and attached to the IAM Node Group Role.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*

Onboarding Docker Hub Container Registry

To configure container registry scanning of a Docker Hub Container Registry environment, you need to onboard the environment to CloudGuard.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment (Kubernetes or ECS Scanner).
- If required, configure a CA certificate for the registry - see *"Configuring CA Certificate" on page 268*.
- For authentication with the Docker Hub container registry, create an access token and set its access permissions to **Read & Write**. For more information, see the Docker documentation.

Onboarding

To onboard a Docker Hub Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in Kubernetes cluster scanning environment, you can open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

- For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
- For AWS, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **DockerHub Registry**.

g. **Registry URI** - Use a URI in the format `registry_url/<namespace>`.

h. **Authentication Method - Docker Hub Access Token** - For the Kubernetes scanner, enter the details below each method. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

Pull Secret Name - Enter the image pull secret name that you create on the hosting cluster with your credentials.

Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

To create the secret, run:

```
kubectl create secret docker-registry <pull-secret-
name>
--namespace <namespace>
--docker-server=hub.docker.com/<namespace>
--docker-username=<robot-account-user-name>
--docker-password=<robot-account-token>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:

- Select an existing Service Account with its corresponding API Key.
- Enter a Service Account manually.
- Click **Add Service Account** to create a new account.

2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:

- **Secret type**: **Other type of secret**
- **Key**: <registry_URI>
- **Value**: <docker_username>: <docker_access_token>

5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.
   f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
      - **Environment ID**
      - **CloudGuard API Key ID**
      - **CloudGuard API Key Secret**
   g. In AWS, enter these details:
      - **Subnet** - Select a subnet.
      - **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
      - **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.
7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- API Reference Guide

Onboarding Google Container Registry

CloudGuard lets you onboard two types of registries on the Google Cloud Platform:

- Google Container Registry (GCR)
- Google Artifact Registry (GAR)

This topic describes how to onboard your GCR to CloudGuard.

Prerequisites

Before onboarding your Container Registry to CloudGuard, select an authentication method:

- **GCP Service Account Key** - GCP Service Account key file with the relevant permissions required for the relevant GCR.
- **GCP GKE Internal Authentication** - If you have a Google Kubernetes Engine (GKE) cluster, CloudGuard can use the Compute Engine or GKE metadata server to authenticate with GCR.

  To use this option, make sure that the service account of the GKE node pool has permissions to access the container registry.

  This requirement is met by default when you work in the same project and use the default service account. You must set the required permissions for the service account if:

  - The GKE cluster is in a different project
  - The cluster uses a different service account

Notes:

- The option of GKE internal authentication is not supported if GKE is configured with Workload Identity.
- Only GCP Service Account Key authentication is available for onboarding with an AWS ECS scanner.

Onboarding

To onboard a Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.
3. In the Container Registry Onboarding wizard, enter the registry details:

a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard

b. **Environment Description** - Optionally, enter a description.

c. Select an **Organizational Unit**.

d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS, select from the list of all AWS environments onboarded to CloudGuard.

f. **Choose Registry type** - Select **Google Cloud Container Registry (GCR)**.

g. **Registry URI** - Select one of the approved endpoint names of your GCR, based on your region.

h. **Authentication Method** - Select one of the methods:

- **GCP Service Account Key**

  i. Create a service account with the needed permissions (minimal required role: *roles/browser*) and generate a json key for it. To do this, open a Google Cloud Shell terminal and run the commands below. This creates a service account `gcp-svc-acc` with the needed permissions to access GCR and generates the file `gcp-svc-acc-keyfile.json` to use for an image pull secret.

```
gcloud iam service-accounts create gcp-svc-acc
gcloud iam service-accounts keys create gcp-
svc-acc-keyfile.json --iam-account gcp-svc-
acc@<your_project_name>.iam.gserviceaccount.com
gcloud projects add-iam-policy-binding <your_
project_name> --member="serviceAccount:gcp-svc-
acc@<your_project_
name>.iam.gserviceaccount.com" --
role="roles/browser"
gsutil iam ch serviceAccount:gcp-svc-acc@<your_
project_
name>.iam.gserviceaccount.com:objectViewer
gs://artifacts.<your_project_name>.appspot.com/
```

  ii. **Pull Secret Name** - Create the image pull secret on your hosting cluster where the Check Point Image Assurance agents are deployed (can be done later).

  Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

  To create the secret, run:

```
kubectl create secret docker-registry <secret-
name> \
--namespace <cloudguard-namespace> \
--docker-server=<google-registry-uri> \
--docker-username=_json_key \
--docker-password="$(cat gcp-svc-acc-
keyfile.json)"
```

  iii. Enter the pull secret name in the CloudGuard onboarding wizard.

- **GCP GKE Internal Authentication**

  To use this method, make sure the hosting cluster satisfies all the prerequisites in *"Prerequisites" on page 231*. When you select this option, no more configuration is required.

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with a Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with an ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <GCR_URI>
   - **Value**: _json_key: <json-secret>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.

f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:

- **Environment ID**
- **CloudGuard API Key ID**
- **CloudGuard API Key Secret**

Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:

- **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
- **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, in the **Scanners** tab, see the status of the registry and the cluster that scans it.

The related Kubernetes cluster page shows information on the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Onboarding Google Artifact Registry" on page 237*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- Google Cloud documentation: Service account JSON key file

Onboarding Google Artifact Registry

CloudGuard lets you onboard two types of registries on the Google Cloud Platform:

- Google Container Registry (GCR)
- Google Artifact Registry (GAR)

This topic describes how to onboard your GAR to CloudGuard.

Prerequisites

- You must have a Kubernetes cluster onboarded to CloudGuard before you scan your container registry. This hosting cluster environment must have Image Assurance enabled. Otherwise, CloudGuard instructs you to install Image Assurance agents (or to update the agents to the latest version) on that cluster as part of the onboarding process.
- Before onboarding your Container Registry to CloudGuard, select an authentication method:
  - **GCP Service Account Key** - GCP Service Account key file with the relevant permissions required for the relevant GAR.
  - **GCP GKE Internal Authentication** - If you have a Google Kubernetes Engine (GKE) cluster, CloudGuard can use the Compute Engine or GKE metadata server to authenticate with GAR.

    To use this option, make sure that the service account of the GKE node pool has permissions to access the container registry.

    This requirement is met by default when you work in the same project and use the default service account. You must set the required permissions for the service account if:
    - The GKE cluster is in a different project.
    - The cluster uses a different service account.

Note - The option of GKE internal authentication is not supported if GKE is configured with Workload Identity.

Onboarding

To onboard a Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

Alternatively, open the hosting cluster page and click **Scan Registry** on the top menu.

3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes**. Only this scanner type is available for Google Artifact Registry.
   e. Select a Kubernetes cluster on which you can run the registry scanner.
      - Select from the list of onboarded clusters with enabled Image Assurance.
      - For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   f. **Choose Registry type** - Select **Google Cloud Artifact Registry (GAR)**
   g. **Registry URI** - Enter one of the approved endpoint names of your GAR based on your region. See the list of valid regions in the Google documentation.

      Example: `us-central1-docker.pkg.dev`

h. **Authentication Method** - Select one of the methods:

- **GCP Service Account Key**

  i. Create a service account with the needed permissions (minimal required roles: *roles/artifactregistry.reader* and *roles/browser*) and generate a json key for it. To do this, open a Google Cloud Shell terminal and run the commands below. This creates a service account `gcp-svc-acc` with the needed permissions to access GCR and generates the file `gcp-svc-acc-keyfile.json` to use for an image pull secret.

```
gcloud iam service-accounts create gcp-svc-acc

gcloud iam service-accounts keys create gcp-svc-
acc-keyfile.json --iam-account gcp-svc-acc@<your_
project_name>.iam.gserviceaccount.com

gcloud projects add-iam-policy-binding <your_
project_name> --member="serviceAccount:gcp-svc-
acc@<your_project_name>.iam.gserviceaccount.com" -
-role="roles/browser"

gcloud projects add-iam-policy-binding <your_
project_name> --member="serviceAccount:gcp-svc-
acc@<your_project_name>.iam.gserviceaccount.com" -
-role="roles/artifactregistry.reader"
```

  ii. **Pull Secret Name** - Create the image pull secret on your hosting cluster under the same namespace that are used for the Image Assurance agents.

  Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

  To create the secret, run:

```
kubectl create secret docker-registry <secret-
name> --namespace <cloudguard-namespace> --
docker-server=<google-registry-uri> --docker-
username=_json_key --docker-password="$(cat
gcp-svc-acc-keyfile.json)"
```

  iii. Enter the pull secret name in the CloudGuard onboarding wizard.

- **GCP GKE Internal Authentication**

  To use this method, make sure the hosting cluster satisfies all the prerequisites in *"Prerequisites" on page 237*. When you select this option, no more configuration is required.

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

CloudGuard opens the onboarded registry. For onboarding validation, in the **Scanners** tab, see the status of the registry and the cluster that scans it.

The related Kubernetes cluster page shows information on the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Onboarding Google Container Registry" on page 231*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- Google Cloud documentation:
  - Service account JSON key file
  - Artifact Registry regions

Onboarding Harbor Registry

To configure container registry scanning of a Harbor Registry environment, you need to onboard the Harbor Registry to CloudGuard.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment and an applicable authentication method.
- You must provide a Certificate Authority (CA) certificate to CloudGuard resources deployed on your Kubernetes cluster. For more details, see *"Configuring CA Certificate" on page 268*.

Onboarding

To onboard a Harbor Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in the Kubernetes cluster scanning environment, open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS ECS Scanner, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **Harbor**.

g. **Registry URI** - Enter the FQDN of your Harbor server endpoint, **without the protocol (https)**.

h. **Authentication Method - Harbor Credentials** - For the Kubernetes scanner, enter the details below. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

   Pull Secret Name - Create a secret on the cluster with credentials of a user with at least a **Limited Guest** role, for each project to scan. For `Harbor-URL`, use the same Registry URI provided above for onboarding the container registry to CloudGuard, without the protocol (https).

   Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

   To create the secret, run:

```
kubectl create secret docker-registry <secret-name> \
    --namespace <cloudguard-namespace> \
    --docker-server=<Harbor-URL> \
    --docker-username=<username> \
    --docker-password=<password>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.

4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <registry_URI>
   - **Value**: <HARBOR_USERNAME>: <HARBOR_PASSWORD>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.
   f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
      - **Environment ID**
      - **CloudGuard API Key ID**
      - **CloudGuard API Key Secret**

      Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:
      - **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
      - **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*

Onboarding JFrog Artifactory

To configure container registry scanning of a JFrog Artifactory environment, you need to onboard the Artifactory to CloudGuard. It supports the self-hosted Artifactory and the cloud-based solution provided by JFrog. CloudGuard discovers all types of JFrog Artifactory Docker repositories (local, remote, and virtual) and scans images in those repositories.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment and an applicable authentication method.
- CloudGuard uses HTTPS connection to the JFrog Artifactory registry. Note that by default a self-hosted registry is configured for HTTP only. For more information, see JFrog Artifactory Documentation.
- If required, configure a CA certificate for the registry - see *"Configuring CA Certificate" on page 268*.
- For authentication with the JFrog Artifactory Docker repositories, it is necessary to have a JFrog user with *Read* permissions. CloudGuard discovers and scans all repositories to which this user has access.

Onboarding

To onboard a JFrog Artifactory to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in Kubernetes cluster scanning environment, open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS ECS Scanner, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **JFrog Artifactory**.

g. **Registry URI** - Enter the FQDN of your registry endpoint, **without the protocol (https)**:

   - For Artifactory instances hosted in the JFrog Cloud, use a URI in the format `<company>.jfrog.io`.
   - For self-hosted Artifactory instances, use a URI with a resolvable host name or the IP of your JFrog Artifactory instance.

h. **Authentication Method** - With at least *read* permissions, select one of the authentication methods below.

For the Kubernetes scanner, enter the details below each method. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

- **JFrog Artifactory Basic Authentication**

  Pull Secret Name - Enter the image pull secret name that you create on the hosting cluster with your credentials.

  Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

  To create the secret, run:

```
kubectl create secret docker-registry <secret-name> \
    --namespace <cloudguard-namespace> \
    --docker-server=<artifactory-registry-URI>\
    --docker-username=<username> \
    --docker-password=<password>
```

- **JFrog Artifactory Access Token**

  Enter the token that you generate per user or with Admin permission:

  a. In the JFrog platform, navigate to **Platform Configuration** and go to **User Management**.

  b. Configure a **Scoped** Access Token with the **Reference Token** option.

  c. Configure a secret on your cluster with the received Reference Token. For this, run:

```
kubectl create secret docker-registry <secret-
name> \
    --namespace <cloudguard_namespace> \
    --docker-server=<artifactory-registry-URI>\
    --docker-username=jfrog \
    --docker-password=<reference-token>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with an existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:
    - **Secret type**: **Other type of secret**
    - **Key**: <registry_URI>
    - **Value**: <JFROG_USERNAME>: <JFROG_PASSWORD>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
    a. On the **Stacks** page, click **Create stack**.
    b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
    c. For **Template source**, select **Amazon S3 URL**.
    d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
    e. In **Step 2 Specify stack details**, enter a name for the stack.

f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:

- **Environment ID**
- **CloudGuard API Key ID**
- **CloudGuard API Key Secret**

Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:

- **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
- **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- How to configure a Reverse Proxy in JFrog
- How to generate a Scoped Token in JFrog
- API Reference Guide

Onboarding Sonatype Nexus Registry

To configure container registry scanning of a Sonatype Nexus environment, you need to onboard the environment to CloudGuard. CloudGuard discovers only the *hosted* type of Sonatype Nexus Docker repositories and scans images in these repositories only.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment and an applicable authentication method.
- CloudGuard uses HTTPS connection to the Sonatype Nexus registry.
- You must provide a Certificate Authority (CA) certificate to CloudGuard resources deployed on your Kubernetes cluster or AWS ECS environment. For more details, see *"Configuring CA Certificate" on page 268*.
- For authentication with the Sonatype Nexus Docker repositories, it is necessary to have a Sonatype user with *Read* permissions. CloudGuard discovers and scans all repositories to which this user has access.

Onboarding

To onboard a Sonatype Nexus Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in Kubernetes cluster scanning environment, open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS ECS Scanner, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **Nexus**.

g. **Registry URI** - Enter the FQDN of your Nexus server endpoint, **without the protocol (https)**.

h. **Authentication Method - Nexus Basic Authentication** - For the Kubernetes scanner, enter the details below. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

   Pull Secret Name - Enter the image pull secret name that you create on the hosting cluster with your credentials.

   Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

   To create the secret, run:

```
kubectl create secret docker-registry <secret-name> \
    --namespace <cloudguard-namespace> \
    --docker-server=<nexus_registry_URI> \
    --docker-username=<username> \
    --docker-password=<password>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.

4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <registry_URI>
   - **Value**: <NEXUS_USERNAME>: <NEXUS_PASSWORD>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.
   f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
      - **Environment ID**
      - **CloudGuard API Key ID**
      - **CloudGuard API Key Secret**

      Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:
      - **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
      - **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- API Reference Guide

Onboarding GitHub Container Registry

To configure container registry scanning of a GitHub Container Registry environment, you need to onboard the environment to CloudGuard.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment.
- If required, configure a CA certificate for the registry - see *"Configuring CA Certificate" on page 268*.
- For authentication with the GitHub container registry, it is necessary to have a GitHub user with *Read* permissions. CloudGuard discovers and scans all repositories to which this user has access.

Onboarding

To onboard a GitHub Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in Kubernetes cluster scanning environment, you can open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

   - For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
   - For AWS ECS Scanner, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **GitHub Container Registry**.

g. **Registry URI** - This field is not available for editing.

h. **Authentication Method - GitHub Container Registry Personal Access Token** - For the Kubernetes scanner, enter the details below. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

   i. **Create a personal access token (classic)** - see the GitHub Container Registry Documentation

   ii. **Pull Secret Name** - Enter the image pull secret name that you create on the hosting cluster with your credentials.

   Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

   To create the secret, run:

```
kubectl create secret docker-registry <secret-name> \
  --namespace <cloudguard-namespace> \
  --docker-server=ghcr.io \
  --docker-username=github \
  --docker-password=<personal-access-token>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.

4. Click **Store a new secret** to create an image pull secret with:
   - **Secret type**: **Other type of secret**
   - **Key**: <registry_URI>
   - **Value**: <GITHUB_USERNAME>: <GITHUB_PASSWORD>
5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In the CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process**:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.
   f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
      - **Environment ID**
      - **CloudGuard API Key ID**
      - **CloudGuard API Key Secret**

      Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:
      - **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
      - **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- API Reference Guide

Onboarding Quay.io Container Registry

To configure container registry scanning of a Quay.io Container Registry environment, you need to onboard the environment to CloudGuard. It supports the SaaS solution provided by Quay.io for private repositories in an organization with the quay.io/organization pattern.

Prerequisites

- Before onboarding your Container Registry for scanning, select a type of hosting environment (Kubernetes or ECS Scanner).
- For authentication with the Quay.io container registry, create a *robot account*. For more information, see the Quay documentation. CloudGuard discovers and scans all repositories to which this user has the *Read* access.

Onboarding

To onboard a Quay Container Registry to CloudGuard:

STEP 1 - Registry Configurations

1. In the CloudGuard portal, navigate to **Asset** > **Environments**.
2. From the top menu, select **Add** > **Container Registry** and follow the setup steps.

   Alternatively, in Kubernetes cluster scanning environment, you can open the hosting cluster page and click **Scan Registry** on the top menu.
3. In the Container Registry Onboarding wizard, enter the registry details:
   a. **Environment Name** - Enter a new name for the registry or use the default name. This name allows you to identify the registry later in CloudGuard.
   b. **Environment Description** - Optionally, enter a description.
   c. Select an **Organizational Unit**.
   d. Select the type of environment to host your scanner - **Kubernetes** or AWS **ECS Scanner**.

e. Select a Kubernetes cluster or an AWS environment on which you can run the registry scanner:

- For Kubernetes, select from the list of clusters with enabled Image Assurance. For a new cluster, click **Onboard a new Kubernetes Cluster** and see *"Onboarding Kubernetes Clusters" on page 188*. In this case, you quit the registry onboarding and, after onboarding a new cluster, you need to start the registry onboarding from the beginning.
- For AWS, select from the list of all AWS environments onboarded to CloudGuard.

f. **Registry Type** - Select **Quay.io Container Registry**.

g. **Registry URI** - Use a URI in the format `quay.io/<organization>`.

h. **Authentication Method - Quay Access Token** - For the Kubernetes scanner, enter the details below each method. For the AWS ECS scanner, only select the method and enter the details later in Step 3.

Pull Secret Name - Enter the image pull secret name that you create on the hosting cluster with your credentials.

Make sure that the `<secret-name>` is a valid Kubernetes name. For more details, see the Kubernetes Documentation.

To create the secret, run:

```
kubectl create secret docker-registry <pull-secret-
name>
--namespace <namespace>
--docker-server=quay.io/<organization>
--docker-username=<robot-account-user-name>
--docker-password=<robot-account-token>
```

4. Click **Next** to continue with Cluster Configurations.

STEP 2 - Cluster Configurations

In this step, you configure the CloudGuard Service Account credentials if in Step 1 you selected to onboard with a new cluster or with the existing cluster that requires an agent update.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

1. Configure a Service Account by one of these methods:

- Select an existing Service Account with its corresponding API Key.
- Enter a Service Account manually.
- Click **Add Service Account** to create a new account.

2. Click **Next** to continue to the next step.

For Onboarding with the Kubernetes Scanner

STEP 3 - Instructions

This step appears when you select to associate the registry with a new cluster or with an existing cluster that requires an agent update. CloudGuard instructs you how to install Image Assurance agents or to update them to the latest version on the cluster.

For onboarding with the hosting cluster that has updated agents, this step is skipped.

CloudGuard shows the details of your new registry and its related cluster.

1. Follow the on-screen instructions to copy the Helm commands and run them on your cluster with Helm 3.
2. Click **Next**.

STEP 4 - Onboarding Summary

CloudGuard shows the full details of your new registry and its related cluster. If your registry onboarding includes onboarding or updating the cluster, this page shows the cluster onboarding summary. The cluster deployment takes several minutes, and you can see its progress in the Cluster and Registry Status.

For more information on the cluster onboarding summary, see *"STEP 4 - Onboarding Summary" on page 189*.

- Wait for the deployment completion based on the Cluster and Agent Status or click **Finish** to skip the process.

For Onboarding with the ECS Scanner

STEP 3 - AWS Configurations

Follow the on-screen instructions to use the provided CloudFormation Template and launch the CFT for the ECS scanner.

1. Select to use a new ECS cluster or an existing one.
2. Use the URL to review the CloudFormation Template.
3. Open the AWS Secrets Manager and click **Secrets**.
4. Click **Store a new secret** to create an image pull secret with:

- **Secret type**: **Other type of secret**
- **Key**: <registry_URI>
- **Value**: <QUAY_USERNAME>: <QUAY_PASSWORD>

5. Open the image pull secret and copy **Secret ARN** from **Secret details**. You need this ARN in step **6g**.
6. In CloudGuard wizard, click the link in step 4 to start the **CloudFormation Stack Creation Process** in your AWS account:
   a. On the **Stacks** page, click **Create stack**.
   b. In **Step 1 Create stack**, for **Prepare template**, select **Choose an existing template**.
   c. For **Template source**, select **Amazon S3 URL**.
   d. In the **Amazon S3 URL** field, paste the URL you copied in step 2 and click **Next**.
   e. In **Step 2 Specify stack details**, enter a name for the stack.
   f. In **Parameters > CloudGuard**, paste these details copied from step 5 of the CloudGuard wizard:
      - **Environment ID**
      - **CloudGuard API Key ID**
      - **CloudGuard API Key Secret**

      Optionally, you can configure a proxy server and enter these details for the proxy address and proxy bypass:
      - **Optional - HTTPS Proxy address** - Enter an HTTPS address for a network proxy server
      - **Optional - Proxy bypass list** - Enter one or more addresses that should bypass the proxy

g. In AWS, enter these details:

- **Subnet** - Select a subnet.
- **Optional - Registry Secret ARN** - Enter the ARN of the secret created in step 3.
- **Optional - Custom CA Certificates ARN** - see *"Certificate for AWS ECS Scanner" on page 268*.

7. After the creation of the stack, click **Finish**.

CloudGuard opens the onboarded registry. For onboarding validation, see the **Scanners** tab that shows the status of the registry and its scanning environment (cluster or AWS ECS).

For registries with the Kubernetes scanner, the related Kubernetes cluster page shows information about the registries that the cluster scans, in the list on **Blades** > **Image Assurance** > **Image Scan Engine agent**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Container Registry Scanning" on page 464*
- *"Image Assurance" on page 434*
- *"Onboarding Kubernetes Clusters" on page 188*
- API Reference Guide

Configuring CA Certificate

To scan your container registries, CloudGuard uses a bundle of public Certificate Authorities (CA) for authentication. The CloudGuard scanning agents accept certificates signed only by these trusted CAs. If you use a self-hosted registry, the CA that signs your registry certificate can be absent in the bundle of the trusted CAs. Follow the steps below to provide the CA to CloudGuard.

Certificate for Kubernetes Scanner

To configure the Kubernetes cluster:

1. **Obtain a registry CA certificate in Base64 format**

   Use one of the available methods; for example, download the certificate with your web browser from the registry's website.

2. Create a configmap registry `registry-ca-bundle` in the CloudGuard namespace (by default, *checkpoint*). Use the `registry.cer` file as the key. Use the value as the CA certificate file content.

   ```
   kubectl create configmap registry-ca-bundle --from-
   file=registry.cer=<certificate_file_path> --namespace
   <cloudguard_namespace>
   ```

Notes:

- For clusters that scan multiple registries with different CA certificates, the certificate file must contain a bundle of certificates of all your registries. You can add CA bundles of all the relevant CAs.
- If the registry CA is replaced, update the configmap and restart the Imagescan pods.

Certificate for AWS ECS Scanner

Configuring the certificate includes two stages:

1. Prepare a file that contains a bundle of CA certificates and store its content as a secret.
2. Add the secret to the CloudFormation stack of the scanner.

Note - For AWS ECS scanners that scan multiple registries with different CA certificates, the certificate file must contain a bundle of certificates of all your registries. You can add CA bundles of all the relevant CAs to your secret's content.

To configure a secret:

1. Open the AWS Secrets Manager, go to **Secrets**, and click **Store a new secret** to create it:
   - **Secret type**: **Other type of secret**
   - In **Key/value pairs**, select **Plaintext**.
   - Paste the contents of one or more certificates in text format as in the example below:

```
JFrog Registry CA
==================================
-----BEGIN CERTIFICATE-----
<PEM-Encoded-Certificate-Content>
-----END CERTIFICATE-----
Harbor Registry CA
==================================
-----BEGIN CERTIFICATE-----
<PEM-Encoded-Certificate-Content>
-----END CERTIFICATE-----
```

2. Configure the rest of the parameters based on your preferences and click **Store**.
3. Open the new secret and copy **Secret ARN** from **Secret details**. You need this ARN for the next stage.

To add the secret to the stack in AWS CFT:

1. Edit the stack used for the ECS scanner.
2. In the **Specify stack details**, go to **Parameters > AWS > Optional - Custom CA Certificate ARN**.
3. Paste the **Secret ARN** you copied in the previous stage.

For more information about the stack creation, see **Step 3 - AWS Configurations** in the onboarding wizard for your container registry.

Protected Assets

This page shows a summary of your environments onboarded to CloudGuard. These assets can include, for example, compute services (such as EC2s, Lambdas, and containers), database services (such as RDS, SQL DB, and BigQuery), and more. After onboarding your account, CloudGuard fetches information about these assets from the environment and presents it in the portal. In addition, CloudGuard monitors the security posture of these assets with the Compliance Engine. CloudGuard can fully protect environments that support full protection, such as AWS. CloudGuard can actively make corrections, for example, apply or change a Security Group policy if its configuration is incorrect.

Benefits

CloudGuard presents one view of your cloud assets, on all platforms, from which you can search or filter for specific assets of interest and see details about their security posture.

For some asset types, you can apply Security Group or IAM policies directly from the CloudGuard portal.

Use Cases

Here are some typical use cases for the CloudGuard Protected Assets.

- Find assets matching specific criteria across all accounts and platforms - see *"Filter and Search" on page 863*
- Review attributes and status for an asset - see *"Viewing your assets" on the next page*
- Review and change the security policies for an asset
- Export inventory information to files - see *"Exporting protected asset information" on page 272*

Protected Assets Table

You can filter or search the protected assets table by asset type, region, VPC, and other conditions. By default, the page shows the assets grouping by Environment. For more information on grouping, see *"Group Arrangement" on page 122*.

Organize the table columns as necessary and adjust these parameters:

- **Visibility** - To select which columns to see in the table, click **Customize** on the right. Click a parameter to add its column to the table or search for a parameter name in the internal search bar.
- **Position** - To change the column's location, click the column header and drag it to a specific location.

- **Width** - To change the column width, move the right separator line of its header in the desired direction. To adjust the width by the longest column value, double-click the right separator.
- **Sorting** - To switch between the default, ascending or descending order of the entries, click the column header.

To restore the default settings of the table, click **Reset Columns** on the top right.

To make the columns fit the screen, click **Autofit Columns** on the top right.

Select an asset from the list to see more details. The number of details depends on the type of asset. For some assets, you can see flow logs. If your environment supports full protection mode and is managed by CloudGuard in this mode, you can change the network security settings.

You cannot set other details for your assets here; this is done in your cloud account on the cloud platform.

Actions

Viewing your assets

The primary page shows assets that are protected by CloudGuard. Use the filter to filter the list or search for assets by name in the search box.

In addition, you can see the Dashboard of your protected assets, for more details see *"Dashboards" on page 86*. The dashboard has widgets that show the distribution of your assets based on different parameters, such as region, type, environment, etc. You can change the dashboard to include specific widgets and create new custom dashboards.

Click one of the assets to see its details. For more information, see *"Asset Details" on page 273*.

Changing details (instances only)

You can change details for assets that are instances (EC2s on AWS or virtual machines on Azure or Google) if the assets are in *"Full Protection" on page 373* mode by CloudGuard.

For AWS instances:

You can add Security Group or NACL policies to AWS instances.

1. Select an instance-type asset from the list to show details for it. You can change network settings, in the Network Security or IAM Policies tabs.
2. Click **+ Attach** to attach a security group or NACL to the instance from those already configured. To configure a new security group or NACL, go to *"Security Groups" on page 366*).

3. Select the group or NACL and click **Attach**.

For Azure:

You can change the rules for Security Group applied to virtual machines. You cannot add or remove the Security Group itself.

1. Click the instance from the list.
2. Click the Subnet NSG Policy that is necessary to change and click **Edit Mode** (the security group must be set to Manage, not Read Only, to do this).
3. Click **Edit** to change a firewall rule or **Delete** to delete it. See *"Modify an Azure Network Security Group" on page 374* for details about how to change Network Security Groups (NSGs).

Viewing Flow Logs

Some assets configuration allows CloudGuard access to Flow Logs. These are marked with  . Click this icon to show the Flow Logs. See *"VPC Flow Logs" on page 382* for details about controlling this view.

Exporting protected asset information

You can export information for protected assets to a CSV file, if your filtered view has less than 10,000 results.

1. To select a view of the protected assets of interest, use the filter.
2. Click **Export** in the top right and then select if to export the basic filtered view or the view of all assets.

Asset Details

This page contains all the details of a specific asset in your onboarded environment. Use tabs on the left to navigate between different types of information about the asset. The number and type of the asset tabs can differ based on the asset type.

Main Details

The top of the screen shows the basic asset details, such as name, type, or ID. These parameters depend on the asset type. For some assets, this part shows information related to the asset's risks: business priority, risk score, or exposure.

Symbols

Based on the asset type, permissions, and connected services, asset details can have these symbols:

Icon	Meaning	Explanation
	Risk score	Read more in *"Risk Calculation" on page 101*
	IAM sensitivity	Read more in *"Entitlement Map" on page 387*
	Business priority	Read more in *"Business Priority" on page 118*
	Network exposure	Read more in *"Network Exposure" on page 105* and *"IAM Exposure" on page 112*
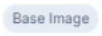	Base Image	Read more in *"Base Image" on page 426*
	Vendor Image	Read more in *"Vendor Image" on page 428*

Overview

The overview page displays important high-level information on the asset.

Risk Management

For assets supported in Risk Management, the Risk Management section shows the data considered for the calculation of the asset risk score, as well as the Context Graph and top remediation actions to mitigate the asset's risk.

Context Graph

The Context Graph is a graphical representation of the asset exposure to the Internet and to other assets in the network. It is applicable to AWS EC2 instances, Lambda functions, RDS, ECS service, and Azure Virtual Machines.

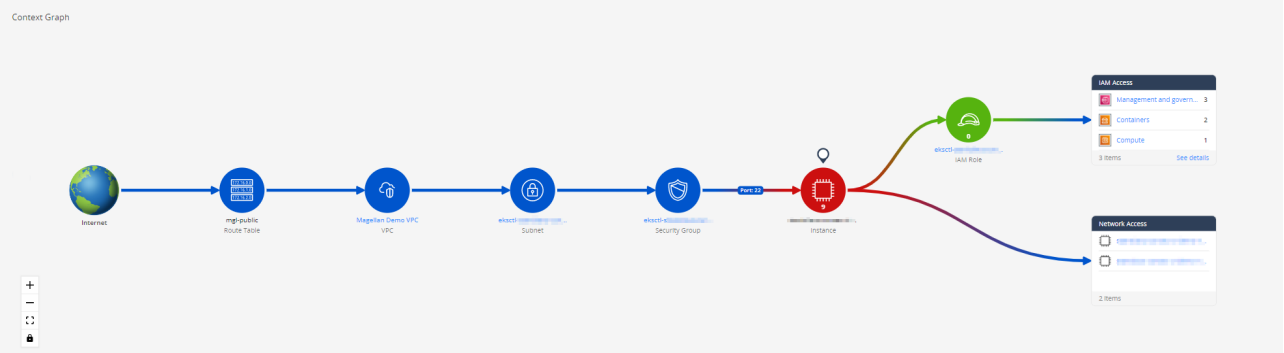

The Context Graph presents the configuration blocks that determine if the asset is exposed to the Internet or not (see *"Network Exposure" on page 105*).

When you put the cursor on a graph node, its tooltip shows important high-level information about the asset.

The Context Graph presents the potential impact that the asset exploit can have on the cloud environment, from a network and IAM perspective. The IAM impact is determined by CIEM, which analyzes the permissions granted to the asset. The network impact is based on an analysis of security groups.

The node color on the graph aligns with the risk level. Assets in blue are not supported in Risk Management and therefore do not have a risk score.

The assets with the highest business priority (*"Business Priority" on page 118*), Crown Jewels, have a little crown symbol.

Open Ports

For AWS EC2 instances, the Context Graph shows the open ports of the security group connected to the instance and their categories, such as *Web Server*, *Database*, or *Checkpoint Service*.

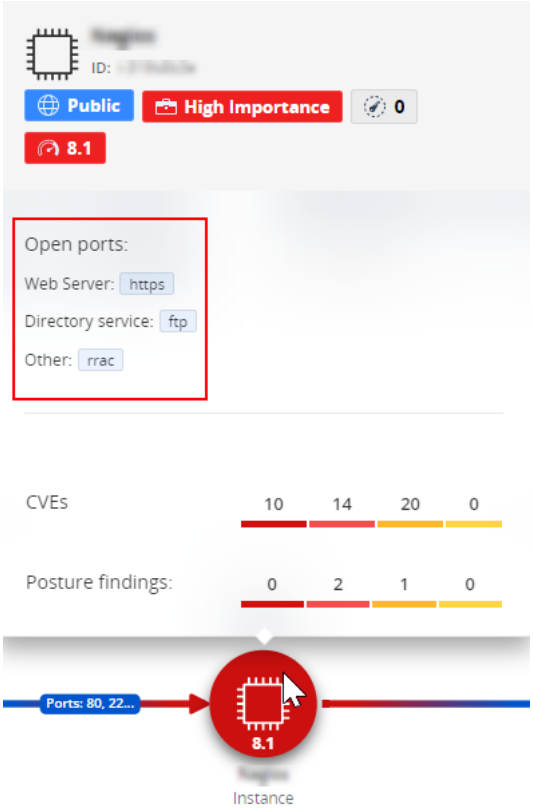

For the port details, put the cursor on the port number to see its tooltip.

Top 5 Remediations

This section shows the top 5 actions that you need to do to reduce the risk score of the asset. Click one of them to open a relevant tab with more details.

Events

This widget shows the number of critical and high-severity security events and posture findings the asset has.

Images

- **Policy status** - Shows if the image meets your organization standards (compliant / non-compliant)
- **Risk score** - Calculated risk score of the image
- **Vulnerability by severity** - Statistics for severity level of Image Assurance findings
- **Vulnerability by category** - Statistics for categories of Image Assurance findings
- **Workloads** - The list of workloads that use this image

- Kubernetes image shows workloads in the environment that uses it
- Container Registry image shows workloads from all Kubernetes environments

- **Scan status** - Shows the results of scanning the image by CloudGuard agents. For more information, see *"Image Scan Status" on page 429*.

SBOM

Software Bill of Materials (SBOM) is a list of all licenses that govern the components, the versions of the components used in the code, and their patch status. This list helps you quickly identify any security or license risk.

Supported Assets

CloudGuard shows SBOM for:

- Virtual Machines
- EC2 instances
- Container images
- Function apps

Note - SBOM is available only for assets scanned with *"Agentless Workload Posture" on page 489* and *"Scan Engine V2" on page 467*.

SBOM Export

On the asset page, go to the **SBOM** tab to see information about all the packages in the applicable asset.

To see the information in a file, click the **SBOM Export** button and select the export format. It can take some time for CloudGuard to generate the requested file.

Vulnerabilities

This page shows the CVEs, threats, and secrets received through the asset scanning. The **Remediation Summary** page contains information that a specific scanner finds.

Scanner examples:

- *"Agentless Workload Posture" on page 489*
- *"Container Registry Scanning" on page 464*

To search for a specific CVE, click **Search CVE** and go to the *"Vulnerability Search" on page 447* page.

Permissions

For more information, see *"Entitlement Map" on page 387*.

Posture Findings

This page presents the part of the Posture Findings related to this asset.

For more information, see *"All Events" on page 120*.

Threat & Security Events

This page presents the part of the Threat & Security Events related to this asset.

For more information, see *"All Events" on page 120*.

Properties

This page shows the asset properties based on the asset type.

Traffic Activity

This page shows part of the traffic logs for the asset.

For more information, see *"Traffic Explorer" on page 359*.

Account Activity

This page shows part of the account activity logs for the asset.

For more information, see *"Activity Explorer" on page 399*.

Permissions

CloudGuard requires permission to access your environment to read information about assets. These permissions are typically granted with an IAM role, with specific policies, which you must configure in your account and grant to CloudGuard.

AWS Policies and Permissions

CloudGuard uses AWS policies to manage your environments and periodically updates permissions for AWS account entities.

The policies give CloudGuard permission to manage specific entities, such as Security Groups and instances, on your AWS environment. The permissions type depends on your environment's selected mode (Monitor or Full Protection) of your environment.

Policies

These are the AWS policies that CloudGuard uses:

- **Mandatory**:
  - **SecurityAudit** (managed by AWS) for proper CloudGuard functionality.
  - **CloudGuard-readonly-policy** created during the onboarding procedure is required for different CloudGuard features, such as Posture Management and Network Security.

    This policy contains specific permissions to fetch information from AWS and use it in CloudGuard. If one of these permissions is not explicitly added to the policy, then information for that specific service becomes unavailable in CloudGuard. This does not affect CloudGuard functionality related to other services that are explicitly included in the policy.

    Best Practice - Check Point recommends to use the latest version of the *CloudGuard-readonly-policy* available for download from GitHub.
- **Optional:**
  - **AmazonInspectorReadOnlyAccess** (managed by AWS) enables CloudGuard to fetch the AWS Inspector information.
  - **ReadOnlyAccess** (managed by AWS) grants CloudGuard reading permissions to support new services in the future.
  - **CloudGuard-write-policy** created during the onboarding or update permissions procedure enables CloudGuard to manage your AWS account in the Full-Protection mode. It contains permissions for CloudGuard to manage Network Security.

    Best Practice - Check Point recommends to use the latest version of the *CloudGuard-write-policy* available for download from GitHub.

AWS Permissions used by CloudGuard

The table below shows the AWS permissions included in the CloudGuard Read, Write, and IAM Policies, as they are used by each CloudGuard module.

In addition, CloudGuard uses the AWS SecurityAudit policy, and the permissions included in this policy.

AWS Permission	CloudGuard Mode	CSPM	Network Security	IAM Safety
ec2:AuthorizeSecurityGroupIngress	Read-Only, Full	No	Yes	No
ec2:CreateSecurityGroup	Read-Only, Full	No	Yes	No
ec2:DeleteSecurityGroup	Read-Only, Full	No	Yes	No
ec2:RevokeSecurityGroupEgress	Read-Only, Full	No	Yes	No
ec2:RevokeSecurityGroupIngress	Read-Only, Full	No	Yes	No
ec2:ModifyNetworkInterfaceAttribute	Read-Only, Full	No	Yes	No
ec2:CreateTags	Read-Only, Full	No	Yes	No
ec2:DeleteTags	Read-Only, Full	No	Yes	No
dynamodb:DescribeTable	Full	Yes	No	No
elasticfilesystem:Describe*	Full	Yes	No	No
elasticache:ListTagsForResource	Full	Yes	No	No
firehose:Describe*	Full	Yes	No	No
firehose:List*	Full	Yes	No	No
guardduty:Get*	Full	Yes	No	No
guardduty:List*	Full	Yes	No	No
kinesis:List*	Full	Yes	No	No
kinesis:Describe*	Full	Yes	No	No
kinesisvideo:Describe*	Full	Yes	No	No
kinesisvideo:List*	Full	Yes	No	No

AWS Permission	CloudGuard Mode	CSPM	Network Security	IAM Safety
logs:Describe*	Full	Yes	No	No
logs:Get*	Full	Yes	No	No
logs:FilterLogEvents	Full	Yes	No	No
lambda:List*	Full	Yes	No	No
s3:List*	Full	Yes	No	No
sns:ListSubscriptions	Full	Yes	No	No
sns:ListSubscriptionsByTopic	Full	Yes	No	No
waf-regional:ListResourcesForWebACL	Full	Yes	No	No
iam:Get*	-	No	Yes	No
iam:List*	-	No	Yes	No
iam:AttachRolePolicy	-	No	Yes	No
iam:DetachRolePolicy	-	No	Yes	No
iam:AddUserToGroup	-	No	Yes	No
iam:RemoveUserFromGroup	-	No	Yes	No

Updating AWS Permissions

When onboarding your AWS account, CloudGuard receives permissions for specific entities in your AWS environment. It is necessary to update these permissions, at intervals, to ensure that CloudGuard obtains up-to-date information about these entities. Missing permissions for an entity in your environment cause CloudGuard's inability to manage or monitor the entity. Nevertheless, this does not affect other entities, if CloudGuard has the correct permissions for them. For troubleshooting steps after onboarding, see *"Troubleshooting AWS Onboarding" on page 167*.

Those permissions that are irrelevant to you, you can ignore. If after it the permissions become relevant, restore them.

You can select the method of the CloudGuard permissions update during the onboarding procedure only. When an environment is onboarded to CloudGuard with a specific method of permissions updating or deletion, you cannot change the method from the CloudGuard portal. On the Welcome page of the Environment Onboarding wizard, below CFT Permissions Management:

- Select **Allow CloudGuard to update and delete its CloudFormation stack resources** if you agree to start the procedure of permissions update automatically, from CloudGuard. See *"Updating Permissions Automatically" on the next page*.
- Do not select **Allow CloudGuard to update and delete its CloudFormation stack resources** if you do not agree, and, when a permissions update is required, do it manually in the AWS portal. See *"Updating Permissions Manually" on page 284*.

CloudGuard requires specific permissions in AWS defined in the AWS policies listed above, see *"Policies" on page 279*.

Mandatory Policies:

- **SecurityAudit** policy managed by AWS
- **CloudGuard-readonly-policy** created during the onboarding procedure.

Optional Policies:

- **AmazonInspectorReadOnlyAccess** managed by AWS is required only if your AWS environment uses the Inspector
- **ReadOnlyAccess** managed by AWS grants CloudGuard reading permissions to support new services in the future.
- **CloudGuard-write-policy** created during the onboarding or update permissions procedure is required for Full Protection (Read/Write) mode.

Ignoring and Restoring Permissions

Ignore irrelevant permissions so they do not affect the environment status in CloudGuard.

Ignoring Permissions

To ignore permissions that are missing for CloudGuard:

1. Go to **Assets** > **Environments**.
2. Search for an environment that requires permissions to update and click its name.
3. In the **Missing Permissions** message, click **Show more** to see the missing permissions table.
4. Select one or more permissions and do one of these:
   - Click **Ignore** on the top menu.
   - In the **Ignore** column, set the toggle to **ON** individually for each permission.

When all missing permissions are ignored or updated, the environment status becomes validated.

Restoring Ignored Permissions

To restore permissions:

1. Go to **Assets** > **Environments**.
2. Open the environment.
3. In the **Missing permissions** message, click **Show more** to see the missing permissions table.
4. Select one or more permissions and do one of these:
   - Click **Restore** on the top menu.
   - In the **Ignore** column, set the toggle to **OFF** individually for each permission.

Reviewing Permissions

You can review the list of affected cloud resources and associated fail messages.

1. In the Missing Permissions table, select a Resource and click **Show Entities**.
2. Compare the list of affected resources with the total population of resources of that type in the affected cloud environment (for example, using the **Protected Assets** page).
3. If the list of affected resources represents the entire population of resources of that type in the affected cloud environment, then a problem can be on the environment level (such as a missing permission in the IAM role).
4. If the list of affected resources is less than the entire population of resources of that type in the affected cloud environment, then the source of the problem must necessarily be specific to the individual affected cloud resources (for example, resource-level IAM "deny" policies, "ghost" resources deleted incorrectly or incompletely and so continue to trigger permission errors, "cross-account" resource deployment or resource sharing/reference issues, etc.).

Updating Permissions Automatically

You can update your account permissions automatically if you agreed to start this procedure during your account onboarding. This option is available for environments onboarded with a Unified procedure, if you selected **Allow CloudGuard to update and delete its CloudFormation stack resources** when onboarding your account. For more details, see *"Unified Onboarding of AWS Environments" on page 54*..

Automatic Updating

You can update the permissions remotely from the CloudGuard portal.

1. Go to **Assets** > **Environments**.
2. Search for an environment that requires permissions to update. The alert icon in the environment Status shows missing permissions.
3. Click the environment name.
4. Click **Update Permissions** at the top.

   The Update Permissions window opens with the details of your AWS account and the number of current and available versions of permissions.
5. As an alternative, click the link to view the changes in GitHub that opens in a new browser tab.
6. Click **Update** to start the automatic procedure.
7. After you complete the steps on the AWS account, click **Validate** in the Validate Permissions window in CloudGuard.

When CloudGuard updates the environment permissions, its status changes to approved. During this process, CloudGuard considers all missing permissions, including ignored permissions.

Updating Permissions Manually

- To update permissions for newly onboarded AWS accounts, without the option to update and delete CloudFormation stack resources (you did not select **Allow CloudGuard to update and delete its CloudFormation stack resources** during the onboarding), see *"Local Updating" below*
- To update permissions for the old AWS accounts (onboarded before March 2022), see *"Updating Permissions in Old Accounts" on the next page*

Local Updating

You can update the permissions locally, on the AWS portal.

1. Go to **Assets** > **Environments**.
2. Search for an environment that requires permissions update. The alert icon in the environment Status shows missing permissions.
3. Click the environment name.
4. Click **Validate Permissions** at the top.
5. The Validate Permissions window opens with instructions to follow to update the permissions in your AWS account.
6. After you complete the steps on the AWS account, click **Validate** in the Validate Permissions window in CloudGuard.

When CloudGuard validates the environment permissions, its status changes to approved. During this process, CloudGuard considers all missing permissions, including ignored permissions.

Updating Permissions in Old Accounts

Follow these steps to update permissions in AWS accounts onboarded before March 2022.

Notification of missing permissions

CloudGuard fetches information about your cloud assets from your environments, across all regions. If access is denied for any asset, CloudGuard retries several times, after which it marks the specific entity as missing permissions. CloudGuard notifies you of missing permissions on the **Environments** page.

When you click the environment, the missing permissions notifications inform you of permissions for the specified Role.

Click **Show more** to see details for the missing permissions and the number of affected entities.

There are two options to resolve missing permissions:

- **VALIDATE PERMISSIONS** - this option resets the mechanism and tries to validate the permissions. If this succeeds, the warning is removed. If not, it suggests to run the Permissions Wizard to add the missing permissions.
- **RUN PERMISSIONS WIZARD** - this option opens the Permissions Wizard that guides you to add the missing permissions to the policies. See the explanation below.

Updating Permissions with the Permissions Wizard

This wizard guides you to add missing permissions to the AWS policies used by CloudGuard.

Before you start, select the operation mode for your CloudGuard environment - **Monitor** or **Full Protection** (see *"AWS Security Group Management Considerations" on page 371*).

Note - If you select Full Protection for the environment, this does not set your Security Groups tobee fully managed. Security Groups can be individually set as Read-Only or Full Protection. See *"Full Protection Mode" on page 373*.

1. Click **Validate Permissions**.
2. In the window that opens, click **permission wizard**. A new instance of the CloudGuard portal opens with the selection of the operation mode for the account.
3. Follow the wizard instructions. If the policy exists, the data you provide updates it; if it does not exist, the wizard creates a new one.

For example, in step 4, search for **CloudGuard-readonly-policy** and answer *Yes* or *No*; the answer shows instructions to update the policy (**Yes**) or how to create a new policy (**No**).

4. In the last window, click **Finish**. After about 30 minutes, the changes are applied.

UnauthorizedOperation Exception

The message "**Error: UnauthorizedOperation: You are not authorized to perform this operation**" in one of your environments means that something happened to a valid policy and CloudGuard cannot use it.

The primary reasons for this are:

1. The mandatory policies **SecurityAudit** or **CloudGuard-readonly-policy** are detached from the role.
2. The role is deleted, or the External ID is changed.
3. There is a global policy that denies some permissions that CloudGuard uses. (AWS Organizations check organization policies).

To solve this issue, follow these steps:

1. Update your permissions by the instructions above.
2. If required, use a new Role and click **Edit Credentials** on the environment page to update the new Role details.
3. Fill in the new Role ARN and the External ID (this must have a value, which you can create. It must be the same as the value given in the Role external ID).
4. Review your global policies that can affect the Role connection and make sure that there is no `Deny` for EC2* in each of the global policies.
5. If these steps do not resolve the problem, contact *Check Point Support Center*.

More Links

- *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287*
- *"Troubleshooting AWS Onboarding" on page 167*

AWS Resources and Permissions for Serverless Runtime Protection

When Serverless Protection is applied to the Lambda functions in your AWS account, some resources are created in the AWS account and CloudGuard is granted permissions to access the resources. This is in addition to the information accessed by, and permissions granted to CloudGuard when the account is onboarded.

Serverless protection is enabled for a specific user's AWS account by launching a CFT (CloudFormation Template) stack, which is downloaded from CloudGuard to the user's AWS account. This stack creates resources and permissions that CloudGuard uses.

In addition, if Serverless Runtime Protection is enabled for specific functions in an account, a Serverless Runtime Protection Lambda Layer is deployed for each function.

CloudFormation Template

Each data center has its CloudFormation Template in YAML syntax, which you can download from the designated data center locations:

Data Center	Region	CloudFormation Template Link
United States	us-east-1	https://magnatar-protego.s3.amazonaws.com/magnatar-unified-cross-account-template.yaml
Ireland	eu-west-1	https://723885542676-protego.s3.amazonaws.com/723885542676-unified-cross-account-template.yaml
India	ap-south-1	https://guru-protego.s3.amazonaws.com/guru-unified-cross-account-template.yaml
Singapore (Dome9 accounts only)	ap-southeast-1	https://uranus-protego.s3.amazonaws.com/uranus-unified-cross-account-template.yaml
Australia	ap-southeast-2	https://neptune-protego.s3.amazonaws.com/neptune-unified-cross-account-template.yaml
Canada	ca-central-1	https://polaris-protego.s3.ca-central-1.amazonaws.com/polaris-unified-cross-account-template.yaml

Deployed Resources

This section describes the resources that CloudGuard deploys when you enable serverless protection for an AWS account and when you apply Serverless Runtime Protection for specific functions.

CFT Stack Resources

CloudGuard applies protection to serverless functions with a CFT stack. CloudGuard deploys this stack which runs in your AWS account.

This stack creates resources on your account, and CloudGuard uses the resources when it scans functions for issues and monitors function runtime activity.

Functions

- code scanning functions - These are deployed for each supported runtime environment

IAM roles and policies

- cross-account role - Allows CloudGuard to access a user's AWS account and read information about functions. The specific permissions are listed in the table below.

IAM Permission	Use
cloudwatch: GetMetricData	Collect statistical information such as number of invocations, their durations, and errors
cloudwatch: GetMetricStatistics	Collect statistical information such as number of invocations, their durations, and errors
lambda:ListVersionsByFunction	Explore functions in the account and get data about functions
lambda: ListAliases	Explore functions in the account and get data about functions
lambda: ListFunctions	Explore functions in the account and get data about functions
lambda: ListTags	Get the functions tags
lambda:GetLayerVersion	Function code analysis: scan any layer used by the function
lambda:ListEventSourceMappings	Continuous Scanning and Analysis
lambda:GetFunction	Get information about the function

IAM Permission	Use
lambda:GetFunctionConfiguration	Get the functions configuration for serverless function, to update the function code
lambda:GetPolicy	Get the function policy
lambda: GetFunctionUrlConfig	Continuous Scanning and Analysis
iam: ListRolePolicies	Describe the set of permissions configured to the function, used in order to provide security insights
iam: ListAttachedRolePolicies	Continuous Scanning and Analysis
iam: GetRolePolicy	Continuous Scanning and Analysis
iam: GetPolicyVersion	Continuous Scanning and Analysis
iam: GetPolicy	Get information about the policies
iam: GetRole	Continuous Scanning and Analysis
iam:SimulatePrincipalPolicy	Check the account required permissions
events:ListRuleNamesByTarget	-
sns:ListSubscriptionsByTopic	-
ec2:DescribeRegions	Get all enabled customer regions
s3:GetBucketNotification	Serverless Function Runtime
s3:GetBucketLocation	Serverless Function Runtime
s3:GetBucketAcl	Serverless Function Runtime
s3:GetBucketPolicy	Serverless Function Runtime

- execution role - For code scanning functions
- execution role - For the Serverless Runtime Protection Log Sender function, used for Serverless Runtime Protection.
- LogSenderRole - CloudGuard resource

IAM Permission	Use
sqs:GetQueueUrl	Serverless Function Runtime
sqs:SendMessage	Serverless Function Runtime

- FSPInjectorRole - CloudGuard resource

IAM Permission	Use
lambda:ListLayerVersions	Serverless Function Runtime Instrumentation
lambda:ListLayers	Serverless Function Runtime Instrumentation
lambda:UpdateFunctionCode	Serverless Function Runtime Instrumentation
lambda:UpdateFunctionConfiguration	Serverless Function Runtime Instrumentation
logs:GetQueryResults	Serverless Function Runtime Instrumentation
logs:StartQuery	Serverless Function Runtime Instrumentation

- CodeAnalysisRole - CloudGuard resource

IAM Permission	Use
lambda:ListLayerVersions	Continuous Scanning and Analysis, Serverless Function Runtime Server
lambda:ListLayers	Continuous Scanning and Analysis

Log Groups

- log groups - For each code scanning lambda function - used to forward results to CloudGuard backend.

S3 Bucket

- S3 bucket - Used to store Serverless Runtime Protection policy for function, and connect it with Serverless Runtime Protection layer, for runtime protection monitoring.

Lambda Layer

CloudGuard deploys a lambda layer on Serverless Runtime Protection, for each function that has this protection enabled.

Scanning Resources

CloudGuard uses these resources in the user's AWS account when it scans serverless functions.

Resource	Description
Functions	A code analysis function is deployed for each supported runtime (Python, Java, C#, Node.js)
Log Groups	A log group is created for each code scanning function
IAM Roles	CodeAnalysis Execution role - Used by functions; these roles have *Allow* permissions for these actions: *logs:CreateLogStream, logs:PutLogEvents, lambda:GetFunction, lambda:ListLayers, lambda:GetLayerVersion, lambda:ListLayerVersions*

Serverless Runtime Protection Resources

CloudGuard uses these resources in your AWS account, to monitor the runtime of activities of serverless functions (if runtime protection is enabled for the function).

Resource	Description
Lambda Layer	For each function in an AWS makes sure which Serverless Runtime Protection is enabled, CloudGuard deploys a Layer. This layer monitors function activities, and enforces the runtime policy.
Functions	Serverless Runtime Protection Log Sender - This function is deployed on-demand from CloudGuard in a specific region of the user's AWS account. It is not deployed by the CFT stack.
IAM Roles	An IAM role is used by the Serverless Runtime Protection Log Sender function, to send log group to the CloudGuard backend; these roles have *Allow* permissions for these actions: *lambda:CreateFunction, lambda:DeleteFunction, lambda:AddPermission, logs:CreateLogGroup, logs:PutRetentionPolicy, logs:DeleteLogGroup, logs:CreateLogStream, logs:PutLogEvents*
S3 Buckets	An S3 bucket is created by the CFT stack, for each account. A folder is created in the bucket for each function that has Serverless Runtime Protection enabled.

Resource	Description
Log Groups	A log group is created for the Log Sender function, which sends runtime information to the CloudGuard backend.

Azure Roles and Permissions

This topic describes the Azure applications and roles that CloudGuard uses to manage your accounts.

The applications and the roles granted give CloudGuard permission to manage specific entities (such as Security Groups, Instances, etc.) in your Azure account.

Roles

The roles depend on if the account is managed as **Read-Only** or **Manage**.

You must create a new Web App/API application (and name it *CloudGuard-Connect*, for example)

- **Read-Only**

  You must add this Access Control role to the Web App/API application, in your subscription: Reader.

- **Manage**

  You must add these Access Control roles to the Web App/API application, in your subscription:

  - Reader
  - Network Contributor

Permissions

Important - To ensure the maximal security of your assets, CloudGuard adheres to the principle of least privileges and requires only minimal set of access permissions. If you want to allow CloudGuard more access, grant access to additional resources and services. For troubleshooting information, see *"Troubleshooting Azure Onboarding" on page 176*.

An administrator consent is necessary to add the API application permissions below:

- `Directory.Read.All`, which includes and can be replaced by these permissions:
  - `User.Read.All`
  - `Group.Read.All`
  - `Application.Read.All`
- `Reports.Read.All`, which is required for Security Center-related entities, such as `DefenderServerVulnAssmt`
- `Policy.Read.All` used for fetching AD access policies, such as:

    - `ADAuthorizationPolicy`
    - `ADCondAccessPolicy`
    - `ADSecurityDefaults`
- `AccessReview.Read.All` used for fetching AD-level review policies, such as:
    - `ADAccessReviewsScheduleDefinition`
    - `ADCondAccessNamedLocation`
- `Audit.Log.Read.All`

For more information about used permissions, see Microsoft Graph permissions reference.

More Links

- *"Onboarding Azure Organizations" on page 170*
- *"Onboarding an Azure Subscription" on page 61*
- Microsoft Graph permissions reference

GCP Permissions and Roles

This topic describes the GCP APIs and roles that CloudGuard uses to manage your account.

The APIs and roles allow CloudGuard to manage specific entities (such as Security Groups, Instances, etc.) in your GCP account.

APIs

You must enable the Compute Engine API and the Cloud Resource Manager API, and create a new service account for CloudGuard. CloudGuard uses this service account to connect to your GCP account.

Important - To ensure the maximal security of your assets, CloudGuard adheres to the principle of least privileges and requires only minimal set of access permissions. If you want to allow CloudGuard more access, grant access to additional resources and services. For troubleshooting information, see *"Troubleshooting GCP Onboarding" on page 185*.

In addition, you can optionally enable these APIs:

- GKE API - for GKE entities, such as GkeCluster
- KMS API - for KMS entities, such as KmsKeyRing
- IAM API - for IAM entities, such as GcpIamGroup, and GcpIamUser
- BigQuery API - for the BigQuery entity
- Admin SDK - for IAM entities like users or groups
- Kubernetes API
- App Engine Admin API
- Cloud Functions API
- Cloud SQL Admin API
- Cloud BigTable Admin API
- Cloud Pub/Sub API
- Cloud Memorystore Redis
- Service Usage API
- Cloud Filestore API
- API Keys API
- Cloud Logging API
- Cloud DNS API

- Cloud Asset API
- Essential Contacts API
- Access Approval API

Roles

In addition, you must add these roles for the service account:

- Viewer (in Project)
- Security Reviewer (in IAM)

Organizational Units

Overview

You can organize your environments in CloudGuard into Organizational Units. Organizational Units are user-defined groupings of accounts. An Organizational Unit could depict, for example, the accounts for a business unit in an enterprise, or a geographical location. You can associate your accounts with an Organizational Unit, with accounts from different cloud providers. In addition, you can create Organizational Units in existing Organizational Units, creating a logical hierarchy.

Initially, your account has a root entity that includes all environments that have been onboarded to CloudGuard. This root entity is not an Organizational Unit, and it serves only as a starting point for creating your own Organizational Units. From there, you can create more Organizational Units and associate environments with them (they are moved from the root). An account can correspond with only one Organizational Unit at a time, but one Organizational Unit can be a sub-unit of a different one. Onboarded AWS Organizations appear under the root entity as its children.

You can label Organizational Units with a name, but sub-Organizational Units of the same parent cannot have the same name.

You can delete Organizational Units. All environments related to it and its sub-Organizational Units are moved to the 'root' unit, and all sub-Organizational Units are deleted with it.

Benefits

- See your accounts based on logical groupings, for example, business units or geographical regions.
- Better visibility of your account inventory by seeing them grouped logically and hierarchically (with collapsible views).
- Define & apply tailored compliance policies for groupings that are logical for your enterprise.
- Apply user access (RBAC) policies to your accounts based on enterprise logical groupings.

Use Cases

- Streamline the view of environments and assets - see *"Viewing OUs" on the next page*
- Apply a Continuous Posture policy to a business unit - *"Continuous Posture" on page 315*

Actions

Viewing OUs

1. Navigate to the **Organizational Units** page in the **Assets** menu. This shows your Organizational Units. For each, the number of environments corresponding to it is shown, broken down based on the cloud providers. In addition, Sub-Organizational Units are shown. You can use the Filters pane, on the left, to filter the list.
2. Click right or down arrows to expand or close the hierarchy of OUs.

Creating an OU

1. Navigate to the **Organizational Units** page.
2. Click **CREATE OU**.
3. Enter a name for the OU and select its location in the hierarchy of OUs, then click **Create**.
4. As an alternative, create a new OU as a sub-OU for an existing OU. Select the existing OU and click **Create sub OU**.
5. Enter a name for the OU and click **Add**.

Moving OUs

You can change the location of an Organizational Unit in the hierarchy of OUs.

1. Put the mouse on the OU to be moved and click **Move**.
2. Select the new OU below which the OU is moved and click **Move**.

Associating Environments with an OU

When you have created Organizational Units, you can associate environments with them. You can associate accounts with an OU, with accounts from different cloud providers. An environment can correspond to only one OU (or to the root).

1. Navigate to the **Environments** page in the **Assets** menu. This shows your environments onboarded to CloudGuard (from all providers).
2. Select one or more environments.
3. Click **Associate To OU**.
4. Select the Organizational Unit and click **Associate**.

   The Organizational Unit for the environments is updated.

Removing (disassociating) Environments from an OU

You can remove (disassociate) environments from an OU. You can do this by associating the environments with a different OU or with a root. Follow the steps in the procedure above.

Deleting OUs

When you delete an OU, the accounts related to it are moved to the root.

In addition, all sub-OUs for the Organizational Unit are deleted. You cannot delete the root.

1. Navigate to the **Organizational Units** page.
2. Put the cursor on the OU to be deleted and click **Delete**.

Custom Resources

You can create and manage named lists in CloudGuard and use them in place of the full list of items. For example, you can define a named list of IP addresses, and refer to the list (by its name) in a Security Group rule definition. Similarly, you can refer to a list of names in a GSL rule.

User-Managed Lists

You can create two types of lists:

- **Generic List** contains text values. For example, it can be a list of instance types, OS types, or network names.

  You can create a list of entities and refer to its name in a GSL rule statement. It is not necessary to have the full list in the rule. This makes the rule shorter, and you can use the same list in many rules. Changes to a list affect all the rules that use it when CloudGuard runs the rule the next time.

- **IP List** contains IP addresses or CIDR ranges.

  You can create a list of IP addresses or ranges and refer to them in Security Group rules. The same list can be used in many rules, for example, a list of public IP addresses. When there are changes in the IP addresses, update the list, and the rules that use it are updated automatically.

CloudGuard-Managed Lists

For AWS, Azure, Alibaba Cloud, GCP, and OCI environments, CloudGuard automatically creates a Generic list of all onboarded environments, one list for each cloud platform. When you onboard a new environment to CloudGuard, it updates the existing list of onboarded environments with the new entry. You cannot change or delete the lists managed by CloudGuard.

Creating a User-Managed List

1. In the **Assets** menu, navigate to the **Custom Resources** page.
2. Click **Add List**.
3. Enter a name without spaces and, optionally, a description for the list.
4. Select the list type – *Generic* or *IP*.
5. Enter a value for each item of the list and click **Add** to add it. Alternatively, click **Upload CSV** to upload a CSV file with a list of values and optional comments.
6. Click **Save**.

Using Lists

- Refer to the CloudGuard-managed list of onboarded environments in the GSL Builder to run a GSL rule on all environments.
- Refer to a managed list in GSL as *$<list-name>*. For example, the GSL snippet

  ```
  ... VirtualMachine should have operatingSystem in
  ($AzureOSTypes)
  ```

  refers to the list *AzureOSTypes*.

  Note - List names do not read apostrophes (') around them in functions such as *in()*. The GSL builder editor automatically includes apostrophes around some values, so you have to remove these manually (in the GSL text editor).

- Use a list when you define a rule for a Security Group.

Cloud Security Posture Management (CSPM)

CloudGuardCloud Security Posture Management (CSPM) checks your cloud environments' compliance with industry standards and best practices or your organization's security policies. Its engine uses the rules that you define or sets of rules (rulesets) developed by CloudGuard and available out-of-the-box. CloudGuard provides a comprehensive set of rulesets that have many of the same standards, such as PCI-DSS and HIPAA, for cloud security, which you can run immediately on your environments. In addition, you can build and test new rules or change existing rules with an intuitive graphical rule builder, to tailor policies to your organization's specific needs and compliance goals.

Posture Management enables you to manage resources across multiple clouds, flows, and settings on one management platform. You can check your environments and receive notifications when it detects issues. Detailed results of tests and summary reports are available for your review.

Posture Management accesses your environments directly through cloud platform APIs and CloudGuard policies that you set up on these environments. It works with all cloud providers, and you can check compliance even when your cloud presence is distributed on multiple cloud platforms.

Benefits

- At-a-glance dashboard view of organizational compliance across the full cloud presence, on all cloud platforms
- Check compliance with cloud security standards
- Clear reports to indicate non-compliant issues
- Easily build modified rules based on requirements with the graphical GSL builder
- Preconfigured (built-in) rulesets developed by CloudGuard have a wide range of standards and best practices

Use Cases

- Enforce environments compliance with standards - see *"Rules and Rulesets" on page 309*
- Enforce compliance with organizational policies across the estate - see *"Configuring CloudGuard Policies" on page 78*

- Review the security and compliance posture across the estate with a unified dashboard - see *"Dashboards" on page 86*
- Analyze compliance of a proposed cloud design (CloudFormation Template) before actual deployment - see *"Onboarding AWS Environments" on page 144*
- Customize the Posture Management dashboard based on your needs, to put effort into the more sensitive and interesting environments
- Review latest assessment results and apply remediation - *"Automatic Remediation with CloudBots" on page 317*
- Review assessments on a specific environment from a specific point in time - *"Assessment History" on page 321*
- Create customized compliance or organizational policy rules - see *"Rules and Rulesets" on page 309*

CloudGuard GSL (Governance Specification Language)

Rules used by Posture Management are defined with the CloudGuard Governance Specification Language (GSL). This is an intuitive user-readable language that describes the test. For example, the rule

```
S3Bucket should have logging.enabled=true
```

checks that logging is enabled for AWS S3 buckets.

See *"Governance Specification Language (GSL)" on page 326* for details and examples of the GSL syntax. See *"GSL Builder" on page 350* to learn how to build a rule with the graphical interface.

Cloud Entity Domain Model

CloudGuardPosture Management is based on Governance Specification Language (GSL), which defines the syntax for compliance rules. In addition, it includes cloud entities, which are the targets to which the rules apply. These entities represent the real entities in the supported cloud platforms, such as instances or S3 buckets.

Entities have attributes, some unique to specific entities. Some attributes are simple, for example, strings or numbers, while some are compound: usually, sub-entities that are related to the entity. For example, the VPC in which an instance is located. In addition, the attributes can contain list attributes.

GSL includes entities for the cloud platforms that CloudGuard supports. Each platform lists its supported entities.

Views

Posture Management includes these pages:

- Rulesets - Shows your rulesets and rules, preconfigured rulesets, and custom ones that you define.
- Continuous Posture - Lists the policies that continuously assess your environments for compliance.
- Remediation - Lists the environments where automatic remediation with CloudBots is enabled.
- Exclusions - Lists the environments that have rule exclusions.
- Assessment History - Shows a list of previously run assessments, with summary details for each. You can filter the view by account, rulesets, and time to show specific assessments of interest.
- GSL Builder - Lets you build and test GSL rules.
- Posture Overview - Presents a summary view of the compliance assessments run on your environments.

Actions

Running an Assessment

Run a ruleset on a selected environment.

1. Navigate to the **Rulesets** tab in the **CSPM** menu.
2. Select a ruleset to run.
3. Click **Run Assessment**. The **Run Assessment** window opens with a list of available environments.
4. Select the environment, region, and VPC, on which the policy is run, and click **RUN**. The assessment period usually takes from a few seconds to a few minutes, based on the complexity of the ruleset and the number of rules. When it completes, the results appear.

   You can see details for each. They include the number of entities tested (Tested), the number that is included in the scope of the rule (Relevant), the number of entities that are excluded (Excluded), if **Show Exclusions** is selected, and the number of failed tests (Non-Compliant).
5. Click **Expand** to show more details, the rule's details, and a list of the failed entities.

Viewing Assessment History

See *"Assessment History" on page 321*.

Creating Exclusions from Assessment

You can create an exclusion directly from the assessment based on the assessment details. This method of exclusion creation is faster because CloudGuard enters several fields for you.

1. Navigate to **CSPM** and select **Assessment History** from the menu.
2. Use the Filter bar to search for the assessment that can be a basis for your new exclusion.
3. Click the assessment to see its details.
4. Click **Expand** to show more details.
5. Below **Findings**, click the **Exclude finding** icon ( ) in the **Actions** column on the right.

   The **Create New Exclusion** window opens, and the most applicable fields in it are already filled.
6. Edit the fields as you wish based on the steps in *"Configuring CloudGuard Exclusions" on page 80*.
7. Enter a comment for the exclusion.
8. Click **Save**.

Getting Started with Posture Management Policy

When you define a new Posture Management policy, CloudGuard starts to continuously assess the applicable environments with the selected rulesets and notifies you of rules that failed with the notification that you select. All Posture Management policies configured on your CloudGuard account have its *"Continuous Posture" on page 315*.

To set up the first CSPM policy:

1. Navigate to the **Continuous Posture** page, in the **CSPM** menu.
2. Click **Add Policy** to add a new policy. Select a new **Environment Policy**.
3. Select a cloud platform, for example, Microsoft Azure, and click **Next**.
4. Select one or more environments and click **Next**.
5. For the initial configuration, use the configured CloudGuard-managed rulesets. From the list, select one or more rulesets for the policy, select the **Latest** ruleset version, and click **Next**.
6. To add a new notification, click **Add Notification**.
7. In the **Create New Notification** window, enter the notification name and, optionally, a description. For this initial policy, you can use the default settings. Make sure that the **Alert console** is selected. This option allows you to see all findings on the **Events > Posture Findings** page.
8. Click **Save**.
9. Select the notification for the association.
10. Click **Save**.

Your first policy appears on the Continuous Posture page.

More Links

- *"Configuring CloudGuard Policies" on page 78*
- *"Continuous Posture" on page 315*
- *"Rules and Rulesets" on page 309*
- *"Notifications" on page 852*

Kubernetes Posture Management

When you onboard a Kubernetes cluster to CloudGuard, it immediately starts to apply Posture Management rules to the cluster. It can examine your clusters deployed at various cloud providers, as well as clusters located on premises. See *"Onboarding Kubernetes Clusters" on page 188* for details on how to onboard the clusters.

For more details on CloudGuard CSPM, see *"Cloud Security Posture Management (CSPM)" on page 302*.

CloudGuard Workload Protection - Kubernetes Posture Management

Kubernetes Rules and Rulesets

To examine your Kubernetes clusters, CloudGuard uses rulesets as for all other onboarded environments. For your posture management, you can use general or vendor-specific rulesets.

Best Practice - Check Point recommends using rulesets developed for dedicated cloud providers:

- For EKS clusters, apply the **CIS Amazon Elastic Kubernetes Service (EKS) Benchmark** ruleset
- For GKE clusters, apply the **CIS Google Kubernetes Engine (GKE) Benchmark** ruleset
- For Microsoft AKS clusters, apply the **CIS Microsoft Kubernetes Engine (AKS) Benchmark** ruleset
- For OpenShift clusters, apply **CIS OpenShift Container Platform Benchmark** ruleset

For other platforms, use the latest CIS Kubernetes Benchmark and Kubernetes CloudGuard Best Practices rulesets.

All available rulesets are shown on the **Posture Management** > **Policy** > **Rulesets** page. Filter the list for **Platform: Kubernetes** and **Type: CloudGuard Managed**.

For more information on CloudGuard rulesets, see *"Rules and Rulesets" on page 309*.

Kubernetes Posture Findings

The CloudGuard Compliance engine generates Kubernetes posture findings that show on the **Events** > **Posture Findings** page. For more details, see *"All Events" on page 120*.

To send email notifications on findings filtered by Kubernetes labels, see *"Notifications" on page 852*.

More Links

- *"Kubernetes Containers" on page 415*
- *"Admission Control" on page 476*
- *"Image Assurance" on page 434*
- *"Intelligence for Kubernetes Containers" on page 645*
- *"Kubernetes Runtime Protection" on page 549*

For Kubernetes terminology, see the Glossary in the Kubernetes documentation.

Rules and Rulesets

All CloudGuard components, Posture Management, Intelligence, and Image Assurance, use a combination of rulesets to test your environments. Rulesets contain rules, which are individual tests of function in your environment. For example, a rule can test if a password policy is enforced.

Rulesets Management

There are two types of rulesets management:

- **CloudGuard-Managed Rulesets** - CloudGuard includes a set of built-in rulesets developed by its research team. These rules test your environments for compliance with best practices and with the same cloud security standards, such as PCI-DSS, HIPAA, and CIS Foundations for AWS, Azure, GCP, and Kubernetes. They include remediation steps that you can apply to your environment.

  Periodically, CloudGuard research team update the rulesets based on the recent changes, for example, add rules, delete rules, etc. From January 2024, CloudGuard-managed rulesets have designated versions. For more information about versions, see *"Viewing a Ruleset Version" on page 313*.

  Note - When remediation steps are applied to an environment, and CloudGuard is updated (time may change, based on internal sync intervals). Run the assessment again to verify the remedy.

- **Customer-Managed Rulesets** - Although the CloudGuard-Managed rulesets cannot be changed, you can clone them to make a copy and then change the copy. You can find your modified rulesets if you filter them by the Customer Managed type.

Severity Levels

While in CloudGuard each finding has its triggers and conditions, the information below describes the general criteria and implications to reflect finding risk.

CloudGuard assigns each finding one of five severity levels:

- **Informational** - There is no security or infrastructure risk. Administrator awareness is recommended.
- **Low** - There is no security or infrastructure risk. The response is based on best practices.
- **Medium** - A possible security risk exists. Action is required in reasonable time.
- **High** - This may lead to a possible risk. Immediate action is required.
- **Critical** - An asset is compromised. Immediate action is required.

Severity Criteria and Implications

Three criteria below define and affect these severity levels:

1. **Infrastructure exposure:** - If a finding shows an infrastructure exposure that is not necessary, which can provide attackers a possible ground for exploitation.
    - **None** - There is no risk of infrastructure exposure.
    - **May lead** - Some conditions can cause infrastructure exposure.
    - **Exists** - There is an infrastructure exposure.
2. **Information disclosure:** If a finding describes an information disclosure, which can lead to sensitive data exfiltration and can be used maliciously.
    - **None** - There is no risk of information disclosure.
    - **May lead** - Some conditions can cause infrastructure exposure.
    - **Exists** - There is an information disclosure.
3. **Possible impairment:** - If a finding describes a lead to infrastructure or information impairment, in terms of security, misconfiguration, or maintenance.
    - **None** - There is no risk of impairment
    - **May lead** - Some conditions can cause infrastructure exposure or information impairment.
    - **Exists** - There is a lead for an infrastructure or information impairment

In addition, each severity level correlates with two implication levels:

1. Level of required action:
    - **None** - No action is required.
    - **Advised** - The response is based on best practices.
    - **Not immediate** - Action is required in a reasonable time.
    - **Immediate** - Immediate action is required.
2. Compromised assets:
    - **None** - The asset is not compromised.
    - **Compromised** - The asset is vulnerable.

Severity Matrix

The table below shows the relationship between each severity level and the mentioned criteria and implications. For each finding, its severity is defined by the highest severity that meets a minimum of one criterion.

Severity	Criterion			Implication	
	Infrastructure Exposure	Information Disclosure	Possible Impairment	Compromised Asset	Action Level
Informational	None	None	None	None	None
Low	None	None	None	None	Advise
Medium	May lead	None	None	None	Not immediate
High	Exists	May lead	May lead	None	Immediate
Critical	-	Exists	Exists	Compromised	Immediate

Malicious IP Classification

For rules that identify malicious IPs, CloudGuard uses Check Point's ThreatCloud technology. The table below explains the meaning of each IP category.

Class	Description
Unclassified	The service could not classify the IP. There is not sufficient data about this resource.
Adware	The IP domains operate in the gray areas of the law, collecting private data on the users, and show unwanted content or a website that contains sub-application to download.
Volatile	The IP domains contain malicious software, for example, hacking websites.
Benign	Legitimate IP that is not malicious.
CnC Server	Command and control of malware.
Compromised Server	Legitimate IP that was hacked and operates a malicious function.
Phishing	The IP domains attempts to get sensitive information such as usernames, passwords, and credit card details (and sometimes, indirectly, money), frequently for malicious reasons, by masquerading as a trustworthy entity in electronic communication.
Infection Source	The IP domains can infect their visitors with malware.
Web Hosting	The IP domains allow you to rent out space for websites to have your business in.

Class	Description
File Hosting	The IP domains allow you to rent out space for storage to have your business in.
Parked	The IP domains permanently do not have content. Possibly, they contain advertising content on pages that have been registered but do not (at this time) have initial content.
Scanner	The IP is a known Internet scanner.
Anonymizer	The IP is a known Tor (The Onion Router) anonymity proxy server.
Cryptominer	The IP domains are used for crypto mining.
Spam	The IP domains are used for spam.
Compromised Host	The victim's IP.

CloudGuard Rules Repository

CloudGuard Compliance Engine is an end-to-end security and compliance solution for assessment, remediation, and continuous security compliance enforcement. The CloudGuard GSL (Governance Specification Language) is a syntax to configure cloud security and compliance rules that can be applied in assessments for your environments with the CloudGuard Compliance Engine.

The Cloud Security Posture Repository is a shared security and compliance knowledge platform for AWS, Azure, GCP, and Kubernetes. It provides an evolving set of security and compliance best practices, curated and developed by CloudGuard. The controls include risk and remediation details needed for security governance and compliance with public cloud environments.

Actions

Adding a Ruleset

Add a new ruleset. When you have a ruleset, you can add rules to it. Then, the rules can be applied to a VPC in one of your environments or to a CloudFormation Template.

1. Navigate to the **Rulesets** page in the **CSPM** menu.
2. Click **Add Ruleset** to create a new ruleset. Enter a name for the ruleset and, optionally, a description, and select the cloud provider on which it is applied.

Cloning a Ruleset

You can copy an existing ruleset. The copy contains the same rules. This is useful to change or extend rules in a CloudGuard-managed ruleset that you cannot edit.

1. Navigate to the **Rulesets** page.
2. Click the ruleset that you want to copy and open its details.
3. Click **Clone**.
4. In the **Clone <name> ruleset** window, enter a name for the new ruleset and its description.

Viewing a Ruleset Version

For CloudGuard-managed rulesets, you can select a particular version to adhere to or automatically receive updated versions of the rulesets.

1. Navigate to the **Rulesets** page.
2. Click **Add Filter** and select **Type** > **CloudGuard Managed**.
3. Click a ruleset for which you want to see the available versions.
4. From the list on the top bar, see:
   - **Latest** to automatically use the most recent ruleset version.
   - a particular version that you can select to adhere to in the future.

Note - If you change the ruleset version from the ruleset page, it does not affect the policies. To adhere to a specific ruleset version, select it when you edit or add a new *"Continuous Posture" on page 315* policy.

Adding Rules to a Ruleset

Add rules to a ruleset. You can add rules to custom rulesets (new policies that you add), but not to preconfigured rulesets.

1. Navigate to the **Rulesets** page and select the ruleset.
2. Click **New Rule** to add a rule to the policy. This opens the online GSL rule builder (see *"Governance Specification Language (GSL)" on page 326*).
3. Enter a name for the rule and, optionally, a description, remediation (corrective steps), compliance sections that the rule covers, and a severity level (that is, the severity or effect of non-compliance with this rule).

4. Enter the rule in the **GSL Editor** box with GSL syntax, then click **Test** to check the rule. When you finish, click the **Done** button. The rule appears from the list of rules for the policy. You can enter the rule as text, in the **Free text** mode, or graphically, in the **Builder** mode.
5. Optionally, add *"Automatic Remediation with CloudBots" on page 317* tags in the **Compliance Section** of the rule. These tags are used only if the rule is used in a *"Continuous Posture" on page 315* policy. They show a remedial CloudGuard CloudBot to be run if the rule fails in an assessment. The tag has the form:

   ```
   AUTO: ec2_stop_instance
   ```

   The prefix 'AUTO' indicates that this is an auto-remediation tag. The expression that follows the tag is the name of a remediation bot (for example, 'ec2_stop_instance') followed, optionally, by parameters. You can add more than one tag for a rule, in which all the remediation actions are done if the rule fails.
6. Add more rules as needed.

Modifying Rules

You can change existing rules in a custom ruleset. You can configure them with the graphical Rule Builder, in the same procedure that you create new rules. CloudGuard stores rules in JSON format, so you can edit rules for a policy by editing the JSON block.

1. Navigate to the **Rulesets** page and select the ruleset.
2. Click the rule you wish to edit. This opens the Rule Builder. From there you can change the rule, edit the text or use the graphical Builder.
3. Change the text of the rule as necessary and then click **Done**.

Continuous Posture

A CloudGuard Continuous Posture is a compliance ruleset, associated with an environment and a notification. CloudGuard continuously assesses the environments in your compliance policies with the selected rulesets and notifies you of rules that failed with the Notification that you select. You can receive findings as email reports, messages to SNS topics, or events sent to HTTP endpoints (webhooks).

CloudGuard does not send notifications for issues already discovered in assessment done before by the same policy. You receive a notification only the first time a rule fails, but not after the next assessments. If the issue is remedied, and the rule passes in the next assessment, a 'pass' notification is sent to SNS and HTTP endpoints, but not to email notifications. In email reports, it does not show in the list of failed rules.

To set up a Continuous Posture:

1. Navigate to the **Continuous Posture** page in the **CSPM** menu. This page shows a list of policy associations.
2. Click **Add Filter** to search from the list of policies by platform, OU/environment, ruleset, and notification.
3. Click **Add Policy** to add a different policy. Select one of these options:
   - **Cloud Platform Policy** - CloudGuard automatically selects all onboarded environments on the selected platform. In addition, when you onboard a new environment that belongs to this platform to CloudGuard, this policy automatically applies to all new environments.
   - **Environment Policy** - Select one or more environments for the policy.
   - **Organizational Unit Policy** - Select an Organizational Unit, and the policy applies to all existing and newly onboarded environments in this Unit.
4. For a new Cloud Platform Policy:
   a. Select a cloud platform and click **Next**.
   b. Select one or more rulesets for the policy, select the ruleset version, and click **Next**. For *"Adding a Ruleset" on page 312*, open the **Rulesets** page of the **CSPM** menu.
   c. Select notifications for the policy. To add a new notification, click **Add Notification**. When you add a new notification to the Cloud Platform policy, you can select an Executive Summary report as a scheduled report. For more details on notifications, see *"How to Configure a Notification" on page 853*.
   d. Click **Save**.
5. For a new Environment Policy:

a. Select a cloud platform and click **Next**.

b. Select one or more environments and click **Next**.

c. Select one or more rulesets for the policy, select the ruleset version, and click **Next**. For *"Adding a Ruleset" on page 312*, open the **Rulesets** page of the **CSPM** menu.

d. Select notifications for the policy. To add a new notification, click **Add Notification**. For more details on notifications, see *"How to Configure a Notification" on page 853*.

e. Click **Save**.

6. For a new Organizational Unit Policy:

a. Select an Organizational Unit and click **Next**.

b. Select one or more of the rulesets for the policy, select the ruleset version, and click **Next**. For *"Adding a Ruleset" on page 312*, open the **Rulesets** page of the **CSPM** menu.

c. Select notifications for the policy. To add a new notification, click **Add Notification**. For more details on Notifications, see *"How to Configure a Notification" on page 853*.

d. Click **Save**.

Note - When you change the ruleset version for existing policies (with the **Edit** option), it affects the policies that use the ruleset. This can cause closing findings or opening new findings in the assessment results.

Automatic Remediation with CloudBots

CloudBots automatically correct compliance issues discovered in your cloud environments by CloudGuard compliance checks. You can configure your CloudGuard account to use CloudGuard CloudBots.

You can configure remediation steps in different contexts of the CloudGuard portal.

You must deploy CloudGuard CloudBots in the cloud environments to apply remediation steps.

CloudBots

CloudGuard CloudBots are small programs or scripts in Python that operate on the account or cloud asset to correct missing or misconfigured settings. For example, they can close Security Groups that are widely open. CloudGuard invokes CloudBots when compliance rules fail.

CloudBots work with rules invoked from Continuous Posture, Intelligence, and CIEM policies.

For some rules, CloudGuard recommends you use one of the preconfigured CloudBots or create your CloudBot if there is a rule violation. For other rules, you can use only your custom CloudBots.

CloudBots provide:

- Reduction of risks related to misconfigurations in your cloud environment, to comply with the compliance industry standards.
- Reduced workload on the enterprise cloud IT team by performing remedial actions on misconfigured cloud assets and environments automatically.
- Reduced response time to remedy a problem to decrease the window of exposure to risk because of the misconfiguration.
- As CloudBots work with continuous posture assessments, your cloud environments are repeatedly assessed, so any changes (because of accidental or not approved access to the cloud assets) are detected and corrected almost immediately.
- Reliable application of the same correction to misconfigurations of the same type. Correcting an environment policy misconfiguration is the same for all environments. In addition, a full audit trace can be kept of all actions, so you know about the applied changes.

The CloudGuard portal provides multiple options to configure remediation:

- Go to **CSPM** > **Remediation** and click **Create New Remediation**. For more information, see *"Adding Remediation" on page 319*

- Go to **CSPM** > **Rulesets**, select a ruleset, select a rule, and click **Add CloudBot** below **Automated Remediation**.
- Go to **CSPM** > **Assessment History** and select an assessment result. In one of the failed rules, click **Expand** to see the list of findings and click **Configure remediation** ( ) for a failed entity. For more information, see *"Creating Remediations" on page 323*
- Go to **CIEM** > **Remediation** and click **Create New Remediation**. For more information, see *"Remediation" on page 398*
- Go to **CDR** > **Threat Monitoring** > **Remediation** and click **Create New Remediation**. For more information, see *"Remediation" on page 632*
- Go to **Events** > **Posture Findings** or **Threat & Security Events**, select an event, and click **Fix it**. For more information, see *"Applying a CloudBot immediately (Fix it)" on page 127*

Onboarding CloudBots

To apply remediation steps, you must onboard CloudGuard CloudBots in the cloud environments. For manual deployment of the CloudBots, see https://cloudbots.dome9.com/.

For deployment through the CloudGuard portal, see below.

To onboard CloudBots through CloudGuard:

1. In CloudGuard, open the **Environments** page from the **Assets** menu.
2. Select an environment to be protected with CloudBots.
3. In the **CloudBots** column, click **Enable CloudBots** to start the remediation onboarding wizard.

   As an alternative, you can click and open the environment page. From the top menu, select **Add CloudBots**.
4. Follow the on-screen instructions to complete the wizard.

For AWS environments

a. Create an AWS Lambda function that runs the bots and an SNS topic to trigger the Lambda function.

b. Select a region enabled on your account.

c. Deploy a CFT on your AWS account.

d. Click **Check Now** to make sure the CFT deployment is successful.

For Azure subscriptions

a. Create an Azure FunctionApp function that runs the bots.

b. Deploy an ARM template on your Azure subscription.

c. Paste the function URL into the wizard and click **Next**.

Remediation

Adding Remediation

You can add a remediation for a specific rule in a ruleset or all rules in a ruleset. You can limit remediation to specific environments or entities.

To add a remediation for a specific rule:

1. Navigate to **CSPM** > **Rulesets**.
2. Open the ruleset that contains the rule to which to apply a remediation.
3. Use the *"Filter and Search" on page 863* toolbar to find the rule.
4. Click **Add** to add a predefined CloudBot recommended by CloudGuard. If no recommendation exists, click **Add CloudBot** to create a new custom CloudBot and add it.

   The **Edit Remediation** window opens with the selected rule and ruleset.
5. Select the remediation parameters. You can combine the options, so the remediation applies to the combination of all the selected options.

   a. **Environment** that applies the remediation to rules in the selected ruleset only when the ruleset is applied to the selected environments. If you do not select an environment, CloudGuard applies the CloudBot to all the available environments where the selected rule is triggered.

   b. **Entity**, by its entity ID (optional, if missing, all entities are implied); this selects all rules that contain the selected entities.
6. For rules that recommend remediation, the CloudBot appears in the field. For rules without recommendations, select the CloudBot from the list. If the CloudBot is not in the list, select **Custom**, and then add the name of the CloudBot, along with the runtime arguments. The CloudBot must be deployed in the selected environment, in the same folder as the other bots.
7. Add a comment (mandatory field) and click **Save**.

To add a remediation for all rules in a ruleset:

1. Navigate to the **Remediation** page in the **CSPM** menu.
2. Click **Create New Remediation** in the top right.
3. Select the rules for which the remediation applies, from the given options. You can combine the options, so the remediation applies to the combination of all the selected options.
    a. a **Ruleset** (mandatory)
    b. a specific **Rule** in the ruleset (optional, if missing, all rules are implied)
    c. a specific **Environment** that applies the remediation to rules in the selected ruleset only when the ruleset is applied to the selected environments
    d. a specific **Entity**, by its entity ID (optional, if missing, all entities are implied); this selects all rules involving the selected entities
4. Select the CloudBot, from the list. If the CloudBot does not show, select **Custom**, then add the name of the CloudBot, along with the runtime arguments. The CloudBot must be deployed in the selected environment, in the same folder as the other bots.
5. Add a comment (mandatory) and click **Save**.

Deleting Remediation

1. Navigate to **CSPM** > **Remediation**.
2. Select one or more remediations to delete and click **Delete Selected**.

More Links

"All Events" on page 120

"Assessment History" on page 321

Assessment History

You can see a list of run assessments, with summary details for each assessment. Go to **Posture Management** > **Assessment History**. You can filter the view by environment, rulesets, triggering event, and time, to show specific assessments of interest.

For each assessment, the list shows:

- *Date* - Date the assessment run
- *Environment* - Name of assessed environment
- *Result* - Test score
- *Number of failed rules* (critical and high severity) - How many rules with critical and high severity levels failed
- *Number of failed and excluded tests* - How many tests failed or were not included.
- *Triggered by* - Event that triggered the assessment. The source can be:
  - *Manual* - For assessments run from the Rulesets page
  - *Policy* - For assessments made in *"Continuous Posture" on page 315*
  - *System* - For assessments defined on the CloudGuard dashboard, run hourly
- *Assessment Profile* - Applied ruleset
- *Assessment Identifier* - By default, the environment name

Generating an Executive Summary Report

Executive Summary Report allows you to see the status of your environments and assets based on the results of the last assessment. This report is for a specific ruleset and its assessment results in multiple environments on one cloud platform. It shows the environments with the highest number of severity findings. In addition, it shows the distribution of assets that passed or failed the test score, and the number of failed tests sorted by the rule severity.

To create the executive summary report:

1. Navigate to **CSPM** > **Assessment History** to show the list of assessments.
2. On the top right, click **Export** > **Executive Summary Report**.
3. In the window that opens, select a ruleset that is necessary for your report. Use the filter, search bar, or vendor grouping to find the applicable ruleset.

4. Click **Export**. CloudGuard creates the report in HTML format for the selected ruleset and all environments assessed with it. From the browser, print or save the page to PDF.
5. Click **Back** on your browser window to go back to CloudGuard.

To send the executive summary report by email:

1. Navigate to **CSPM** > **Assessment History** to show the list of assessments.
2. On the top right, click **Export** > **Executive Summary Report to Email**.
3. In the window that opens, select a ruleset that is necessary for your report. Use the filter, search bar, or vendor grouping to find the applicable ruleset.
4. In the **E-mail** field, enter one or more emails separated by a comma.
5. Click **Export**. CloudGuard sends a message with the link to the report in CSV format to the supplied email addresses.

To send the scheduled executive summary report in a Cloud Platform policy:

1. Create a new Cloud Platform policy.
2. Select the platform and the rulesets.
3. Add a new Notification. With only this type of policy, the Notification allows you to send the scheduled executive summary report to the selected email targets.
4. Select Scheduled report and, from the list, select **Executive summary report**. Configure the remaining parameters as usual - see *"How to Configure a Notification" on page 853*.
5. Click **Save** to save the notification.
6. Click **Save** to save the policy.

Viewing Assessment History

1. Navigate to **CSPM** > **Assessment History**.
2. On the top bar, set a filter or select a time frame. The time selector allows you to set the period back from the current time (4 h, 24 h, 7 d) or start and end dates for a custom time range. CloudGuard shows the assessments that align with your criteria.
3. Select an assessment to see its details.

Viewing Assessment Details

Results for assessments show the percentage of passed tests from the total number of tests run. A test is the application of a policy rule on a cloud entity. For example, applying a rule on an ES2 instance or S3 bucket is a test. The same rule applied to many entities results in many tests, each with its result.

1. Navigate to **CSPM** > **Assessment History** and select an assessment from the list.
2. The assessment results show all the failed rules in the assigned policy.

Creating Exclusions from Assessment

You can create an exclusion directly from the assessment based on the assessment details. This procedure is faster because CloudGuard enters some fields automatically.

1. Navigate to **CSPM** > **Assessment History**.
2. Use the Filter bar to search for the assessment that can be a basis for your new exclusion.
3. Click the assessment to see its details.
4. Select a rule and click **Expand** to show more details.
5. Below **Findings**, filter the entities by ID, name, environment, region, or network to select an entity whose finding is necessary to exclude.
6. Opposite the entity, click the **Exclude finding** icon ( ) in the **Actions** column on the right.

   The **Create New Exclusion** window opens, where most applicable fields are entered automatically.
7. Edit the fields as necessary based on the steps in *"Configuring CloudGuard Exclusions" on page 80*.
8. Enter a comment for the exclusion.
9. Click **Save**.

Creating Remediations

You can create remediation directly from the assessment based on the assessment details. This procedure is faster because CloudGuard enters some fields automatically.

1. Navigate to **CSPM** > **Assessment History**.
2. Use the Filter bar to search for the assessment that can be a basis for your new exclusion.
3. Click the assessment to see its details.

4. Select a rule and click **Expand** to show more details.
5. Below **Findings**, filter the entities by ID, name, environment, region, or network to select an entity whose finding is necessary to remediate.
6. Opposite the entity, click the **Configure remediation** icon ( ) in the **Actions** column on the right.

   The **Create New Remediation** window opens, where most applicable fields are entered automatically.
7. Edit the fields as necessary based on the steps in *"Adding Remediation" on page 319*.
8. Enter a comment for the remediation.
9. Click **Save**.

Generating a Tested Entities Report

Tested Entities Report allows you to see a specific assessment's status in detail. The report includes the results of passed and failed entities.

To create the Tested Entities Report:

1. Navigate to **CSPM** > **Assessment History** to show the list of assessments.
2. In the window that opens, drill down in the assessment that is necessary for your report. Use the filter, search bar, or vendor grouping to find the assessment.
3. Click the **Export** down arrow and select **Tested Entities Report**. CloudGuard creates the report in `.CSV` format for the selected assessment.
4. To go back to the **Assessment History** page, click the back arrow on your browser or close the **Assessment** tab.

Generating a List of Passed Entities

To export a list of passed entities to .CSV:

1. Navigate to **CSPM** > **Assessment History** to show the list of assessments.
2. In the window that opens, drill down in the assessment that is necessary for your report. Use the filter, search bar, or vendor grouping to find the assessment.
3. Click the **Export** down arrow and select **Export to CSV - Passed entities**. CloudGuard creates the report in `.CSV` format for the selected assessment.
4. To go back to the **Assessment History** page, click the back arrow on your browse or close the **Assessment** tab.

Generating a List of Failed Entities

To export a list of failed entities to `.csv`:

1. Navigate to **CSPM** > **Assessment History** to show the list of assessments.
2. In the window that opens, drill down in the assessment that is necessary for your report. Use the filter, search bar, or vendor grouping to find the assessment.
3. Click the report **Export** down arrow and select **Export to CSV - Failed entities**. CloudGuard creates the report in `.csv` format for the selected assessment.
4. To go back to the **Assessment History** page, click the back arrow on your browser or close the **Assessment** tab.

Governance Specification Language (GSL)

CloudGuard Governance Specification Language (GSL) is a syntax to define posture management rules, which can be included in rulesets in the CloudGuard Posture Management. GSL has a core language augmented by a set of functions that add domain-specific functionality for different cloud providers (AWS, Azure, GCP, Alibaba Cloud, and Kubernetes). These functions include IP addresses and networking, cloud entities such as instances, strings matching, date & time, etc.

Use Case

GSL is used to create compliance rules and Intelligence queries that you can run on your cloud environments to assess them. CloudGuard provides a graphical interface called GSL Builder or GSL Editor that helps you build the rule in the applicable context. For more information, see *"GSL Builder" on page 350*.

Rule Syntax

GSL rules have the form:

```
<Target> should <Condition>
```

The **Target** is the cloud entity type that the rule checks, for example, instance or SecurityGroup. Each rule can check only one target.

You can qualify the target to match a smaller set of entities that you add with the *where* keyword.

The formal definition of the target is:

```
<entity_type> where <expression>
```

Where:

- **entity_type** - The pivot cloud entity type that CloudGuard checks, for example, Instance, SecurityGroup, ELB.
- **expression** - Can be any complex GSL expression to match the target entity type.

Examples:

```
SecurityGroup where name='default' should...
```

```
Instance where (tags contain [key='env' and value='prod'] and not name='test') should ...
```

The **Condition** is the actual essence of the rule. It contains the attributes of the target that CloudGuard examines. The condition can be a complex GSL expression that tests some attributes related to the target.

Context Travel

When you use a command that references an array (`with, contain, contain-any, contain-all, contain-none, groupBy`, and so on), you execute an action that searches in the array's internal objects.

Each time you use array commands, you create a new context and delve deeper into the array. In a sense, you travel deeper into the child objects of the array which themselves can also be arrays. This is called *child context travel*.

You can use commands to refer back to the parent context and return to the parent array.

To do this, you can use some types of commands. Each command has a text representation and a symbol representation.

Command Symbol	Command Text	Description
$	**this()**	Reference the current array context level. The current array you have reached in the GSL query
~	**root()**	Reference the topmost level of the array context. The root level of any GSL query is the asset itself.
^	**parent(n)**	Reference the *n* level above the current array context.

To travel more than one context level, type more than one symbol (**^^** or **^^^**) or use its text equivalents, such as **parent(2)** or **parent(3)**.

Examples

1. The *$* refers to each role in the roles array.

```
ServiceAccount should not have roles contain-any [ $ like
'roles/iam.serviceAccount%']
```

2. This compares the tags value with the value in the ‘instanceType’ field.

```
Instance should have tags contain [ value like ~.instanceType
]
```

3. In this example, the ‘^.name’ equals to ‘IamUser.name’.

```
IamUser should have firstAccessKey with [ isActive = 'true'
and ^.name like '%d%']
```

Expressions

Wildcards

The *%* wildcard indicates zero or more characters.

For example, *like '%er'* matches the strings 'smarter', 'lover', 'her', etc.

The % wildcard can appear a number of times in an expression. For example, like *'%DB%'* matches 'MongoDB', 'DB123' and 'prodDB_111'.

Wildcards are used with *Like* and *Unlike*.

Examples:

1. This checks the region of an instance, and the name of the region must contain the characters “eu_”. Strings like ‘eu_central_1’ or ‘eu_west_2’ can match.

```
Instance should have region like 'eu_%'
```

2. This rule checks that ‘*’ character does not exist in the domain name.

```
AcmCertificate should not have domainValidationOptions
contain [ domainName like '%*%']
```

Comparison Expressions

These expressions are of the form:

```
<property_name> <comparison_operator> <expected_value>
```

Where:

- **property_name** - The applicable property of the entity being tested.
- **comparison_operator** - One of the operators: = , < , >, !=, <=, >= , like.
- **expected_value** - An expression you want to test with (can be a string, number, a different property, or a function result).

Examples:

```
Instance should have inboundRules.length < 10
```

```
Instance where image='ami-1234' should...
```

```
Instance where name like '%db%' should...
```

```
KMS where isCustomerManaged=true and deletionDate<=0 and
isSymmetricKey=true should have rotationStatus=true
```

The "like" Comparison Operator

The GSL *like* operator is almost the same as to 'like' syntax in SQL. It allows you to match text with wildcards (%). You can *like* on string values.

The like comparisons are case-insensitive, meaning that *'a' like 'A'* is true.

Example:

```
Instance where name like '%db%' should...
```

```
EcsService should not have role.combinedPolicies contain [name like '%admin%']
```

The "unlike" Comparison Operator

The *unlike* operator is the opposite of *like*.

The 'have' Operator

The keyword *have* is not mandatory. It makes a rule more readable. The GSL parser removes it when processing the rule.

```
Instance should image='ami-1234'
```

```
ELB should have accessLog.enabled=true
```

It is useful when you write rules that are also human-readable.

Inclusion Expressions (in)

The 'in' built-in function verifies that a given object's property value is found in a list of user-defined values.

Syntax:

```
<property_name> in (val1, val2, val3...)
```

Examples:

```
Instance should have image in ('ami-111','ami-222')
```

```
Instance should have region in ('us_east1', 'us_west_2')
```

```
Instance nics.length should be in (1,2)
```

```
IamUser should have mfaType in('Virtual','Hardware')
```

Contains Expressions

This rule makes that a property of the type 'list' contains a specific item. The syntax allows you to build complex expressions to find or match an item.

Syntax:

```
<property_name> <containment-operator> [<expression>]
```

Where:

- **property_name** - The property of the entity being tested.
- **containment-operator** - Is one of the following:
  - **contain** or **contain-any** - Are syntactically the same and return a Boolean True or False result.

    These options allow you to build rules that look correct in English. If any of the collection elements satisfies the query expression, then the full 'contain' statement is considered true.

    Examples:

    ```
    VMInstance should not have serviceAccounts contain
    [isDefaultServiceAccount=true]
    ```

    ```
    Region should have accessAnalyzers contain-any
    [status='ACTIVE']
    ```

  - **contain-all** - Requires all elements in the collection to satisfy the expression and returns a Boolean True or False result.

    Example:

    ```
    ApiGateway should have resources contain-all [ (methods
    contain-all [ apiKeyRequired=true ]) or (methods
    isEmpty()) ] or authorizers isEmpty() = false
    ```

  - **contain-none** - Requires none of the elements in the collection to satisfy the expression, and returns a Boolean True or False result.

    Example:

    ```
    SQLServer should have firewallRules contain-none [
    startIpAddress='0.0.0.0' and
    endIpAddress='255.255.255.255'
    ```

- **with** - Is almost the same as **contain**, but returns a list of elements matching the expression (where the number of items in the list can be counted).

- **Expression** - In the square brackets block `[ ]` can be any complex GSL expression. It defines the query to test each collection **item** with.

  If the list is blank, the expression returns as *False*.

Examples:

Iterate each rule in the SecurityGroups' inboundRules collection. Each rule is matched with the properties of port and scope:

```
SecurityGroup should have inboundRules with [port=22 and
scope='0.0.0.0/0']
```

This example is a **nested expression** - the outer 'with' iterates each nic in the nics collection, and the inner 'contain-any' iterates each security group:

```
Instance should not have nics with [securityGroups contain-any
[name='default']]
```

This example counts the number of IAM users with Admin access, and checks that it is not more than 3.

```
List<IAMUser> should have items with [combinedPolicies contain
[id='arn:aws:iam::aws:policy/AdministratorAccess']] count() <=3
```

Existence Expressions

These expressions make sure that a property exists and that it is not empty.

CloudGuard can identify if a property type is empty, including empty collections, empty objects, null, 0, or empty strings.

Syntax:

```
<property_name>
```

Examples:

- These are equivalent:

`Instance should have vpc`
`Instance should vpc`
`Instance should have not vpc isEmpty()`

- These are equivalent:

```
SecurityGroup where outboundRules...
SecurityGroup where outboundRules.length >0
```

Regular Expressions

These are regular expressions that match the text in an entity name or element. The expression 'regexMatch' can be applied to every element, including arrays and objects.

Syntax:

```
<property_name> regexMatch /<regex pattern string>/
```

Examples:

```
S3bucket should have name regexMatch /john.*/
StorageAccount should have networkRuleSet.bypass regexMatch
/.*AzureServices.*/
SystemManagerParameter where name regexMatch /
(pass|user|login|pwd|key|secret)/ should have
parameterType='SecureString'
```

Complex Logical Expressions

You can make complex expressions from multiple simple ones with the operators **and**, **or**, **not**, (<expr>)

AND

Make sure that the two sides of the **and** keyword are true.

Syntax:

```
<expr> and <expr>
```

Examples:

- The instance should satisfy the existence of property (name), and the existence of an element in a collection (tags). It could also be written as two separate rules.

  ```
  Instance should have name and tags with [key='owner']
  ```

- A rule that cannot be logically split as each element is validated considering the two conditions:

```
SecurityGroup should not have inboundRules with [port=22 and
scope='0.0.0.0/0']
```

OR

Make sure that the each side of the **or** keyword is a true expression.

Syntax:

```
<expr> or <expr>
```

Examples:

```
Instance should have name or tags with [key='owner']
```

```
AcmCertificate should not have status='FAILED' or
status='VALIDATION_TIMED_OUT'
```

NOT

Satisfies (returns true) if the following expression is false, and vice versa.

Syntax:

```
not <expression>
```

Examples:

```
SecurityGroup where name='default' should not have inboundRules
```

```
VMInstance should not have canIpForward=true
```

Parenthesis

Use parenthesis "(", ")" for complex expressions that contain multiple **and** and **or** expressions, to remove ambiguities.

Syntax:

```
(<expression>)
```

Examples:

```
SecurityGroup where name='default' should not have (inboundRules
or outboundRules)
```

```
S3Bucket should not have policy.Statement with [Effect='Allow'
and (Principal='*' or Principal.AWS='*') and Condition isEmpty
()]
```

Note - Without the parenthesis, the first rule is *ambiguous*:

```
SecurityGroup should not have inboundRules or
outboundRules
```

Because it is not understandable if you mean:

- SecurityGroup should not (inboundRules or outboundRules)
- SecurityGroup should (not inboundRules) or (outboundRules)

The parser uses one of the options, but it is recommended not to rely on an undocumented default implementation, and as an alternative, use '( )' to make the expression unambiguous.

List Expressions

This expression enumerates a list of items (defined by the following clauses in the GSL rule). The expression can be used with the *items* and *with* keywords, after which additional expressions can then be added. The target of the operation is the full list of items. If one or more items in a list fail a rule condition, the result of the rule evaluation is 'fail'. But, in an assessment, this is considered as one failure.

List expressions can also be used with the *groupBy* keyword to evaluate the number of items in an enumerated list. Grouping aggregates items (in a list-type data entity) by a specific attribute, which can then be used in logical expressions.

Syntax:

```
List<entity_type> should have/not have items with <expr>
```

Examples:

- This enumerates a list of instances, and checks that their number is less than a limit.

  ```
  List<Instance> should have items length() < 50
  ```

- The *with* keyword is used to filter the list for a specific name before it is enumerated, and then the result is checked with a condition.

  ```
  List<Instance> should have items with [name like 'db']
  length() < LIMIT
  ```

- This enumerates the Security Groups in each VPC, and checks that there are less than 100 in each (the AWS limit).

  ```
  List<SecurityGroup> should have items groupBy [vpc.id]
  contain-all [values length() < 100]
  ```

Data Types

GSL has different syntax for strings (textual values) and numeric values.

Text

Strings are surrounded by single quotation marks. For example, 'my string value'.

Examples:

```
IamGroup should not have managedPolicies with [name like
'AdministratorAccess']
```

```
VpcEndpoint should have tags contain [key like '%Name%']
```

Numbers

Numbers are written without quotes.

Example:

```
IamServerCertificate should not have expiration before(7,
'days')
```

Lists

List entities have a length property (*entity*.length). For example, inboundRules, in the example above. As an alternative, use the length() function.

Examples:

```
Instance should have inboundRules length() < 10
```

```
Instance should have inboundRules.length < 10
```

Functions

The core GSL syntax is enriched by internal functions that provide domain-specific functionality in multiple areas such as IP addresses, dates, and string matching.

Syntax:

```
<property_name> <function_name> (<param1>,<param2>...)
```

Where:

- **property_name** is the property/object we wish to operate on (equivalent to functions in object-oriented languages)

- **function_name** is the name of the functions from the above list params the required parameters based on the type of the function, separated by

General Functions

isEmpty

Checks if the object/property is empty based on the property type. It returns 'true' for empty collections and empty objects (and as null values, 0 empty strings)

Parameters:

None

Examples:

These are equivalent:

```
Instance should not have vpc isEmpty()
```

```
Instance should have vpc
```

```
EcrRepository should not have encryptionConfiguration.kmsKey
isEmpty()
```

length

Returns the length of a list. Follow this function with a comparator: >, <, =, >=, <=, !=.

Parameters:

None

Examples:

```
SecurityGroup should not have inboundRules length()>5
```

```
VMInstance should have labels length() > 0
```

in

Returns 'true' if the attribute value is contained in the provided list of values.

Parameters:

List of values to match.

Examples:

```
SecurityGroup should have region in ('us_east_1', 'us_east_2',
'us_west_1')
```

```
PostgreSQL should have logsConfiguration contain [ name='log_
retention_days' and value in ('4', '5', '6', '7')]
```

split

Splits a string based on a provided separator and converts it into an array

Parameters:

Position	Description	Values
1	Separator	The separator to use when splitting the string

Example:

Make sure that the instance name format is x-y-z (more tests could be applied on x, y and z)

```
Instance should have name split('-') length() = 3
```

join

Concatenates multiple strings into one string. This function can be useful to dynamically concatenate entity attributes.

Parameters:

Position	Description	Values
1	Separator	The separator to use when concatenating the strings
...	Strings	The strings to concatenate

Examples:

This example makes sure that the instance name is in the format "`PROD-<Instance ID>`" (for example, `PROD-i-a0b01c01`).

```
Instance should have name like join('-', 'PROD', id)
```

This example is for the instance names of the format like t2.micro-us_east_1.

```
Instance should have tags contain [ value like join('-',
~.instanceType ,~.region)]
```

inThisAccount

Returns true if the attribute value contains the environment ID or is equal to it.

Parameters:

None

Example:

```
Instance where vpc.vpcPeeringConnections length()>0 should have
vpc.vpcPeeringConnections contain-none [ not targetVpc.ownerId
inThisAccount() ]
```

containSecrets

Returns 'true' if a target property has values identified as secrets. The function works with entities, inner object, arrays, and single-string properties.

Parameters:

None

Examples:

1. All the instance fields and subfields should not have secrets.

   ```
   Instance should not containSecrets()
   ```

2. EcsTask should not have an inner vpc object that has fields containing secrets.

   ```
   EcsTask should not have vpc containSecrets()
   ```

3. None of the tags of the instance should contain secrets.

   ```
   Instance should not have tags containSecrets()
   ```

4. IAM role should not have external findings with resource names that contain exposed secrets.

   ```
   Iam should not have externalFindings.findings with [
   resourceName containSecrets() ]
   ```

Networking Functions - General

isPublicCIDR() / isPrivateCIDR()

These two functions check if an IP address or a CIDR is private or public. They work in the context of properties that represent IPv4 CIDR address ranges, for example, `scope`, `source`, `destination`, `prefix`, and `addressRange`.

Examples:

1. AWS in the context of `scope`:

```
SecurityGroup should not have inboundRules contain [ scope
isPublicCIDR() ]
```

2. AWS in the context of `destinationCidrBlock`:

```
RouteTable should not have routes contain [
destinationCidrBlock isPublicCIDR() ]
```

3. Azure in the context of `sourceAddressPrefixes`:

```
NetworkSecurityGroup should not have inboundSecurityRules
contain [ sourceAddressPrefixes contain [ $ isPublicCIDR()
]]
```

4. Azure in the context of `addressRange`:

```
Subnet should not have addressRange isPrivateCIDR()
```

Parameters:

None

numberOfHosts

Counts the number of possible IP addresses in the provided CIDR. It removes the broadcast address

Parameters:

None

Example:

```
SecurityGroup should not have inboundRules with [port=22 and
scope numberOfHosts() > 256]
```

containedInNetworks

Operating on an IP address or a CIDR network and testing if it is fully contained in any of the provided networks.

Parameters:

List of CIDRs representing networks to test

Example:

This rule checks that there is no SSH connection from non-local or 'friendly' networks: '1.2.3.4/10','5.6.7.8/24'

```
SecurityGroup should not have inboundRules with [port=22 and
scope isCIDR() and not scope containedInNetworks
('10.0.0.0/8','172.16.0.0/12','192.168.0.0/16','1.2.3.4
/10','5.6.7.8/24')]
```

overlapWithNetworks

Operating on an IP address or a CIDR network and testing if it has some overlap with any of the provided networks.

Parameters:

List of CIDRs representing networks to test

Example:

In this example, '10.1.2.0/24','192.168.100.0/24' (internal) networks belong to a third party peered with our VPC. Make sure that you do not allow these networks to SSH into the servers.

```
SecurityGroup should not have inboundRules with [port=22 and
scope isCIDR() and scope overlapWithNetworks
('10.1.2.0/24','192.168.100.0/24')]
```

isCIDR

Check if a field is in CIDR format.

Parameters:

None

Example:

```
SecurityGroup should not have inboundRules with [scope isCIDR()]
```

isSecurityGroupReference

Checks if the attribute references a Security Group.

Some attributes can contain multiple types of objects (that is, a Security Group or a CIDR). This function determines if the value is a Security Group.

Parameters:

None

Example:

This example makes sure that security groups should only allow traffic from a different security group, and not a CIDR

```
SecurityGroup should have inboundRules with [scope
isSecurityGroupReference()]
```

Networking Functions for AWS NACL and MS Azure NSG

AWS NACL and MS Azure NSGs have different firewall semantics.

The firewall rules are in an order and may contain explicit 'DROP'. This makes the order of the rules very important.

These functions operate on a list of rules.

allowedHostsForPort

Returns the number of IP addresses (hosts) that can be connected to the applicable port protected by this firewall.

Parameters

Position	Description	Values
1	Port number (mandatory)	any valid port number, such as 22, 80, and so on

Example:

This rule prevents having overly permissive port 22 (that is allowed to more than 256 addresses). This is a rule for Azure Virtual Machines.

```
VirtualMachine should not have nics with
[networkSecurityGroup.inboundRules allowedHostsForPort(22) >
256]
```

allowedPublicHostsForPort

Returns the number of public IP addresses (hosts) that can be connected to the applicable port protected by this firewall.

Parameters:

Position	Description	Values
1	Port number (mandatory)	any valid port number, such as 22, 80, and so on

Example:

This rule prevents having overly permissive port 22 (that is allowed to more than 256 addresses). This is a rule for Azure Virtual Machines.

```
VirtualMachine should not have nics with
[networkSecurityGroup.inboundRules allowedPublicHostsForPort(22)
> 256]
```

areCIDRsAllowedForPort

Returns 'true' if the provided IP addresses / CIDRs are allowed to connect to the specified port.

This method returns 'true' if one or some of the addresses in the provided CIDRs are allowed to connect.

Parameters:

Position	Description	Values
1	Port number (mandatory)	a valid port number, such as 22, 80, and so on
2	CIDR to test (a minimum of 1 is mandatory)	CIDR syntax like 1.2.3.4/24
...	other optional CIDRs	CIDR syntax like 1.2.3.4/24

Example:

This rule prevents SSH access from not approved third-party networks such as `1.2.3.0/24', '5.5.6.7/32.`

```
VirtualMachine should not have nics with
[networkSecurityGroup.inboundRules areCIDRsAllowedForPort
(22,'1.2.3.0/24','5.5.6.7/32')]
```

isPortPrivate

Returns 'true' if the specified port can only be accessed by internal IP addresses

This method returns 'true' if one or some of the addresses in the provided CIDRs are allowed to connect.

Parameters:

Position	Description	Values
1	Port number (mandatory)	Any valid port number, such as 22, 80 and more

Position	Description	Values
...	Additional CIDR to mark as private (optional)	CIDR syntax like 1.2.3.4/24

Example:

This rule prevents SSH access from public IP addresses.

```
VirtualMachine should have nics contain-all
[networkSecurityGroup.inboundRules isPortPrivate(22)]
```

Resource Functions

These functions return values for properties of assets. Especially, they can return secondary values for assets (for example, the value of a rule address for an SG assigned to an EC2 instance).

getResource

1. Returns the first instance of a resource, based on a field and value.

   Parameters:

Position	Description	Value
1	resourceType	An asset property, for example, 'subnet', or 'vpc'
2	resourceFieldValue	The value of a resource property
3	resourceFieldName	The name of the resource property

   Example:

   ```
   Instance should have getResource('Subnet', subnet_name,
   'name') with [cidr = '172.31.32.0/20']
   ```

2. Returns the first instance of a resource, given the type and an id.

   This is almost the same as the instance above, only the field is assumed to be the ID of the specified asset type.

   Parameters:

Position	Description	Value
1	resourceType	An asset property, for example, 'subnet', or 'vpc'
2	resourceId	The resource id to be returned, for example, subnet_id

Example:

```
Instance should have getResource('Subnet', subnet_id) with
[cidr = '172.31.32.0/20']
```

```
SecurityGroup should have getResource('VPC', property_with_
the_vpc_id) getValue(property_in_VPC_entity) = true
```

3. For AWS KMS entities only, returns the last instance that matches one of the resources below and an ID.

   Parameters:

Position	Description	Value
1	resourceType	A constant value: **KMS**
2	KMSId	The resource id to be returned, whatever of these matches: ■ arn ■ id ■ aliases.arn ■ multiRegionConfiguration.primaryKey.arn ■ multiRegionConfiguration.replicaKeys.arn

   Example:

```
SecretManager should have getResource('KMS', kmsKeyId)
getValue('isCustomerManaged') = true
```

getResources

This function returns a list of instances of a resource, based on a field and a value. This is almost the same as *getResource* above but returns a list of all instances matching the field and value.

Parameters:

Position	Description	Value
1	resourceType	An asset property, for example, 'subnet', or 'vpc'
2	resourceFieldValue	The value of a resource property (optional value)
3	resourceFieldName	The name of the resource property (optional value)

Example:

```
Instance should have getResources('Subnet', subnet_name, 'name')
contain-all [cidr = '172.31.32.0/20']
```

getValue

Returns the value of an object, given its path.

Parameters:

Position	Description	Value
1	path	The path to a resource, for example, 'vpc.cidr'

Examples:

```
Subnet should have nacl getValue('vpc.cloudAccountId') =
'********-****-435a-959e-ae3dab323de5'
```

You can use getValue(n) with numeric arguments to index lists.

```
'arn:aws:iam::123123123123:role/name' split(':') getValue(0) =
'aws'
```

getValues

Returns an array of values that matches the provided path parameter.

Parameters:

Position	Description	Value
1	path	The path to a resource, for example, 'vpc.cidr'

Example:

```
HDInsight should have properties.computeProfile.roles contain [
getResource('Subnet', virtualNetworkProfile.subnet) contain[
routeTableData.routes contain [ nextHopIpAddress in(getResources
('Firewall') getValues('ipConfigurations.privateIPAddress') ) ]
] ]
```

portsContainedInRange

Supports AWS NACL, AWS Security Group, MS Azure Network Security Group, and GCP Security Group assets.

Verifies that a minimum of one port, which is part of the port list, also exists in the application configuration range.

Parameters:

Position	Description	Value
1	Port number or GenericList	GenericList or any valid port number, such as 22, 80 etc...

Examples:

```
SecurityGroup should have inboundRules contain [$
portsContainedInRange($portList)]
```

```
SecurityGroup should have inboundRules contain [$
portsContainedInRange(15, 25, 35)]
```

Time Functions

before

Test if the date property in the UNIX time format is before, that is, smaller than the desired relative time.

Time is relative to real-time evaluation time rather than the date of rule authorship and can be defined with each of the supported time units.

Parameters:

Name	Description	Values
count	The number of the units; negative numbers resolve into past dates	Any positive or negative number such as -5, 10, ...
unit	The measurement unit	A string - 'minutes', 'hours', 'days', 'months'

Examples:

This verifies that no instance was launched more than three months ago.

```
Instance should not have launchTime before(-3,'months')
```

This verifies that no certificate of an ALB-secured listener expires in seven days.

```
ApplicationLoadBalancer should not have listeners contain [
certificates contain [ expiration before(7, 'days') ] ]
```

after

Test if the date property in the UNIX time format is equal to or is greater than the desired relative time.

Time is relative to real-time / evaluation time rather than the date of writing the rule and can be defined with each of the supported time units.

Note - Unlike **before()**, this function also includes the desired time.

Parameters:

Name	Description	Values
count	The number of the units; negative numbers resolve into past dates	Any positive or negative number such as -5,10, and so on
unit	The measurement unit	A string - 'minutes', 'hours', 'days', 'months'

Examples:

This verifies that user with an enabled password used it during the last 45 days.

```
IamUser where passwordEnabled='true' should have
passwordLastUsed after(-45, 'days')
```

Do an action for the instances that were launched in the last two weeks.

```
Instance where launchTime after(-2,'weeks') should...
```

Best Practice - To use fixed (non-relative) time, you can use a standard numeric comparison with the UNIX date/time format.
For example, to do something with all instances that were created before 1/1/2016 (which translates into the UNIX time 1451606400), use:

```
Instance where launchTime < 1451606400 should ...
```

dateDifference

Returns the time difference between two dates, in time units of your selection (*outputUnit*).

Parameters:

Position	Description	Values
secondDate	The other date to compare to	date (string)
outputUnit	The desired time unit of the output	'seconds', 'minutes', 'hours', 'days', 'months', 'years' (string)

Date formats:

- Unix time (time in seconds)
- YYYY-MM-DDTHH:mm:ss.SSSZ
- YYYY-MM-DDTHH:mm:ss.SSSSSSSZ
- YYYY-MM-DDTHH:mm
- YYYY-MM-DD
- MM-DD-YYYY
- MM/DD/YYYY HH:mm:ss
- Month (as string) DD, YYYY HH:mm:ss

Tip:

To create a date in local time, include the time explicitly as in '2019-06-11T00:00'.

When you create a date without specifying the time, you get the date set in UTC.

Examples:

Check if the creation date (*createDate*) of a user (*IamUser*) is equal to or less than 90 days from the input date ('2022-01-17T11:09:00.000Z').

```
IamUser should have createDate dateDifference('2022-01-
17T11:09:00.000Z','days') <= 90
```

Check if the time difference between the time of creation (*createdDateTime*) of a user (*User*) and the input date is equal to 10 minutes.

```
User should have createdDateTime dateDifference('01/14/2022
10:05:00', 'minutes') = 10
```

Check if the *updateDate* of a *RamUser* is equal to or less than 20 days from the input date.

```
RamUser should have updateDate dateDifference('October 13, 2021
11:13:00', 'days') <= 20
```

GSL Builder

The GSL Builder is a sandbox that helps you write and test GSL rules. Some parts of the GSL Builder appear as a GSL Editor when you start to create a new rule for your environment. The GSL Builder provides an interactive graphical interface for all supported entities. CloudGuard constantly updates the list of entities and attributes.

Building a New Rule

To build a GSL rule, follow the steps below:

Step 1 - Select a cloud platform

In CloudGuard, you can build GSL posture management rules for multiple cloud platforms.

1. Navigate to the **GSL Builder** page in the **CSPM** menu.
2. Select a cloud platform to run the rule on: AWS, Azure, or other.

The platforms with the **Preview** tag are at Early Availability.

All GSL rules and rulesets are platform-specific, so you can run a rule only on the platform it is created for.

Platforms can be static or dynamic:

- Static platforms - AWS, Azure, GCP, Alibaba Cloud, Kubernetes, Image Assurance, and Source Code Assurance. CloudGuard recognizes all supported resources (entities, services, and so on) on static platforms.
- Dynamic platforms - Terraform, AWS CloudFormation. CloudGuard cannot predict which resources an IaC plan contains until you upload the plan files. Then after CloudGuard reads and analyzes the script, it builds a set of entities based on data that the script provides. To build rules on a dynamic platform, see Compliance Assessment of Infrastructure as Code.

Because of different naming conventions, the entity's names on the static and dynamic platforms are not the same.

Step 2 - Select the editor mode

1. Select the mode:
   - **Builder** - Enter the new rule interactively with hints.
   - **Free text** - Enter the rule as text.
2. Optionally, change **Builder** to **Free text** to make small changes before you test the rule.

Important - Make sure not to change back from **Free text** to **Builder**. If you do, CloudGuard erases the expression in the field.

Step 3 - Select the entity

For each cloud platform, CloudGuard provides a set of applicable entities. You build a rule with one entity, which is the rule Target, and a combination of the entity attributes, which are the rule Condition.

Note - The **New** tag indicates that recently CloudGuard started to support the entity.

1. Select an entity from the list.

   CloudGuard shows possible actions below, in the interactive section. The attributes (properties) for this type of entity appear on the right below **Context Preview**.
2. Select one of the actions. CloudGuard shows applicable operators and properties in the interactive section.
3. Select an operator or a property to continue to build a rule. Use the context preview on the right to expand properties and see their structure.

   Note - Some asset properties have the indication *External*, for example, **ExternalObject** or **ExternalArray**. These properties are brought by CloudGuard from another entity to help you write GSL rules easier.

   Note - Context preview contains information on static elements only. Information on dynamic elements, such as objects, is not available.
4. To continue with the rule, add more operators, properties, and functions as the context suggests. To learn more about the GSL syntax, see *"Governance Specification Language (GSL)" on page 326*.
5. To delete one or more elements in the rule, put the cursor on these elements, and click (this gives approval to the deletion).

Step 4 - Test the rule

1. Below **Test Rule**, select one or more environments to test the rule on and click **TEST**.

   CloudGuard runs the rule and shows the **Result Details**, almost the same as the Assessment Results (see *"Running an Assessment" on page 304*). The **Test History** shows the ten last instances when the rule was run.

2. Click **clear result** to erase the test result.

3. Click **clear tests history** to erase the history.

After a successful test, select **Free text**, copy the GSL expression, and paste it into the GSL section of a new rule (see *"Adding Rules to a Ruleset" on page 313*).

Actions

You can export information about protected assets to a CSV file.

To export protected asset information:

1. Click **Export** in the upper right and select the detailed view by asset type.
2. Select an applicable platform, environment or OU, and the asset type.
3. Click **Done**.

Network Security

This section describes how to manage network security groups with CloudGuard and how to control access to your protected cloud assets with short-term leases.

Configuration Explorer

CloudGuard Configuration Explorer gives graphical visualizations of the network security of your cloud environment. It shows the hierarchy and structure and your cloud assets and their interconnectivity. These views are arranged to show the level of exposure to the external world. From this, you can identify assets that are misconfigured in the network, or overly exposed. In addition, you can drill down from these views to see details in CloudGuard for the Security Groups or assets and make corrections directly in CloudGuard.

You can use Configuration Explorer to analyze your cloud network for toxic combinations, such as access to sensitive components from the Internet. Or you can troubleshoot it for connectivity issues such as blocked paths to components.

Benefits

- Logical visualization of inbound traffic to your VPC and its components, and the cloud perimeter
- Visualize complex networks (for example, with many instances, cross-VPC, cross-region)
- Easily identify toxic combinations, blocked paths
- Agent-less & automated information gathering from Cloud environments
- Automatically classifies protected cloud assets based on the level of exposure to the outer world
- Real-time topology map of security groups, and the interrelationships between security policies
- Visualization of traffic flow and dropped traffic between cloud assets - security groups, instances, etc.
- Real-time topology view of cloud assets
- Visibility into the interplay between security policies for multi-tier applications and the effective security posture in a cloud environment
- Similar cross-cloud security visualization experience
- Contextual VPC Flow Logs
- Visualize virtual networks connectivity

Configuration Explorer Views

Configuration Explorer offers different views of your Security Groups, each highlighting different aspects of the Security Groups.

The following sections tell how to select and then visualize a cloud environment in Configuration Explorer, how to use the different views, and the actions you can do to see more information.

Security Group View

This view shows the relationship between the Security Groups in your network. They are grouped logically based on exposure to the Internet. Their interconnections are shown.

The steps below tell how to select a network and open this view, and then how to navigate and use the view.

Select a cloud network

1. Go to **Network Security** > **Configuration Explorer**. A list of your cloud accounts appears on the left.
2. Select an account from the list. A list of regions is shown, in which you have cloud environments. The numbers in brackets show the number of assets in the environment.
3. Select a region.

   The VPCs in the region are shown as nodes (circles). For an AWS account, the connections between them show peering connections between VPCs.
4. Click a node. This represents a cloud network. The pane on the right shows the type and number of assets in the network.
5. Click the **Legend** button on the bottom left. The node color designates the VPC location:
   - Orange: VPC from the selected region
   - Blue: VPC from another region
   - Green: VPC from other Cloud Account
   - Gray: VPC from external CloudGuard Account

The Security Group view

In this step, the network selected in the previous step is shown in the Security Group view. This view is available for AWS cloud accounts only.

1. Go to **Network Security** > **Configuration Explorer** > select a VPC (in the previous section) and then select **Security Group** from the menu bar on the top right.

   This shows the Security Group view of the environment.

   This view shows the following:

- The view is divided into logical zones that show the level of exposure to the external world, from the External zone (red, at left), the most exposed, to the Internal zone (green, at right), the least exposed. Security Group nodes are located in the view based on their level of exposure.
- Each Security Group is shown as a node. Security Groups that are managed by CloudGuard in Full-Protection are shown like this:

  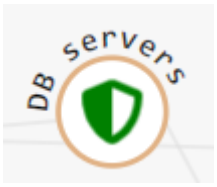

  Security Groups that are managed as Read-Only are shown like this:

  

- In addition, Sources are shown as nodes, with their IP address. These are typically in the External zone (external sources) and the Internal zone (instances).

  

- Lines between sources and Security Groups show that the address is controlled by the Security Group, that is, that a rule in the Security Group affects the address.
- Lines between Security Groups show that one Security Group affects the other (by a rule).

2. Click on a Security Group node. Other Security Groups affected by it are highlighted, with the direction of the arrow that shows if the other group affects this Security Group (the arrow points to the selected Security Group), or is affected by it (the arrow points from it).

   The pane on the right shows details for the Security Group. This shows the following:

   - Sources of inbound network traffic (an external source, or a different Security Group)

- Targets
- Assets controlled by rules in the Security Group
- Inbound and outbound rules

Click the link symbol for a Security Group source or target, to open the CloudGuard page for it.

3. Click on an external source node (on the left, in the External zone). The node is highlighted in the view, and the Security Groups that affect this source are highlighted. The detail pane on the right shows the IP address for the source.

Asset View

This view shows your cloud assets, such as instances and database servers, and the connections between them. Each node in this view shows an asset. They are grouped logically based on their exposure to the internet. Their interconnections are shown.

This view is available for all cloud providers.

The Asset View

1. Go to **Network Security** > **Configuration Explorer** > Select a VPC (*"Select a cloud network" on page 355*), and then select **Asset View** from the menu bar on the top right.

   This shows the Asset View of the selected network.

   This view, like the Security Group view, is divided into zones based on the level of exposure.

   The view has these elements:

   - Each node is an asset.

     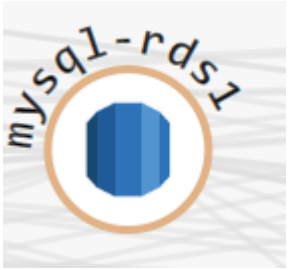

   - Lines between nodes are network connections.

2. Click a node. The view highlights the connections to other assets, with the direction of the arrow showing the direction of the connection.

   The pane on the right shows details for the asset, including the source and target connections (assets or Security Groups that the asset can receive from or send to)

Click a source to see more detail about it on the CloudGuard Protected Asset page.

Click the **Flow Logs** link at the bottom of the detail pane to see VPC Flow Logs, filtered for the selected asset.

Effective Policy Grouping

The Effective Policy grouping in the Asset view groups nodes (assets) together if they are affected by the same security groups.

The pane on the right shows the grouped assets.

Show Peered VPCs

You can see assets in peered VPCs in the Asset View. Move the **Peer VPC** switch at the top of the graph, to enable this.

Navigation and Controls

You can use the following controls from the menu bar to change the Security Group or Asset views.

Button	Description
	Zoom the view in or out.
	Expand or close groups in the view (based on the selected grouping).
	Group external sources or Security Groups based on the selected parameter (affects the same assets).
	Search for elements by name. While you enter text in the text box, the elements with the name that match the text show from the list below the text box. In addition, same elements become highlighted in the visualization map. When you select an element in the search list, the same element is selected it in the map.

Traffic Explorer

Traffic Explorer helps you visualize events of interest in the network traffic of your environments and Kubernetes clusters. It gathers and presents information from environment logs and workload network logs, enriched with information from more sources such as threat intelligence feeds, IP reputation databases, and geolocation databases.

You can find when services, applications, or databases are exposed to the Internet and if there are possible data exfiltration attempts.

Traffic Explorer has three primary views:

- Graph view
- Logs view
- Statistics view

Graph View

The Graph view shows network traffic in your cloud environment or Kubernetes cluster, based on the collected flow logs.

Graph view elements

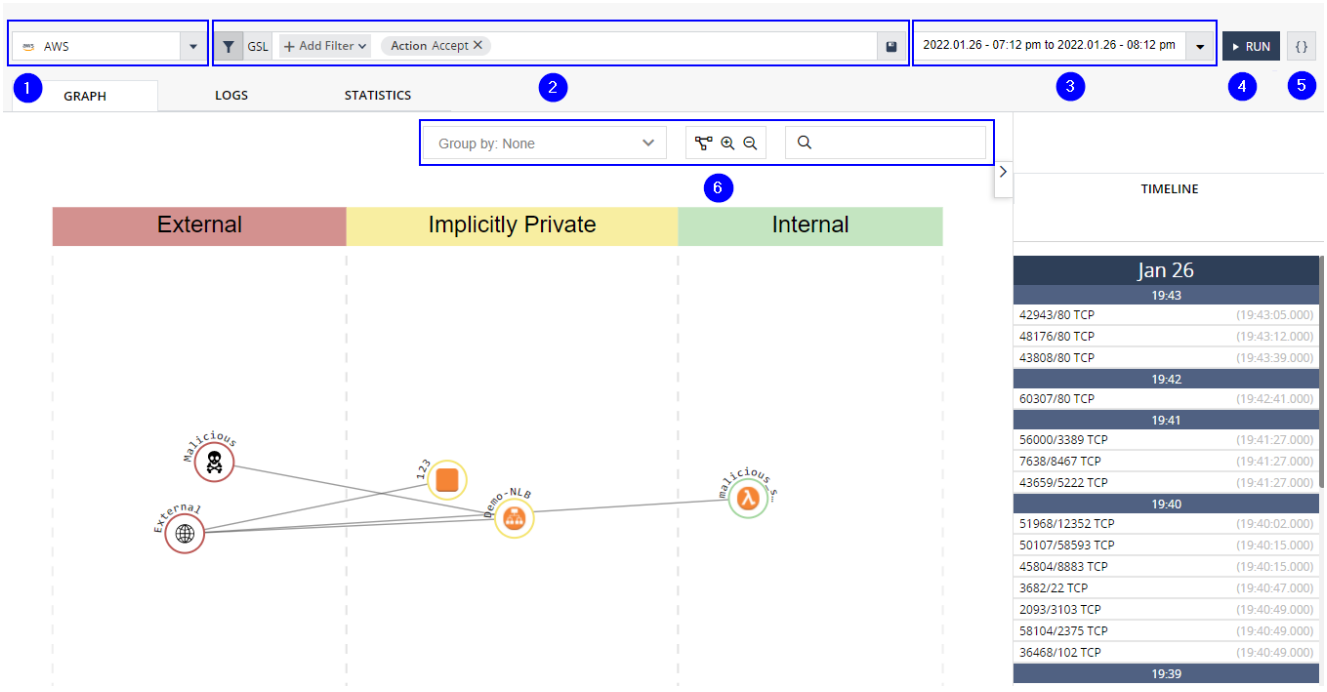

Item	Description
1	Environment selector

Item	Description
2	GSL query
3	Time frame menu
4	Run button
5	Query menu
6	Group and Zoom buttons

Toolbar

The toolbar at the top of the page contains these buttons:

- The Environments list **(1)** on the top left contains only environments with Traffic Activity enabled (onboarded to Intelligence with Flow Logs).
- The GSL query **(2)** enables you to search for network resources or network flows. You can specify in the query the details of specific packets, bytes, source, or destination to monitor traffic and interconnectivity of the resources that belong to your environment or cluster. Edit the query text directly in the box or open a graphic query editor.

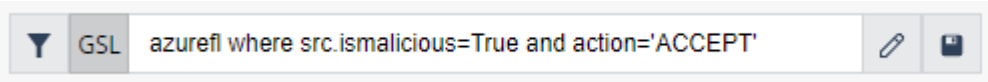

- The time frame **(3)** is the period back from the current time (15 min, 1 h, 24 h, 7 d) or start and end dates for a specific time range. To change the time frame for the view, select a new value and click **Run (4)** to run a new query.

Note - The graph shows only actual traffic between entities. Entities without network activity during the selected time frame are not seen.

- The **Queries** icon  **(5)** allows you to select a query from an applicable category.

The central part shows a graph of the environment entities and the network traffic between them, based on the query and the time frame. Each graph node represents a cloud asset or a Kubernetes entity (pod, node, or service). The entities are grouped into zones (for example, External, DMZ, and Internal) according to the exposure of the entity to the Internet. External entities are exposed, so they have Internet addresses, while Internal entities have no exposure to the Internet.

Note - Kubernetes pods that run on a host network do not appear as different entities, so their traffic combines with the traffic of the applicable node.

Groups

Based on the environment platform, you can group the visible assets by common characteristics:

- **Effective Policy** (AWS, Azure) - Unifies assets with the same policy.
- **Name** (AWS, Azure) - Unifies assets below the same name.
- **VPC** (AWS, Azure) - Unifies assets that belong to the same VPC.
- **Services** (Kubernetes) - Unifies assets by services or workloads (as Controller groups). For example, if pods are part of a DaemonSet or a Service, they appear as a group with an option to ungroup.

  Pods without an owner reference, and cluster nodes, appear below the **No Services** Group.
- **Namespace** (Kubernetes) - Unifies all pods or services in the same namespaces.

Click one of the groups. The pane on the right shows a list of entities in the group. Groups of size 1 appear as regular graph nodes.

Zoom Controls

You can control the view with these tools:

- **Zoom**: Select a point in the center section of the view and use your mouse scroll wheel to zoom the display in or out. Alternatively, use the zoom controls at the top of the view ()
- Select an entity in the view to show its details in the right pane. Many details are links to more information.

Click in the central area of the view (not on an entity) to go back to the previous view, of all entities.

Logs View

You can examine the actual log information for a selected entity. The information is based on the flow logs and enriched by Intelligence with more contextual information.

To open the Logs view, do one of these:

- Navigate to the **Traffic Explorer** page in the **Network Security** menu and click the **Logs** tab.
- Navigate to the **Network Traffic** page in the **Events** menu.

Click an entity in the table to see its details.

Click entities in the details pane to add them to the query and narrow the query to specific items of interest.

Statistics View

You can see the statistics for the environment network traffic. Some traffic statistics are based on the network logs that match the GSL query in the specified time frame.

You can add more filter options to the query from the statistics elements to focus on specific results. Click an element and select the logic with which to add it to the query (AND, OR, NOT). You can add multiple filter additions.

Actions

Viewing Traffic Explorer Graph

This view shows network logs filtered by a query.

1. Navigate to the **Traffic Explorer** page in the **Network Security** menu. The Graph view opens with the default query applied to the flow logs of an environment.
2. From the **Select Environment** list on the top left, select the name of your environment.
3. Click  on the toolbar to open the Query menu.
4. Select a query from the list.
5. The page updates to apply the selected query to the log data for the environment.
6. You can change the query for this view at any time. Open the query menu and select a new query. You can also enter a new query directly in the query box or build one interactively with the editor.

Filtering the Traffic Explorer

You can filter the Traffic Explorer views to focus on events of interest.

Graph filter

Select filter options in the filter bar. You can select more than one option; all are applied to the view.

- **Classifications** - A swimlane zone in the view (for example, External, DMZ, or Private)
- **Action** - A specific traffic action (ACCEPT, REJECT)
- **Malicious** - select malicious or non-malicious sources of traffic
- **Region** (AWS, Azure) - A specific cloud region
- **VPC** (AWS, Azure) - A specific cloud VPC in the environment

- **Source Asset Type** - Select specific asset types for traffic source
- **Destination Asset Type** - select specific asset types for a traffic destination

After you make changes to the filter, click **Run** to re-run the query and update the graph.

Statistics filter

You can add more terms to the query based on statistics results in the Statistics view.

1. Click a detail on a statistics widget.
2. Select the logic to add this to the query (AND, OR, NOT). The selected property adds to the query string.

More Links

- *"Activity Explorer" on page 399*
- *"Cloud Detection and Response (CDR)" on page 565*
- *"Intelligence for Kubernetes Containers" on page 645*

Manage IP Addresses

In this section of the CloudGuard portal, you can see the IP addresses allocated to your cloud resources and which rules (security groups) reference or affect them. This includes internal and external (elastic) IP addresses. In addition, you can label addresses and classify them based on their logical location in the network (external, internal, DMZ). You see these labels and classifications when you view your VPC with the Configuration Explorer (see *"Configuration Explorer" on page 354*).

You can configure lists of IP addresses. You can apply a Security Group policy to the list as an alternative to applying it to each IP address.

Benefits

In the portal, you can see all your IP addresses at a glance and see the rules that reference each address.

When you attach a label to an address, you make it easier to identify in the Configuration Explorer.

Similarly, you can easily group addresses into lists from the portal and apply security policies on lists. This is a simple way to apply policies and decrease possible errors.

Use Cases

Typical use cases for IP addresses management from the CloudGuard portal are:

- Identify IP addresses used in a VPC or find to which VPC an address is related, see *"View IP addresses" below*
- Associate IP addresses with a list, see *"Define IP Lists" on the next page*
- Review the security rules associated with an IP address or group of addresses, see *"Define IP Lists" on the next page*

Actions

View IP addresses

You can see a list of all your IP addresses across all platforms and environments and details for a selected address.

1. Go to **Network Security** > select **IP Address**. A list of all IP addresses shows for all your VPCs on each of your environments.
2. To show more details, click an IP Address.

Classify an IP address

CloudGuard classifies IP addresses as Internal (to the VGC network), External (access from the Internet through an Internet gateway), or DMZ (partial access). This classification is shown in the Configuration Explorer for your VGC network (see *"Configuration Explorer" on page 354*). You can classify each of your IP addresses.

1. Select an IP address from the list, to open the details view for it.
2. Click **Classify IP**.
3. In the window, enter a name for the address that appears later from the list and in the Configuration Explorer and select a classification.
4. If the address already has a classification, click **Edit Classification** to change it or **Delete Classification** to delete it.

   In the Configuration Explorer with your VPC, you can see the IP address in the selected classification with the assigned name.

Define IP Lists

IP lists are groups of IP addresses. You can apply a security group (rule) to an IP list as an alternative to applying the rule to each IP address.

1. Go to **Network Security** > select **IP Lists**. A list of all IP Lists shows.
2. To add a new list, click **Add**. Or select one of the lists shown on the left to show a list of the IP addresses included in it (on the right).
3. Enter new IP addresses in the box, and click **Add**. When done, click **Save**.
4. Click **Delete** to remove IP addresses from a list.

   Click **Delete IP List** to delete a list. This deletes the list but not the IP addresses.

Security Groups

The Security Group section shows the security groups for all your CloudGuard-managed environments. If use CloudGuard to fully manage your environments, then you can actively manage your security groups from CloudGuard: configure new security groups, change them, and apply them to your environments.

In addition, you can review your security groups for all your cloud accounts, on all platforms, in one place. For CloudGuard-managed accounts, you can also apply changes in one place to all accounts. When you configure and apply changes in one place, you make sure that your security groups are harmonious and comprehensive across your cloud presence.

When you make changes in environments through the CloudGuard portal, you decrease the chance of accidental or non-approved changes to your security groups. In addition, if unwanted changes are accidentally made through the cloud account, CloudGuard detects this activity. If CloudGuard fully manages the security groups, it automatically rolls them back to the earlier authorized settings made on CloudGuard and notifies you about it.

Use Cases

Typical use cases to illustrate how you can control your Security Groups from one central location:

- Search for Security Groups - To quickly search for specific security policies across all of your cloud presence, see *"Viewing your Security Groups" below*.
- Review security posture - To analyze your security position effectively by reviewing all your policies in one view, see *"Continuous Posture" on page 315*.
- Apply equal changes - If you expand or change your cloud presence, or add more services or regions, you can change the security policies harmoniously for all regions from one portal - see *"Configuring CloudGuard Policies" on page 78*.
- React to anomalous behavior - If changes are made to one of your cloud accounts, accidentally or maliciously, you immediately receive a notification, which allows you to do corrective actions - see *"Notifications" on page 852*.

Actions

Viewing your Security Groups

To open the primary Security Group page, navigate to **Network Security** > **Security Groups**.

The primary page shows a list of all your managed security groups, which includes all your CloudGuard-managed environments on all cloud platforms. Use the search box or filter options to filter the list.

Generating Security Group Reports

You can export your security groups to a `CSV` file.

To create a security group report:

1. Navigate to **Network Security** > **Security Groups** to show the list of security groups.
2. Filter the list to show only applicable security groups.
3. On the top right, click **Export** and select one of these options:
   - **Export to CSV - Filtered Security Groups** - Download the file with the applied filter criteria.
   - **Export to CSV - All Security Groups** - Download the file with all security groups in your account.
   - **By OU\Environment - Detailed** - Select a platform and environment or Organizational Unit, and receive an email. The email contains a link to the protected assets page. From this page, you can save the downloaded file to your computer.

More Links:

- *"AWS Security Groups" on page 368*
- *"Full Protection Mode" on page 373*
- *"Azure Network Security Groups" on page 374*

AWS Security Groups

This section describes how to create and change AWS Security Groups in the CloudGuard console. The account for which the Security Groups are created must be in Full Protection mode (this allows Security Groups to be managed by CloudGuard as an alternative to the AWS console).

Creating a new AWS Security Group

You can create a new AWS Security Group for a VPC (your account must be Full-Protection to do this).

1. In the CloudGuard portal, navigate to **Network Security** > **Policy** > **Security Groups**.
2. In the filter field, select AWS accounts.
3. To add an account to a Security Group, click the  icon opposite the account.
4. Enter a name and description for the new security group and click **Add**.
5. Add Inbound and Outbound Services to the group.
    a. Select the details for the Service.
    b. In addition, set the Port Behavior as Open or Limited. For Limited behavior, add the source IP addresses to be accepted as individual IP addresses, IP Lists, or a different AWS Security Group.
    c. Click **Create Service**.
6. Add tags to the service, which allows it to be searched. Enter a Key (name) and a Value for the tag, then click **Create**.

Changing details for an AWS Security Group

You can change details for Security Groups (the Security Group must be in **Full Protection** mode to allow this).

1. Click the link for the AWS Security Group you want to change.
2. Click  to add a new service (inbound or outbound), **Edit** to change details for a service in the Security Group, or **Delete** to delete it.
3. You can change the name and description of the service. You can add sources to the services.

Cloning a Security Group

You can clone an existing Security Group to make a copy of it. The copy has the same configurations (services, and more). You can select to apply the new Security Group to the same VPC, or a different one.

1. Click the link for the AWS Security Group to clone, and then click **Clone**.
2. Enter a name and description for the new Security Group. If it necessary to associate it with different VPCs, select **Other VPCs**.
3. Select the Account, Region, and VPC from the lists, and then click **Add** to associate the Security Group with a VPC. You can associate it with more than one VPC.

Setting an AWS Security Group to Full Protection

You can change the protection mode for each AWS Security Group (independently) to Full Protection (or change it to Read-Only). In this mode, you can make changes to the Security Group only in the CloudGuard Console, and not on in the AWS console. Any changes made in the AWS console, or elsewhere, are found by CloudGuard and reverted to the definition in CloudGuard

You can set a Security Group to Full Protection mode only if the AWS account is managed by CloudGuard in Full Protection mode. If the account is managed as Read-Only, you can update it to Full Protection.

1. Navigate to the **Security Groups** page. This shows a list of the Security Groups in your AWS environments.
2. Click on the Security Group to which to apply Full Protection.
3. Move the toggle in the top right to enable Full Protection.
4. Click **Switch** to confirm.

In addition, you can do it on the **Environments** page, see below for more details.

Selecting the Default Protection Mode for New Security Groups in AWS Environments

You can select the Protection Mode that CloudGuard applies to new security groups detected in accounts. CloudGuard defines and applies Security Groups in AWS for each region separately.

You can select from these options:

- **Read-Only** - CloudGuard includes new Security Groups in Read-Only mode, without changes to the rules
- **Full Protection** - CloudGuard includes new Security Groups in Full Protection mode, without changes to the rules

- **Region Lock** - CloudGuard includes new Security Groups in Full Protection mode and clears all inbound and outbound rules

You can set or change the Protection Mode for existing Security Groups, in all regions, for all of your AWS accounts.

To set or change the Protection mode:

1. Navigate to **Assets** > **Environments** and select an environment from the list. The **Network** tab shows the regions for the environment and the number of Security Groups defined for each region.
2. Click one of the regions. This shows a list of the Security Groups defined for the region.
3. Select the Protection mode to apply by default to new Security Groups in the region.
4. Select the Protection mode for each of the existing Security Groups in the region. Click **select entire region** to apply the mode to all Security Groups in the region.

   Note - The account must have a *CloudGuard-write-policy* to apply Full Protection to a Security Group (see *"Setting an AWS Security Group to Full Protection" on the previous page*).
5. Click **Save**.

AWS Security Group Management Considerations

Guidelines for managing AWS Security Groups from the CloudGuard portal:

- When a server instance is launched in AWS, a security group association is assumed. If the Administrator does not assign a security group to a new instance, it is placed in the default security group and uses its policy settings.
- AWS instances belong to one of two supported security group types: EC2-Classic or EC2-VPC. An AWS account can launch instances into both EC2-Classic and EC2-VPC, or only into EC2-VPC by region.
- Security Group rule definitions let specific sources reach an AWS instance using a specific protocol. *Inbound* rules identify the sources that can reach an instance with a given protocol (TCP protocol, UDP, or ICMP) and destination port.

Example: A rule could allow IP address `203.0.113.1` (the source) to reach the instances on TCP port 22 (the protocol and destination port).

- AWS Security Group rules are permissive in nature. When multiple Security Groups are applied to an instance, the rules from each Security Group are effectively aggregated to create a larger set of rules.
- In the case of internal referencing, an Administrator defines the Security Group as a source security group in the inbound security group rules. This enables additional instances to send traffic to instances in the source group.

Amazon VPCs and CloudGuard Service Functionality

A VPC is a virtual private cloud in Amazon Web Services, a private network that closely resembles classic virtual private networks (VPN). A VPC benefits from a scalable infrastructure. Protection of VPC subnet resources is achieved through the application of multiple security layers that contain security groups and network access control lists.

VPC can assign *persistent* private and multiple IP addresses to instances. This lets an Administrator to stop and start instances again and again without reassigning IP addresses. Network interfaces are defined independently and attached to specific instances.

An additional VPC feature is the power to change an instance's Security Group membership on the fly. An instance can be switched to a different Security Group while it is running. Instances can also run on single-tenant hardware.

For more information, see the Amazon Virtual Private Cloud User Guide.

AWS Security Group Management Modes: Full Protection or Read-Only

In CloudGuard, Amazon AWS Security Groups can be managed in one of two modes: Full Protection or Read-Only. Full Protection provides the CloudGuard administrator with full control of AWS security policy definition, access leases, and the ability to interact with dynamic policy objects.

In Full Protection mode, an AWS Security Group can only be managed from CloudGuard. Attempts to change a security group from the AWS environment (such as the AWS console) are detected by CloudGuard and trigger a CloudGuard Tamper Protection message. CloudGuard overrides the change that is mad and reverts to the definition of the Security Group defined in CloudGuard.

In Read-Only mode, Security Groups are defined and modified in the AWS environment, but you can monitor changes in CloudGuard with alerts, and a full audits trail. Use this mode initially as you plan a transition from managing your cloud environment in AWS to managing it in CloudGuard. In addition, it is the recommended mode of operation for Security Groups that are automated/managed by other tools (such as *AWS OpsWorks*).

The following table summarizes the differences between Read-Only and Full Protection modes:

Mode	Policy Visualization	Alerts & Audits	Tamper Protection	Policy Editing	Access Leases
Monitor	Yes	Yes	No	No	No
Full Protection	Yes	Yes	No	No	No

When a Security Group is switched to Full Protection mode, CloudGuard normalizes the rules in the group. Rules for IP address ranges that are fully included in the range of a different rule, and with the same ports are removed.

For example, the rule to allow inbound traffic on port 22 to address `192.168.10.10` is fully included in the rule to allow inbound traffic on port 22 to the address range `192.168.0.0/16` and would be removed.

More Links

- AWS EC2 User Guide
- Amazon Virtual Private Cloud User Guide

Full Protection Mode

In CloudGuard, there are two modes to manage Amazon AWS Security Groups:

- Full Protection
- Read-Only

Full Protection provides the CloudGuard administrator with full control of AWS security policy definition, access leases, and can interact with dynamic policy objects.

Full Protection

In Full Protection mode, you can manage an AWS Security Group only through CloudGuard. CloudGuard detects attempts to change a security group from the AWS environment (such as the AWS console), which starts **Tamper Protection** and can send an alert/notification. CloudGuard overrides the change that is made and reverts to the definition of the Security Group defined in CloudGuard.

The alerts and notifications initiated from Tamper Protection occur when you start Full Protection for the necessary regions in your cloud account. CloudGuard locks down the configuration of the security groups in that region to make sure that the security group stays correctly configured.

Configure a Security Group that has Tamper Protection enabled

To make a change in a Security Group that has Tamper Protection enabled, the change is made in CloudGuard.

1. Navigate to the Security Groups page in the Network Security menu.
2. Select the Security Group to be modified.
3. Make the necessary changes to the Security Group (for example, add or change Inbound or Outbound services). See *"AWS Security Groups" on page 368* for details on how to create or change Security Groups.
4. Save the changes.

Troubleshooting Workflow

When you receive an alert or notification in regards to Tamper Protection, this is seen in the CloudGuard Audit Log. To view and verify the action of CloudGuard Tamper Protection and its associated information, you can navigate to the System Audit Log and view the CloudTrail details.

1. From the menu, select **Events** > **Operational** > **System Audit Logs**.
2. Filter events by the Event Name **Security group tamper detected and handled**.
3. Select the event to see its details.

Azure Network Security Groups

This topic describes how to create and change Network Security Groups for an Azure account in CloudGuard. The account must be in Manage mode.

You can create NSGs for each region or resource group in your Azure account.

Create an Azure Network Security Group (NSG)

1. In the CloudGuard portal, navigate to **Network Security** > **Policy** > **Security Groups**.
2. Select the Azure account and click +.
3. Enter a name and description for the Security Group.

   The new NSG is created with default rules:

Set an Azure Environment to Managed Mode

This procedure describes how to set an Azure environment in CloudGuard to Managed mode. You have to start with *"Onboarding Azure Subscriptions" on page 169* to CloudGuard.

In Managed mode, you can manage the Security Groups for the account from CloudGuard.

1. In the **Assets** menu, navigate to the **Environments** page.
2. Select the Azure environment.
3. In the toolbar, move the switch from **Read only** to **Managed**.

   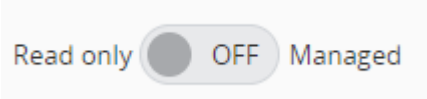

4. A confirmation message opens. Click **Switch**.
5. Click **OK** to affirm the change.

   Note - You can switch the environment back to Read-Only. In this mode, you cannot set Security Groups from CloudGuard.

Modify an Azure Network Security Group

You can change details for an Azure NSG in CloudGuard. The NSG must be in Manage mode. You can add, remove, or change rules for the NSG.

1. Navigate to the **Security Groups** page in **Network Security**. It shows your Security Groups, for all your environments.
2. Click the Azure NSG of interest from the list.

3. Click **Edit Mode**.
4. Select **Click to add new rule**.
5. Enter details for the rule.

   For example, add an SSH rule:

   Set the parameters for the Security Group:

   Service Type - Contains a list of predefined services, and type selection automatically fills most of the required fields.

   Action - Deny or Allow - Type of access to apply if the rule matches.

   Priority - Rules are checked in order of priority. When a rule applies, no more rules are tested for matching.

   Protocol - TCP, UDP, or *.

   Destination Port Range - Destination port range to match the rule.

   Destination Type - Source address prefix or tag to match the rule.

   Name - Name for the rule.

   For more information, see https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-nsg.
6. When the NSG contains several rules, you can drag the new rule and place it between other rules.

   Note - You can **Drag or Click to add new rule** between rules to create a rule directly at that location.
7. Click **Save Changes**.

Apply Tamper Protection to an Azure NSG

You can apply Tamper Protection to an Azure Security Group. Tamper Protection detects not approved changes made to the Security Group, that is, changes not made in CloudGuard, and resets them to the settings you configure in CloudGuard.

You can only apply Tamper Protection to Azure NSGs in an account that is Managed.

1. Navigate to the **Security Groups** page in **Network Security**. It lists your Security Groups, for all your environments.
2. Click the Azure NSG of interest from the list.
3. Move **Tamper Protection** to On.
4. In the confirmation message, click **Confirm**.

Dynamic Access Leasing

Overview

Dynamic Access Leasing is a CloudGuard feature that controls access to protected resources on AWS accounts. Access is given to specific users for a limited time to resources through specific Service Groups (as in SSH or Remote Terminal).

Dynamic access leasing allows AWS cloud servers and other resources to be almost hermetically closed, opening tiny security "holes" for management activities only when necessary.

Note - Dynamic Access is available for AWS environments only.

Access Lease

An Access Lease is a grant of access to specific Services Groups on an AWS cloud entity, for a limited period. Leases can be assigned to any of the following recipients:

- **Yourself** - The lease is for you, to access a selected service on a cloud entity, for a specific period, from the same device from which you are currently connected to CloudGuard.
- **Specific IP/CIDR** - The lease is for a specific IP address (or CIDR), to access a cloud entity, for a specific period.
- **An email recipient** - The lease is for an email recipient not necessarily a CloudGuard user), to access a cloud entity from the device on which the recipient opens the email.

A lease is for one-time access during a specific period. An expired lease cannot be extended, but it can be renewed by sending a new invitation. In addition, it is possible to terminate a current lease. The option to **Terminate Access** option appears for each lease in the Active Access Leases list.

How it Works

- Configure the AWS account and the Security Groups to be fully protected by CloudGuard
- CloudGuard admin users create Leases that, when activated, provide access to an AWS cloud resource (such as an EC2) through a specific Security Group, for a limited time period
- Recipients activate Leases by clicking on a link; access to the cloud resource is from the same host (IP) from which the link was activated, and for the specific service(s) or port(s) specified in the lease

- Recipient receives an email with a link to activate the Lease. Activation of the lease triggers the creation of one temporary Security Group Inbound Access Rule for each Inbound port or continuous port range that is selected for Dynamic Access.
- At the end of the time period, access to the cloud entity is blocked

The configuration and management of access leases for access to a security group service is the primary function here. The administrative user can get leases for themselves and assign them to other users.

Main Features

- Access to cloud services is usually **blocked** and only opened as necessary for limited periods to specific individuals.
- Access to some cloud services or cloud resources with one lease.
- Full audit trail of all access and changes to the cloud resources (see *"System Audit Logs" on page 134*).
- Admin users make decisions about which services to manage with IAM Safety.

Note - Access is for specific services that are attached to security groups. Gaining access to that service means that a user can interact with all servers in the selected security group.

Prerequisites

AWS account must be Fully Protected by CloudGuard (see *"Unified Onboarding of AWS Environments" on page 54*).

The Security Groups on the AWS account to be managed by IAM Safety must be managed by CloudGuard and configured not to be open to everyone

The Security Group in which the lease is established must **not** already have the maximum number of Inbound Access Rules permitted by AWS at the time when the lease is activated. For CloudGuard Customers for whom the default AWS "soft limit" of 50 Inbound Access Rules apply for each Security Group, this means that one Dynamic Access lease for one protocol/port or continuous port range can be activated on any appropriately configured Security Group that contains 49 or fewer Inbound Access Rules.

Access Groups

You can configure an Access Group to grant access to a number of services or ports with one lease. This is useful if activities are done on the group of services or ports together. In this case, the access lease specifies an access group as an alternative to a specific service or port.

Methods of creating leases

An admin user can create access leases from the following applications:

- the CloudGuard portal (admin user)
- the CloudGuard mobile app (for users with Dome9 accounts)
- the CloudGuard Chrome add-on

Google Chrome Add-on for Dynamic Access

The CloudGuard Chrome extension allows CloudGuard users to create dynamic access leases on-demand from their Chrome browser without signing-in to the CloudGuard console.

Use Cases

- User access a resource in a cloud VPC, for example, troubleshooting an issue, see *"Getting Access (How to Set Up a Lease)" below*.
- Configuration of an Access Group for IAM Safety, see *"Setting up an Access Group" on the next page*.

Actions

Getting Access (How to Set Up a Lease)

Configuration of access leases begins with the assignment of leases. It all starts on the **Access Leases** page with the **Get Access** option.

To assign leases:

1. Navigate to the **Access Leases** page in the **Network Security** menu, and select the **Get Access** tab. This tab shows a list of your AWS cloud VPCs and the services groups for each that are fully protected by CloudGuard and, for each, the services that they control. These services can be accessed with Access Leases. Use the filter and search fields to filter the list or search for a specific asset or service.
2. Click **Get Access** next to the service you wish to access to create a lease to access the service. The default lease period is set in **Settings** (**Configuration** > **Access Leases**).
3. To open a lease for a different period, click  and select the access period (1, 5, or 10 hours).

   When the lease is opened, you can access the cloud asset with the service you selected (for example, SSH) from the same device from which you opened the lease.

Sending an Access Invitation to a Different User

An administrator can send an invitation to external (non-CloudGuard) users through the Get Access/Send Invitation option. This is useful for inviting contractors, support personnel, and more who do not have a CloudGuard account. The specified user is then notified of a pending invitation by e-mail that includes a lease activation link. Selecting the link initiates the lease.

Note - Access invitations are marked *pending* until accepted and activated, and can be terminated before they are activated.

1. Navigate to the **Access Leases** page.
2. Click  next to the service for which you wish to open a lease, and select **Send Invitation**.
3. Select the lease period and procedure of delivery for the invitation. You can send an email through CloudGuard to the recipient, with a link to activate the lease, or you can copy the link, and send it on your own (for example, by private email, or messaging).
4. The recipient receives an email (or message) with a link to activate the lease. The lease is activated when the link is followed. The user can access the cloud asset with the selected service during the lease from the device from which the lease was activated (that is, from which the link was followed). At the end of the lease period, access is closed.

Viewing Active Leases

1. Navigate to the **Access Leases** page, and select the **Active Leases** tab. This shows a list of active leases and leases which have not been activated.
2. To terminate immediately an active lease, click **Terminate Lease**.

Setting up an Access Group

Access Groups are groups of services. They can be from different service groups, and for different VPCs. Select an Access Group when creating a lease, to open access to all the services in the group with one lease.

1. Navigate to the **Access Leases** page.
2. Select the check box next to the services to be grouped, and then click **Save as Access Group**.
3. Enter a name for the group, and then click **SAVE**. You can select to make the group public (with access to other CloudGuard users) or private (only you can access it). When Access Groups are configured, a new tab, **Access Groups**, appears on the Dynamic Access primary page. This tab shows a list of access groups.

Creating a Lease for an Access Group

You can create Access Leases for services in an Access Group. These leases can only be assigned to you.

1. Navigate to the **Access Leases** page, and select the **Access Groups** tab (this tab only appears if there are Access Groups configured). The tab shows a list of all the Access Groups.
2. Select the Access Group to be used in the lease from the list of groups on the left. The services in the group are shown on the right.
3. Click **Get Access for All n Services** to create a lease for the group of services. As an alternative, click  and select a different lease period.

Updating an Access Group

You can modify the composition of an Access Group. This affects new leases that use the group.

1. Navigate to the **Access Leases** page, and select the **Access Groups** tab.
2. Click the edit icon next to the group to be modified. The **Get Access** tab is opened, showing the list of services. The services in the group are selected. Select or clear Access Groups from the list, to modify the composition of the group, and then click **Update Group**.

Configuring the CloudGuard Chrome extension

Follow these instructions to install the CloudGuard extension for the Chrome browser.

1. Visit the Chrome web store and search for CloudGuard extension or go directly to CloudGuard Chrome Extension.
2. On the extension's details page, click **Add to Chrome** and then click **Add extension** in the dialog.
3. A CloudGuard icon ( ) appears in the browser menu bar. Click the icon to open the extension.

Using the CloudGuard Chrome extension to create a lease

Use the CloudGuard Chrome extension to create a lease. The lease can only include Access Groups.

1. Click the CloudGuard icon ( ) in your Chrome browser menu bar.
2. Select an Access Group for the lease. A lease is created (for the default period). A validation message shows.

Using the CloudGuard mobile app to create a lease

You can create leases with the CloudGuard mobile app, if you have a Dome9 account. You must install the app and pair it with your CloudGuard account first (see *"CloudGuard Mobile Application" on page 841*)

1. Open the CloudGuard mobile app and select **Dynamic Access** from the primary menu.
2. Tap on the Access Group to create a lease for it.

VPC Flow Logs

You can see the traffic into and out of, and in, your Amazon Virtual Private Cloud (AWS VPC) in CloudGuard. You can select traffic for any of your VPCs and then filter for specific flow items of interest. CloudGuard extracts this information from the cloud platform and enriches it with contextual information, such as source and target names if they are labeled. You can export the displayed information to a file.

In addition, you can see VPC flows from the Configuration Explorer (see *"Configuration Explorer" on page 354*).

Note - Configuration Explorer is available for AWS VPCs only.

Benefits

- Console view of all VPC networks and flows on all cloud providers, all accounts, and regions.
- See flow in network context (in Configuration Explorer, for AWS only).
- Variety of filter and search options to narrow the scope, and look for specific flows of interest.

Use Cases

Here are some typical use cases for viewing VPC Flow Logs:

- Analyze incidents with network traffic in the VPC, see *"Traffic Explorer" on page 359*.
- Filter traffic for specific network elements, see *"Filter the flow list for specific detail" on the next page*.

Actions

Setup your AWS environment for VPC Flow Logs

Your AWS environment must be configured for VPC Flow Logs to view them on the CloudGuard portal. This is done on the AWS console, in the VPC Dashboard.

1. **Create a VPC flow log on AWS for our VPC**. Follow the steps described in https://aws.amazon.com/blogs/aws/vpc-flow-logs-log-and-view-network-traffic-flows/ to enable Flow Logs for a specific VPC. This step must be done for each VPC for which you wish to view flow logs.

   Set the filter on the flow logs to capture all traffic for Accepted and Rejected.

2. **Enable the IAM policy for the CloudGuard user on AWS** (this is applicable for AWS for accounts that were added before September 2015). On the AWS console, select the IAM Dashboard.

    a. In the AWS IAM Dashboard, select **Roles** (on the left) and select the *CloudGuard-Connect* role.

    b. Check that the *CloudGuard-readonly-policy* appears in the **Permissions** tab for this role. If the role or the policy do not appear, the AWS account has not been fully onboarded to CloudGuard - check or repeat the procedure in *"Unified Onboarding of AWS Environments" on page 54*.

View a VPC flow log

View flows for any of your VPCs, in any of your cloud accounts.

1. Select the VPC (account, region, assets), and the period (back from the present time, or click the **Custom Date** link to select a specific date & time).
2. A list of entries for the selected VPC is shown. Each entry represents a flow.
3. Put the cursor over an entry for more details.

    These are the filter options:

Icon	Action
	Show IP address
	Show the geolocation, hostname, and network, of the host
	Filter for this value
	Not this value (for example, other than)

Filter the flow list for specific detail

You can filter the flow list to show entries of interest. The filter options are at the top of the list

Filter options:

- Select the VPC & instance - This is the primary filter.
- Select specific values for one of the columns (click on the terms, or enter as free text).

- Add terms to build up the filter. As you add terms, the list of flows is incrementally filtered (the result is the AND of all selections):
- Filter on a specific value(s) of a field: press the filter icon next to the value to filter for entries with this value.

Cloud Infrastructure Entitlement Management (CIEM)

The goal of **Cloud Infrastructure Entitlement Management** (CIEM, formerly called Identity) is to reduce your attack surface by ensuring that cloud entitlements or permissions respect the principle of *least privilege*. This means that identities are only granted the smallest set of permissions to do their tasks.

In addition, CIEM provides in-depth visibility into permissions granted to cloud entities and calculates which permissions are effective.

Use Cases and Prerequisites

To use CIEM, you must finish onboarding a cloud environment to CloudGuard with all necessary permissions. There are additional prerequisites for some use cases of CIEM.

	Use Case	Prerequisite for the Use Case
1	Gives visibility into cloud entitlements and effective permissions.	No additional prerequisites are required.
2	Automatically identifies overprivileged AWS Lambda functions and provides a least-privilege suggestion for remediation.	You must enable Serverless Risk Assessment for the AWS environment. For more information, see *"Serverless Risk Assessment" on page 525*
3	Automatically identifies overprivileged cloud identities based on actual use of permissions. Provide least-privilege suggestions for remediation.	You must onboard your account to Intelligence Account Activity. For details about onboarding, see *"Intelligence Onboarding and Offboarding" on page 570*.

CIEM Dashboard

The CIEM **Dashboard** screen shows you some aggregated data on your cloud assets protected with CIEM and on CIEM Findings. The time filter allows you to review data over different periods.

- **Inactive Users & Roles** - Shows *Inactivity*, which is one type of misconfiguration that CIEM finds and highlights. Those listed must be removed, deactivated, and not stay on the list of users and roles.
- **Severely Overprivileged Entities** - Shows entities that are overprivileged with a severity of High or Critical for each asset type.
- **Findings by Severity** - Findings that CloudGuard flags as CIEM related.
- **Top Entities** - Shows top entities based on the number of alerts and includes a selectable time filter.
- **Trend Over Time** - Shows the number of daily alerts.
- **Labels** - Show the categories of issues that CIEM finds and show the ranking of findings for each label.

  The time filter is shared between the four widgets in the **Findings** section.

Entitlement Map

The Entitlement Map visually presents the permissions granted to the cloud identity and shows all the elements that contribute to the set of permissions.

It is applicable only for:

- AWS IAM roles, IAM users, EC2 instances, and Lambda functions
- Azure users, groups, user-assigned managed identities, virtual machines (VM), app registrations, function apps

To see the Entitlement Map for a specific entity:

1. In the CloudGuard portal, go to **Assets** > **Protected Assets**.
2. Filter the view by the Asset Type and select one of the supported entities.
3. Select an entity to open its details page.
4. Select **Permissions**. The Entitlement Map opens to visualize the asset entitlements and effective policy.

In addition, you can see the entitlement map from the CIEM findings, see *"Findings" on page 392* for details.

Entitlement Map in AWS

Policy Types

For AWS identities, the Entitlement Map shows these policy types: Identity-based, Permissions boundaries, and Organizations SCPs (service control policy). For more information about these policies, see AWS Identity and Access Management.

Identity-based policies are Inline or Managed. Managed policies are managed by AWS or by the customer.

This is an example of an IAM role. For an EC2 instance or a Lambda, there is one node with the role attached to them.

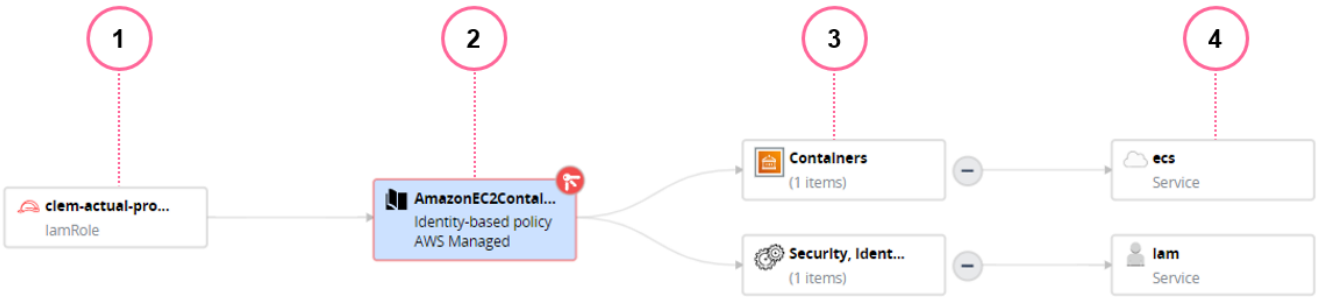

Number	Description
1	Name of the context asset.
2	Policy (or policies) attached to the Role. Below the policy's name is the type of policy. For this example, the policy is "Identity-based".
3	AWS Service Category allowed by the policy. In the example above, the policy gives access to Containers and Security, identity and compliance.
4	List of services and resources that the policy gives permission to.

- When you select **Detailed View** and then click a policy, it opens in the right pane in JSON format. Or select **Consolidated View** and click the Consolidated permissions node to see the details of the permissions presented in a table. For more information, see *"Consolidated View" on page 390*.
- When you click on a Service or Resource, it highlights it in the policy's details pane.
- The map shows only services and resources linked to an `Allow` effect in the policy. `Deny` effect statements show in the policy's detail pane.
- The default view shows all the policies, and the alternate view is called **Consolidated Permissions**.

Policy Sources

Use the Entitlement Map to see how policies are assigned to specific IAM users, including through direct assignment, group assignment, and through trust relationships.

Three types of policy sources:

Policy	Description	Example
IAM User Group	Attached through an IAM user group	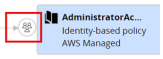
Trust Relationship	Attached through a role trust relationship	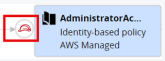
Direct	Attached directly to the user	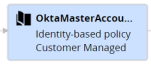

Policy by trust relationships with other IAM roles

Policies obtained through a *trust* relationship are indicated as such by an icon. For policies obtained through more than one trust relationship, CloudGuard gives the shortest path.

When you open a policy's details pane, it shows the **Original** and **Effective** set of permissions. If the policy is overpermissive, then **Suggested** also shows.

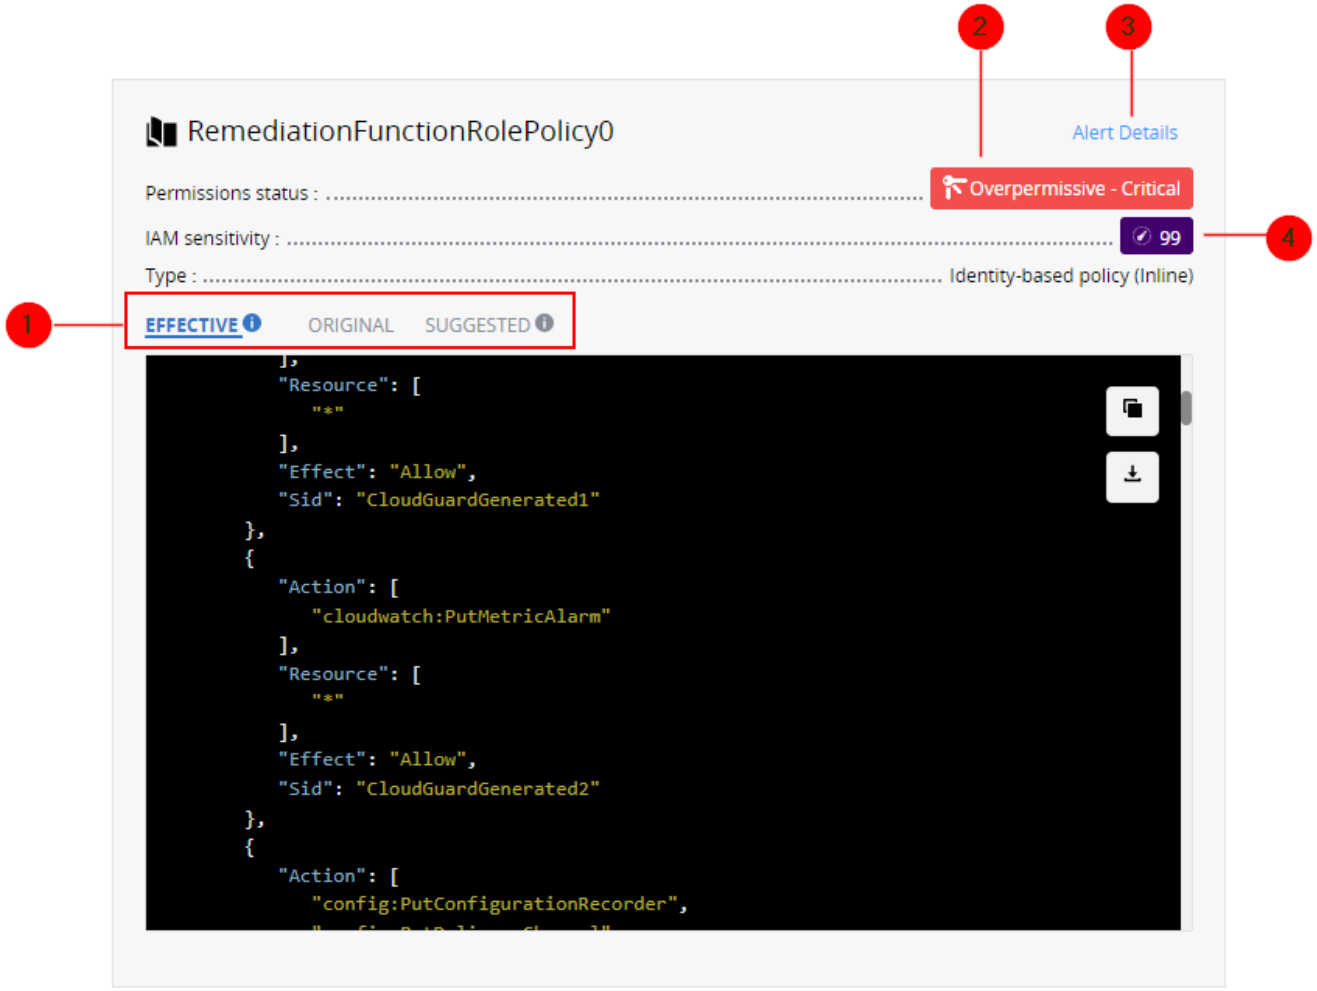

Number	Description
1	■ The *Original Policy* shows the policy as defined in a cloud platform. ■ The *Effective Policy* considers permission boundaries and SCPs and shows the *Effective* set of permissions. ■ The *Suggested Policy* only shows if the policy is overpermissive. It shows the set of permissions that must be granted by the policy so that it complies with the least-privilege principle.
2	**Alert Details** - Only relevant when the entity is overprivileged and within overpermissive policies (the link redirects to the original finding).

Number	Description
3	**Permission status** - The severity of the finding depends on the sensitivity of the excessive permissions. ■ Valid - Shows when the policy is not overpermissive. ■ Overpermissive - Shows the severity of the overpermissive policy. ■ N/A - The entity is in an account that is not onboarded to Intelligence Account Activity, as such, CIEM cannot analyze the permissions that are in use. In addition, it can be that the role is idle for the last 90 days.
4	**IAM sensitivity** - A number in the range from 0 to 100 that CIEM calculates. The **Sensitivity** score represents the possible damage that IAM permissions can cause to the cloud environment.

Entitlement Map in Azure

For Azure identities, the Entitlement Map shows Role Definitions. These entities are granted permissions through Role Assignment. In addition to assigning a role, Role Assignment defines the scope of the permissions, that is, for which resources these permissions are valid.

Azure VMs can obtain permissions in two ways: through a user-assigned managed identity or through a system-assigned managed identity. For more information about managed identities, see Managed Identity Types.

Some Azure identities, such as Users, Groups, and Service Principals (user-assigned or system-assigned identities) can be part of an Azure Group and obtain the Role Definition through this Group.

The **Detailed View** of the map shows the identity and all its paths to obtain permissions. You can see its user-assigned and system-assigned identities that obtain their role definition through the role assignment. The icon to the left of the role definition shows the role assignment and its ID. Below you see the scope of the role assignment.

On the graph, each role definition connects to a list of services grouped by categories.

The same information appears on the right pane of the detailed view. Click the role assignment button to navigate to its asset page. All assets on the right pane are clickable.

Consolidated View

The **Consolidated View** of the map combines all the permissions effectively given to the identity under one title called **Consolidated Permissions**. It lets you understand the overall effective permissions granted to the identity.

The table in the right pane shows the effective set of actions per Service that the identity can do.

Note - If an *Action* is on different resources, it may show several times.

- To search the table, use the search bar. The search is a free text search.
- To expand the table, click  .
- The Access Level is not available when an action is not specific and contains a set of actions. Such action cannot have a description otherwise available as a tooltip.
- When you click a Service node, the corresponding row in the table is highlighted.

Limitations

CIEM does not support:

- Azure Locks. For more information about locks, see https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/lock-resources?tabs=json
- Azure deny assignments. For more information about these assignments, see https://learn.microsoft.com/en-us/azure/role-based-access-control/deny-assignments
- Azure Policy. For more information, see https://learn.microsoft.com/en-us/azure/governance/policy/overview

In the entitlement map, you can see the role definitions obtained through RBAC role assignments on the specific subscription onboarded to CloudGuard.

In Azure, Microsoft Entra ID roles are used to manage Microsoft Entra ID resources in a directory, for example, create or edit users, assign administrative roles to others, reset user passwords, manage user licenses, or manage domains. Those roles are handled on the tenant or Entra ID level. As a result, permissions granted through Microsoft Entra ID roles are not part of the entitlement map and of the effective permissions of Microsoft Entra ID identities.

Findings

The **Findings** screen shows these types of findings related to IAM and permission management:

- Posture Findings
- Security Events
- Overprivileged Entity findings from Cloud Infrastructure Entitlement Management (CIEM)

Findings use **labels** to show insights into the category of the specific finding, such as inactive entities, risky permissions, password policy, and misconfigurations. You can group and filter the findings by labels.

CIEM overpriviliged findings are supported for AWS and for Azure.

To see findings in CIEM:

In the CloudGuard portal, from the left toolbar, click **CIEM** > **Findings**.

CIEM Suggestions

CIEM compares IAM permissions *granted* to IAM permissions *in use* to find identities that received unnecessary permissions. For these entities, CloudGuard creates a finding with the source **CIEM** and the label **Overprivileged Entity**. The finding contains suggestions to grant the least privilege for the entity.

Prerequisites

To allow CIEM to analyze the permissions in use, you must onboard your account to Intelligence Account Activity. For more information, see *"Intelligence Onboarding and Offboarding" on page 570*

Note - In an Azure environment, you must onboard Azure Activity logs to use this feature.

To get least privilege suggestions for AWS lambda functions, you must enable Serverless Risk Assesment for the AWS environment. For more information, see *"Serverless Risk Assessment" on page 525*

Supported Entities

CIEM policy suggestions support these entities:

Cloud Provider	Supported Entities
AWS	■ IAM Users ■ IAM Roles ■ Lamdba functions
Azure	■ User-Assigned Managed Identity ■ App Registration

How CloudGuard Makes Policy Suggestions

After you onboard an environment to Intelligence Account Activity, CIEM does a full scan of your identities and their permission usage after 24 hours. CIEM scans again after 7 days, 30 days, 60 days, and 90 days. Then, CIEM does a full scan every 90 days for roles that it already scanned and that were not modified.

For newly created roles or for roles that had permissions modified, CIEM does a new permission assessment after 3 hours, 24 hours, and 1 week.

For identities whose behavior is harder to predict (such as IAM roles used for SAML federation and some IAM users), CIEM uses machine learning to build a baseline of the behavior of an entity. When CIEM creates a profile of the entity's behavior, it starts to provide suggestions to reduce the access rights for entities with a minimum of 75 days of activity.

Alert Details

The details of the alert appear in the **CIEM Findings** screen.

To see CIEM Overprivileged alert details:

1. Go to **CIEM** > **Findings**.
2. Filter by the label **Overprivileged entity**.
3. Click an alert to see its details.

The **Permissions** tab of the finding shows the Entitlement Map. The Entitlement Map shows the asset entitlements and the effective policy. For more information about the Entitlement Map, see *"Entitlement Map" on page 387*.

Remediation

The Remediation section shows options to remediate findings in your cloud platform. Select **one** of these options:

Remediation for AWS

Select **one** of these and do the suggested actions in AWS:

- **Option A: Update existing policies with suggestion** - Shows the set of permissions that must be attributed to the entity for it to comply with the principle of least privilege. This suggestion is provided for each policy. To see the suggested policy, click **Show**.

  This option requires you to edit the policies used by the entity based on CloudGuard's suggestion. If a managed policy cannot be updated because it is used by a different entity or is managed by AWS, you can create a new policy with CloudGuard's suggestion and use it in place of the existing policy.

- **Option B: Add suggested Permission Boundary to the entity** - Creates a permission boundary policy for the entity. The permission boundary policy is a type of policy that sets the maximum permissions an entity can have. This option allows for easy roll-back because you do not make changes to the existing policies.

  AWS allows one permissions boundary for each entity. If the entity already has a permission boundary, change the permission boundary to match CloudGuard's suggestion.

  Note -
  CloudGuard recommends a permission boundary based on actual usage of permissions. It is possible for CloudGuard to recommend a policy that is longer than the maximum permitted length of a policy in AWS (see AWS documentation). In this case, do **one** of these:
  - Use asterisks (`*`) to shorten the policy when this fits your use case.
  - Select **Option A: Update existing policies with suggestion** instead of **Option B**.

- **Redundant permissions** - Shows the permissions that CIEM recommends for you to change with the policy name, the original permissions (as they appear in the policy), and the suggested permissions. A suggestion of empty brackets (`[]`) means to remove a full statement from the policy.

  To see redundant permissions, click  . The **Redundant Permissions** window opens with the original permissions, for each policy, and a corresponding recommendation. The permissions are color-coded to indicate their sensitivity, which is also written as part of the original permissions. Use the filter options to display permissions with specific sensitivity levels.

Remediation for Azure

Select **one** of these and do the suggested actions in Azure:

- **Option A: Update/replace role definitions with corresponding suggestion** - Shows the set of permissions that must be attributed to the entity for it to comply with the principle of least privilege. This suggestion is provided for each role definition. To see CloudGuard's suggestion, click **Show**.

This option requires you to edit the role definitions used by the entity based on CloudGuard's suggestion. If a role definition cannot be updated because it is managed by Azure, you can create a new Azure role definition with CloudGuard's suggestion and use it in place of the existing role definition.

- **Option B: Create a single role definition with suggestion and use in place of all existing ones** - Shows the set of permissions that must be attributed to the entity for it to comply with the principle of least privilege. In Azure, remove all existing role definitions and replaces them with CloudGuard's suggested role definition.
- **Redundant permissions** - Shows the permissions that CIEM recommends to change. Shows the role definition, the original permissions (as they appear in the role definition), and the suggested permissions.

  To see redundant permissions, click . The **Redundant Permissions** window opens with the original permissions, for each role definition, and a corresponding recommendation. The permissions are color-coded to indicate their sensitivity, which is also written as part of the original permission. Use the filter options to display permissions with specific sensitivity levels.

Data Events (for AWS)

Because CIEM is based on AWS CloudTrail, it can suggest remediations only for actions that are tracked in CloudTrail. This is done to make sure that CIEM does not block your operations.

For example, if you do not log data events in CloudTrail, CIEM does not suggest remediation for actions that are in the category of data events.

Policy suggestions depend on the status of data events in your CloudTrail:

- Data events are disabled on the account or for a specific resource - No information is logged in CloudTrail - No policy can be suggested for data events.
- Data events are enabled only for Write (or Read) events - No information is logged in relation to Read (or Write) events - No policy can be suggested for Read (or Write) data events.
- Data events are enabled on all resources for Write and Read events - Information is logged in CloudTrail - Policy suggestion can be to remove the event if it does not show in the log or to keep the event if it shows in the log.

Note - Data events are not supported in policy suggestions based on machine learning for these entities:

- roles used for SAML federation
- IAM users

Other Findings

Lambda Functions - Overprivileged Lambda functions are identified through code analysis, which is a feature of Serverless Risk Assessment. For more information, see *"Serverless Risk Assessment" on page 525*.

For Permissive Role findings for Lambda functions, see *"Finding Types" on page 526*.

Finding Severity

The severity of the finding is dynamic. It is based on an IAM Sensitivity score that CIEM calculates. The Sensitivity score represents the possible damage that IAM permissions may cause to the cloud environment.

After you make the suggested changes to the permissions, CloudGuard updates the finding accordingly.

CIEM Policies, Exclusions, and Remediation

Policies

To receive notifications about CIEM events, it is necessary to create a policy. You define the CIEM policy like other CloudGuard policies. For more information, see *"Configuring CloudGuard Policies" on page 78*.

The findings that the CIEM policy discovers already exist in the system, so no new findings are created.

Start to configure the policy from **CIEM** > **Policies** > **Add Policy** > **Environment Policy**.

You can select these types of notifications in your CIEM policy:

- **Immediate Notification** - Send findings immediately to one or more of the given destinations:
  - *Separate Message* - By email, to a list of email recipients.
  - *SNS notifications* - To an AWS SNS topic; enter the ARN for the AWS SNS topic and select the format for the notification.
  - *HTTP Endpoint* - To an HTTP endpoint for third-party applications.
- **Security Management Systems** - Send notifications to a security management system.
- **Issue Management Systems** - Send notifications to an external ticketing system, such as *PageDuty*.

For more information about the options, see *"How to Configure a Notification" on page 853*.

Exclusions

You define the CIEM exclusions like other CloudGuard exclusions. For more information, see *"Configuring CloudGuard Exclusions" on page 80*.

Start to configure the exclusion from **CIEM** > **Exclusions** > **Create New Exclusion**.

The CIEM exclusions always apply to the preselected Entitlement Management ruleset. The exclusion is based on these parameters:

- **Environment or Organization unit**
- **Date range**
- **Entity**
- **Account number**
- **Alerts severity**

To create an exclusion from a finding, see *"Creating exclusion for findings" on page 126*.

To create an exclusion from scratch, see *"Creating an Exclusion" on page 81*.

Remediation

You define remediations for CIEM like other remediations for Posture Management. For more information, see *"Automatic Remediation with CloudBots" on page 317*.

Start to configure the remediation from **CIEM** > **Remediation** > **Create New Remediation**.

The CIEM remediation always applies to the preselected Entitlement Management ruleset.

To create a remediation, see *"Adding Remediation" on page 319*.

The applied CloudBot is *iam_entity_create_and_attach_permission_boundary*. The default policy name is *CIEMSuggestion*, but you can enter another permission boundary (policy) name. This policy is a permissions boundary policy added by the CloudBot from the finding as described in Option B.

Activity Explorer

Activity Explorer is a tool to search for and see events of interest in your environment's network traffic or event activities. It gathers and presents information from logs for the environments, enriched with information from more sources such as Threat Intelligence feeds, IP reputation databases, and geolocation databases.

Benefits

- Quickly identify unwanted network traffic from unknown or suspicious sources
- Identify gaps in cloud security settings or misconfigurations
- Monitor and analyze user activity on your cloud environments for unusual behavior

Activity Explorer Views

Intelligence combines cloud assets and configuration information with real-time monitoring data from a variety of cloud platform sources and current threat intelligence feeds, IP reputation, and geolocation databases. This results in enhanced visualization that highlights suspicious traffic from legitimate traffic. For example, sources of network traffic from other cloud elements are shown based on type, and malicious external sources are marked as such.

The Activity Explorer provides visualization of event activities in your environment. You can view activities on all of your assets or filter the view for specific assets or activities. You can use this to identify anomalous activities from unwanted and potentially malicious sources or unexpected activities from trusted sources.

Actions

Viewing Activity Explorer Graph

This view shows events in your environment, based on logs, such as CloudTrail for AWS, or Log Search for Azure.

1. Navigate to **CIEM** > **Visibility** > **Activity Explorer**.
2. The Activity Explorer page opens and shows all events for all assets in the selected account (during the hour before).
3. To open the **Query** menu, click  , in the top left.
4. Select a Query from the list of queries.
5. Click **Select this query**.

The page is refreshed, with the selected query applied to the log data for the account.

This view has the following elements:

- The GSL query **(2)** enables you to search for network resources or network flows. You can specify in the query the details of specific packets, bytes, source, or destination to monitor traffic and interconnectivity of the resources that belong to your environment or cluster. Edit the query text directly in the box or open a graphic query editor.

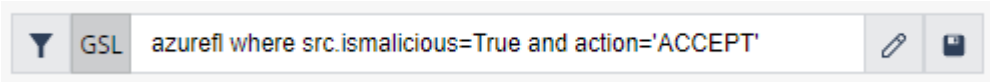

- The time frame **(3)** is the period back from the current time (15 min, 1 h, 24 h, 7 d) or start and end dates for a specific time range. To change the time frame for the view, select a new value and click **Run (4)** to run a new query.

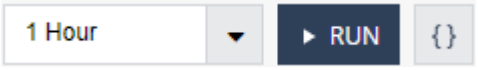

Note - The graph shows only actual traffic between entities. Entities without network activity during the selected time frame are not seen.

- The **Queries** icon  **(5)** allows you to select a query from an applicable category.

The central part of the view shows a graph of entities in the environment, and the account activities based on the selected query, and the time frame. The entities are grouped in zones that show the direction of activity, Identity, Issuer, and Target Asset. The Identity zone shows entities that did or initiated the activity, and the Target Asset zone shows the assets on which the activity was done. The Issuer zone shows the assets that were used to gain access to the Target, such as a role or token.

The default view groups assets by types (these are the large circles with numbers that show the number of assets). Click  to ungroup assets.

Click on a node in the graph to highlight the events relating to it.

The pane at the right shows details for the node in the INFO tab.

Click on an identity to add it to the query (for example, you can narrow the query to this asset only by adding it as an AND clause to the current query. Or do not include it by adding it as an AND NOT clause).

The target assets list is organized into Top and Least type lists. Select one of the asset types in the list to add it to the query (as above).

Activity Explorer Graph Controls

You can control the view with these controls:

- **zoom**: Select a point in the center section of the view, and use your mouse scroll wheel to zoom the display in or out, or use the zoom controls at the top of the view ()
- Select an entity in the view, to show **details**, in the pane on the right.

  There are three tabs of details:

  - **Info** - Shows the list of Target Assets with details for each.
  - **Statistics** - Groups the entities into Identities and Target Assets.
  - **Timeline** - Shows a list of activities for the selected asset, organized as a timeline (this is only shown if a specific asset is selected, not a group of assets).

  Click in the central part of the view (not on an entity), to go back to the previous view, of all entities.
- **Group** entities in clusters based on the entity types with the **Group by** selector at the top of the view. And then you select if to expand or collapse the groups with the expand/collapse symbol next to the group in the view or at the top of the view() .

Viewing Activity Explorer Logs

You can view Activity Explorer logs for specific events.

Select an event in the Activity Explorer graph, and then click **OPEN LOGS** in the detail pane on the right.

Logs for the selected entity show in a new tab.

Click a record in the logs to show more detail, for example, details for the Issuer, Identity, and Target Assets involved in the specific activity.

Click entities in the detail pane to add them to the query more phrases to narrow the query to specific items of interest.

Filter Views

Filtering Activity Explorer graphs

You can filter the Activity Explorer view to focus on events of interest.

Activity Explorer filters

Select filter options from the following:

- **Status** - The event status (success or failure)
- **Interactive** - The event involved user interaction on the cloud platform console
- **Assets** - Select specific asset types

After making changes to these filters, click 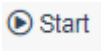

to run the query again and update the graph.

Statistics filters

In the Traffic Activity view, you can add more terms to the query based on results in the statistics pane (on the right).

1. Click a detail in the statistics pane.
2. Select the logic with which to add this to the query (AND, OR, NOT). The selected property is added to the query string.

In addition, in the Activity Explorer view, you can select details from the detail pane (on the right).

1. Click on an entity in the detail pane.
2. Select the logic with which to add this to the query (AND, OR, NOT). The selected property is added to the query string.

More Links

- *"Traffic Explorer" on page 359*
- *"Cloud Detection and Response (CDR)" on page 565*
- *"Intelligence for Kubernetes Containers" on page 645*

IAM Safety

Overview

CloudGuard IAM Safety controls access to services on AWS environments by IAM users, and requires that these users have explicitly given permission from a CloudGuard account administrator to access these services. This hardens the AWS account console and restricts users from making that are not approved or accidental changes to account settings without the knowledge and authorization of an administrator. Users can continue to access the account to view settings without restrictions (based on their AWS permissions).

For IAM users to access protected services, they must have an authorization window opened for them for the service. The window can be opened for them by a CloudGuard admin user, on the CloudGuard portal, or, for users with Dome9 accounts, by the IAM users themselves with the CloudGuard Mobile App. The authorization window is for a limited period. During this time, the IAM user can access the protected AWS services. At the end of the window, access to the services is blocked.

In addition, all actions taken by IAM users on protected services are logged, and appear in the CloudGuard Audit Trail.

How it Works

CloudGuard IAM Safety protects AWS services or specific actions for these services. To set up IAM Safety on CloudGuard, you configure a CloudGuard IAM policy on your AWS account which grants CloudGuard permissions to control select AWS services. You included in this policy the AWS services or actions that are protected by CloudGuard (AWS actions or services that are not selected can be accessed by IAM users based on their AWS permissions and are not restricted or protected by CloudGuard).

After the policy has been applied to the AWS account, you use CloudGuard to e xplicitly apply protection to the IAM users of the AWS account for the protected services you selected. This means to access the protected services or actions on AWS they are given explicit access permission from a CloudGuard admin user. This is called *elevation*. It is for a limited time set at the time it is granted. During this time, the IAM user can access the service based on their AWS IAM role. At the end of the time, they are blocked from access.

In addition, you can apply IAM Safety to IAM Roles. In this case, all IAM users with this role can access protected AWS services when the role is elevated.

Use these methods to elevate an IAM user or role:

- By a CloudGuard super user from the CloudGuard portal.(for users with Infinity Portal or Dome9 accounts)
- They can elevate themselves if they are in addition a CloudGuard user, have installed the CloudGuard Mobile app, and associated it with a protected account. (for users with Dome9 accounts only)

Note - IAM users of a protected account, who do not have protection applied to them, are not restricted by CloudGuard from accessing services in the account (based on their AWS permissions only). To protect an AWS account, it is important to immediately apply protection to all IAM Users and Roles the account is protected.

Considerations

CloudGuard recommends through categories of actions and services to be protected by IAM Safety. These are grouped as Templates when you set up the IAM Safety, and cover Computing, Networking, Security & Identity, Storage, and Database actions. Check Point recommends to lock down services/actions that are not done frequently, or are irrevocable when done, or both. For example, IAM, Route53, KMS, services, or actions such as changing S3 bucket permissions, deleting buckets, or deleting EBS snapshots.

Prerequisites

The AWS account with the services that you wish to protect with IAM Safety must be onboarded to CloudGuard. See *"Onboarding AWS Environments" on page 144*.

CloudGuard users must be associated with a protected AWS account to grant access to themselves or other users. This is done by invitation from a CloudGuard admin user.

If it is necessary for CloudGuard users to use the CloudGuard Mobile app to elevate themselves to access AWS-protected services, they must install the application and then pair it with their CloudGuard(Dome9) account.

Protected vs Protected with Elevation

Two procedures used to protect an AWS service:

Protected - Protected AWS IAM users cannot do protected actions on these AWS services in any circumstances. Users can only do these actions if the CloudGuard protection is permanently removed from the service.

Protected with Elevation - CloudGuard users (who are associated with the protected account) can elevate themselves or other IAM users to access protected services for limited periods.

Tamper Protection

IAM users or roles that are protected with IAM Safety are protected against tampering. These users and roles are included in restricted groups or policies in AWS (as part of the procedure CloudGuard implements the protection). If someone tries to remove a user or role from these groups or policies on the AWS console (and not through CloudGuard) it is detected by CloudGuard (and logged in the Audit trail) and rolled back.

Benefits

- Reduce not approved or accidental access to AWS accounts to change settings or entities
- Control who can make changes to AWS accounts settings
- Must have more authorization (the mobile app on the Dome9 user's device) to grant access
- Access permissions are temporary, and are automatically removed at the end of the authorization window
- Full audit trail of access to sensitive services

Use Case

An AWS IAM user account must have:

- To change settings on the AWS account, see *"Add an AWS environment to be Protected by CloudGuard IAM Safety" below*.
- Add/change cloud entities associated with the account or the account's VPCs.

Actions

Add an AWS environment to be Protected by CloudGuard IAM Safety

To set up your CloudGuard account to manage IAM user access to an AWS account, you must configure a policy in the AWS account. This policy lists the AWS services and actions that are protected. When this policy is in place, access to these services is blocked to all IAM users and only permitted to specific users when authorization is given.

1. In the primary menu, navigate to **CIEM** > **IAM Safety** > **Accounts**.
2. Select the AWS services and actions to be managed by your CloudGuard account from the list of services. The list of services expands, to show specific actions. As an alternative, select one or more templates (aggregate groups of services) at the top. After making your selections, click **Copy to Clipboard**. Click **Next**.
3. Follow the steps described in the next screen, to create a new policy and group on your AWS account, which permits your CloudGuard account to manage AWS IAM users. Copy the Policy and Group ARNs from the AWS console, paste them in the applicable places on this screen, and then click **Next**.

   Note - Review carefully the services and actions that you select for protection before proceeding. When you complete the policy setup for these services, there is no simple procedure to make changes to it.
4. Select the AWS account to be managed by CloudGuard, and then click **Next**.

5. Connect the IAM Safety policy with the account. Follow the on-screen instructions and then click **Next**.
6. CloudGuard connects to your AWS account and tries to gain control of the selected services. If this is successful, the confirmation message appears.

Protect an IAM User or Role

After the AWS account has been protected with CloudGuard IAM Safety, you can apply CloudGuard protection to IAM users of the account, so that they can access the protected services. These users are called 'Protected' users. Applying protection to them does not grant them access, but allows temporary access to be granted to them with an 'elevation' (or authorization).

IAM Users and Roles can be protected. If a role is protected, any IAM user with this role can access protected services if the role is elevated.

Note - Until you apply protection to an IAM user, the user can access AWS services (and protected services) without restriction. It is important to apply protection to all IAM users immediately after configuring CloudGuard IAM Safety on the account.

1. Navigate to the IAM Safety page and select the **IAM Users** tab. This shows a list of the IAM (AWS) users of the AWS account. The protection status of each user is shown (initially all are Not Protected).
2. Select a user or users to protect and click **Protect All**.
3. Select the type of protection to apply to the user, then click **Save**. **Protected** restricts the user from accessing protected AWS services. **Protected With Elevation** restricts the user from accessing protected services, but allows the user to be elevated, to access services. In addition, select the CloudGuard users that can elevate these users. This can be a group of users.
4. Click **Save**.

Apply protection to IAM Roles in the same procedure. Select the **IAM Roles** tab.

Select the roles to be protected, then click **Protect All**.

Unprotect an IAM User or Role

A CloudGuard super user can remove protection from an IAM User for an AWS account. When protection is removed, this user can access protected services on the account without any CloudGuard restriction or control (or Tamper Protection). In addition, actions by this user are audited by CloudGuard.

1. Navigate to the IAM Users tab.
2. Select  opposite the User or Role.

Invite (Add) CloudGuard Users to Protected Accounts

A CloudGuard account administrator invites other CloudGuard users to a protected account. These users can then elevate IAM users to access the protected AWS account.

1. Navigate to the **Users** page in the **Settings** menu.
2. Select the user from the list and click **Invite User** from the menu bar.

   An email invitation is sent to the user.
3. The invited users receive an email to join IAM Safety. To join, they must click the link.

The invited user with Dome9 accounts can optionally install the CloudGuard mobile app (see *"CloudGuard Mobile Application" on page 841*), to elevate IAM users from the app.

Open an Authorization Window for an IAM User or Role (Elevate a User)

A CloudGuard user, related to a protected IAM user or role, can elevate them, to access the protected services. This can be done from the CloudGuard portal, or on the CloudGuard Mobile app for users with Dome9 accounts.

The IAM user must be protected by IAM Safety with **Protect With Elevation** protection.

The elevation is for a limited period, during which the elevated user can access the protected AWS services.

Elevate with the CloudGuard portal

CloudGuard super users can elevate IAM users from the CloudGuard console app.

1. Navigate to the IAM Users tab.
2. Select the user or users to be elevated from the list of IAM Users (the user must be Protected). Click **Elevate** opposite the user to elevate the user for 15 minutes, or select a specific elevation period from the drop-down list.
3. To elevate a number or users, check the box adjacent to each one, then select the elevation period.
4. If the intended user is not yet protected, press **Protect** to include them in CloudGuard protection, and select the Protected With Elevation option, after which they can be elevated.

Elevate with the CloudGuard Mobile App

CloudGuard Dome9 account users can elevate themselves with the CloudGuard Mobile app.

1. Open the mobile app, and select **IAM Safety** from the primary menu.
2. Tap on a Role or User from the list, to grant an authorization window to access the AWS service. The period of the window is indicated. The size of the authorization window can be configured on the Settings page of the app.

IAM Reports

The IAM report gives you a dashboard view of the CloudGuard IAM users for your cloud accounts. This shows summary statistics for these users, the number that uses MFA, the rate of password or access key rotation, and the use of IAM Safety leases.

Put the cursor on one of the summaries for the details (the numbers) and click it to open a detailed report on the Policy or Credential report page.

Two more reports can be reached from the dashboard:

- The IAM Policy Report
- The Credentials Report

IAM Policy Report

This shows the IAM policies that have been defined in your cloud accounts. These policies give permissions to these users, so this view shows you which permissions have been granted. IAM policies grant permissions to IAM users or roles (which are then assumed by users). The view shows the specific cloud service (such as EC2, RDS, etc) affected by the policy, and the IAM entity (user or role) give permission.

You can filter the view for specific values in the columns:

- **Filters** - Saves a set of filter settings to use again.
- **Group By** - Aggregates entries into groups.
- **Export to CSV** - Exports the details as a file.

Credentials report

This report shows details for the IAM user accounts on your cloud account. It includes if the user is enabled for *"IAM Safety" on page 403* access, if MFA is enabled, and the date the password was last changed. To filter the view for specific columns, use the same filters as the IAM Policy Report.

Workload Protection

This section describes how to use CloudGuard to manage and protect workloads in your AWS and Microsoft Azure environments and to secure containers in clusters.

CloudGuard provides these kinds of protection:

- **Risk Assessment** (Proact) - Runs risk assessments on the AWS Lambda (serverless) functions in AWS environments that are onboarded to CloudGuard. This includes, for example, identification of overly permissive IAM roles, scanning for vulnerabilities, and hard-coded credentials. The result of the assessment is shown for each asset. This type of protection is available for AWS environments only.
- **Runtime Protection** - Monitors AWS serverless functions at runtime, checks inputs and runtime behavior, and generates notifications for suspicious behavior. In addition, you can apply Runtime Protection to AWS functions at the CI/CD stage on their deployment in your environment. For Kubernetes clusters, Runtime Protection monitors the kernel system calls done by workload containers. You can optionally configure CloudGuard to block unwanted, malicious, or anomalous activity that it discovers. This type of protection is available for AWS environments and Kubernetes clusters.
- **Admission Control** - Monitors your clusters and enforces a security baseline on a namespace or cluster. It can detect if your clusters do not comply with the common practices of having good labels, annotations, resource limits, or other settings.
- **Image Assurance** - Helps you analyze Kubernetes images at each stage of their life cycle to make sure that they are clean. The Image Assurance agents continuously check the clusters and registries for all images. If the agent identifies an unknown image, it scans and analyzes the image to find vulnerabilities, exploits, malware, viruses, trojans, credential leakage, and other malicious threats. This type of protection is available only for container images.
- **Agentless Workload Posture** (AWP) - Provides continuous security assessment of your workloads without the need to install agents in each virtual machine.

More Links

- *"Container Assets" on page 412*
- *"Kubernetes Containers" on page 415*
- *"Serverless Risk Assessment" on page 525*
- *"Kubernetes Posture Management" on page 307*
- *"Runtime Protection" on page 529*
- *"Admission Control" on page 476*
- *"Image Assurance" on page 434*

- *"Agentless Workload Posture" on page 489*
- *"Intelligence for Kubernetes Containers" on page 645*
- *"Onboarding Kubernetes Clusters" on page 188*
- *"Onboarding Container Registries" on page 204*

Container Assets

CloudGuard provides workload protection for these container assets:

- Container Environments
  - Clusters
  - Registries
- Container Workloads (Kubernetes pods, ReplicaSets, Deployments, DaemonSets, CronJobs, etc.)
- Container Images
  - Kubernetes images
  - ShiftLeft images
  - Container Registry images

Container Environments

The (Container) Environments page shows your environments onboarded to CloudGuard:

- Kubernetes clusters
- ShiftLeft environments
- Container registries

as well as assets that are part of other managing environments (for example, AWS or Azure) but not protected by CloudGuard:

- AWS EKS clusters
- Azure Kubernetes Service clusters
- ECR
- ACR

For each environment and asset, you can enable available features and see the status of the features that are already enabled.

Use Cases

Typical use cases to illustrate the control of Container Environments from one central location.

- Review feature status
- Review agent status
- Identify unsecured assets and onboard them (set protection)

Features to Onboard

- Posture Management - This feature is enabled by default when you onboard an environment to CloudGuard, and it cannot be disabled. To learn more, see *"Kubernetes Posture Management" on page 307*
- Admission Control - To learn more, see *"Admission Control" on page 476*
- Image Assurance - To learn more, see *"Image Assurance" on page 434*
- Threat Intelligence - To learn more, see *"Intelligence for Kubernetes Containers" on page 645*
- Runtime Protection - To learn more, see *"Kubernetes Runtime Protection" on page 549*

Viewing Unsecured Environments

With multiple environments onboarded to CloudGuard, it is sometimes hard to monitor which clusters and container registries are onboarded and which are not. The Container Environments page provides information about these assets at a glance. It supports EKS, AKS, ECR, and ACR.

To see unprotected clusters and registries:

1. Navigate to **Workload Protection** > **Container Assets** > **Environments**. This shows a list of environments added to CloudGuard.
2. Click the header of the **Onboarding Time** column to adjust its order and scroll the table to see the environments with a **Click To Onboard** link in this cell. These assets are onboarded to CloudGuard as part of other environments and are not secured with container security features.
3. As an alternative, filter the environments by **Status: Unsecured**.

To enable protection:

- Use the **Click to Onboard** link to open the onboarding wizard.
  - For more information about onboarding clusters, see *"Onboarding Kubernetes Clusters" on page 188*.
  - For more information about onboarding container registries, see *"Onboarding Container Registries" on page 204*.

More Links

- *"Kubernetes Containers" on page 415*
- *"Create a ShiftLeft Environment and Service Account" on page 469*
- *"Container Registry Scanning" on page 464*

Kubernetes Containers

Kubernetes is an open-source container orchestration system for automating the deployment, scaling, and management of containerized applications. It operates with a range of container tools and runs containers in a cluster with images built with Docker, OCI, or Kaniko. It groups containers that make up an application into logical units for easy management and discovery.

Before you can use Kubernetes Containers features in CloudGuard, your Kubernetes cluster must already be onboarded to CloudGuard. See *"Onboarding Kubernetes Clusters" on page 188* for details on how to do this.

Supported Versions

Name	Version
Kubernetes	■ Version 1.21[1] and higher (with managed and unmanaged distributions)[2]
Kubernetes-based Container Orchestration Platforms	■ Red Hat OpenShift v4.6 and higher (Runtime Protection: nodes running Red Hat Enterprise Linux CoreOS) ■ VMware Tanzu TKG v1.2 and higher, TKGI v1.10 and higher
Container Runtime	■ Docker v1.32 and higher ■ containerd 1.1 and higher ■ CRI-O 1.16 and higher

Name	Version
Node Operating System	■ All blades, except for Runtime Protection and Flow Logs: • Any Linux ■ Flow Logs: • Any Linux with kernel 4.1 and higher ■ Runtime Protection - Linux distros with kernel 4.14 and higher, as follows: • Ubuntu • Debian • CentOS • RHEL including these Cloud Vendor Operating Systems: • Amazon Linux 2, Amazon Linux 2023 • Red Hat Enterprise Linux CoreOS • AWS Bottlerocket • Google Container-Optimized OS (Autopilot clusters are not supported) **Note** - Some Linux distributions may require additional setup – see *"Prerequisites" on page 549*)
Node architecture	■ AMD64 ■ ARM64 (Runtime Protection is not supported)

[1] Kubernetes versions from 1.16 to 1.20 are supported only with Helm deployment instructions, with the regular `helm upgrade --install` command.

[2] CloudGuard does not support hybrid clusters with multiple (mixed) container runtimes. You cannot change the container runtime after the service is onboarded in the cluster. For this, upgrade the solution.

Notes:

- CloudGuard for Kubernetes containers is available on platforms such as: vanilla Kubernetes, AKS on Azure, EKS on AWS, GKE on GCP, OpenShift, Tanzu, Rancher (k3s), and others.
- CloudGuard automatically detects the container runtime on onboarding with Helm.
- CloudGuard automatically detects OpenShift and Tanzu platforms and adjusts the Helm deployment accordingly.

Version Deprecation

Deprecated Kubernetes versions are not supported by cloud vendors and do not get important security updates.

Important - Clusters with deprecated versions can be at risk.

Platform	Service	Supported versions
Microsoft Azure	AKS	Link
Google Cloud Platform	GKE	Link
Amazon Web Services	EKS	Link
Oracle Cloud Infrastructure	OKE	Link
Red Hat	OpenShift	Link
Kubernetes	Kubernetes	Link

Requirements

Permission Requirements

The CloudGuard agent requires Kubernetes permissions for:

Verb	Group	Resource	Scope
get, list	-	pods services nodes nodes/proxy serviceaccounts namespaces resourcequotas	Cluster-wide
	apps	daemonsets deployments replicasets statefulsets	Cluster-wide
	networking.k8s.io	networkpolicies ingresses	Cluster-wide
	extensions	ingresses	Cluster-wide
	policy	podsecuritypolicies	Cluster-wide
	rbac.authorization.k8s.io	roles rolebindings clusterroles clusterrolebindings	Cluster-wide
	batch	cronjobs	Cluster-wide
	-	pods secrets configmaps	Agent namespace (default: *checkpoint*)
patch	admissionregistration.k8s.io	validatingwebhookconfigurations	Cluster-wide
all (*)	*.cloudguard.checkpoint.com	all (*)	Agent namespace (default: *checkpoint*)

The CloudGuard agent requires these Kubernetes permissions for OpenShift Kubernetes clusters:

Verb	Group	Resource	Scope
get, list	config.openshift.io	clusteroperators (resourceName: openshift-apiserver)	Cluster-wide
	operator.openshift.io	openshiftapiservers, kuberapiservers (resourceName: cluster)	Cluster-wide
	security.openshift.io	securitycontextconstraints	Cluster-wide
get -	-	configmaps (resourceName: config)	openshift-kube-controller-manager, openshift-apiserver, openshift-kube-apiserver
		configmaps (resourceName: kube-scheduler-pod)	openshift-kube-scheduler

Note - For Kubernetes clusters on OpenShift, CloudGuard creates additional roles in OpenShift namespaces. The role bindings bind these roles to the CloudGuard service account (in CloudGuard namespace used for the agent).

Pod and Container Requirements

Linux allows assigning specific capabilities to processes, thus restricting the processes to only the minimum privileges required to perform their tasks and reducing the risk of security breaches.

The CloudGuard agents require these Linux kernel capabilities:

Agent Name	Container Capabilities	Host Network	Privileged
Runtime Protection (runtime-daemon)	▪ SYS_RESOURCE ▪ SYS_ADMIN ▪ SYS_NICE ▪ SYS_PTRACE ▪ FOWNER ▪ SYS_PACCT ▪ NET_ADMIN ▪ NET_RAW ▪ AUDIT_READ ▪ AUDIT_WRITE ▪ AUDIT_CONTROL	required	OpenShift, CRI-O, CoreOS, Bottlerocket: required
Flow Logs(flowlogs-daemon)	▪ SYS_RESOURCE ▪ SYS_ADMIN ▪ NET_ADMIN	required	OpenShift: required
Image Assurance (imagescan-daemon)	▪ NET_BIND_SERVICE	not required	OpenShift and CRI-O: required

Connectivity Requirements

CloudGuard agents must have connectivity to these domains:

Blade or Agent	Address
CloudGuard Image Assurance Image Scan	.dome9.com **Notes**: ▪ The domain `.dome9.com` represents Check Point's Network Domain Objects. ▪ For Image Assurance agents ver. 2.28.0 and lower, add the `.checkpoint.com` domain.
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query
Container Registry	https://quay.io/checkpoint

Instead of the domain objects, you can use the region-specific URLs for your Data Center location from the table below. Add these endpoints to the allowlist.

Blade or Agent	Address
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query
Container Registry	https://quay.io/checkpoint

For Image Scan agents **ver 2.28.0 and lower**, you must use additional endpoints. To learn more about agent's version, see *"Agent Version Life Cycle" on page 196*.

Blade or Agent	Address
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com

United States (US)

Blade or Agent	Address
CloudGuard	https://api-cpx.dome9.com https://api.dome9.com
Threat Intelligence	https://validator-prod-k8s.s3.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://us-gw.sg.iaas.checkpoint.com

Europe (EU)

Blade or Agent	Address
CloudGuard	https://api-cpx.eu1.dome9.com https://api.eu1.dome9.com
Threat Intelligence	https://validator-prod-533924475734-k8s.s3.eu-west-1.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://eu-gw.sg.iaas.checkpoint.com

Australia (AU)

Blade or Agent	Address
CloudGuard	https://api-cpx.ap2.dome9.com https://api.ap2.dome9.com
Threat Intelligence	https://validator-prod-583664506098-k8s.s3.ap-southeast-2.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://au-gw.sg.iaas.checkpoint.com

Canada (CA)

Blade or Agent	Address
CloudGuard	https://api-cpx.cace1.dome9.com https://api.cace1.dome9.com
Threat Intelligence	https://validator-prod-052001227150-k8s.ca-central-1.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://ca-gw.sg.iaas.checkpoint.com

India (IN)

Blade or Agent	Address
CloudGuard	https://api-cpx.ap3.dome9.com https://api.ap3.dome9.com
Threat Intelligence	https://validator-prod-573281234161-k8s.s3.ap-south-1.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://in-gw.sg.iaas.checkpoint.com

Singapore (SG) - for Dome9 accounts only

Blade or Agent	Address
CloudGuard	https://api-cpx.ap1.dome9.com https://api.ap1.dome9.com
Threat Intelligence	https://validator-prod-155213570047-k8s.s3.ap-southeast-1.amazonaws.com
Image Scan (agent ver. 2.28.0 and lower)	https://sg-gw.sg.iaas.checkpoint.com

If the CloudGuard pod image is uploaded to a private repository, connectivity to Container Registry is not necessary. In this case, the Helm chart parameter `image.repository` must be changed to indicate the location of the image. For more information about how to set this parameter, see https://github.com/CheckPointSW/charts/tree/master/checkpoint/cloudguard.

Resource Requirements

Each CloudGuard feature can have pods that run in daemonsets and pods that run in deployments. For the pods in daemonsets, resources are shown in the table below **per node**. For the pods in deployments, resources are below **per cluster**. The pods that run in deployments can certainly run on different nodes.

You can find the default values of requests and limits in the defaults.yaml on the Helm Chart.

Basic Feature	Per cluster or node	CPU (millicores)		Memory (MiB)	
		requests	limits	requests	limits
Posture Management	per cluster	100	200	50	50
Image Assurance	per cluster	200	1050	200	2600
	per node - Docker	50	50	50	50
	per node - containerd	200	250	150	150
	per node - CRI-O*	200	300	250	250

Basic Feature	Per cluster or node	CPU (millicores)		Memory (MiB)	
		requests	limits	requests	limits
Admission Control	per cluster	1150	1350	330	450

* OpenShift uses the CRI-O runtime

Note - Image Assurance scanning engine (imagescan-engine) requires free ephemeral storage, which size is double maximal size of the scanned image. Maximal image size means the uncompressed, `docker save` output image size.

Premium Feature	Per cluster or node	CPU (millicores)		Memory (MiB)	
		requests	limits	requests	limits
Runtime Protection	per cluster	50	50	30	50
	per node*	200	400	300	800
Flow Logs	per node	100	200	30	100

* Large Nodes (above 8 vCPUs) may require additional resources in case of pod restarts.

More Links

- *"Admission Control" on page 476*
- *"Image Assurance" on page 434*
- *"Kubernetes Posture Management" on page 307*
- *"Intelligence for Kubernetes Containers" on page 645*
- *"Kubernetes Runtime Protection" on page 549*

For Kubernetes terminology, see the Glossary in the Kubernetes documentation.

Images

To see your workload images, you must onboard the environment that contains these images to CloudGuard. See *"Onboarding Cloud Environments" on page 53* to onboard your environment.

CloudGuard supports images built with Docker, OCI, and Kaniko.

When you enable Image Assurance on your cloud environments, you can see all the images that run on these environments and their scan status on **Workload Protection** > **Containers Assets** > **Images**. CloudGuard does not show images that have not been run on an onboarded workload. CloudGuard considers a Kubernetes or ECS image *running* when a workload with this image is running in the relevant Kubernetes or ECS environment. CloudGuard considers a Container Registry or ShiftLeft image *running*, when it is running in some Kubernetes or ECS environment in the tenant.

After onboarding, the images start to appear on the page with the **Scanned** scan status. In addition to the regular scans, you can schedule on-demand scanning of Kubernetes and Container Registry images whose status is other than **Scanned**.

Scanning Time Frames

Scanning time frames can be different:

- After onboarding, CloudGuard shows the list of discovered Kubernetes images. Because scan results are shared between environments, some images can already be scanned in other Kubernetes or Container Registry environments, so they appear first on the **Images** page.
- Kubernetes images are scanned gradually. The first scanned images are shown several minutes after they appear in the portal.
- New images added to the registry take up to 12 hours to be scanned, based on the registry configuration. The scan period is configurable on the container registry page.
- For images, the on-demand scanning request schedules the image for scanning in several minutes with priority over other regular images. Actual scanning can start later if multiple images are prioritized.
- For environments, the on-demand scanning request schedules their images for scanning within several minutes.

Image Parameters

Use the **Platform** filter to show available images by group:

- Container Registry image
- Kubernetes image
- ShiftLeft image

The Images page allows you to see immediately the vulnerability level and risk score of the scanned images:

- **Risk** - Image risk score from 0 to 10 based on the Common Vulnerability Scoring System (CVSS).
- **Registry** - Registry that stores this image.
- **Is Running** - The check mark indicates images corresponding with currently running workloads. The empty space indicates inactive images.
- **CVEs** - Summary of the CVEs by severity.
- **Scan Status** - See *"Image Scan Status" on page 429*.
- **Last Running Date** - Indicates when a related workload was last seen running. An empty cell means that a related workload has not been seen.

Click the image name to see more details about its status, properties, and posture findings.

Base Image

While scanning your workloads, CloudGuard sends notifications about all findings: CVEs, secrets, or other vulnerabilities. If you want to focus only on the findings relevant to remediate, you can exclude vulnerabilities inherited from the base image. For example, if you build your image on an Ubuntu image, you can create a rule that excludes the base image vulnerabilities from your findings and see only those findings created by your dependencies.

In CloudGuard, the base image is an entity that was used as a basis to create your images or other base images. When you decide which images to mark as base images, select those extended by other images in use.

Important - The base image feature requires upgrading your Image Assurance agent to version 2.37 or higher.

Base Image Repository

Best Practice - Check Point recommends to set up a separate repository for base images that you use. CloudGuard automatically recognizes all images added to the repository as base images.

To set a repository as a Base Image repository, create a Base Image rule.

To create a Base Image Rule:

1. Navigate to **Workload Protection** > **Container Assets** > **Images** and open your image.
2. On the top right, click the menu ⋮ and select **Add Base Image rule**.

Or you can open **Workload Protection** > **Vulnerabilities** > **Base Image Rules** and click **Add**.

3. In the **Add Base Image rule** window, enter the details:
   a. **Name** (mandatory) - Enter a distinctive name for the rule, for example, *My Org Nginx Rule*.
   b. **Registry Environment** (mandatory) - Select one or more container registries where you apply the rule.
   c. **Repository** (mandatory) - Enter a repository to contain the base images. For example, if your image URL is `myrepo.com/this/is/my/imageName:11`,
      i. registry - `myrepo.com`
      ii. repository - `this/is/my/imageName`
      iii. tag - `11`
   d. **Description**
4. Click **Save**.

The image located in the repository obtains the *Base Image* indication in its details, and, in the **Entity Viewer**, the image belongs to the **BaseImage** group.

Images extended from the base image have the indication in the asset details that they are **Based On** the base image, with a link to the base image. The link opens the **Images** page filtered by the base image's SHA256, which shows all its copies in this account.

Important - An image set as a base image is considered as such in all environments (for example, container registries, clusters, or ShiftLeft environments) onboarded to the CloudGuard account.

To sort out base image vulnerabilities:

1. Navigate to **Workload Protection** > **Container Assets** > **Images** and open your image.
2. Go to the **Vulnerabilities** tab. It shows the list of the found CVEs.
3. Group the CVEs by **Base image**.
   - The CVEs in the Base image group are found in the packages installed on the base image.
   - The CVEs in another group are found in the packages installed only on your image, so you can remediate them.
4. Go to the **Threats** or **Secrets** tab. In the same manner, sort out the threats and secrets that come from the base image or from the resources in your image.

CloudGuard scans only the latest (most recently used) images from the repository. On the asset page, you can configure the maximum number of these images for each base repository. For more details, see *"Configuring Scanning of Registries" on page 466*.

Posture Findings

You can create a policy that allows you to see only those findings that are found in your image but not in the base image. For this, create a rule that excludes all findings originating in the base image. Use the entity property called `baseImages`.

Example 1

The rule is triggered when the package has CVEs with a High severity level *and* while it originates outside the base image.

```
Package where baseImages isEmpty() should not have cves contain-
any [severity='High']
```

Vendor Image

Similar to the base images, CloudGuard indicates vendor images. These images are created and maintained by cloud vendors, such as AWS, Azure, or Google Cloud Platform, so you can exclude findings related to these images. CloudGuard regularly updates the list of vendor images. If a vendor image is missing in the image list, contact *Check Point Support*.

Posture Findings

You can create a policy that allows you not to include those findings that are found in the vendor image. For this, create a rule that excludes all findings originated in the vendor image. Use the entity property called `imageGroups` under `scannedAssets`.

Example 2

The rule is triggered when the package has CVEs with High severity level *and* while it does not originate in the vendor image.

```
Package where scannedAsset.imageGroups contain-none
[name='Vendor Image'] should not have cves contain-any
[severity='High']
```

Layers

CloudGuard shows the layers that comprise the image. The layers are ordered based on their creation date. When CloudGuard discovers a vulnerability, it can show at which layer the vulnerability was introduced and present the summary of all layer commands. You can filter the layers by **Layer ID** or by **Layer Command**. To learn more details, click the **Layers** panel and open the **Vulnerabilities** tab.

The **Vulnerabilities** page shows the image vulnerabilities (CVEs and Threats) grouped by the layers and sorted by severity from **Critical** to **Low**. For example, see CVE statistics for each layer and then expand it to see the CVEs found in this layer.

Example

This example shows how to resolve the image vulnerability:

1. Open the image **Overview** and review the image **Layers**.
2. Notice at which layer the vulnerability entered the image and see its layer command.
3. Navigate to the **Vulnerabilities** page. The CVEs are arranged by the **Layer Command**.
4. Expand the layer command to see all CVEs found.
5. Open the CVE that you want to resolve, for example, one of the critical severity, and find its **Remediation**.
6. Upgrade the package to the recommended version.

After the upgrade, the image has a new layer where the relevant vulnerabilities are resolved.

Image Scan Status

See the table below for all statuses.

Scan Status	Description	Corrective Action
Scanned	The image is successfully scanned.	
Pending Scan	▪ The image awaits to be scheduled for a scan. Applicable to Fargate images: ▪ No matching image scans are found for the Fargate image.	
Partial	Scan results are partial; the image will be scheduled for rescanning.	

Scan Status	Description	Corrective Action
Unsupported OS	The image operating system is not supported (for example, Windows is not supported).	
Unmatched	Applicable for ECS images: No matching image scans were found for the ECS task image.	
Not an image	An artifact found in the registry is not an image (for example, Helm chart).	
Network Error	Unable to create a connection to scanning services, possibly because of a firewall or a proxy.	Verify your firewall/proxy configuration to make sure it does not block access to the required CloudGuard URLs. See the **Connectivity Requirements** section in *"Kubernetes Containers" on page 415*.
Unauthorized	Failed on one of these: 1. Failed to authenticate with CloudGuard. 2. Failed to authenticate with the container registry (for example, because of expired credentials). 3. Failed to verify CloudGuard certificate, possibly because of the firewall/proxy.	Verify your firewall/proxy configuration to make sure it does not block access to the required CloudGuard URLs. See the **Connectivity Requirements** section in *"Kubernetes Containers" on page 415*. If the image is from a container registry environment, follow the procedure for Error **2** of *"Error Messages in Agent Status" on page 458*.

Scan Status	Description	Corrective Action
Insufficient resources	■ The image is too large to be scanned. or ■ No space left on your host machine.	The maximum allowed image size is 20 GB. If you need to scan larger images, contact *Check Point Support Center*. If the image size is less than 20 GB, examine the space left on your cluster machine.
Timeout	Timeout on pulling the image to be scanned.	Examine your network connectivity on the cluster and try to increase the image pull timeouts by setting the environment variables. See the **Central Agent Environment Variables** section in *"Image Assurance Troubleshooting" on page 452*.
Internal Error	An unknown error has occurred. The image will be rescheduled for a scan.	■ Review the imagescan-engine logs and identify the engine container reporting errors and the node running it. ■ Examine the container metrics. If it reaches memory limits, increase the limits. If the node's memory utilization is high, increase the number of memory requests of the container. ■ Examine the free disk space of the node. For ECS scanning environments, examine the ephemeral storage of the task (the default is 20 GB). If the problem continues, contact *Check Point Support Center*.

Inactive Images

CloudGuard deletes inactive images in a specific period.

- **Kubernetes Images** - CloudGuard considers a Kubernetes image inactive if none of its corresponding workloads are running. You can set the period after which CloudGuard deletes inactive images (by default, 7 days).
- **Container Registry Images** - A container registry image is live (active) if at least one Kubernetes container corresponding to this image is running in your CloudGuard account. CloudGuard deletes inactive container registry images during the 24 hours (not immediately) after they were deleted from the registry. You cannot set the period for the deletion of these images.

- **ShiftLeft Images** - CloudGuard considers a ShiftLeft image inactive if none of its corresponding workloads are running. You can set the period for the image deletion after the last scanning of this image (by default, 30 days).

You can set the lifetime for inactive Kubernetes and ShiftLeft images in the *"Workloads Settings" on page 861*.

On-Demand Image Scanning

In addition to the regular scans, you can schedule on-demand scanning of these:

- Kubernetes environments and images
- Container Registry environments and images
- AWS environments and images

Inactive (non-running) images cannot be requested for scan:

- If an environment is requested for scan, its inactive images are not considered.
- If an image is requested for scan, the scanning process is not triggered.

Scanning Failed Images

This process starts the scanning of all images in an environment that are not in the **Scanned** status.

To start the image scan on demand:

1. Navigate to **Assets** > **Environments** and select a cluster or a container registry.
2. Click to open the environment page.
3. Click **Retry Failed Scans**.

Scanning Individual Images

This process schedules an image for scanning regardless of its status.

To start the image scan on demand, do one of these:

- On the image level:
  1. Navigate to **Workload Protection** > **Containers Assets** > **Images** and select an image.
  2. Click **Request Scan**.
- Or on the environment level:
  1. Navigate to **Assets** > **Environments** and select a cluster or a container registry.
  2. Click to open the environment page.
  3. Open the Images tab.
  4. In the image row, see its Scan Status, then click the menu ⋮ and select **Request Scan**.

For on-demand image scanning with API, see Workload Image Assurance in the API Reference Guide.

Limitations

- Images used by short-lived pods may not be visible to Image Assurance.
- The **Request Scan** usage is limited to 200 requests in an hour.
- Requests for a scan of *inactive* images are not available.
- On-demand scanning is not supported for ShiftLeft images and environments.

Image Assurance

CloudGuard Image Assurance analyzes container images for vulnerabilities at each stage of their life cycle to make sure they meet your organizational policies.

The Image Assurance agents continuously check Kubernetes clusters and registries to scan the discovered container images. If an agent identifies an unknown image, it scans and analyzes the image for vulnerabilities, exploits, malware, viruses, trojans, credential leakage, and other malicious threats. In the Kubernetes clusters, only images of the running workloads are scanned.

CloudGuard Workload Protection - Image Assurance

Before you can see Vulnerabilities, you must onboard your Kubernetes cluster or AWS account to CloudGuard. See *"Onboarding Kubernetes Clusters" on page 188* and *"Onboarding AWS Environments" on page 144*.

How Image Assurance Works

What to scan	How to scan	Prerequisites
Kubernetes clusters	Deploy Image Assurance on the Kubernetes cluster	Onboard a Kubernetes cluster
Container registry	Deploy Image Assurance on a hosting Kubernetes cluster	Onboard a Kubernetes cluster
	Use a CloudFormation Template to deploy the ECS scanner resources on your AWS account	Onboard an AWS account
AWS ECS tasks container images	Deploy Image Assurance on a hosting Kubernetes node and scan the ECR that hosts container images used in the applicable ECS cluster	Onboard a Kubernetes cluster Onboard the related ECR
	Use a CloudFormation Template to deploy the ECS scanner resources on your AWS account and scan the ECR that hosts container images used in the applicable ECS cluster	Onboard an AWS account Onboard the related ECR

Resources

Image Assurance scanner uses these resources:

- **ImageScan List** - A single-replica Deployment that sends CloudGuard container image lists. The lists are collected from the image-scan-daemon pods and from the connected Container registries.
- **ImageScan Engine** - A single-replica Deployment that analyzes and scans container images. The agent sends CloudGuard the necessary information to complete the scan. For more information about the agent's version, see *"Agent Version Life Cycle" on page 196*.
- **ImageScan Daemon** (for Kubernetes images only) - A DaemonSet that provides a list of local images (on each node) and the content of the requested images.

CPU

When the ImageScan Engine pod scans images, it can consume more than one CPU. In a stable state, when only new images are scanned most of the time, the Engine pod consumes a very low CPU.

Reduction of the values of the requests and limits for CPU can have an opposite effect on the scan time.

Supported Packages

Image Assurance and the CI/CD tool support these types of packages:

- Distro package managers (Alpine, Debian, Ubuntu, RHEL, and CentOS)
- .Net languages (C#, C++, F#, VB)
- Node.js packages
- Python packages (requirments.txt)
- Ruby gems
- Java artifacts (JAR files)
- Go packages

Image Assurance on AWS Fargate

Because AWS Fargate is managed by AWS, it does not allow to install CloudGuard agents on its nodes. This means that Image Assurance cannot scan images on Fargate nodes.

To obtain this information, CloudGuard shares data from other scanners available on the same account and shows it for Fargate images with this note: "This image is hosted on Fargate node and cannot be scanned/rescanned from the node. Scan results depend on correlation with other scanners (Kubernetes and registry)". If no matching image scans are found for the image, its scan status is **Pending Scan**, until the image appears in an environment where it can be scanned.

To see the Fargate image status:

1. Navigate to **Assets** > **Environments** and select an environment that contains the Fargate image.
2. Go to the **Images** tab. A Fargate image has the above note in its details, and the option of **Request Rescan** (from ⋮ menu) is not available for it.
3. See the image status in the **Scan Status** column.

Image Assurance on GKE Clusters

Image Assurance is supported on Google Kubernetes Engine (GKE) with Image Streaming. To be able to scan Kubernetes images in GKE clusters, you need the Image Assurance agent v2.30.0 and higher included in Helm chart v2.30.0 and higher.

Image Streaming is enabled by default in GKE Autopilot clusters. In GKE Standard clusters, you have to enable it explicitly to use Image Assurance.

GKE Supported Versions:

- GKE version 1.28.9 and higher
- GKE version 1.28.8 using nodes with images `1.28.8-gke.10950000` and higher

More Links

- *"Kubernetes Containers" on page 415*
- *"Vulnerability Policies (Image Assurance)" on page 444*
- *"Vulnerability Findings (Image Assurance)" on page 439*
- *"Configuring CloudGuard Exclusions" on page 80*
- *"Image Assurance Troubleshooting" on page 452*
- *"Image Scan Findings" on page 449*
- *"Images" on page 425*
- *"Image Admission" on page 486*
- *"Container Registry Scanning" on page 464*

Getting Started with Image Assurance Policy

When you enable Image Assurance, CloudGuard protects your cluster with an automatically created default ruleset and default policy. For configuration of the default policy, see *"Vulnerability Policies (Image Assurance)" on page 444*.

Image Assurance policy apples to Kubernetes image assurance, container registries, ECS images and ShiftLeft image scans.

Configuring an Image Assurance Policy

Follow the steps below to change the default policy or create a new policy.

Step 1. Changing the Image Assurance managed ruleset

1. Navigate to **Workload Protection > Vulnerabilities > Rulesets**.
2. Click the **Container Image Assurance** ruleset.
3. Click **Clone**.
4. Give the ruleset a name and a description.

Step 2. Creating a New Image Assurance Ruleset

1. Navigate to **Workload Protection > Vulnerabilities > Rulesets**.
2. Click **Add Ruleset**.
3. Enter a name and a description for the ruleset.

Step 3. Adding a Rule to a Ruleset

1. Navigate to **Workload Protection > Vulnerabilities > Rulesets**.
2. Click the thumbnail of the ruleset to which you add a rule.
3. Click **New Rule** to open the rule editor page.
4. Click the **GSL** text box to open the GSL Rule Editor. For more information, see *"GSL Builder" on page 350*.
5. Click **Verify** to verify the rule.
6. Click **Done**.
7. Enter or update other available fields (Title, Description, etc.)
8. Click **Save** to save the new rule.

Step 4. Removing the default policy associated with a Kubernetes cluster

1. Navigate to **Workload Protection** > **Vulnerabilities > Policies**.
2. Find the Kubernetes cluster association.
3. Click the menu above the table and select **Unassociate**.

Step 5. Adding a new Image Assurance Policy

You can configure more than one policy on a cluster.

1. Navigate to **Workload Protection > Vulnerabilities > Policies**.
2. Click **Add Policy**.
3. Select **Environment Policy** if you want to apply this policy to a cluster, or **Organizational Unit Policy** if you want to apply it to the unit to which the cluster belongs.
4. Select the cluster or Organizational Unit to apply the policy.
5. Click **Next**.
6. Select the ruleset to bind to the cluster or OU. You can select multiple rulesets.
7. (Optional) For Kubernetes environments, enable **Admission Control** (**Image Admission)** in Detection or Prevention mode. For more details, see *"Image Admission" on page 486*.
8. Select the notification that appears when this policy is violated. You can configure a new notification or select an existing one.

   Note - To see findings in the CloudGuard portal, make sure the **Alerts Console** option is selected.

9. Click **Save**.

More Links

- *"Image Assurance" on page 434*
- *"Vulnerability Policies (Image Assurance)" on page 444*
- *"Image Assurance Troubleshooting" on page 452*
- *"Image Admission" on page 486*

Vulnerability Findings (Image Assurance)

CloudGuard creates Vulnerability findings for Container images based on the assigned policy.

CloudGuard automatically creates a policy with a default Image Assurance ruleset for applicable clusters. If the default policy is sufficient, no more actions are necessary. If the onboarded environment is part of an Organizational Unit with an Image Assurance policy, no default policy is associated with the environment.

Viewing one-image findings

To see the findings in the CloudGuard portal:

1. Navigate to **Workload Protection** > **Containers Assets** > **Images**.
2. Select an image. Use Environment, Asset Type, or other criteria to filter images.
3. Go to the **Events** page and click **Vulnerabilities**. Make sure you set the time selector to **All** to see all findings for the image.

Viewing Vulnerability findings

To see the vulnerability findings for all clusters and images in the account, navigate to **Workload Protection** > **Vulnerabilities** > **Findings**.

CloudGuard creates the findings when it scans the image for the first time. Afterward, the CloudGuard portal checks it (one time) in several hours for changes or newly discovered vulnerabilities.

To see findings for your AWS ECS images, use the filter for the **AWS** Platform and **AwsEcsImage** Entity Type. In addition, see the vulnerabilities in the **AWS ECS image** object.

On this page, use the *"Filter and Search" on page 863* toolbar to select parameters to filter out from the **Findings** table.

Use these preconfigured filters:

- Environment or OU - Select one or more cluster environments or organizational units.
- Severity - Select from the available alert severity objects.
- Ruleset - Select from the available rulesets.

Viewing workloads

To see the workloads that use vulnerable images:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Findings**.
2. Select one of the findings. On the right, the entity card shows information about the image.
3. Click the image link. CloudGuard redirects you to the asset page of the image.
4. The Overview page shows workloads that contain this image. For more information about Overview, see *"Asset Details" on page 273*.

Categories of Findings

Image Assurance finds different types of findings grouped in the categories:

- **CVE** - Common Vulnerabilities and Exposures
- **MaliciousURL**
- **MaliciousIP** - For more details, see *"Malicious IP Classification" on page 311*
- **MaliciousFile** - Malware
- **InsecureCode**
- **InsecureContent** - Credential leakage

  Note - This feature is in Early Availability.

- **ImageScan** - Indicates that the number of issues or severity of the issues found on an image exceeds a preconfigured threshold. See *"Image Scan Findings" on page 449*
- **Package** - Package license, package info, and CVEs

Details of Findings

The fields in Image Assurance findings are almost the same as other fields in the finding details (see *"All Events" on page 120*).

Fields for Kubernetes images:

- **Title** - The specific ID or type for which the finding is created based on the finding category.

  - ImageScan findings have the title with the name of the image
  - Common Vulnerabilities and Exposures (CVE) findings have the title with the CVE ID
- **Description** - The issue description, for example, the CVE description as it appears in the National Vulnerability Database (NVD).
- **Environment** - The Kubernetes cluster that contains the image with the finding.

More Links

- *"Image Assurance" on page 434*
- *"Image Scan Findings" on page 449*
- *"All Events" on page 120*
- *"Asset Details" on page 273*

Vulnerability Exclusions

You can select to exclude a specific vulnerability (CVE, threat, or secret) that appears in a specific package. If you set the vulnerability as not important or relevant, CloudGuard ignores it. Such vulnerability has the **Excluded from Findings** indication on the asset pages. CloudGuard rules do not take the vulnerability into account.

The vulnerability exclusions are applied to the raw data in the package before running an assessment and obtaining findings. Therefore, these exclusions affect findings, toxic combinations, notifications, and more.

To exclude only findings related to vulnerabilities, see *"Configuring CloudGuard Exclusions" on page 80*.

Creating a CVE Exclusion

1. Navigate to **Workloads** > **Vulnerabilities** > **Vulnerability Exclusions**.
2. Click **Add** in the top left.
3. In the **Create new exclusion** window, enter a name for the exclusion. This parameter is mandatory.
4. Select to **Include in Assessment**, if you want CloudGuard to create a finding related to the vulnerability regardless of its being ignored. This option affects only findings; for toxic combinations, risk score, etc., the vulnerability remains ignored.
5. Select the scope of the exclusion application. For **Environments**, select one or more environments to apply the exclusion. For **Organizational Units**, select one or more units.
6. Select one of these options:
   - **CVE** - Enter these details:
     - CVE ID (mandatory)
     - Package name
     - Package version

- **Threat** - Enter these details:
  - Category
  - Path (mandatory) - file path
- **Secret** - Enter the file path (mandatory).

7. **Entity Name** / **Entity ID** - Exclude vulnerabilities that correspond to specific entities. Enter the entity name or ID. You can enter one or more entity names. Start to type the entity name to see and select a matching option. You can include the wildcard '%' in the entity name, to include a group of entities. For example, `%s3%` matches all entities with 's3' in their name.
8. Enter your comment to distinguish between different exclusions. The comment is a mandatory parameter.
9. Select the date range.
10. Click **Save**.

You can also create a CVE exclusion from the CVE page. In the top right corner, click the menu icon ⋮, select **Exclude**, and enter the required parameters.

If CloudGuard finds the excluded CVE in one of the images, the Vulnerabilities page of the image shows the CVE with the **Is Excluded** indication. Click the CVE ID to open its page and learn more details. It also shows information on who excluded the CVE and when.

Editing a Vulnerability Exclusion

1. Navigate to **Workloads** > **Vulnerabilities** > **Vulnerability Exclusions**.
2. Select an exclusion to edit and click **Edit** on the top bar.
3. Modify the exclusions parameters and click **Save**.

Deleting a Vulnerability Exclusion

1. Navigate to **Workloads** > **Vulnerabilities** > **Vulnerability Exclusions**.
2. Select an exclusion to delete and click **Delete** on the top bar.
3. Click **Confirm** to confirm your choice.

CloudGuard rules start to take the CVE in account.

You can also delete a CVE exclusion from the CVE page. In the top right corner, click the menu icon ⋮, select **Remove exclude**, and click **Confirm** to confirm your choice.

Vulnerability Policies (Image Assurance)

On CloudGuard, findings are created when objects violate GSL rules, as a result of the assessment process as explained in *"Image Assurance Assessment" on the next page*.

A ruleset is a list of rules that defines what is considered a violation. A ruleset corresponds to a specific Kubernetes cluster or an Organization Unit with the use of a Policy configuration.

In addition, Policy Configuration includes the notifications sent when findings are created.

On Image Assurance, everything is configured automatically, so the assessment process is invoked automatically when needed. No user action is required. When you onboard a new cluster to CloudGuard (or enable the Image Assurance feature) and associate it with an Organizational Unit, the cluster obtains the Image Assurance policy configured for this Organizational Unit. If no such policy exists, a new policy is created to associate the new cluster with the default ruleset.

You can assign a new policy or edit an existing one as explained in this section.

Image Assurance policy apples to Kubernetes image assurance, container registries, ECS images and ShiftLeft image scans.

Image Assurance Default Policy Configuration

Note - This section provides all the details of the default Image Assurance configuration. No user action is required.

Image Assurance Managed Ruleset

A managed ruleset is a list of rules that CloudGuard creates automatically for each account. These rules represent CloudGuard security recommendations.

You cannot adjust a managed ruleset. As an alternative, duplicate it and use the edited version of this ruleset, as explained below.

To examine the Image Assurance managed ruleset:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Rulesets**.
2. Find a ruleset named **Container Image Assurance**. For this, filter rulesets by CloudGuard-Managed Type.
3. Click the **Container Image Assurance**, **Container Image Assurance 1.0**, or **Workload Vulnerability Default 2.0** ruleset to see the default rules.

Image Assurance Default Policy

The policy allows you to associate a ruleset with a specific Kubernetes cluster, ShiftLeft environment, Container Registry, or an Organization Unit. Additionally, it allows you to configure alerts and notifications on findings.
Without a related policy, CloudGuard cannot run assessments on image scan results and cannot create Image Assurance findings to alert on weaknesses and vulnerabilities found on vulnerable images.

CloudGuard creates an Image Assurance policy by default for each ShiftLeft environment, Container Registry, and Kubernetes cluster with enabled Image Assurance. This happens if the environment is not associated with an Organizational Unit that already has an Image Assurance policy.

To examine the Image Assurance default policy:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Policies**.
2. Optionally, filter the view by the name of your environment.
3. Find your environment association and a link to the default notification.

Image Assurance Default Notification

A notification defines which actions CloudGuard must do when a violation of rules triggers the creation of findings.

The default notification configuration is the **Alerts Console**. It defines that the findings appear on the CloudGuard portal on **Posture Findings** or the **Findings** page of Vulnerabilities. Check Point recommends to have this option always selected.

To examine the default Image Assurance notification:

1. Navigate to **Workload Protection** > **Vulnerabilities** >**Policies**.
2. Click **Default Image Assurance notification** to review the configuration of Image Assurance findings.

To change the default Image Assurance policy or create a new policy, see *"Getting Started with Image Assurance Policy" on page 437*.

Image Assurance Assessment

During the Image Assurance Assessment process, CloudGuard analyzes image scan results and creates Image Assurance findings. The Image Assurance findings are created in this process for each violation of the Image Assurance rules.

CloudGuard runs the assessment process automatically when the image assurance agents identify an unknown image and scan it. The assessment process runs automatically when you change Image Assurance rulesets and policy, or when you make OU hierarchy changes that have an effect on the association between the Image Assurance policy and your environment. During these events, the assessment process runs on all images applicable to the affected environment.

To see the Image Assurance assessment history of your account, navigate to **Workload Protection > Vulnerabilities > Assessment History**.

Vulnerability Search

You can search for a particular CVE found in your scanned assets regardless of policies configured for your environment. The **CVE Search** page shows all the CVEs found in your onboarded environments.

Use Case

CVE Search allows you to find all workload assets affected by a CVE. The search results show all the occurrences of the searched CVE ID in the affected assets, with one result for each occurrence.

Note - If an asset has multiple occurrences of the searched CVE ID, the search results show each of the occurrences.

How it Works

CloudGuard scanners (AWP and Image Assurance) find the CVEs when scanning your environments and assets, for example, images or Virtual Machine instances.

Supported assets:

- Azure Virtual Machines, FunctionApps
- AWS EC2 instances
- Kubernetes images, workloads
- Container Registry images
- ShiftLeft images

Searching

To search for a CVE:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Vulnerability Search**.
2. Enter the CVE ID in the search bar and click **Search**. CloudGuard shows all assets affected by the CVE and their information:
    - Affected asset (for example, a workload that runs the affected image), its type, name, and link to the page
    - Scanned asset (for example, an image scanned by the CloudGuard scanner), its type, and name
    - Package name and version

- Environment where the vulnerability is found
- Remediation for the vulnerability

3. To limit the search results by certain criteria, use the filter. For example, you can show only fixable vulnerabilities or only in a particular environment.
4. Click the CVE ID to see its details. CloudGuard shows this information in a sliding panel:
   - CVE severity (color and symbol)
   - Source
   - Description
   - Remediation (if applicable) - Remediation is the package version that you need to upgrade to.
   - Affected asset details
   - Containing package

To export the search results:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Vulnerability Search**.
2. Enter the CVE ID in the search bar and click **Search**. CloudGuard shows all assets affected by the CVE.
3. On the top bar, click **Export** and select:
   - **Export** to see basic CVE information
   - **Export Extended** to see full CVE information

   CloudGuard shows the results in CSV format.

Known Limitations

The exported results do not include the **Scan Source** parameter.

Image Scan Findings

CloudGuard creates Image Assurance findings for Kubernetes images based on the assigned policy.

The findings of the **ImageScan** category include events related to CVEs, packages, sensitive data, and malware. For them, you can receive an email (or other notification) that contains aggregated information about all these findings.

To do this, create a new Image Assurance policy in three steps:

1. Create a new ruleset.
2. Create a new notification.
3. Configure a new policy.

Step 1. Creating a Ruleset

1. Navigate to **Workload Protection > Vulnerabilities > Rulesets**.
2. Click **Add Ruleset**.
3. Enter a distinctive name, for example, *Aggregated-ImageScan*, and a description for the ruleset.
4. Click **Create**. CloudGuard creates and opens a new ruleset.
5. Click **New Rule** to open the rule editor page.
6. Click the **GSL** text box to open the GSL Rule Editor.
7. Write a rule that contains a criterion for sending a notification, for example, the number of critical- or high-severity findings.

   Best Practice - For this ruleset, you can create only one rule to trigger sending you **only one finding** that aggregates all information about the image vulnerabilities. This includes all vulnerability statistics and remediation actions that you can do to mitigate the risks. Make sure to use **ImageScan** as a rule target.
8. Click **Verify** to verify the rule.
9. Click **Done**.
10. Enter or update other available fields (Title, Description, etc.).
11. Click **Save** to save the new rule.

Examples

The default CloudGuard rulesets **Container Image Assurance** and **Container Image Assurance 1.0** contain an ImageScan rule that you can use as an example.

1. This rule triggers sending a notification when the scan results have at least one critical vulnerability.

```
ImageScan should not have totals.critical > 0
```

2. This (default) rule triggers sending a notification when the scan results have at least one critical vulnerability or no less than 10 high severity vulnerabilities.

```
ImageScan should not have (totals.critical > or totals.high
> 10)
```

3. This rule triggers a notification when the scan results of a specific image (repo:tag) have risk score higher than 8.

```
ImageScan where repo-url='quay.io/checkpoint' and repo-
tag='consec-imagescan-engine:2.21.0' should not have risk-
score > 8
```

Step 2. Creating a Notification

1. Navigate to **Settings** > **Configuration** > **Notifications**. menu. This shows a list of all your Notifications.
2. Configure a new notification as in *"How to Configure a Notification" on page 853*.
3. In the **Add Filter** area, select **ImageScan** as a **Category**. This enables sending notifications only for the ImageScan findings.

   If the **ImageScan** category is not available, see *"Limitations" on the next page*.

Step 3. Configuring a Policy

1. Navigate to **Workload Protection > Vulnerabilities > Policies**.
2. Click **Add Policy**.
3. Select **Environment Policy**.
4. Select the **Kubernetes** platform.
5. Click **Next**.
6. Select a cluster to apply the policy.
7. Click **Next**.
8. Select the ruleset that you created in Step 1.
9. Enable **Admission Control** (**Image Admission**) in Detection or Prevention mode. For more details, see *"Image Admission" on page 486*.

10. Select the notification that you created in Step 2.
11. Click **Save**.

Viewing ImageScan Findings

1. Navigate to **Workload Protection > Vulnerabilities > Findings**.
2. Filter the view by **Category** : **ImageScan**. Click an applicable finding.

   If the **ImageScan** category is not available, see *"Limitations" below*.
3. The finding overview contains information about the image risk score, statistics of findings by severity, and the aggregated remediation (click **Show more** to see it). Full information is available in JSON format through the configured notification.

   Note - In the finding notification, the description and remediation appear separately for the rule and for the ImageScan entity. In the ImageScan entity, CloudGuard generates the finding information:
   - For remediation, it creates an aggregated remediation from all applicable vulnerabilities.
   - For description, it aggregates data with all statistics of the image vulnerabilities.

   In Jira notifications only, the generated data concatenate with the values of **Description** and **Remediation** configured in the rule.
4. The finding description is available as a tooltip when you put the cursor on the finding row and the **Description** column in the findings table.

Limitations

- Sometimes, the ImageScan category is not available in the filter when you create a notification. This happens with newly onboarded environments where CloudGuard has not finished yet to scan images for the first time. Wait approximately 5-10 minutes to let it finish and try again.
- The remediation length is limited to 25,600 symbols. The remediation that exceeds this length is truncated to 25,600 symbols.

More Links

- *"Image Assurance" on page 434*
- *"Vulnerability Policies (Image Assurance)" on page 444*
- *"Rules and Rulesets" on page 309*
- *"Governance Specification Language (GSL)" on page 326*

Image Assurance Troubleshooting

Verify the Agent Installation Status

Installation of the Inventory agent is a basic requirement to run Image Assurance.

- On the Image Assurance page, you can see the status for all nodes, while practically only the engine pod connects with CloudGuard
- Installed pods:
  - An Asset management agent (Inventory deployment) - A basic requirement for Image Assurance.
  - A minimum of one *imagescan-engine* pod.
  - For CloudGuard Helm Chart versions from 2.13.0, an *imagescan-list* pod.
  - Daemonset of *imagescan-daemon* pods - Based on the number of nodes. In addition, it runs on available control plane nodes.

    Example of a cluster with 2 nodes:

```
~$ kubectl -n checkpoint get pods
NAME                                              READY
STATUS    RESTARTS   AGE
asset-mgmt-imagescan-daemon-blzlj                 2/2
Running   0          5m
asset-mgmt-imagescan-daemon-kmftj                 2/2
Running   0          5m
asset-mgmt-imagescan-engine-767fdb686f-rqmlc      1/1
Running   0          5m
asset-mgmt-imagescan-list-5b48574cbd-q7g4p        1/1
Running   0          5m
asset-mgmt-inventory-agent-79764cc64f-dhrpk       1/1
Running   0          5m
```

- Get pod - Make sure that all the containers are ready:
  - *imagescan-engine* (central) - 1 pod - 1/1 ready
  - *imagescan-list* - 1 pod - 1/1 ready
  - *imagescan-daemon* - X pods (where X is the number of nodes) - 2/2 ready for each pod; for Docker container runtime, it is 1/1
- Get logs - Make sure that you do not see errors:

- *inventory-agent* pod - Make sure that there are no errors in logs and that all the entities uploaded successfully.
- *imagescan-engine* (central) - Make sure that there are no errors, and that scans are completed successfully.
- *imagescan-list* - Make sure that there are no errors and that image lists are uploaded successfully.
- *imagescan-daemon* - Make sure that there are no errors and that the imagescan-list and imagescan-engine pods can get image list and image content from these agents.

  Example:

```
# kubectl -n checkpoint logs asset-mgmt-imagescan-engine-
6f7bf8787b-tkw77 engine | grep -i error
```

- Basic use - When Image Assurance is deployed, user interaction is not necessary (only for advanced configuration settings). It scans running images first, from the smallest to the largest.

Central Agent Environment Variables

Name	Workloads	Default Value	Max Value	Comments
LOG_LEVEL	imagescan-engine imagescan-daemon imagescan-list	info	N/A	Possible values: debug, trace, warn, error The same for the *imagescan-daemon* agent
CP_ IMAGESCAN_ INTERNAL_ PROTO	imagescan-engine imagescan-daemon imagescan-list	HTTPS	N/A	If you set it as HTTP, agents use HTTP and not HTTPS for communication in the cluster.
CP_ IMAGESCAN_ SCAN_ TIMEOUT	imagescan-engine	N/A	24h	By default, the scan timeout is set in by the CloudGuard engine. It is possible to override this value. The timeout is in seconds.

Name	Workloads	Default Value	Max Value	Comments
IMAGE_ TRANSFER_ TIMEOUT_ SECONDS	imagescan-engine imagescan-daemon	N/A	24h	Kubernetes image is transferred between *imagescan-daemon* and *imagescan-engine* pods to be scanned by the *imagescan-engine*. This variable configures the timeout for the transfer operation. The timeout is in seconds. **Important** - Define the environment variable for two workloads and use the same value for each of them.

To configure the environment variables:

1. Edit the applicable *imagescan* workload and add one or more applicable variables with valid values. Use one of these methods:

- Edit a running deployment, for example:

```
kubectl edit deployment asset-mgmt-imagescan-image -n
checkpoint
```

- Set the environment variables with the Helm commands.

  Examples

  - To set `CP_IMAGESCAN_INTERNAL_PROTO` to *HTTP*, run:

```
--set-string addons.imageScan.daemon.env[0].name=CP_
IMAGESCAN_INTERNAL_PROTO,addons.imageScan.daemon.env
[0].value="HTTP" \
--set-string addons.imageScan.engine.env[0].name=CP_
IMAGESCAN_INTERNAL_PROTO,addons.imageScan.engine.env
[0].value="HTTP" \
--set-string addons.imageScan.list.env[0].name=CP_
IMAGESCAN_INTERNAL_PROTO,addons.imageScan.list.env
[0].value="HTTP"
```

  - To set `IMAGE_TRANSFER_TIMEOUT_SECONDS` to *1 hour* and `CP_IMAGESCAN_SCAN_TIMEOUT` to *2 hours*, run:

```
--set-string addons.imageScan.daemon.env
[0].name=IMAGE_TRANSFER_TIMEOUT_
SECONDS,addons.imageScan.daemon.env[0].value=3600 \
--set-string addons.imageScan.engine.env
[0].name=IMAGE_TRANSFER_TIMEOUT_
SECONDS,addons.imageScan.engine.env[0].value=3600 \
--set-string addons.imageScan.engine.env[1].name=CP_
IMAGESCAN_SCAN_TIMEOUT,addons.imageScan.engine.env
[1].value=7200
```

Note - When you set manually more than one environment variable, increase accordingly the index of the variable (0 and 1 in the above).

Istio

As Istio adds HTTPS proxies which break mutual TLS between ImageScan engine and ImageScan daemon agents. Change the protocol they use to connect to HTTP. It is done through environment variables passed to ImageScan Engine deployment and ImageScan daemon DaemonSet: CP_IMAGESCAN_INTERNAL_PROTO=HTTP. For this, append the below lines to the Helm installation or upgrade command (index 0 is used assuming no other environment variables are changed):

```
--set-string addons.imageScan.daemon.env[0].name=CP_IMAGESCAN_
INTERNAL_PROTO,addons.imageScan.daemon.env[0].value="HTTP" \
```

```
--set-string addons.imageScan.engine.env[0].name=CP_IMAGESCAN_
INTERNAL_PROTO,addons.imageScan.engine.env[0].value="HTTP" \
```

If the list deployment exists, add this line:

```
--set-string addons.imageScan.list.env[0].name=CP_IMAGESCAN_
INTERNAL_PROTO,addons.imageScan.list.env[0].value="HTTP" \
```

Low Rate of Image Scan

By default, there is only one image scan engine that scans images sequentially. You can increase the rate of image scanning. For this, deploy more image scan engines on your cluster.

- Add this parameter to the Helm command:

  ```
  --set addons.imageScan.engine.replicaCount=<number-of-
  scanners>
  ```

Example

To allow the scanning of three images in parallel and increase the scan rate, increase the number of engines to three:

- Add this line to the Helm install command:

  ```
  --set addons.imageScan.engine.replicaCount=3
  ```

Common Errors

1. When the engine and daemon pods cannot connect and the ImageScan engine reports a timeout error, it can be because of the cluster configuration.

   ```
   handleImageListTask returned error: get image list failed from
   agent at 172.17.0.3 (ubuntu2004): from https://consec1-
   imagescan-daemon:8443/imagelist: Get
   "https://172.17.0.3:8443/imagelist: dial tcp 172.17.0.3:8443:
   i/o timeout
   ```

2. When the engine and daemon pods cannot connect and the ImageScan engine reports a certificate validation error, it can occur one time (for each connection with the ImageScan daemon pod) after upgrading Helm. This occurs because pods are restarted one after the other, and some can use certificates from before for internal communication, and some possibly already use the new ones.

Example

```
get image failed from agent at 192.168.50.164 (ip-192-168-
29-125.ec2.internal): from https://asset-mgmt-imagescan-
daemon:8443/getimage/sha256:ee597f5bb5bc95c01d79a04ed053a388
d05836c96ae3aed117df5b2fea81f6aa: Get
"https://192.168.50.164:8443/getimage/sha256:ee597f5bb5bc95c
01d79a04ed053a388d05836c96ae3aed117df5b2fea81f6aa": x509:
certificate signed by unknown authority (possibly because of
"x509: invalid signature: parent certificate cannot sign
this kind of certificate" while trying to verify candidate
authority certificate "asset-mgmt-imagescan-daemon"
```

3. Sometimes the Image Assurance Agent status shows the error message: "No images found, please set containerRuntime to be docker", although the cluster worked before. It is possible that your cluster has recently changed its containerRuntime.

   Solution: Run the Helm upgrade command again to solve the issue.

Errors Troubleshooting in ImageScan Agent

The environment page of Container Registries and Kubernetes clusters shows information about their agents' status.

Error Messages in Agent Status

The agent status can show these error messages:

1. **Image scan fails. Last scan date is <date> with scan status code: <error>**

   This error message appears when all image scans fail (no successful image scans).

   The handling of this error depends on the specific error that is shown in the agent status description. See *"Image Scan Status" on page 460*.

2. **Unable to access the registry, last successful communication with the registry is at: <date>**

   This error message appears if the ImageScan agent cannot get access to the registry.

   Do these steps:

   a. Make sure that your cluster has outbound network connectivity. Make sure that no proxy/firewall blocks traffic.

b. If your authentication method requires to create a secret, make sure that you did it on the cluster correctly.

   The command to create a secret:

   ```
   kubectl create secret docker-registry <secret_name> \
   --namespace <namespace> \
   --docker-server=<registry_uri> \
   --docker-username=<key> \
   --docker-password=<password>
   ```

   Make sure that:

   - the values in `registry_uri` and `secret_name` match the values in Registry URI and Pull Secret Name in the CloudGuard portal.
   - the `namespace` is the same as configured during the cluster onboarding (by default, *checkpoint*).

   Make sure that your `key` and `password` are:

   - defined for your registry
   - have the correct permissions to get access to the registry
   - not expired

c. If your authentication method does not require a secret (GKE/AWS/Azure internal authentication):

   - Make sure that your cluster has the correct permissions to get access to the registry.
   - Make sure that you assigned the needed roles/permissions for the cluster to get access to your registry. Make sure to follow the steps related to the specific authentication method for your specific registry.
   - For more information, see the documentation of registries - *"Onboarding Container Registries" on page 204*.

d. If none of the steps worked for you, run this command to examine the agent logs for errors:

   - `kubectl -n <namespace> logs deployments/<release-name>-imagescan-list`
   - `kubectl -n <namespace> logs deployments/<release-name>-imagescan-engine`

3. **Failed to get image list from the registry, last image list is received at: <date>**

This error message appears when CloudGuard does not receive the list of registry images in the expected time.

By default, the configuration is to send the registry image list every 12 hours. You can change it with **Advanced Configurations** option of **Scan period in hours** on the registry **Scanners** page.

To solve the issue, do the procedure in Step 2 above.

Note - It can take five minutes at most after solving the problem to see the change in the agent status.

4. **Agent is disconnected**

   This error message appears when the agent has no successful connection with CloudGuard for more than one hour.

   Do these steps:

   a. Make sure that your cluster has outbound network connectivity. Make sure that no proxy/firewall blocks traffic.

   b. Make sure that your CloudGuard credentials on the cluster, API key and API secret, are correct.

      If you use Helm to install CloudGuard resources on your cluster, run this command to see the installation command arguments:

      ```
      helm -n <namespace> get values <release-name>
      ```

      Verify that `credentials.user` and `credentials.secret` are correct.

   c. If none of the steps worked for you, run this command to verify the agent logs for errors:

      - `kubectl -n <namespace> logs deployments/<release-name>-imagescan-list`
      - `kubectl -n <namespace> logs deployments/<release-name>-imagescan-engine`

Note - In the above examples, the default values for `namespace` and `release-name` are *checkpoint* and *asset-mgmt*, respectively.

Image Scan Status

The Scan Status of an image is shown on the Images page (**Workload Protection** > **Container Assets** > **Images**) and on the image details page.

See the table below for all statuses.

Scan Status	Description	Corrective Action
Scanned	The image is successfully scanned.	
Pending Scan	▪ The image awaits to be scheduled for a scan. Applicable to Fargate images: ▪ No matching image scans are found for the Fargate image.	
Partial	Scan results are partial; the image will be scheduled for rescanning.	
Unsupported OS	The image operating system is not supported (for example, Windows is not supported).	
Unmatched	Applicable for ECS images: No matching image scans were found for the ECS task image.	
Not an image	An artifact found in the registry is not an image (for example, Helm chart).	
Network Error	Unable to create a connection to scanning services, possibly because of a firewall or a proxy.	Verify your firewall/proxy configuration to make sure it does not block access to the required CloudGuard URLs. See the **Connectivity Requirements** section in *"Kubernetes Containers" on page 415*.

Scan Status	Description	Corrective Action
Unauthorized	Failed on one of these: 1. Failed to authenticate with CloudGuard. 2. Failed to authenticate with the container registry (for example, because of expired credentials). 3. Failed to verify CloudGuard certificate, possibly because of the firewall/proxy.	Verify your firewall/proxy configuration to make sure it does not block access to the required CloudGuard URLs. See the **Connectivity Requirements** section in *"Kubernetes Containers" on page 415*. If the image is from a container registry environment, follow the procedure for Error **2** of *"Error Messages in Agent Status" on page 458*.
Insufficient resources	▪ The image is too large to be scanned. or ▪ No space left on your host machine.	The maximum allowed image size is 20 GB. If you need to scan larger images, contact *Check Point Support Center*. If the image size is less than 20 GB, examine the space left on your cluster machine.
Timeout	Timeout on pulling the image to be scanned.	Examine your network connectivity on the cluster and try to increase the image pull timeouts by setting the environment variables. See the **Central Agent Environment Variables** section in *"Image Assurance Troubleshooting" on page 452*.

Scan Status	Description	Corrective Action
Internal Error	An unknown error has occurred. The image will be rescheduled for a scan.	■ Review the imagescan-engine logs and identify the engine container reporting errors and the node running it. ■ Examine the container metrics. If it reaches memory limits, increase the limits. If the node's memory utilization is high, increase the number of memory requests of the container. ■ Examine the free disk space of the node. For ECS scanning environments, examine the ephemeral storage of the task (the default is 20 GB). If the problem continues, contact *Check Point Support Center*.

More Links

- *"Container Registry Scanning" on page 464*
- *"Image Assurance Troubleshooting" on page 452*
- *"Asset Details" on page 273*
- *"Images" on page 425*

Container Registry Scanning

With Image Assurance, CloudGuard can scan container images on these private registries:

- Azure Container Registry (ACR)
- AWS Elastic Container Registry (ECR)
- Docker Hub Container Registry
- Google Cloud Container Registry (GCR)
- Google Artifact Registry (GAR)
- Harbor Registry
- JFrog Artifactory
- Nexus
- GitHub Container Registry
- Quay.io Container Registry

Note - GAR repositories can store helm charts in image format together with the actual docker images. If your repositories include helm charts in addition to images, CloudGuard shows them with the *Not an image* scan status.

To onboard your container registry to CloudGuard, see *"Onboarding Container Registries" on page 204*. These are two options to scan your Container Registry in CloudGuard:

- Link it to a Kubernetes cluster that has the ImageScan agents scanning your registry
- Deploy ImageScan with an AWS ECS scanner (available for selected types of registry)

The Image Assurance agents deployed on a cluster scan new images as they appear, on this cluster and on a linked ACR, ECR, or GCR container registry.

Note - Registry scanning requires Image Assurance agent version 2.10.0 or higher included in the CloudGuard Helm chart version 2.11.1 or higher. See *"Upgrading the Agent" on page 196* for more information.

AWS ECS Image Assurance

To launch containers, Amazon ECS uses Docker images in task definitions. The Docker images are commonly hosted in AWS ECR registries.

CloudGuard provides scanning results for the AWS ECS Docker images based on the inventory information of the onboarded AWS environment and ECR scanning. Installation of CloudGuard agents in the AWS ECS clusters is not necessary.

Prerequisites

Before you start, make sure to:

- Onboard the AWS environment to CloudGuard with the relevant ECS clusters
- Configure the ECR Container Registry Scanning for the ECR registry, that is, onboard to CloudGuard the ECR registry that hosts the Docker images of AWS ECS containers. For more details, see *"Onboarding AWS Elastic Container Registry" on page 214*.
- Enable the AWS ECS images scanning for the AWS environment with this API call: `https://api.dome9.com/v2/ecs/configuration/{cloudAccountId}`

Known Limitations

- By default, CloudGuard adds to Protected Assets and scans only 10 recent images of each repository. You can change the default value with the API call (maximal number is 1000 for a JFrog Artifactory and Sonatype Nexus). For more information, see the API Reference Guide.
- Scanning Windows container images is not supported.
- For JFrog Artifactory, it can take about 20 minutes that the images start to show for the first time.
- For JFrog Artifactory and Sonatype Nexus, the maximal number of tags per repository is 1000. Container images from the repositories with more than 1000 tags are neither shown as protected assets, nor scanned. The number is limited due to extensive API calls and performance considerations.
- To receive scanning results, ECS images must be onboarded in the same CloudGuard account as the Private ECR that scans them.
- CloudGuard creates ECS images only for running tasks.

Actions

Showing Scanning Results

To see the results of scanning, in CloudGuard, navigate to **Workload Protection** > **Containers Assets** > **Images**. For more details, see *"Images" on page 425*.

Removing a Container Registry from the Scan List

A container registry must have a Kubernetes cluster linked to it to scan the images.

1. Navigate to **Workload Protection** > **Container Assets** > **Environments** and select your container registry.
2. On the Scanners page, find the scanning environment (cluster).
3. In the Actions column, click **Unlink**.
4. In the confirmation message, click **Save**. CloudGuard informs you that the container registry is unlinked successfully.

The cluster stops to scan the registry, and the registry does not appear anymore in the cluster's list.

Configuring Scanning of Registries

CloudGuard scans only the latest (most recently used) images from the repository. You can configure a maximal number of images that it fetches for each regular or base repository.

1. Navigate to **Workload Protection** > **Container Assets** > **Environments** and select your container registry.
2. On the Scanners page, click **Advanced Configurations**.
3. For **Max images per repository**, enter the number of images that CloudGuard fetches from a repository (by default, **10**).
4. For **Max images per base repository**, enter the number of images that CloudGuard fetches from a base repository (by default, **50**).
5. Click **Confirm**.

More Links

- *"Onboarding Container Registries" on page 204*
- *"Kubernetes Containers" on page 415*
- *"Images" on page 425*
- *"Image Assurance" on page 434*
- API Reference Guide

Scan Engine V2

The Scan Engine is a module that you can run in the command-line shell or in your Continuous Integration / Continuous Deployment (CI/CD) pipeline. This module scans container images for security risks and vulnerabilities and assesses the image compliance with the organization's security policy defined in CloudGuard.

The Scan Engine exists in two versions. This section describes working with version 2.0.0.

To see the Scan Engine version enabled on your CloudGuard account:

1. Navigate to **Settings** > **Configuration** > **Workloads**.
2. On the **Image Assurance** heading, see **Scan Engine Version**. Version 2.0.0 and higher means that your account has V2 support.

Requirements

CloudGuard Account

You need a CloudGuard account to use Scan Engine version 2.0.0. If you do not have one, go to the Infinity Portal to create an account.

Connectivity

The Scan Engine must have connectivity to the `*.dome9.com` domain to properly communicate with the CloudGuard portal.

Instead of the domain address, you can use the region-specific URLs from the table below. Add these endpoints to the allow list.

Region	Region Code	Address
United States (US)	US	https://api.dome9.com
Europe (EU)	EU1	https://api.eu1.dome9.com
Singapore (SG) Dome9 accounts only	AP1	https://api.ap1.dome9.com
Australia (AU)	AP2	https://api.ap2.dome9.com
India (IN)	AP3	https://api.ap3.dome9.com
Canada (CA)	CA	https://api.ca.dome9.com

Workflow

1. Configure an Image Assurance policy for CI/CD pipeline:
    a. Create a ShiftLeft environment. See *"Create a ShiftLeft Environment and Service Account" on page 469*.
    b. Create a new CloudGuard service account or use the credentials of an existing one.
    c. Download and install the Scan Engine. See *"Download the Scan Engine" on page 469*.
    d. Create an Image Assurance policy. See *"Create an Image Assurance Policy" on page 470*.
2. Run the Scan Engine. See *"Running the Scan Engine" on page 473*.
3. View the assessment results in the CloudGuard portal or on the CLI terminal. See *"Viewing Scan Results" on page 475*.

Configuring a Policy for Scan Engine

Create a ShiftLeft Environment and Service Account

1. In the CloudGuard portal, navigate to **Assets** > **Environments** (or **Workload Protection** > **Containers Assets** > **Environments**), click **Add** and select **ShiftLeft Environment**.
2. In the **ShiftLeft Onboarding** wizard that opens, enter these details:
   a. Environment Name
   b. Environment Description (optional)
   c. Organizational Unit (optional)
3. Configure a Service Account by one of these methods:
   - Select an existing Service Account with its corresponding API Key.
   - Enter a Service Account manually.
   - Click **Add Service Account** to create a new account.
4. Click **Next**.

Download the Scan Engine

1. Select one of the operating systems to run the program on:
   - Windows
   - Linux
   - macOS
2. Download the correct version of the program based on your host architecture.
3. Click **Next**.
4. Copy the commands from the wizard and paste them into the terminal. Run the commands to set up the downloaded executable file. For more information, see *"Scan Engine Installation" on page 471*

   Note - If you decide to store your data on a non-US CloudGuard data center, then you must set the environment variable `SHIFTLEFT_REGION` value based on your data center. For region codes, see *"Connectivity" on page 467*.
5. The summary page informs that the environment onboarding is complete. Click **Finish**,

The new environment page opens. It contains the name, the description (if provided), and the CloudGuard ID that CloudGuard assigns to the environment. You must use this ID later when you initiate assessment with the Scan Engine.

The instructions for downloading and setting up the Scan Engine are available on the ShiftLeft environment page. Open the page and click **Read Instructions**.

Create an Image Assurance Policy

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Policies**.
2. Click **Add Policy**.
3. In the **Create New Policy** window, select **ShiftLeft** for **Platform** and click **Next**.
4. For **Environments**, select the environment that you created and click **Next**.
5. For **Rulesets**, select one or more Image Assurance rulesets and click **Next**.
6. For **Notifications**, select one or more notifications, or click **Add Notification** to create a new notification.

   Note - No notifications are required for ShiftLeft, so you can use the default CloudGuard notification.
7. Click **Save**.

The new policy is ready, and you can start *"Scan Engine Installation" on page 471*.

More Links

- *"Assets" on page 138*
- *"Vulnerability Policies (Image Assurance)" on page 444*
- *"Configuring CloudGuard Policies" on page 78*
- *"Rules and Rulesets" on page 309*
- *"Notifications" on page 852*

Scan Engine Installation

The Scan Engine installation depends on your operating system and exists in three options:

- Windows
- Linux
- macOS

You can use one of these methods to install the Scan Engine:

- Download and set up the binary file.
- Use a Docker image that contains the binary file.

Downloading the binary file

Installing the Scan Engine on Windows

1. Download the x64 or 386 standalone binary.
2. (Optional) Save the `shiftleft.exe` file in a directory in your current PATH.
3. Launch a new command terminal and verify that ShiftLeft is properly installed:

   ```
   C:\Downloads\>shiftleft image-scan --version
   ```

   The sample output:

   ```
   0.400.0
   ```

Installing the Scan Engine on Linux

1. Download the x64 or 386 standalone binary.
2. (Optional) Make the file executable and move the file into a directory in your current PATH, for example:

   ```
   chmod +x shiftleft
   mv shiftleft /bin/shiftleft
   ```

3. Launch a new command terminal and verify that the program is properly installed:

   ```
   $shiftleft image-scan --version
   ```

   The sample output:

   ```
   0.400.0
   ```

Installing the Scan Engine on macOS

1. Download the x64 standalone binary.
2. Make the file executable, add a Gatekeeper exception, and move the file into a directory in your current PATH, for example:

```
chmod +x shiftleft
spctl --add shiftleft
sudo mv shiftleft
/usr/local/bin/shiftleft
```

3. Launch a new command terminal and verify that the program is properly installed:

   ```
   user-mbp:~ user$ shiftleft image-scan --version
   ```

   The sample output:

   ```
   0.400.0
   ```

Using a Docker image

Using a Docker image of ShiftLeft provides an easy and flexible way to run scans in containerized environments. The Docker image is available on DockerHub at `checkpoint/shiftleft`.

To use the ShiftLeft image:

1. Run the command:

```
docker run -e CHKP_CLOUDGUARD_ID=<CloudGuard ID> -e CHKP_
CLOUDGUARD_SECRET=<CloudGuard ID Secret> -e SHIFTLEFT_
REGION=<Region> -e ShiftLeftEnvID=<ShiftLeft environment id> -
v <path to image>:/tmp/images/image.tar
checkpoint/shiftleft:latest_v2
```

2. Replace `<Cloudguard ID>`, `<CloudGuard ID Secret>`, `<REGION>`, `<ShiftLeft ENVIRONMENT ID>`, and `<PATH_TO_IMAGE>` with your CloudGuard API key, secret, region, your environment ID, and the local path to your image directory, respectively.

   Make sure to use the `latest_v2` tag to get the most recent features and improvements.

Scan Engine Update

Each time the Scan Engine runs, it automatically checks for updates and installs the latest version available.

Running the Scan Engine

After you configure the Image Assurance policy and install the Scan Engine, you can run it on your CI/CD pipeline.

The Scan Engine scans container images for security risks and vulnerabilities.

Usage

```
shiftleft [-t timeout] image-scan [OPTIONS] -e <ENVIRONMENT_ID> -i
<IMAGE_PATH>
```

Image-Scan Arguments

Argument	Description
-e, `--environment <string>`	CloudGuard ShiftLeft environment ID
-h, `--help`	Show help
-i, `--image <string>`	Path to docker image TAR file
-j, `--json`	JSON output
-o, `--output <string>`	Full CLI output to the provided file path
-t, `--timeout <int>`	Scan timeout in seconds (default: 3600) **Note**: Make sure to use the `-t` flag **before** the `image-scan` expression.
-v, `--version`	Show version

Exit Codes

The exit code of the command is non-zero in case of a policy violation or an error.

Exit Code	Description
0	Image scan succeeded, empty assessment / Assessment passed, no rules failed. Image is compliant.
1	Network error
3	Authentication error
4	Missing arguments
5	Internal error

Exit Code	Description
6	Image scan succeeded, assessment is not empty / Assessment failed at least on one rule. Image is not compliant.
11	Error in getting the assessment result. Try again or contact support.
99	Insufficient memory for image scanning
101	Insufficient disk space for image scanning
253	Scan timeout

Examples

Scanning a Container Image TAR File

To scan a container image `myrepo/myimage:version` for vulnerabilities and security threats, run:

```
docker save -o /home/my_container.tar myrepo/myimage:version
shiftleft image-scan -i /home/mycontainer.tar -e <ENVIRONMENT_
ID>
```

The Scan Engine scans the container TAR file at the provided path `/home/mycontainer.tar`. For the assessment, the policy uses all CloudGuard rulesets attached to the provided environment. The assessment results are published as findings associated with the applicable ShiftLeft image entity in the ShiftLeft environment. Use the Scan Engine flags to print the results as text or JSON.

Scan Timeout

When you scan a large image, you can receive a timeout error from the program.

To solve the issue, increase the default one-hour timeout with the `-t` flag.

For example, if the scan in the above example fails after a one-hour timeout, run this command to start another scan with a two-hour timeout:

```
shiftleft -t 7200 image-scan -i /home/mycontainer.tar -e
<ENVIRONMENT_ID>
```

Viewing Scan Results

Scan results are available as the CLI output in your terminal or in the CloudGuard portal.

CLI Output

The Scan Engine prints the scan results to the CLI. This output can be displayed in a table view or as a JSON file. The results are based on the assessments that consider vulnerabilities found in the scanned image against the attached policy.

The JSON output has the same format as the results of the assessment API. For more information about API, see the API Reference Guide.

When the scan results include more than 2 CVEs, they show the CVEs in short form: two CVEs with the highest severity and the number of the remaining CVEs (for example, CVE-123, CVE-456 and 10 others).

To see the full output:

- use the `-o` flag to print the full output to a file
- use the `-j` flag to see the output in JSON

Output in CloudGuard

The scan results shown in the CloudGuard portal include both the particular vulnerabilities found in the image and their assessment against the attached policy.

Upon scan completion, the Scan Engine sends the results to the associated environment in CloudGuard.

To see the scan results, do one of these:

- Navigate to **Workload Protection** > **Containers Assets** > **Images**. This page shows images scanned in CI/CD pipeline.
- Navigate to **Events** > **Posture Findings** and filter the view for the **ShiftLeft** Platform and **ShiftLeft Image** Entity Type. This page shows posture findings in your images. To learn more about posture findings, see *"All Events" on page 120*.
- Click a ShiftLeft image to open its page. Go to the **Vulnerabilities** tab to see vulnerable entities in the image.

  To learn more about vulnerabilities, see *"Agentless Workload Posture" on page 489*.

Admission Control

CloudGuard Admission Control monitors your clusters and enforces a security baseline across a namespace or cluster. It can detect if your clusters do not comply with standard practices like having good labels, annotations, resource limits, or other settings. If you configure the Admission Control in the Prevention mode, not only does it detect a breach, but it stops the unwanted action.

CloudGuard Workload Protection - Kubernetes Admission Control

Before you can use Admission Control, your Kubernetes cluster must be onboarded to CloudGuard. See *"Onboarding Kubernetes Clusters" on page 188*.

How it Works

When you make a change in an existing workload, Kubernetes must update the workload configuration in etcd. The API request to the Kubernetes API server goes through the admission stages as in the diagram below before it is recorded (persisted) in etcd. CloudGuard's Admission Control utilizes the validating admission webhook to enforce your policies created in the CloudGuard portal or through the CloudGuard API.

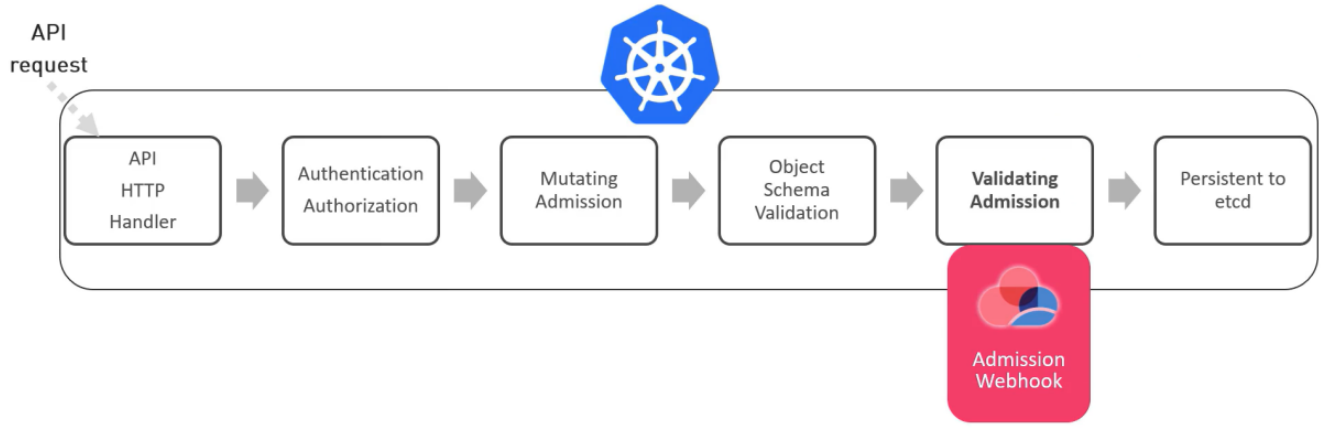

When you enable Admission Control, CloudGuard installs or upgrades Kubernetes agents with these resources:

- **Policy Agent** - One replica Deployment responsible for retrieving configuration and policies (configured by the user) from the CloudGuard backend.
- **Enforcer Agent** - A Deployment with a replica count of two that establishes the admission webhook receives the API calls and validates them based on the enforced policy.
- **Custom Resource Definitions (CRD)** - Kubernetes custom resources (CR) that contain Detection and Prevention rules for each supported object type. The Policy Agent has permission to change the CRs. The Policy and Enforcer Agents have permissions to read the CRs.

Access to the CRDs (together with `read` permissions like `get` and `list`) must be restricted as they contain information on the cluster hardening.

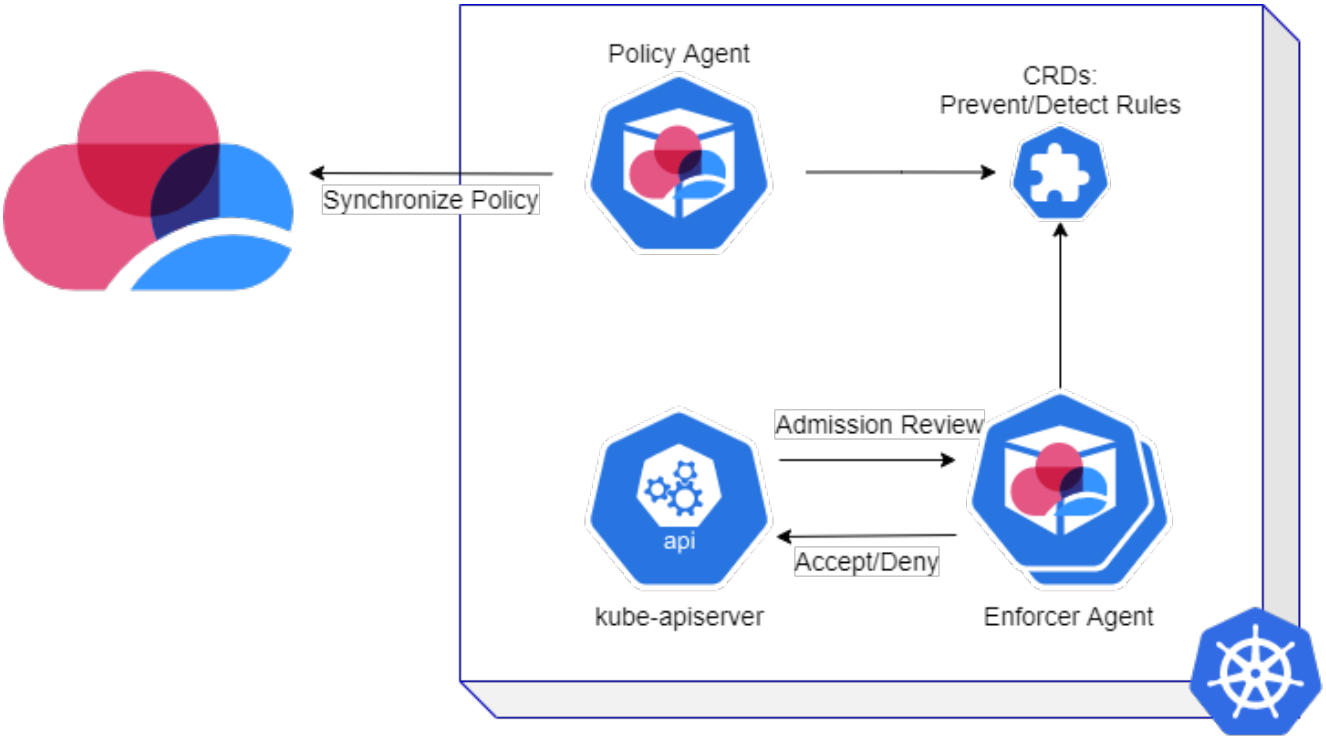

Kubernetes deploys and registers the Enforcer Agent as the Admission Control validation webhook on the supported resources.

The Enforcer Agent does this:

1. Stops every *create*, *update*, *delete*, or *connect* operation triggered on the supported resources by the API calls to the API server.
2. Validates these changes against the Admission Control rules.

If an API call violates the rule, CloudGuard sends an alert (go to **Events** (*"All Events" on page 120*) and filter by Source > **Kubernetes Admission Control**) based on the notification in the policy. Based on your selection of the mode, CloudGuard:

- Blocks the operation for policies in the **Prevention Mode**.
- Allows the operation of policies in the **Detection Mode**.

Admission Control blocks or detects violations on API calls only; it does not delete or block running workloads.

Example

Add this rule:

```
kubernetesPod should not have labels with [key='app' and
value='my_app']
```

Suppose that you already have a workload running with the label app *my_app*. This workload is not blocked, and no alert is triggered. You can use Posture Management (see *"Cloud Security Posture Management (CSPM)" on page 302*) to detect violations on the running state of the cluster.

When there is an attempt to create a workload with the label app *my_app* or to add such a label to a workload, Admission Control detects (or blocks) the operation.

Note - Admission Control monitors (blocks or detects) all operations initiated by users. It ignores system components operated by the Kubernetes control plane.

Supported Resources

CloudGuard Admission Control allows you to enforce your policy on these resources:

- Pods
- Deployments
- ReplicationControllers
- ReplicaSets
- DaemonSets
- StatefulSets
- Jobs
- CronJobs
- Roles
- ClusterRoles
- RoleBindings
- ClusterRoleBindings
- Services
- Ingresses
- ServiceAccounts
- NameSpaces
- ConfigMaps
- NetworkPolicies
- PodSecurityPolicies

Alert Recurrence

CloudGuard Admission Control issues only one alert for each unique incident.

An issue is unique when it meets all these conditions:

- All occurrences violate the same rule, in the same ruleset and policy.
- The same entity (Username or Service Account) relates to all occurrences.
- The target is the same in all occurrences, that is, it has the same entity or the same root owner of the entity, for example, the same deployment for different pods.

Alerts Severity

- Determined by the severity configured on the use case or rule.
- When the Kubernetes server-side dry-run option calls the API, Admission Control stops this event as all other events. If the applicable event is a Kubernetes dry-run, the severity level of the alert is **Informational**.

Kubernetes Definitions

Some Kubernetes resources have an owner resourced:

- **Owner resource** - The parent resource that created a resource.
- **Root Owner resource** - The first resource that led to the creation of a resource.

For example:

- A pod can be created directly or indirectly by a ReplicaSet; in the latter case, the ReplicaSet is the owner of the pod.
- A resource can have a chain of owner resources, for example, a deployment can create a ReplicaSet which can create pods.

Example

Deployment-X creates ReplicaSet-X, which creates two replicas: Pod-X1 and Pod-X2.

Deployment-X is the root owner of Pod-X1 and Pod-X2, and the root owner of ReplicaSet-X.

Deployment-X is the owner of ReplicaSet-X.

ReplicaSet-X is the owner of Pod-X1 and Pod-X2.

Deployment-X has no owner or root owner.

Admission Control Default Policy Configuration

Container Admission Control is a CloudGuard-Managed ruleset that contains the best practice rules for Kubernetes Admission Control. You can find this ruleset if you navigate to **Workload Protection** > **Admission Control Rulesets** and filter on the **CloudGuard-Managed Type**.

The default Admission Control policy uses this ruleset. CloudGuard attaches the default detection mode policy to all new and existing clusters, where you enable Admission Control. The default Admission Control ruleset constantly undergoes updates.

When you onboard a new cluster to CloudGuard (or enable the Admission Control feature) and associate it with an Organizational Unit, the cluster obtains the Admission Control policy configured for this Organizational Unit. If no such policy exists, a new policy is created to associate the new cluster with the default ruleset.

Best Practice - For Prevention policies, clone the default ruleset and use it in your policy.

To provide the security solution, it is sometimes necessary to give CloudGuard agents elevated permissions that must be restricted for most workloads. To address this requirement, the default policy has preconfigured exclusions to streamline the CloudGuard solution.

Configuring Admission Control in CloudGuard

Follow these steps to configure a GSL policy on the cluster:

1. Create an Admission Control Ruleset.
2. Add rules to the Ruleset.
3. Create an Admission Control Policy that binds the Ruleset to the cluster.

For more, see *"Getting Started with Admission Control Policy" on page 481*.

Exclusions

Admission Control exclusions are almost the same as the regular CloudGuard exclusions. For more details, see *"Parameters for Admission Control" on page 83*.

More Links

- *"Kubernetes Containers" on page 415*
- *"Image Assurance" on page 434*
- *"Configuring CloudGuard Exclusions" on page 80*
- Kubernetes documentation:
  - Guide to Kubernetes Admission Controllers
  - Using Admission Controllers

Getting Started with Admission Control Policy

Container Admission Control is a CloudGuard-Managed ruleset that contains the best practice rules for Kubernetes Admission Control. You can find this ruleset if you navigate to **Workload Protection** > **Admission Control Rulesets** and filter on the **CloudGuard-Managed Type**.

The default Admission Control policy uses this ruleset. When you onboard a new cluster to CloudGuard (or enable the Admission Control feature) and associate it with an Organizational Unit, the cluster obtains the Admission Control policy configured for this Organizational Unit. If no such policy exists, a new policy is created to associate the new cluster with the default ruleset.

To provide the security solution, CloudGuard agents sometimes need elevated permissions that must be restricted for most workloads. To address this requirement, the default policy has preconfigured exclusions to streamline the CloudGuard solution.

Configuring Admission Control in CloudGuard

Follow these steps to configure a GSL policy on the cluster:

1. Creating an Admission Control Ruleset.
2. Adding rules to the Ruleset.
3. Creating an Admission Control Policy that binds the Ruleset to the cluster.

Step 1. Creating an Admission Control Ruleset

1. Navigate to **Workload Protection** > **Admission Control Rulesets** and click **Add Ruleset**.
2. In the Create New Ruleset window, enter the Ruleset name and the description.
3. Click **Create**. The new Ruleset page opens.

Step 2. Adding Use Cases or Rules to the Ruleset

Best Practice - The **Use Case** option covers the most popular scenarios for the Admission Control Ruleset creation. Click the **New Rule** button only if you do not find an applicable Use Case.

1. In the newly created Ruleset page, click **New Use Case**. The **Edit Rule GSL** window opens with the default Use Case option selected.
2. Select the desired Use Case from the list. Some parameters appear configured automatically in part of the fields.

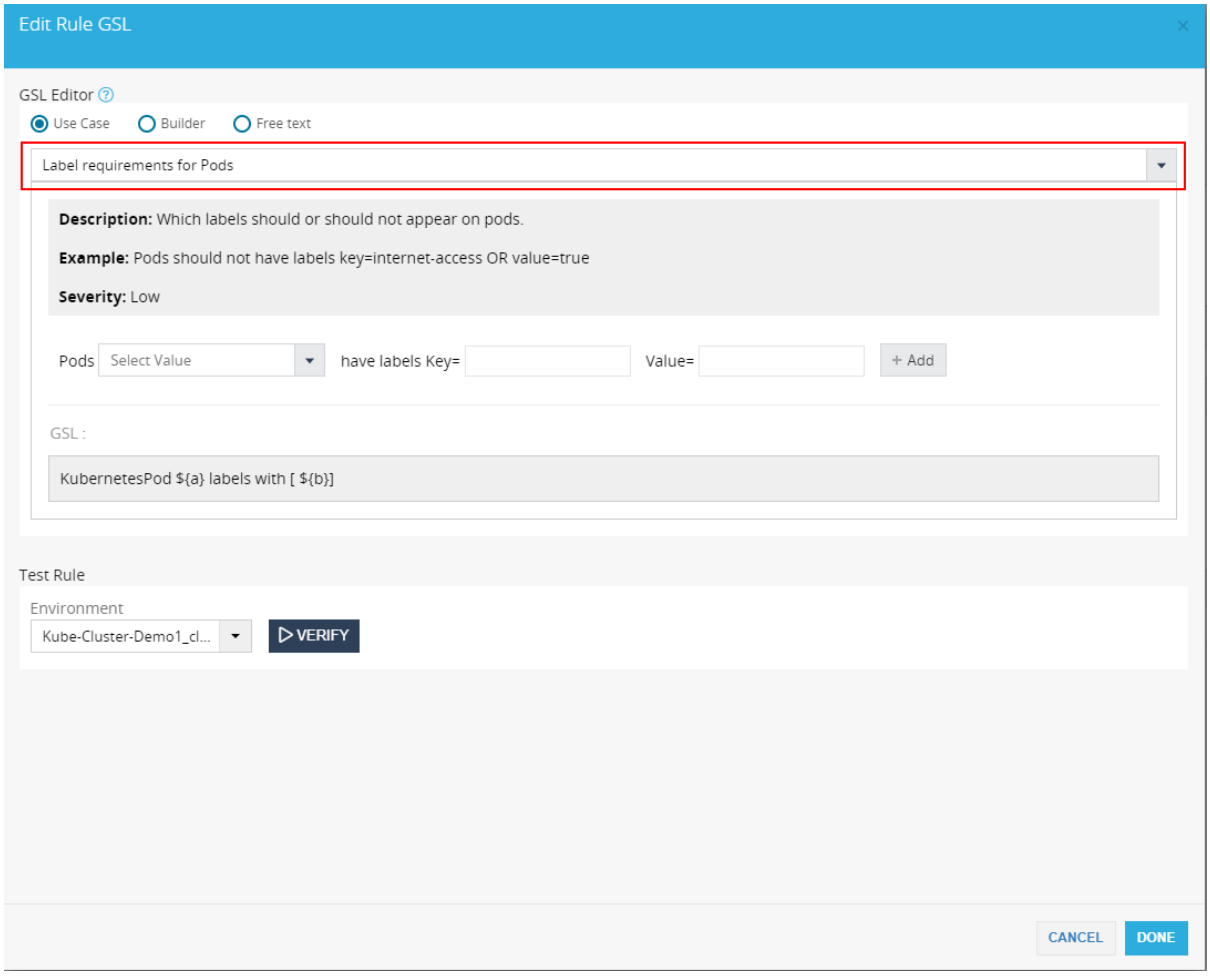

3. Enter missing information in the available fields to configure the GSL.
4. As an alternative, you can configure the GSL with the **Builder** or **Free text** options.

a. To use the GSL Builder, select **Builder**. When you select an object to configure the rule on, the Builder provides a context and helps you set up the rule.

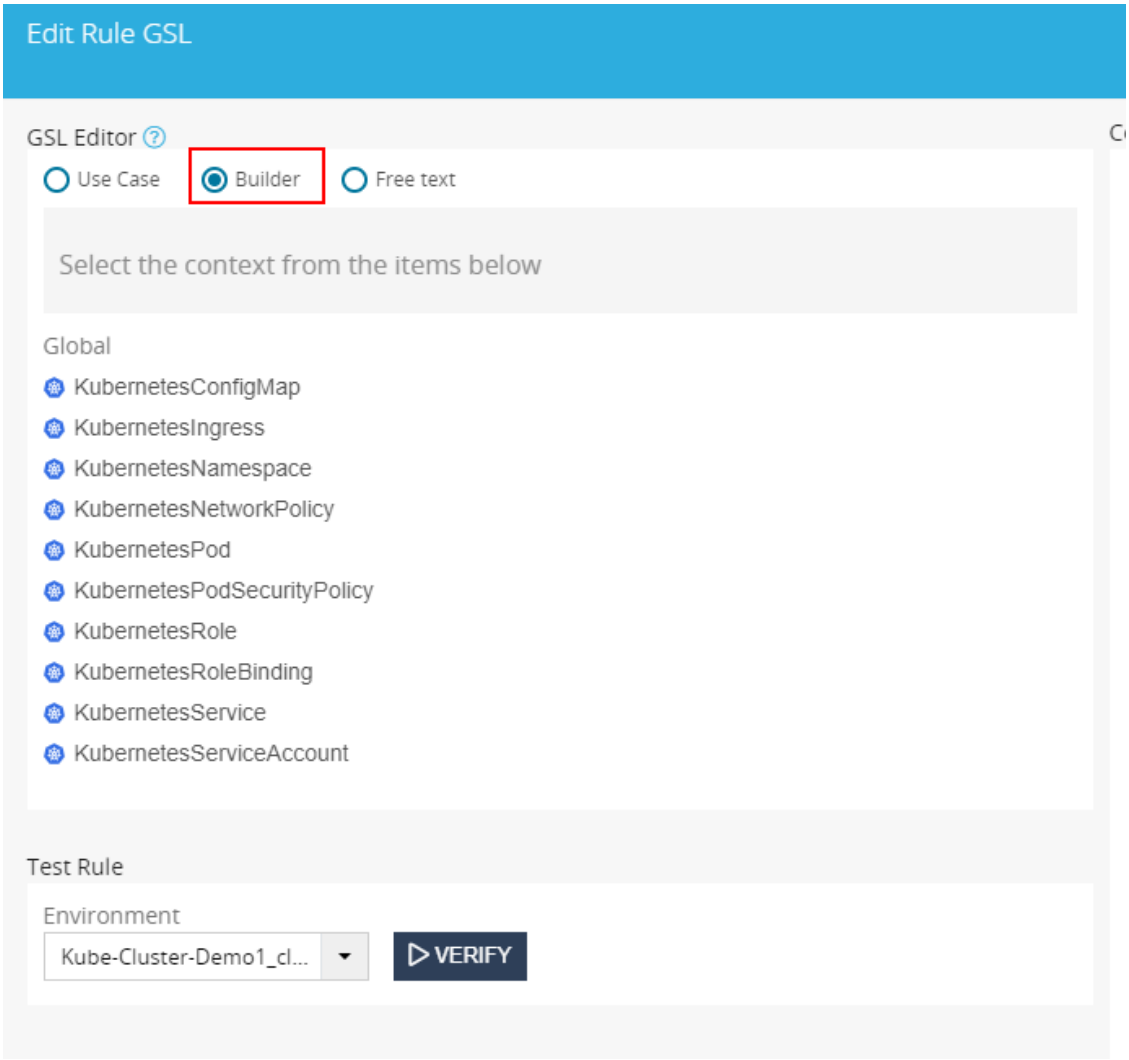

b. If it is necessary for you to configure the rule, select **Free text** and start to write it.

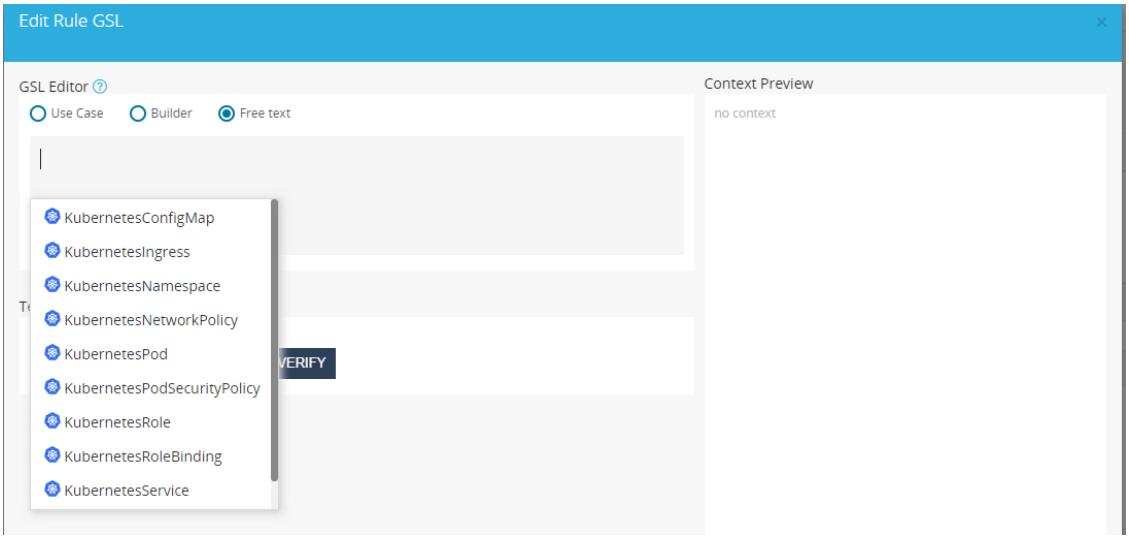

5. Below **Test Rule**, select an **Environment** from the list and click **Verify** to test the rule.
6. If the verification result is correct, click **Done**. The rule page opens.

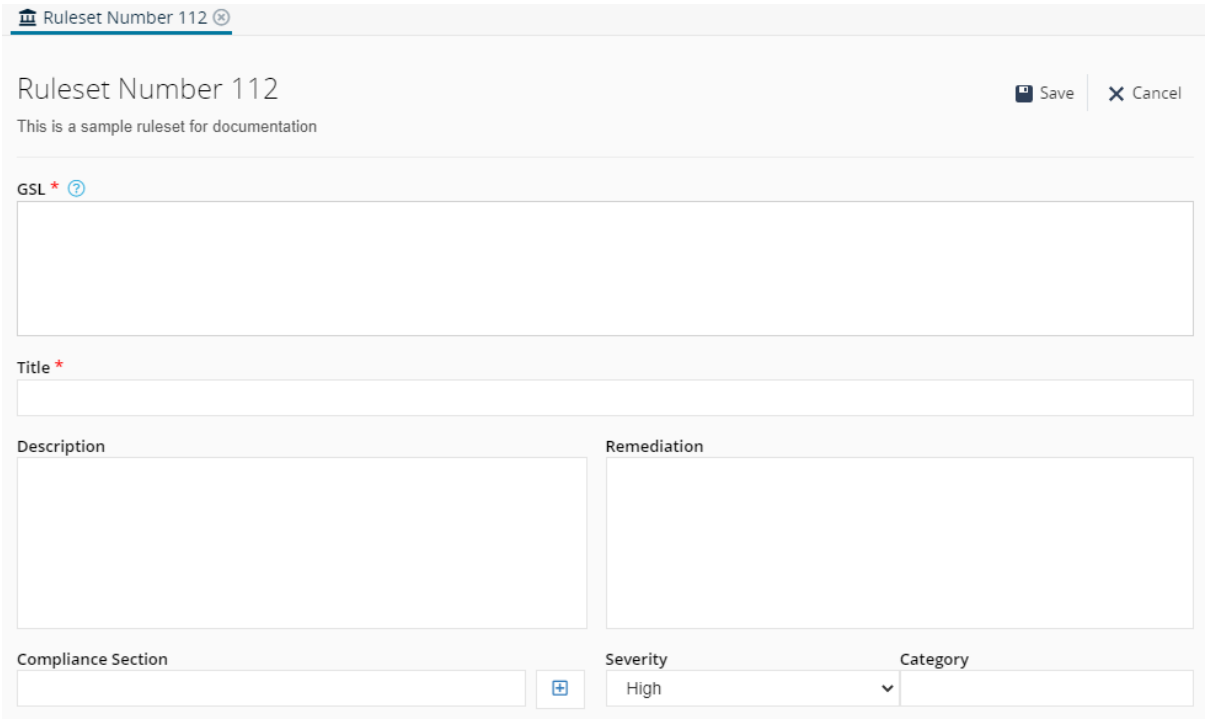

As an alternative, you can open this page when you click **New Rule** on the **Ruleset** page. When you click in the **GSL** field, the **Edit Rule GSL** window opens again.

7. Enter the **Title**, select the **Severity** and, optionally, enter **Description**, **Remediation**, **Compliance Section**, and **Category**.
8. Click **Save** to save the new rule.

Note - Rules on Kubernetes Pod are enforced on all workload resources.

Step 3. Creating an Admission Control Policy

The Admission Control Policy binds the Ruleset to a cluster and configures how to create notifications.

1. Navigate to **Workload Protection** > **Admission Control Policies**. This page shows all Admission Control Policies available for your assets.
2. Click **Add Policy** and select:
   - **Environment Policy** if it is necessary to apply this policy on a cluster
   - **Organizational Unit Policy** if it is necessary to apply it on the OI to which the cluster belongs
3. In the **Create New Policy** window, select one or more Environments or Organizational Units on which the policy applies. Click **Next**.
4. Select one or more Rulesets configured in Step 2 above.
5. Select one of two actions when the Rule(s) is violated:
   - **Detection Mode** - The policy does not block events on the cluster, but only sends an alert notification with severity configured in the violated rule.
   - **Prevention Mode** - The policy blocks an event that violates it on the cluster and sends an alert notification with severity configured in the violated rule.

   Note - You can configure multiple policies on the same cluster with different Action configurations.

   Best Practice - Before you use the **Prevention mode**, validate new policies in the **Detection mode**.
6. Click **Next**.
7. Select one or more Notifications to receive when the policy is violated. You can click **Add Notification** to configure a new notification.
8. Click **Save**. CloudGuard Admission Control protects your cluster.

More Links

- *"Admission Control" on page 476*
- *"Kubernetes Containers" on page 415*

Image Admission

With the Image Assurance policy, CloudGuard makes an assessment of scanned image compliance. While Image Assurance can only detect vulnerabilities in the image, Image Admission can prevent image deployment in a cluster. The image is allowed if it was scanned in one (minimum) of the related environments (ShiftLeft, Container Registry, or Kubernetes cluster) and found compliant.

For Image Admission, images scanned in a ShiftLeft or Container Registry environment are considered a primary source. If no previous image sources are available, .images scanned from a Kubernetes cluster environment are considered a secondary source.

Image Assurance Policy

The Image Assurance policy can have these two actions, disabled by default:

- Image Admission Action of **Detection** or **Prevention** to enforce non-compliant images based on the scan results.
- Image Admission UnScanned Action of **Detection** or **Prevention** to enforce not scanned images.

To receive the correct results of the image scanning, you must use for the Kubernetes environment the same ruleset that was used for the applicable ShiftLeft, Container Registry, or Kubernetes cluster environment.

Note - Image Enforcement is done based on the image name. If the image is scanned on a ShiftLeft environment, make sure that the image is correctly tagged before the scan.

Detect or Prevent Modes

Enforcement

When your container creates a new workload, all workload images are checked if they are compliant or scanned, based on the selected action. When it updates the workload, CloudGuard checks only changed or added images if they are compliant or scanned (for each selected action).

Prevention

When you enable a Prevention policy, the Image Admission Enforcer agent blocks the deployment of workloads with non-compliant or not scanned images. The Kubernetes user receives this error message in their CLI:

```
Error from server: error when creating "deployment.yaml":
admission webhook "cloudguard-enforcer-
webhook.cloudguard.checkpoint.com" denied the request:
[CloudGuard] The request has been blocked because the image
'myregistry.domain.com/my-ubuntu:v1' has not passed compliance
check.
```

Exclusions

Use exclusions to allow registries, specific images, usernames, roles, or namespaces. For more information, see *"Configuring CloudGuard Exclusions" on page 80*.

In exclusions, use "**%**" as a wildcard.

Wildcard Usage Examples

- Username - Write the exclusion as GSL, for example, **%@checkpoint.com**
- Roles - Write the exclusion as GSL, for example, **%myrole%**
- Groups - Write the exclusion as GSL, for example, **%mygroup%**
- Namespace - Write the exclusion as GSL, for example, **mynamespace%**
- Image
  - Registry - **myregistry.domain.com/%**
  - Tag - **myregistry.domain.com/my-ubuntu:%**
  - Specific image (full name) - **myregistry.domain.com/my-ubuntu:v1**

Actions

When you configure Image Admission, you can select one of these options:

- **Detect** to generate an event on the image deployment. No CLI output, no blocked API.
- **Prevent** to block the API call, send a CLI message, and generate an event on the image deployment.
- **Disable** to ignore the deployment.

Enabling Image Admission in the Image Assurance Policy

Based on previous assessments from CI/CD pipeline (ShiftLeft), Registry Scanning, or Kubernetes cluster, Image Admission can detect or block the deployment of the image in a cluster.

To configure Image Admission:

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Policies**.
2. Click **Add Policy** and start to add a policy for a Kubernetes environment as in *"Getting Started with Image Assurance Policy" on page 437*.
3. On the **Image Admission** page of the wizard, select the actions:
   a. For non-compliant images, select **Detect**, **Prevent**, or **Disable**.
   b. For not scanned images, select **Detect**, **Prevent**, or **Disable**.
4. Click **Next** to continue the policy configuration.

Changing the Admission Control Action

1. Navigate to **Workload Protection** > **Vulnerabilities** > **Policies**.
2. Select an existing policy and click **Edit** on the top bar.
3. On the **Image Admission** page of the wizard, change the actions:
   a. For non-compliant images, select **Detect**, **Prevent**, or **Disable**.
   b. For not scanned images, select **Detect**, **Prevent**, or **Disable**.
4. Optionally, you can edit the policy notification.

Creating an Image Assurance Policy with API

```
POST /v2/kubernetes/imageAssurance/policy
```

Use these properties:

- Image Admission Action - `admissionControllerAction`
- Image Admission UnScanned Action - `admissionControlUnScannedAction`

For more information, see the API Reference Guide.

More Links

- *"Kubernetes Containers" on page 415*
- *"Image Assurance" on page 434*
- *"Vulnerability Policies (Image Assurance)" on page 444*
- *"Admission Control" on page 476*
- *"Configuring CloudGuard Exclusions" on page 80*

Agentless Workload Posture

The Agentless Workload Posture (AWP) solution for VM instances and serverless functions (AWS EC2, Azure Virtual Machines, and Function Apps) provides continuous security assessment of your workloads without the need to install agents on each virtual machine. AWP continuously checks your assets for vulnerabilities to ensure the workloads meet your organization's security standards. AWP shows the vulnerabilities of each workload and suggests remediation.

Benefits

- Deep security visibility with seamless deployment
- Continuous scanning for vulnerabilities and secrets
- Automatic update of scanning tools and vulnerability databases
- Feed for CloudGuard Risk Management solution to identify and prioritize risks

Prerequisites

Your account (AWS account or Azure subscription) must be onboarded to CloudGuard before AWP can scan your virtual machines and serverless functions. If your account is not yet onboarded, see instructions in *"Onboarding Cloud Environments" on page 53*.

How AWP Works

AWP focuses on the file system of your workload.

AWP does not install an agent to scan the files on the machine. Instead, it makes snapshots of the virtual machine volumes or disks. AWP uses these snapshots to statically scan your packages, dependencies, and libraries on a dedicated AWP scanner machine. During scanning, AWP checks the VMs for known vulnerabilities (such as Log4j) and hardcoded secrets registered in security databases. The databases are updated daily according to the current security trends.

Azure Function Apps scanning is enabled by default when you select to scan Azure VMs with In-Account or Sub-Account mode. You can disable this option when you start onboarding. After initial scanning upon onboarding, AWP scans Function Apps when it detects changes.

After scanning is done, the CloudGuard portal shows the scan results for each supported entity. If AWP detects a vulnerability or a hardcoded secret, it shows you the vulnerable entity and suggests remediation. Then, by default AWP scans your VMs once every 24 hours. For Function Apps, AWP inspects the `lastModifiedTimeUtc` attribute of the Function App and rescans it when the attribute changes.

You can select one of two modes for AWP:

- **SaaS Mode** - AWP creates the snapshots of your EC2 volumes or VM disks and scans the snapshots on a virtual machine located in CloudGuard's own AWS account or Azure subscription. With this mode, you do not pay for the scans, and CloudGuard fully manages all the required resources.
- **In-Account Mode** - AWP scans data locally, so everything stays in your AWS or Azure account. The only data sent to CloudGuard are the AWP scanner findings. With this mode, you can keep all your data private, but the volumes/disks scanning entails additional costs.

Onboarding AWP

To enable AWP in your environment:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. In the AWP column, click **Enable** for your environment.
3. Follow the instructions on the wizard page that opens. For more details, see:
   a. AWS - *"AWP for AWS Environments" on page 497*
   b. Azure - *"AWP for Azure Environments" on page 510*
4. In the CloudGuard wizard, click **Next**. CloudGuard completes the process to enable AWP scanning.

AWP starts to scan the VMs and functions and shows the first results within several minutes. Depending on the number of assets, the scan can take up to a few hours. The scanned assets appear on the **Protected Assets** page of the CloudGuard portal.

Viewing Results

To see the scan results:

1. In the CloudGuard portal, go to **Assets** > **Protected Assets** and filter the view by the asset type **AWS EC2 Instance**, **Azure Virtual Machine** or **Azure Function App**.
2. Make sure that the Scan Status of the asset that you need is **Scanned**.
   - **In Progress** - The asset is being scanned.
   - **Internal Error** - The asset scan encountered an error.
   - **Pending Scan** - The asset is not scanned yet.
   - **Scanned** - AWP scanned the asset, and the results are available.

- **Skipped** - AWP excluded the asset from scanning (For types of skipped entities, see *"Known Limitations" on the next page*).

3. Click the asset to see its page and go to the **Vulnerabilities** tab that contains the scan results. It shows the most recent scan date and time.

You can search and filter the scan results by appropriate criteria in the Remediation Summary.

See in the tabs these types of vulnerability:

- **CVEs** - Shows scan of packages installed on the EC2, scanning package managers existing on the machine, and all libraries. Results are sorted by severity. Each package contains a list of CVEs found on it, sorted by severity as well. The header shows the file path, so if the package is installed in more than one place, you must apply the remediation for every instance of the CVE. If the issue is fixable, the Remediation section in the header shows the way.

  To search for a specific CVE, click **Search CVE** and go to the *"Vulnerability Search" on page 447* page.

- **Secrets** - Shows insecure or exposed keys, passwords, and where each of them was found. You find the insecure item in the code and delete it.

- **Remediation Summary** - Shows the contents of the three previous pages in one location. For secrets and threats, it directs you to the file. For CVEs, it indicates which package requires an upgrade.

To export the scan results:

Use the CloudGuard API to export the results to a file. For details, see the API Reference Guide.

Viewing AWP Details

To see the AWP configuration details of the onboarded environment, navigate to **Assets** > **Environments**, select the environment, and open the **AWP** tab.

This page shows the scan statistics, analysis, and other configurable parameters.

To configure these parameters for all onboarded environments, use the *"Workloads Settings" on page 861*.

Custom Tags

AWP adds its tags to every resource it dynamically creates in your AWS account or Azure subscription. It also allows you to add custom tags to these resources. The maximum number of custom tags is 20.

To set the scan preferences:

1. On the AWP page, click the Settings icon in the **Scan Settings** section.
2. To set the maximum number of simultaneous scans in the same region, enter a number between 1 and 20 (by default, **20**) in **Max Concurrent Scans per Region**.
3. To set the period between consecutive scans of the same VM, in **Scan Interval in Hours**, enter a number:
   - For the SaaS scan mode, between 24 and 1000 (by default, **24**)
   - For the In-Account scan mode, between 4 and 1000 (by default, **24**)
4. To configure a custom tag, click **Add** and enter the tag's key and value.
5. For AWS accounts with In-Account mode only, to enable the scanning of AWS Marketplace licensed images, select the **Scan licensed images** option.

   Note - Scanning of AWS Marketplace licensed images may result in additional charges based on AWS pricing policies. AWP attempts to scan with the regular machine types. If this fails, it reverts to the original machine types, subject to some cost limitations. Review your cloud storage plan to understand potential costs.
6. Click **Confirm**.

Known Limitations

AWP cannot scan some types of assets and skips them. In the table below, see the reasons for the **Skipped** status.

Virtual Machine Instances

AWS

Skipped	Scanned
AWS Marketplace licensed images on Windows	AWS Marketplace licensed images on Linux, with In-Account mode (when enabled) Linux operating system Windows operating system
Stopped VM	Running VM
One of the VM disks is larger than 1 TB	

Azure

Skipped	Scanned
Subscription scope lock	
AWP resource group scope lock	
Azure China and Azure Gov regions	
	Linux operating system Windows operating system
Azure Server-Side Encryption + Customer-Managed key: ■ SaaS Azure Disk Encryption: ■ SaaS ■ Centralized See also: *"Azure Account Encryption" on page 518*	Azure Server-Side Encryption: ■ Platform-Managed key – SaaS, In-Account, Centralized ■ Customer-Managed key – In-Account Centralized Azure Disk Encryption: ■ In-Account[1] Azure Encryption at Host: all modes
VMs runtime is less than 4 hours	VMs runtime is 4 hours or more
Stopped VM	Running VM
One of the VM disks is larger than 1 TB	

[1] - Mixing ADE with logical volumes is not supported

File System in Scanned Machines

Linux

- Operating System files and directories are mounted on the top level of the file system ('/').
- Data Disk can be partitioned logically or physically.
- Supported file system formats:
  - XFS
  - EXT2
  - EXT3
  - EXT4
  - NTFS
  - BTRFS

Windows

Temporary disks are not mounted.

Function Apps

To scan Azure Function Apps, AWP needs to download the Function App source code. The availability of the source code depends on how the Function App is deployed.

In general, Function Apps are deployed and available on the SCM (Source Code Management) site.

The SCM site can be inaccessible in these cases:

- SCM IP is restricted.
- Linux Function App is deployed with the **Consumption** hosting plan.

In addition to the SCM site, Azure allows Function App deployment with other technologies. For more details, see Azure documentation.

Supported Function Apps

The table below shows the types of Function Apps that AWP can download and scan. AWP uses available methods according to their priority and applicable conditions.

Operating System	Download Source	Priority	Hosting Plan	Usage Conditions
Windows, Linux	External URL	1	All	External URL deployment
Windows	SCM	2	All	SCM is not blocked
Linux	SCM	2	Non-consumption[1]	SCM is not blocked
Linux	Blob - Storage Account	2	Consumption	-
Windows, Linux	File Share - Storage Account	3	Consumption, Premium	1. SCM is blocked 2. The File Share option is enabled at the Function App creation

[1] Non-consumption hosting plans are, for example, Premium and Dedicated.

AWP skips scanning Function Apps in these cases:

- Function App has Container deployment
- Logical-App Function App (no source code)
- The Function App content is not available. This can happen when the SCM site is the only option and it is blocked (Public Networking is entirely disabled). To allow AWP scanning, see *"Blocked SCM Troubleshooting" below*.

Blocked SCM Troubleshooting

For security reasons, it is usually recommended to block public networking for your SCM.

To allow AWP to scan your SCM, apply these restrictions:

1. Enable the inbound networking to the SCM site from the AWP scanner VNet only.
   a. For **Main site**, apply the deny rule to the entire site.
   b. For **Advanced tool site**, apply the allow rule to inbound traffic from the AWP scanner VNet only.
2. AWP scanner VNet is created after first scan of any VM or Function App.

Azure Resources

AWP scanner requires 4 vCPUs on average. Make sure that your regional vCPU quota is sufficient for launching the scanner.

More Links

- *"AWP for AWS Environments" on page 497*
- *"AWP for Azure Environments" on page 510*
- *"Unified Onboarding of AWS Environments" on page 54*
- *"Onboarding Azure Subscriptions" on page 169*
- *" Risk Management" on page 89*
- Settings: *"Agentless Workload Protection" on page 861*
- Azure documentation: Disk Encryption Overview

AWP for AWS Environments

When you enable AWP, it creates a cross-account stack on your AWS account. The cross-account stack deploys in your account these primary resources:

- Cross-account IAM role
- Proxy utility Lambda
- Multi-Region key (for SaaS mode)

The data that the AWP scanner sends to CloudGuard are only CVEs and paths of the secrets.

All resources that AWP creates in your account have the **Owner : CG.AWP** tag.

Onboarding Workflow

To enable AWP on your AWS environment:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. Click **Enable** in the AWP column for your AWS environment.

   Alternatively, open the AWS environment, go to the AWP tab, and click **Enable Agentless Workload Posture**.
3. Follow the instructions on the wizard page that opens.
   a. Select one of the options:
      - **In-Account** to scan one account locally. See *"Independent Accounts" on the next page*
      - **In-Account Sub** to scan one or more accounts on another (centralized) account. See *"Centralized Account and its Sub-Accounts" on page 500*
      - **SaaS** to scan one account remotely on CloudGuard accounts. See *"AWS SaaS Mode" on page 503*
      - **In-Account Centralized** to scan one or more accounts on this (centralized) account. See *"Centralized Account and its Sub-Accounts" on page 500*
   b. Copy the URL generated by CloudGuard and paste it in a new browser tab. When you sign in to your AWS account, you are redirected to the CloudFormation page to create a new CFT stack.
   c. In AWS, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.

d. Click **Create stack**. CloudFormation starts to create the stack. After you create the stack, additional permissions are granted to CloudGuard.

Important - Make sure you do not change the mode (SaaS/In-Account) during the onboarding. For successful onboarding, you must use the same mode that you selected before the stack creation.

4. In the CloudGuard wizard, click **Enable AWP**. CloudGuard completes the process to enable AWP scanning.

Note - If you onboarded your AWS environment to CloudGuard with Terraform (*"Onboarding with Terraform" on page 144*), you can use Terraform to onboard it to AWP. For more information, see the Terraform documentation for AWP.

AWS In-Account Mode

With the In-Account mode, AWP scans data locally, so everything stays in your AWS account. The only data sent to CloudGuard are the AWP scanner findings. With this mode, you can keep all your data private, however, the volumes scanning entails additional costs.

Independent Accounts

In the In-Account mode for independent accounts, AWP generates the scan resources inside the same accounts, whose workloads it scans.

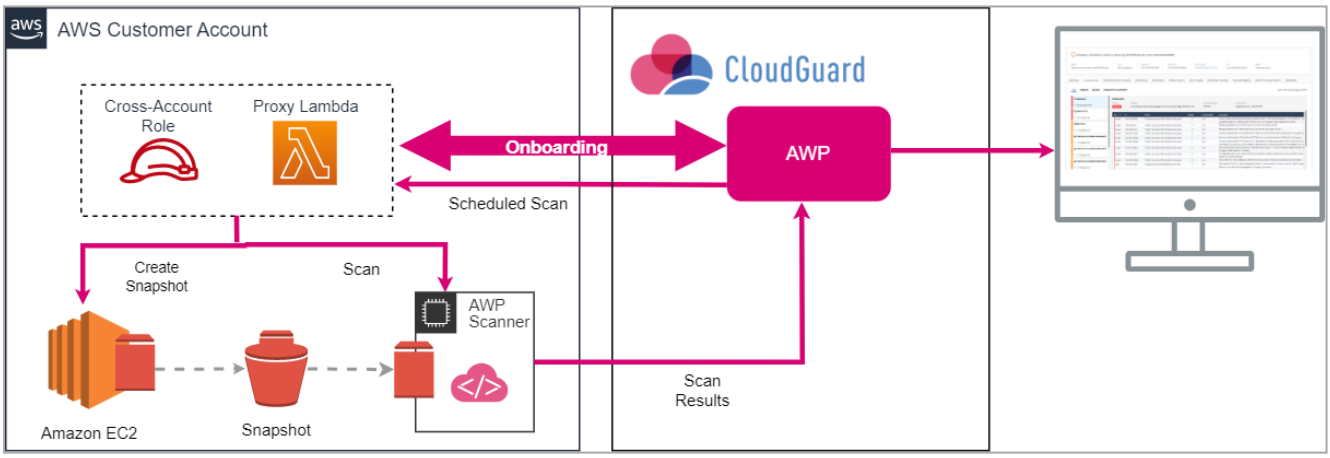

Scanning Workflow

1. With the Cross-account role, AWP gathers information about your instance volumes.
2. With the Cross-account role, AWP remotely invokes the proxy utility Lambda with a request to create snapshots from the instance volumes.
3. When all instance volumes have their equivalent snapshots created, AWP launches an EC2 instance on the customer account and performs the scanning. This instance runs in the **same region** as the original EC2, in a custom VPC that AWP creates for this task.

- The scanner outbound traffic is restricted by a security group rule that limits access exclusively to S3 IP addresses.
- No inbound rules are configured.

To ensure that network traffic remains within the AWS backbone, the VPC utilizes an S3 endpoint.

Note - Make sure that your account does not reach the VPC quota (by default, 5 VPCs for each region) before you enable AWP.

4. When the scan is complete, AWP sends a request to Lambda to delete the snapshots created for the scan.

Onboarding Independent Accounts

1. In the onboarding wizard, select the **In-Account** option for the scan mode.
2. The link in step 2 has some default onboarding parameters. If you need to change them, click **Advanced** and edit:
   - Name of the CloudFormation stack
   - Cross-Account External ID
   - Cross-Account Role Name
   - Optionally, you can enable the scanning of AWS Marketplace licensed images by selecting the **Scan licensed images** option. Due to internal restrictions, AWP can skip some licensed images with an indication **unsupported license**.

     Note - Scanning of AWS Marketplace licensed images may result in additional charges based on AWS pricing policies. Initially, AWP attempts to scan with the regular machine types. If this fails, it reverts to the original machine types, subject to some cost limitations. Review your cloud storage plan to understand potential costs.
3. Copy the URL generated by CloudGuard and paste it into a new browser tab. When you sign in to your AWS account, you are redirected to the CloudFormation page to create a new CFT stack.
4. In AWS, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
5. Click **Create stack**. CloudFormation starts to create the stack. After you create the stack, additional permissions are granted to CloudGuard.

Centralized Account and its Sub-Accounts

When you select to use the In-Account mode for independent accounts, AWP creates multiple resources on your account during its scanning. If you have many accounts, these multiple distributed resources can be impracticable for management and billing. In such cases, you can configure one of your onboarded AWS accounts as a *Centralized account*, where all AWP scans and resources are located. You can configure other AWS accounts in the same organization as *Sub-accounts* to have their scanners and AWP resources located in the centralized account.

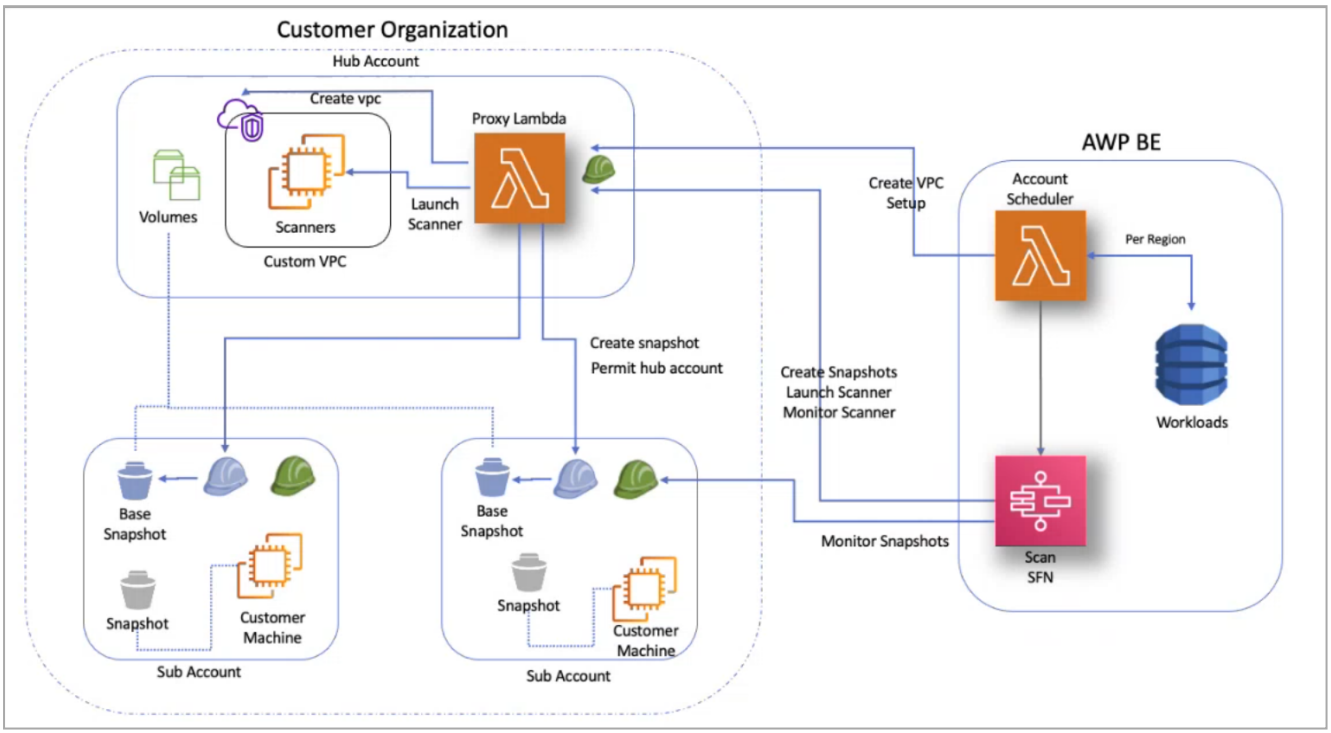

Scanning Workflow

1. With the Cross-account role, AWP uses the proxy utility Lambda to create a custom VPC where the scanners can run.
2. AWP remotely invokes the proxy Lambda with a request to create snapshots on a sub-account. The proxy Lambda assumes the operator role of the sub-account and, in case of the encrypted volume, re-encrypts the snapshots with the CloudGuard multi-Region key.

   Note - The key resources are regional. Sometimes the instance is in a region where the multi-Region key does not have a replica yet. In such a case, AWP creates a replica of the key in the relevant region, with its Cross-account role.
3. AWP monitors the creation of the snapshots with the Reader role until the snapshots are ready.

4. When all sub-account snapshots are created, AWP launches the proxy Lambda on the centralized account and performs the scanning in the created VPC. In the case of the encrypted volume, the AWP scanner has access to the data of the sub-account snapshots because they are encrypted with AWP's own multi-Region key. The EC2 instance runs in the **same region** as the centralized account, in a custom VPC that AWP creates for this task.

   - The scanner outbound traffic is restricted by a security group rule that limits access exclusively to S3 IP addresses.
   - No inbound rules are configured.

   To ensure that network traffic remains within the AWS backbone, the VPC utilizes an S3 endpoint.

   Note - Make sure that your account does not reach the VPC quota (by default, 5 VPCs for each region) before you enable AWP.

5. When the scan is complete, AWP sends a request to Lambda to delete the snapshots created for the scan.

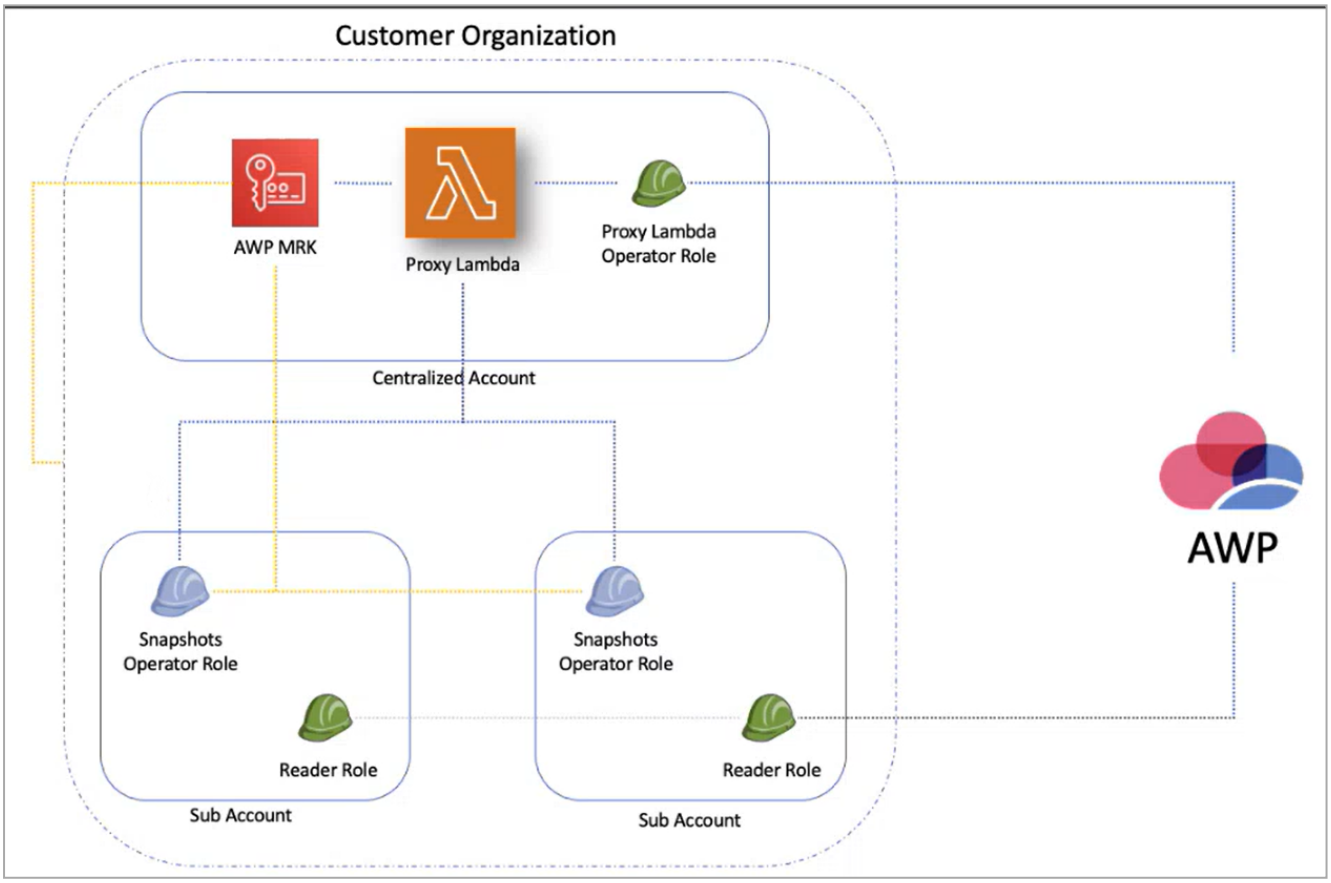

The centralized account serves the entire organization. If the organization has several centralized accounts, each of them has permissions to manage scanners and AWP resources for all sub-accounts in the organization.

Onboarding Centralized and Sub-Accounts

First, you need to onboard a centralized account. Then you can onboard a sub-account which assets are to be scanned on the centralized account.

Centralized Account

1. In the onboarding wizard, select the **In-Account Centralized** option for the scan mode.
2. The link in step 2 has some default onboarding parameters. If you need to change them, click **Advanced** and edit:
   - Name of the CloudFormation stack
   - Cross-Account External ID
   - Cross-Account Role Name
3. Copy the URL generated by CloudGuard and paste it in a new browser tab. When you sign in to your AWS account, you are redirected to the CloudFormation page to create a new CFT stack.
4. In AWS, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
5. Click **Create stack**. CloudFormation starts to create the stack. After you create the stack, additional permissions are granted to CloudGuard.

Note - Make sure to validate in AWS that all regions enabled for sub-accounts are also enabled for the centralized account.

Sub-Account

1. In the onboarding wizard, select the **In-Account Sub** option for the scan mode.
2. Select the centralized account where this sub-account has its scanner and AWP resources.
3. The link in step 3 has some default onboarding parameters. If you need to change them, click **Advanced** and edit:
   - Name of the CloudFormation stack
   - Cross-Account External ID
   - Cross-Account Role Name
4. Copy the URL generated by CloudGuard and paste it in a new browser tab. When you sign in to your AWS account, you are redirected to the CloudFormation page to create a new CFT stack.

5. In AWS, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
6. Click **Create stack**. CloudFormation starts to create the stack. After you create the stack, additional permissions are granted to CloudGuard.

When the onboarding is completed, the AWP tab of the AWS environments shows its mode, its related centralized account and other scan statistics.

Roles and Permissions

Centralized account stack

Role	Permission
Proxy Lambda Operator role	▪ Invoke proxy Lambda
Proxy Lambda role	▪ Assume role to sub-accounts dedicated role ▪ Create VPC setup ▪ Launch scanner

Sub-account stack

Role	Permission
Reader role	▪ Describe instances ▪ Describe volumes ▪ Describe snapshots
Operator role	▪ Create snapshot ▪ Modify snapshot attribute ▪ Copy snapshot ▪ Delete snapshot

AWS SaaS Mode

In the SaaS mode, AWP creates the snapshots of your EC2 volumes and scans the snapshots on a virtual machine located on the CloudGuard's own AWS account. With this mode, you do not pay for the scans, and CloudGuard fully manages all the required resources.

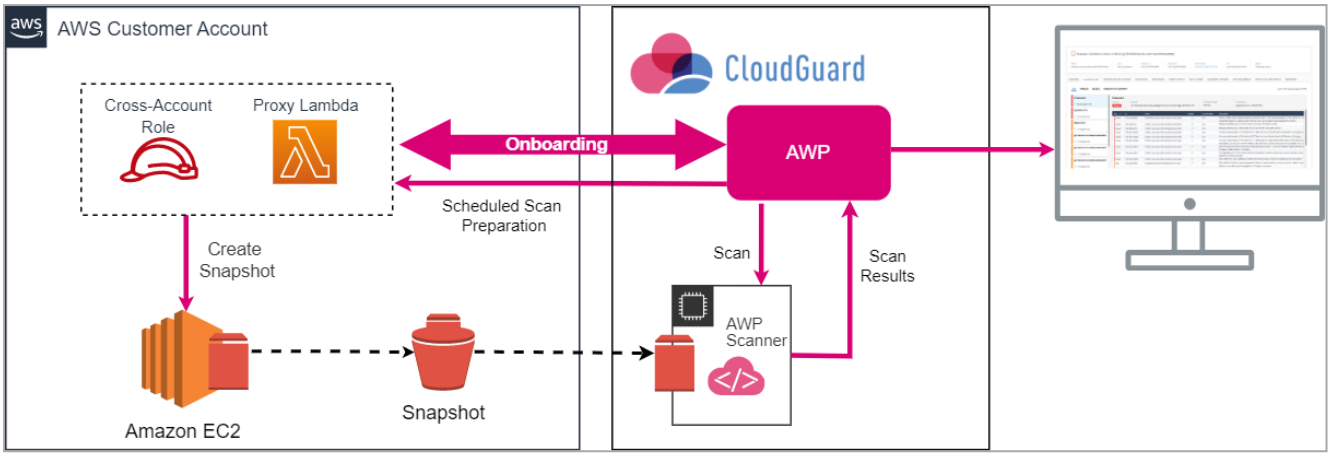

Scanning Workflow

1. With the Cross-account role, AWP gathers information about your instance volumes.
2. With the Cross-account role, AWP remotely invokes the proxy utility Lambda with a request to create snapshots in the customer account from the instance volumes.
3. When all instance volumes have their equivalent snapshots created, a scanner machine launches at the AWP engine with the snapshots attached to it and performs the scanning.
4. When the scan is complete, AWP sends a request to Lambda to delete the snapshots created for the scan.

Scanning Encrypted Volumes in SaaS Mode

For security reasons, AWP does not have access to encrypted volumes in your EC2 instance and cannot scan them. It happens because CloudGuard does not require access to the encryption keys and never obtains them from you as this can compromise your data.

To scan the encrypted volumes securely, CloudGuard re-encrypts the volume data with its own multi-Region key. It installs the key as part of the AWP cross-account stack. Likewise, it installs a proxy utility Lambda function as part of the cross-account stack. With its cross-account role, this Lambda function manages all procedures of the snapshots' creation and re-encryption on remote requests (invocations) from the AWP engine.

This limits access to your account keys only to the proxy Lambda, where you have full visibility and control.

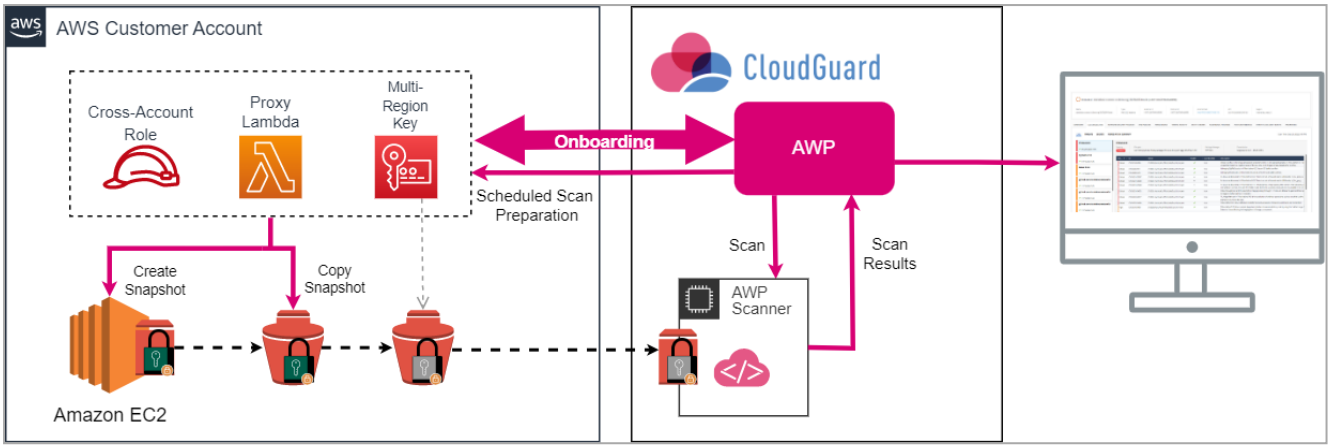

Scanning Workflow

1. With the Cross-account role, AWP gathers information about your instance volumes.
2. With the Cross-account role, AWP remotely invokes the proxy utility Lambda with a request to create snapshots from the instance volumes.

   In the case of encrypted snapshots, AWP sends one more request to Lambda to re-encrypt the snapshots created in the customer account with the CloudGuard multi-Region key.

   Note - The key resources are regional. Sometimes the instance is in a region where the multi-Region key does not have a replica yet. In such a case, AWP creates a replica of the key in the relevant region, with its Cross-account role.
3. When all instance volumes have their equivalent snapshots created, a scanner machine launches at the AWP engine with the snapshots attached to it and performs the scanning.

   For encrypted snapshots, the AWP scanner has access to the snapshots' data because they are encrypted with AWP's own multi-Region key.
4. When the scan is complete, AWP sends a request to Lambda to delete the snapshots created for the scan.

Onboarding Accounts in SaaS Mode

1. In the onboarding wizard, select the **SaaS** option for the scan mode.
2. The link in step 2 has some default onboarding parameters. If you need to change them, click **Advanced** and edit:
   - Name of the CloudFormation stack
   - Cross-Account External ID
   - Cross-Account Role Name

3. Copy the URL generated by CloudGuard and paste it in a new browser tab. When you sign in to your AWS account, you are redirected to the CloudFormation page to create a new CFT stack.
4. In AWS, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
5. Click **Create stack**. CloudFormation starts to create the stack. After you create the stack, additional permissions are granted to CloudGuard.

Actions

Ignoring an Instance Scan

If you need to deliberately skip scanning an instance, set a tag for the instance on the AWS console.

To set a tag for the AWP scanner to ignore an EC2 instance:

1. In the AWS console, navigate to **EC2** > **Instances** and select your instance.
2. On the Tags tab, click **Manage tags**.
3. In the **Manage tags** window, click **Add tag** and add a new tag with the key `CG_AWP_SKIP_SCAN`.

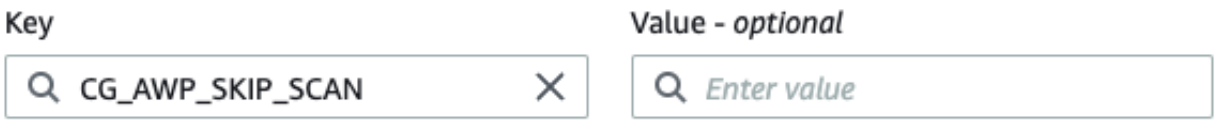

Offboarding AWP

When you disable AWP for your environment, you must delete the CloudFormation stack created on your account. This process removes the stack created with the CloudFormation Template.

Disabling AWP in independent accounts in SaaS or In-Account mode

1. Sign in to your AWS account.
2. Delete the AWP CloudFormation stack.
3. Use an API call to remove AWP from your environment. For more information, see Delete Agentless Account.

Disabling AWP in Sub-Accounts

1. Sign in to your AWS account.
2. Use an API call to remove AWP from your environment. For more information, see Delete Agentless Account.
3. Delete the AWP CloudFormation stack.

Disabling AWP in Centralized Accounts

1. Make sure that the centralized account does not have connected sub-accounts. If it has, first remove AWP from the connected sub-accounts.
2. Sign in to your AWS account.
3. Delete the AWP CloudFormation stack.
4. Use an API call to remove AWP from your environment. For more information, see Delete Agentless Account.

Switching between AWP Modes

You cannot instantly switch the AWP mode from SaaS to In-Account and in reverse. For this, you must offboard AWP and then onboard it again with another mode.

To change the AWP mode:

1. Remove AWP from your AWS account.
2. Onboard AWP on the account with another mode.

Creating a Dedicated VPC

During each scanning process, AWP creates a custom VPC in the region where your EC2 instance runs. You can use your own VPC (not recommended), for accounts onboarded to AWP with In-Account or In-Account Centralized mode.

To use your dedicated VPC, create all the required resources and follow these general instructions:

- **VPC** - Create a VPC in each region that has workloads to be scanned.

  Note - The VPC and other resources created for a Centralized account are used for all its sub-accounts.

- **Custom tag** - Add a custom tag with the `CG_AWP_NETWORK` key (key: `CG_AWP_NETWORK` value: any) to the VPC and related resources (subnet, route table, and security group).

- **S3 Bucket access** - Make sure you have access to S3 buckets and to AWS regional S3 IPs.
- **Route table** - Create a route table and attach it to the VPC.
- **Network ACL** (Access Control List) - Create a network ACL for the VPC.
- **Subnets** - Create the necessary subnets in the VPC, minimum one subnet. One subnet for each Availability Zone (AZ) is recommended. The number and configuration of subnets depend on your needs. Associate each subnet with the route table created earlier.
- **Security Group** - Create a security group in the VPC and add a new rule to allow outbound traffic to the specified S3 prefix list.
- **Internet Gateway** or **S3 Gateway Endpoint** - Based on your needs, create an Internet Gateway and attach it to the VPC, or an S3 Gateway Endpoint.

  If you select to work with S3 Endpoint, make sure:

  - You have access to these S3 buckets:
    - `arn:{CloudGuard_aws_account}:s3:::agentless-*`
    - `arn:{CloudGuard_aws_account}:s3:::agentless-*/*`
  - The VPC has permissions to AWP bucket with these actions:
    - `s3:GetObject`
    - `s3:PutObject`
    - `s3:DeleteObject`

Customer VPC API

You can indicate that your VPC is managed by you during the onboarding process or update the account settings after it is onboarded.

To add an independent account:

POST V2/workload/agentless/aws/accounts/{accountNumber}/enable

```
{
      "scanMode": "inAccount",
       "agentlessAccountSettings": {
            "inAccountScannerVPC": "ManagedByCustomer"
          }
}
```

For more details, see the CloudGuard API Reference Guide

To add a centralized account:

POST V2/workload/agentless/aws/accounts/{accountNumber}/enableCentralizedAccount

```
{
      "agentlessAccountSettings": {
           "inAccountScannerVPC": "ManagedByCustomer"
         }
}
```

For more details, see the CloudGuard API Reference Guide

To update account settings:

PATCH V2/workload/agentless/{provider}/accounts/{accountNumber}/settings

```
{
      "inAccountScannerVPC": "ManagedByCustomer"
}
```

For more details, see the CloudGuard API Reference Guide

More Links

- *"Agentless Workload Posture" on page 489*
- *"AWP for Azure Environments" on page 510*

AWP for Azure Environments

When you enable AWP, it adds permissions to the App Registration created during your environment onboarding to CloudGuard. These permissions allow AWP to manage the necessary resources for the scan.

The data that the AWP scanner sends to CloudGuard are only CVEs and paths of the secrets.

All resources that AWP creates in your subscription have the **Owner: CG.AWP** tag.

Onboarding Workflow

To enable AWP on your Azure environment:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. Click **Enable** in the AWP column for your Azure environment.
3. Follow the instructions on the wizard page that opens.
    a. Select the scan mode: **SaaS** or **In-Account**. For more details, see *"Azure SaaS Mode" on the next page* and *"Azure In-Account Mode" on page 512*.
        i. For In-Account mode, select enabling AWP for an independent account or centralized configuration. For more details, see *"Independent Accounts" on page 512*.
        ii. For Centralized In-Account configuration, enable AWP separately for a centralized account and for each Sub-account. For more details, see *"Centralized Account and its Sub-Accounts" on page 513*.
    b. Click the link to open your account in Azure Cloud Shell or use AZ CLI in your terminal.
    c. Copy the script created by the AWP engine. You can download the script to review or edit it based on your needs.
    d. Run the script in the shell or terminal.

    AWP creates the required resources and roles in your Azure subscription to allow the AWP scanner to run.

    Important - Make sure you do not change the mode (SaaS/In-Account) during the onboarding. For successful onboarding, you must use the same mode that you select in Step **3a**.
4. In the CloudGuard wizard, click **Enable AWP**. CloudGuard completes the process to enable AWP scanning.

Caution - When you enable AWP, make sure there is no lock on your Azure subscription or on the AWP resource group. AWP cannot delete locked resources, which causes additional costs.

Onboarding script

Based on the selected mode, CloudGuard enables different flags for the command that generates an onboarding script. You do not have to copy and paste the command as is. Instead, you can use the flags to change the command and run it edited.

Flag	Description
-t, `--additional_tags`	(Optional) Additional tags to apply to resources created during the onboarding process; format: 'Key1=Value1,Key2=Value2'. These tags are not related to the custom tags created during the scan.

Azure SaaS Mode

In Azure SaaS mode, AWP copies customer disks to CloudGuard's snapshots and disks. They are connected to the scanner VMs that display the toxic combinations found on the copied disks. During the workload scanning, no resources are created on the customer's side.

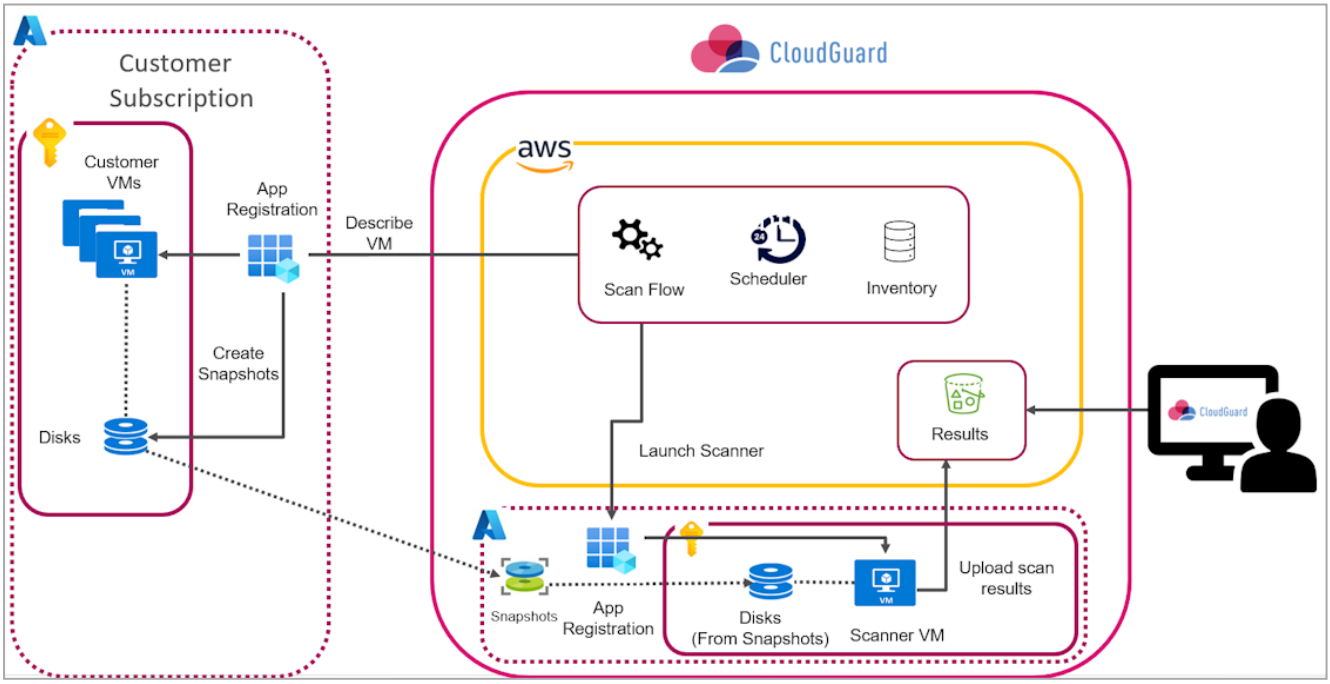

Note - AWP does not support scanning of Function Apps in SaaS mode.

Resources and Permissions for Azure SaaS Mode

In this mode, all scan resources are created on the CloudGuard AWP side, including snapshots, disks, and scanner VMs. AWP creates a single custom role called *CloudGuard AWP VM Data Share* which includes actions to permit reading customer's disks.

AWP creates the AWP-Data-Share custom role with the permissions to:

- describe the account VM configuration to get the disk IDs
- read disk data and generate its snapshot

CloudGuard AWP VM Data Share role is assigned to a subscription that is onboarded.

Note - In Azure SaaS mode, AWP skips encrypted volumes with a customer-managed key.

Azure In-Account Mode

In the In-Account mode, all scan resources (snapshots, disks, VNets, and scanner VMs) are created on the customer's side, therefore the workload data always remains inside the customer tenant perimeter.

Independent Accounts

In the In-Account mode for independent accounts, AWP generates the scan resources inside the same subscriptions, whose workloads it scans.

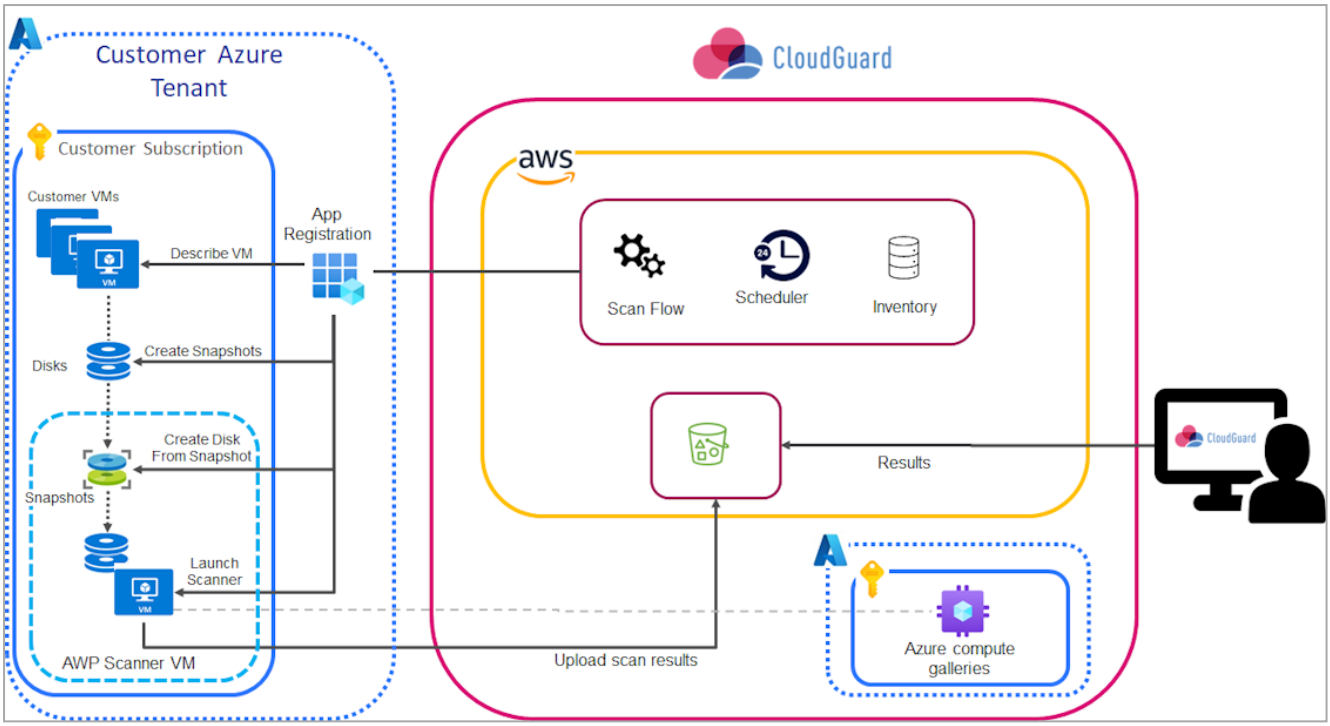

To scan Function Apps, the scanner Virtual Machine launches in the customer subscription. To reach the Function App resource and download its content in runtime, it needs a User-Assigned Managed Identity. This resource belongs to a resource group. It can be granted with permissions in the same way as an App Registration (RBAC).

Centralized Account and its Sub-Accounts

When you select to use the In-Account mode for independent accounts, AWP creates multiple resources on your subscription during its scanning. If you have many subscriptions, these multiple distributed resources can be impracticable for management and billing. In such cases, you can configure one of your onboarded Azure subscriptions as a *Centralized account*, where all AWP scans and resources are located. You can configure other Azure subscriptions on the same tenant as *Sub-accounts* to have their scanners and AWP resources located in the centralized account.

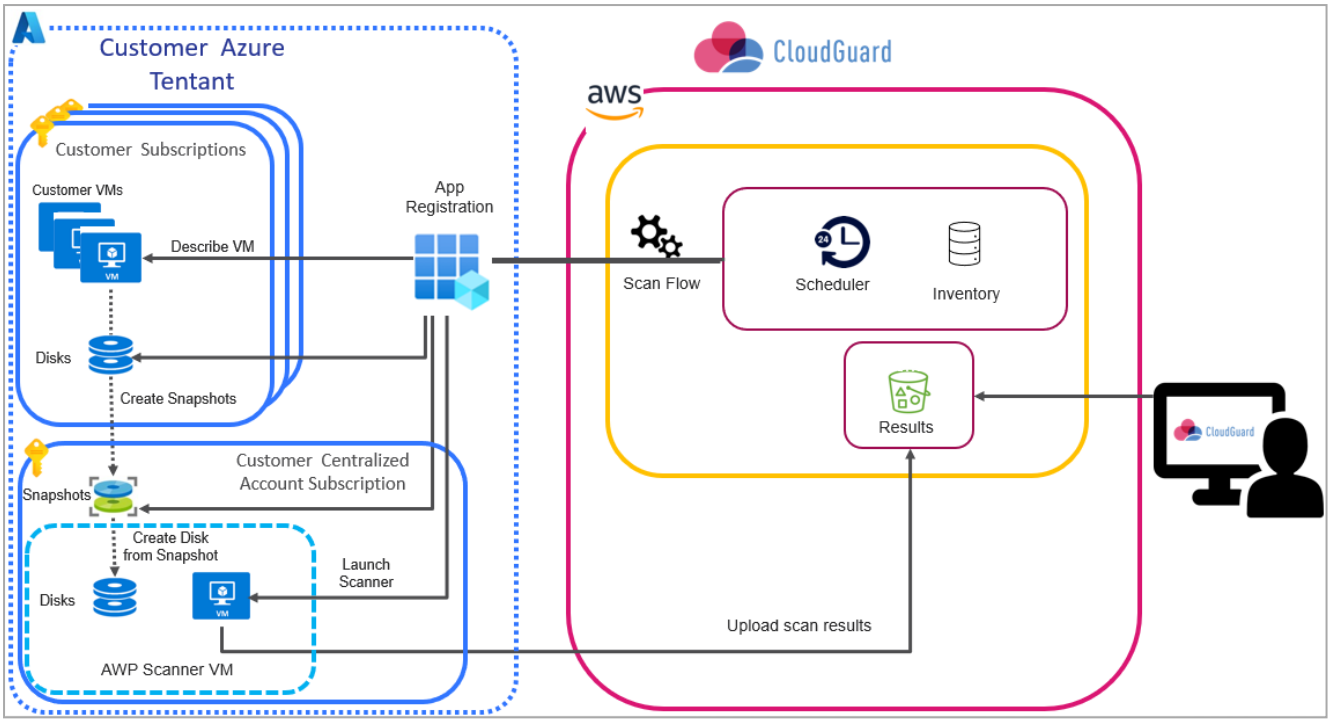

Prerequisites

- For the centralized account, you can select every Azure subscription onboarded to CloudGuard but not onboarded to AWP.
- For a sub-account, make sure it has the required permissions to create resources on the centralized account - see *"Resources and Permissions for Azure In-Account Mode" below*.

Resources and Permissions for Azure In-Account Mode

Virtual Machines

AWP creates two primary custom roles in the customer Azure tenant and assigns them based on the scan mode selected during the subscription onboarding. These roles are:

1. **CloudGuard AWP VM Scan Operator** role is assigned to a subscription that hosts the scans, for example, in the Standard In-Account or Centralized Account modes. It provides permissions to:
    a. Read, Create or Delete Snapshots
    b. Read, Create or Delete Disks
    c. Create network components such as VNets and Security Groups
    d. Create and Delete Scanner VMs

2. **CloudGuard AWP VM Data Share** role is assigned to a subscription whose workloads are to be scanned, for example, in the Standard In-Account or Sub-Account modes. It provides permission to read customer disks.

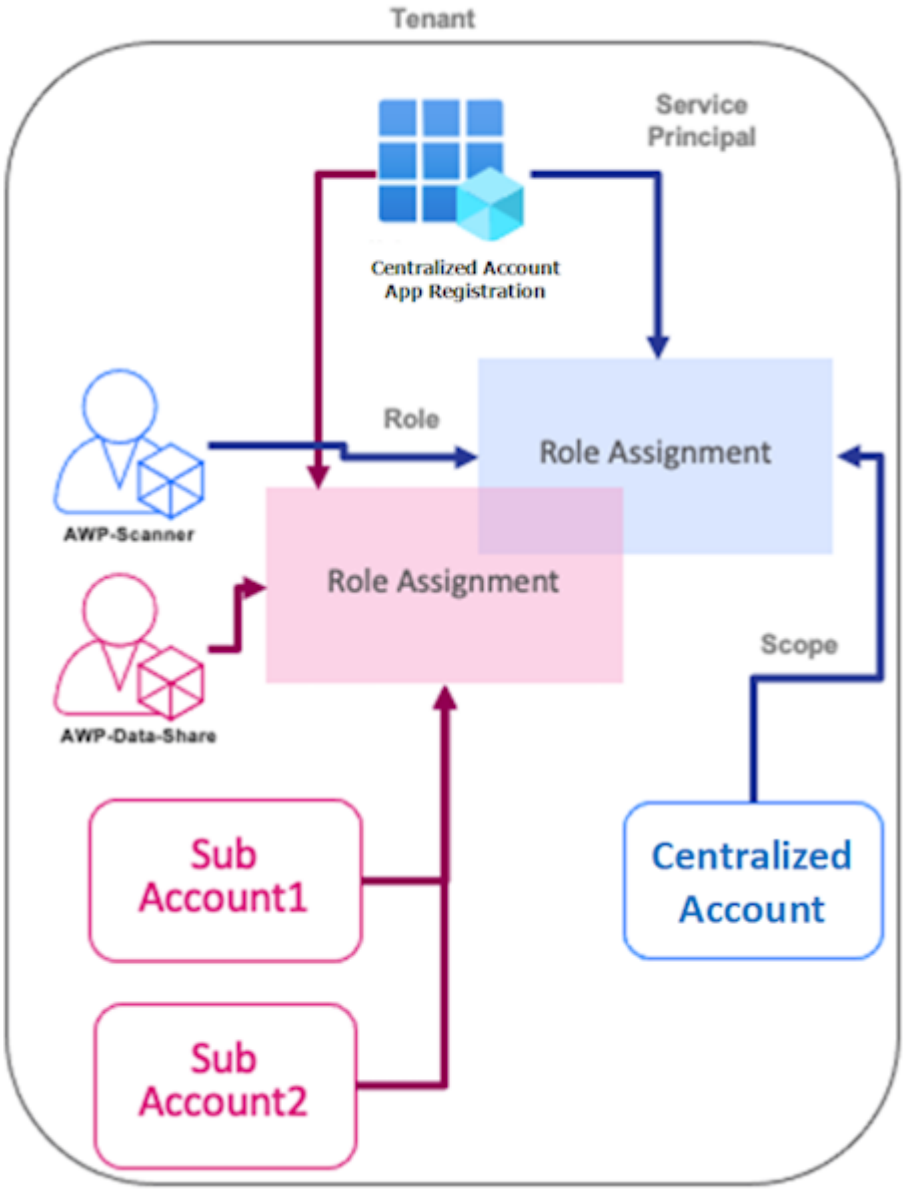

For Standard In-Account mode:

1. Create an AWP dedicated resource group.
2. Assign the Scan-Operator and Data-Share custom roles.

For Centralized Account:

1. Create an AWP dedicated resource group for common resources, such as VNet, to be in use for various Sub-Accounts' scans.
2. Assign the Scan-Operator custom role.

For Sub Account:

1. Create an AWP dedicated resource group for the subscription in the Centralized account subscription, where all scan resources for the sub-account workloads are located.
2. Assign the Data-Share custom role with the scope of the sub-account subscription.

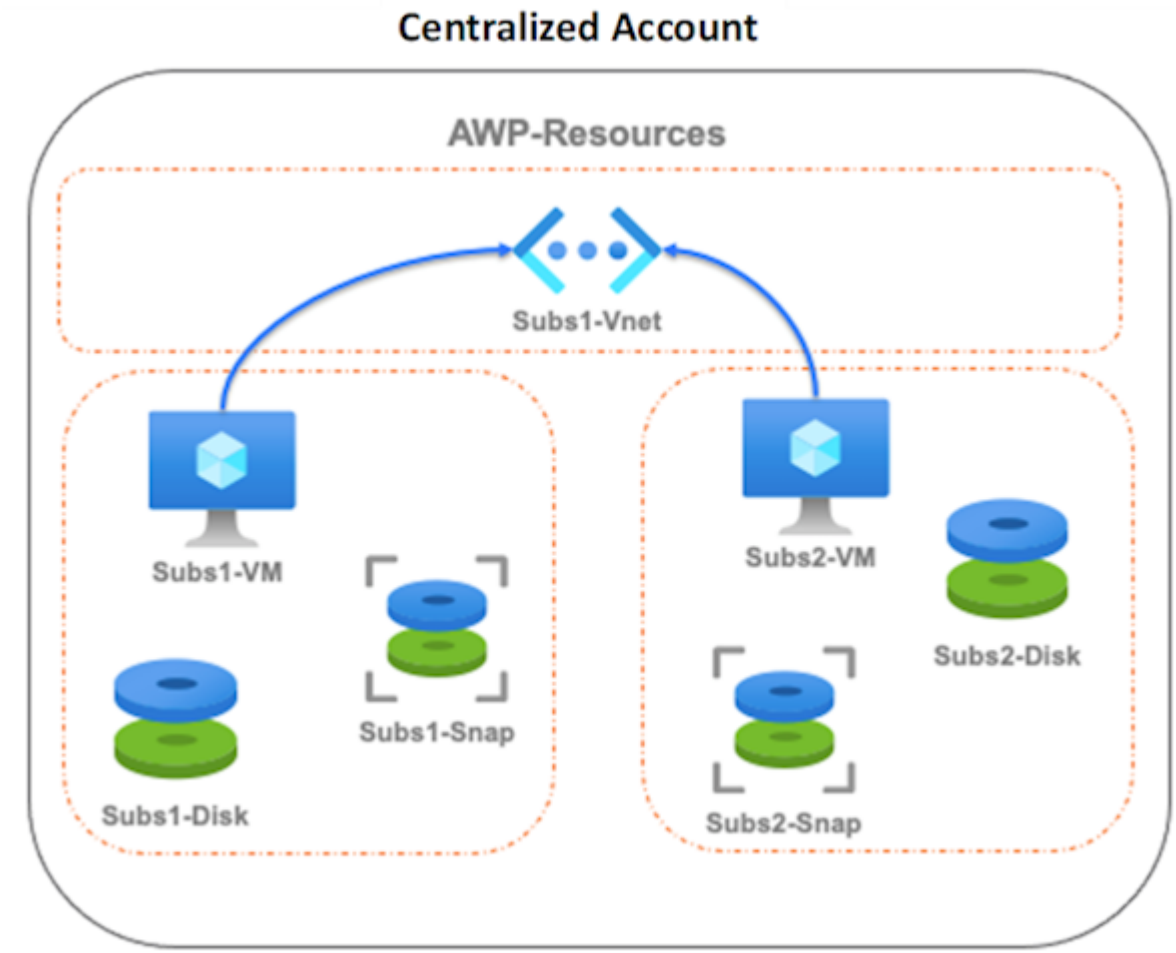

Function Apps

AWP uses a user-assigned managed identity attached to the Function App scanner. This identity is the most safe way to launch a scanner capable of obtaining a Function-App source code, because it does not need transferring secrets or credentials.

When you click **Enable AWP**, CloudGuard onboarding script creates its resources and custom roles for a specific task of Function App scanning:

- Resources
  - **CloudGuardAWPScannerManagedIdentity** – A user-assigned managed identity that enables the Function App scanner to download the Function App source code.
- Custom Roles

- **CloudGuardAWPFunctionAppScanOperator** - A role assigned to the *CloudGuard-Connect* app registration that allows to launch Function App scanner attached to the managed identity.
- **CloudGuardAWPFunctionAppScanner** - A role assigned to the managed identity to enable the Function App scanner to access the Function App source code.

The diagram below shows an In-Account deployment that enables the Function-App scanning.

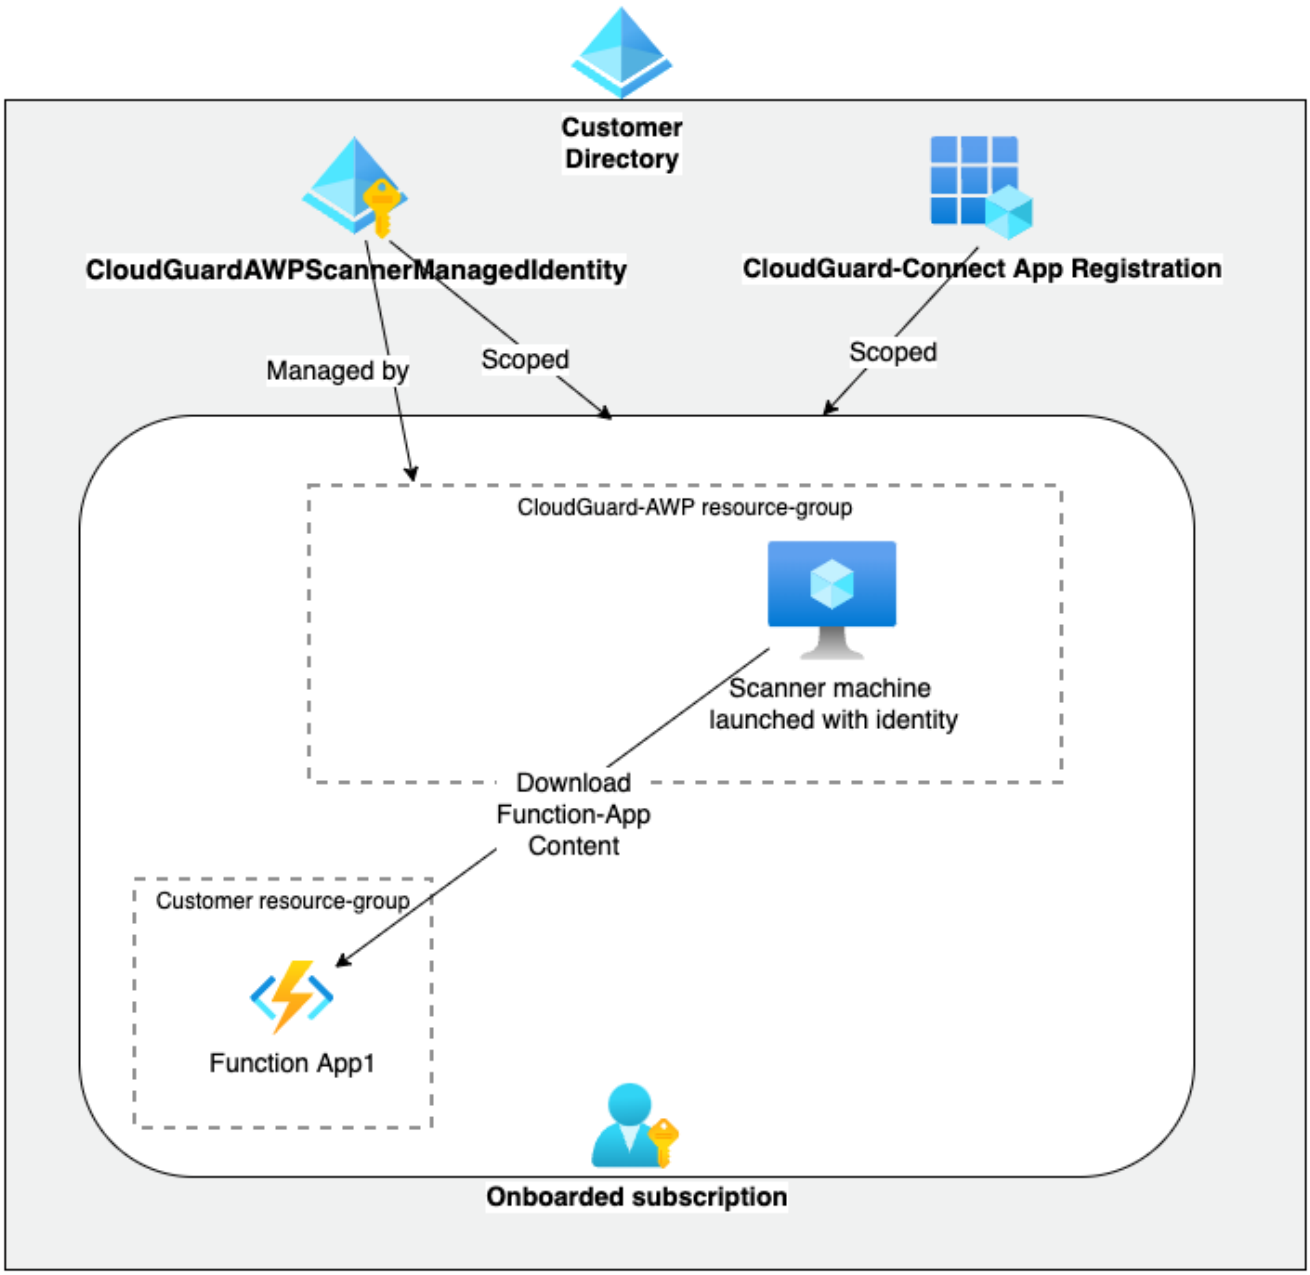

The diagram below shows a Centralized-Sub-Account deployment that enables the Function-App scanning in a Centralized account.

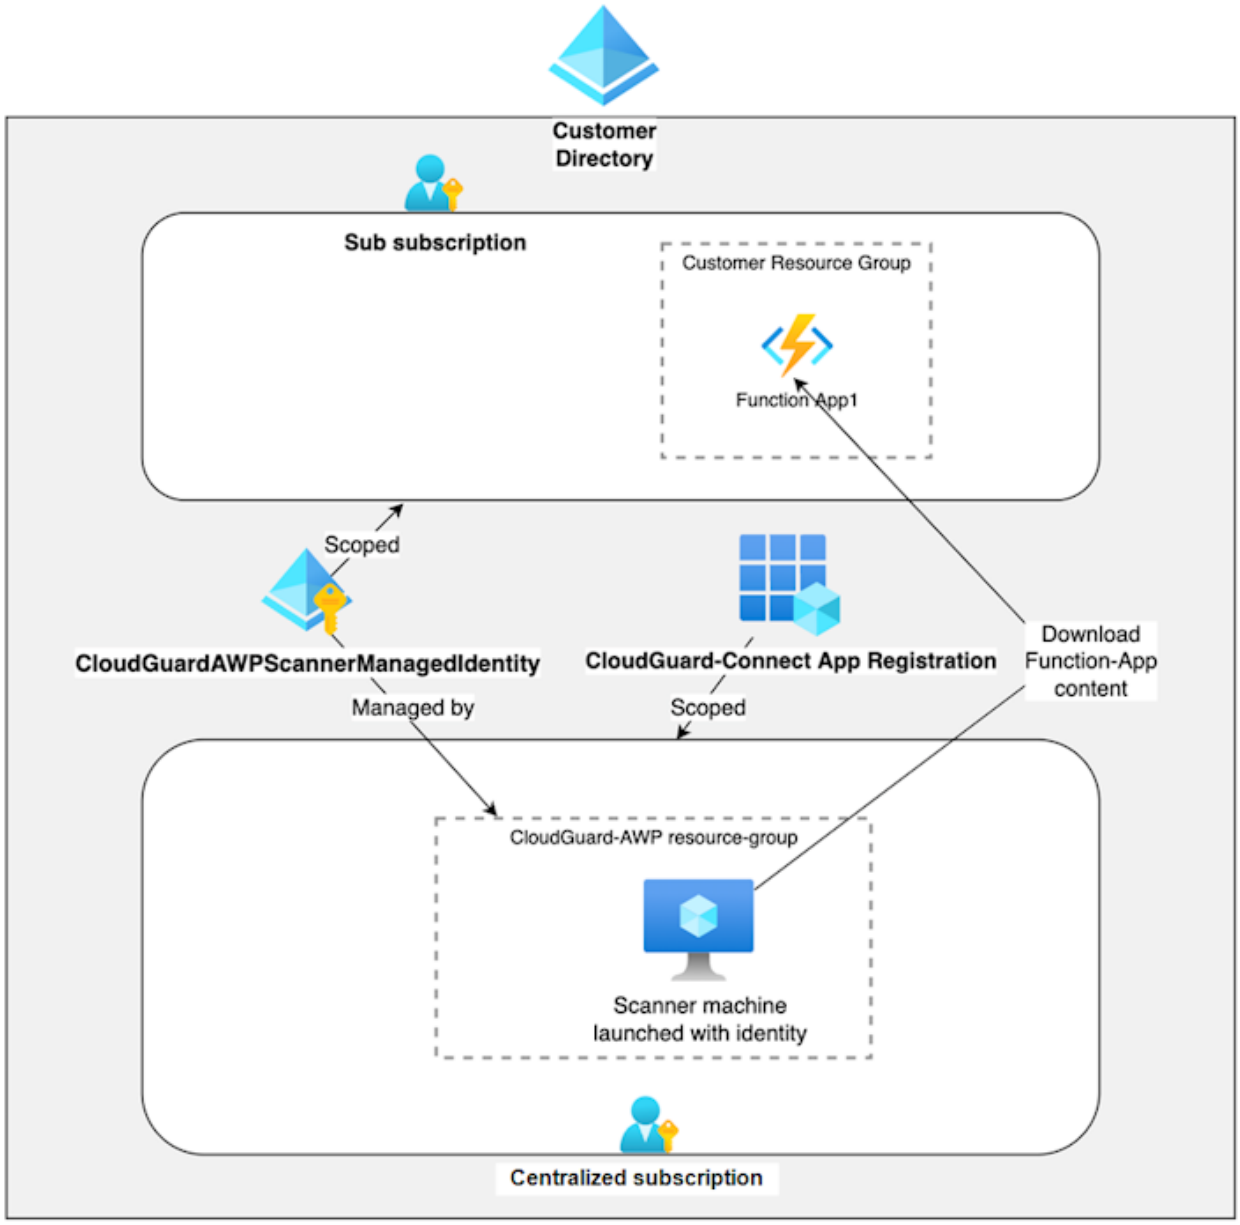

Azure Account Encryption

AWP supports several types of Azure encryption:

- Encryption at Host
- Azure Disk Storage Server-Side Encryption
- Azure Disk Encryption

Encryption Type	Key Type	SaaS		In-Account		Centralized	
		Linux	Windows	Linux	Windows	Linux	Windows

Encryption Type	Key Type	SaaS		In-Account		Centralized	
Encryption at Host	N/A	Yes	Yes	Yes	Yes	Yes	Yes
Server-Side Encryption	Platform-Managed key	Yes	Yes	Yes	Yes	Yes	Yes
	Customer-Managed key	No	No	Yes	Yes	Yes	Yes
Azure Disk Encryption	Customer-Managed key	No	No	Yes	Yes	No	No

For more information about Azure encryption, see the Azure documentation.

Encryption at Host

Encryption at host (EAH) adds one more level to Azure security while enhancing Azure Disk Storage Server-Side Encryption. It ensures that all temp disks and disk caches are encrypted at rest and flow encrypted to the Storage clusters.

AWP uses encryption at host whenever possible. Thus, if your subscription has EAH registered, AWP scanners are launched with EAH activated for In-Account and Centralized scan modes. For SaaS scan mode, AWP scanners are launched with EAH by default.

Note - For VMs encrypted with Azure disk encryption, AWP scanners are launched without EAH.

Server-Side Encryption

If you use server-side encryption (SSE) with customer-managed keys (CMK) for managed disks, AWP can scan the disks with the same encryption method (SSE + CMK) in Azure Centralized Account mode.

For each region with a VM in one of the sub-accounts, AWP dynamically creates a new CMK in the AWP Key Vault and a dedicated DES (Disk Encryption Set). The AWP scanner uses these resources to scan the customer's VM. All resources are created in the centralized account, with one set for each region.

Infrastructure Preparation

When one of the sub-accounts has a VM in a new region, AWP prepares the scanning infrastructure and creates in the region:

- Key Vault
- RSA Key in the Key Vault (CMK-HUB)
- Dedicated DES (Disk Encryption Set) associated with the CMK-HUB.

Scanning

To scan the disks:

1. AWP makes a snapshot from the sub-account disk and stores it encrypted with CMK-HUB in the centralized account.
2. AWP creates a Disk encrypted with CMK-HUB from the snapshot and launches the scanner.

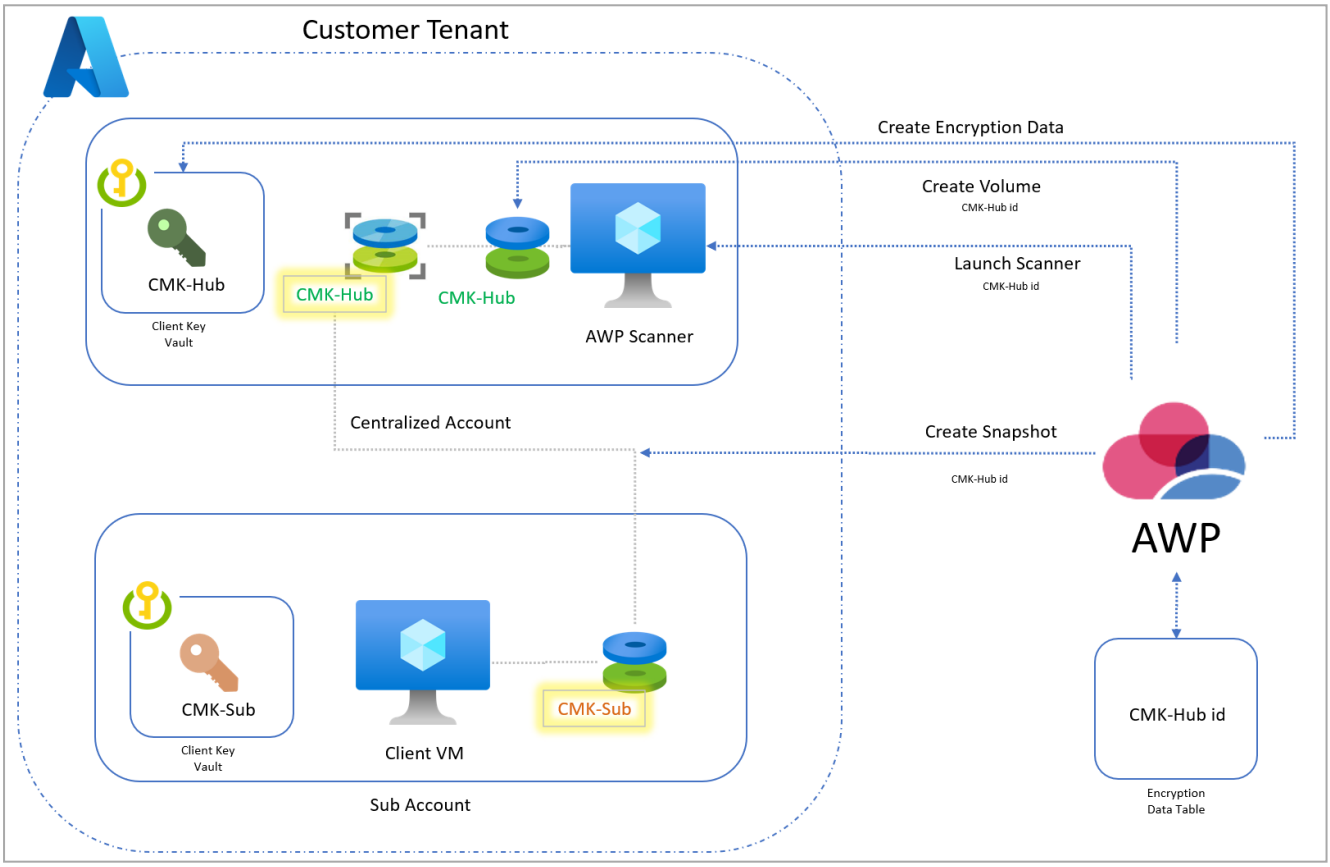

Note - CloudGuard skips scanning of VM instances with CMK-encrypted disks if you do not enable this option explicitly.

Scanning of workloads with multiple disks and mixed encryption is done based on the particular disk's encryption method:

- The disks encrypted with PMK (Platform-Managed key) are scanned with the centralized account's PMK.
- The disks encrypted with CMK are scanned with the dynamically generated AWP CMK.

To enable scanning of disks with SSE CMK encryption:

1. During the subscription onboarding to AWP (see *"Onboarding Workflow" on page 510*), in step 2, select the **Advanced** option, then select **Use SSE CMK encrypted disks scan**.
2. Optionally, edit the names of these roles:
   a. CloudGuard AWP Crypto Creator
   b. CloudGuard AWP Disk Encryptor

   CloudGuard adds the applicable parameters to the onboarding script.
3. Continue onboarding as usual.

To enable scanning of disks with SSE CMK encryption through API:

POST /workload/agentless/azure/accounts/<<ACCOUNT_ID>>/enableCentralizedAccount

```
{
   "agentlessAccountSettings": {
   "sseCmkEncryptedDisksScan": true  //default: false
   }
}
```

For more details, see the API Reference Guide

Scanning Workflow

To prepare your VM for AWP scanning:

1. With the App Registration, AWP gathers information about your VM disks.
2. With the App Registration, AWP remotely creates snapshots from the VM disks.
3. When all disks have their equivalent snapshots created, a scanner machine launches with the disks created from snapshots and performs the scanning.

   For **In-Account mode**, this VM runs in the **same region** as the original VM, in a custom VNet that AWP creates for this task.

   - The scanner outbound traffic is restricted by a security group rule that limits access exclusively to:
     - Azure AppService
     - Azure ManagementAPI
     - Azure Storage

    - Azure ResourceManager
  - No inbound rules are configured.

  To ensure that network traffic remains within the Azure backbone, the VNet utilizes Microsoft.Storage and Microsoft.Web service endpoints.

4. When the scan is complete, AWP sends a request to delete all the resources created for the scan.

To prepare your Function App for AWP scanning:

The AWP scanner downloads code.

Ignoring a VM Scan

If you need to deliberately skip scanning a VM, set a tag for the VM in the Azure portal.

To set a tag for the AWP scanner to ignore a Virtual Machine:

1. In the Azure portal, open your Virtual Machine.
2. Navigate to **Tags** and add a new tag with the name `CG_AWP_SKIP_SCAN` and the value `ANY`.

Offboarding AWP

To remove AWP from your Azure subscription, you must complete these two major steps:

I. Running a script that deletes all the resources that AWP created in your Azure account.

II. Deleting your account from the AWP database.

Before you remove AWP from a centralized account, make sure to remove it from all connected sub-accounts.

Running the script

Use the CloudGuard API to generate the offboarding script.

1. Generate the script with the **Get Azure Offboarding Data** API call. For more details, see the API Reference Guide
2. Open your account in Azure Cloud Shell or use AZ CLI in your terminal.
3. Paste the script created by the API engine. You can download the script to review or edit it based on your needs.
4. Run the script in the shell or terminal.

Deleting AWP from your account

Use the CloudGuard API to remove AWP. In the case of deleting AWP from a sub-account, this also deletes the resources associated with it in the centralized account.

1. Remove the account with the **Disable Azure Account** API call. For more details, see the API Reference Guide

Note - Deleting your Azure account from CloudGuard removes the account from AWP, which does not scan it anymore.

Creating a Dedicated VNet

During each scanning process, AWP creates a custom VNet in the region where your VM runs. You can use your own VNet (not recommended), for subscriptions onboarded to AWP with In-Account or In-Account Centralized mode.

To use your dedicated VNet, create all the required resources and follow these general instructions:

- **Virtual Network (VNet)** - Create a VNet in each region that has workloads to be scanned.

  Note - The VNet and other resources created for a Centralized subscription are used for all its sub-accounts.

- **Network Security Group** - Create a Network Security Group with specific security rules that allow outbound traffic to regional Azure Storage, Azure App Service, Azure API Management, and Azure Resource Manager. No inbound is required. If you do not include Function App scans, include only Azure Storage and omit other rules.
- **Subnet** - Create the necessary subnets in the VNet. Associate each subnet with the Network Security Group created earlier.
- **Custom tag** - Add a custom tag with the `CG_AWP_NETWORK` key (key: `CG_AWP_NETWORK` value: any) to the VNet and related resources (route table).
- **Storage account access** - Make sure you have access to a storage account.

Note - When you switch the VNet mode from `ManagedByAWP` to `ManagedByCustomer`, make sure to delete the existing VNet created by AWP in each applicable region.

Customer VNet API

You can indicate that your VNet is managed by you during the onboarding process or update the account settings after it is onboarded.

To add an independent account:

POST V2/workload/agentless/azure/accounts/{accountNumber}/enable

```
{
      "scanMode": "inAccount",
       "agentlessAccountSettings": {
            "inAccountScannerVPC": "ManagedByCustomer"
          }
}
```

For more details, see the CloudGuard API Reference Guide

To add a centralized account:

POST V2/workload/agentless/azure/accounts/{accountNumber}/enableCentralizedAccount

```
{
      "agentlessAccountSettings": {
           "inAccountScannerVPC": "ManagedByCustomer"
      }
}
```

For more details, see the CloudGuard API Reference Guide

To update account settings:

PATCH V2/workload/agentless/{provider}/accounts/{accountNumber}/settings

```
{
       "inAccountScannerVPC": "ManagedByCustomer"
}
```

For more details, see the CloudGuard API Reference Guide

More Links

- *"Agentless Workload Posture" on page 489*
- *"AWP for AWS Environments" on page 497*

Serverless Risk Assessment

CloudGuard Proact Serverless protection evaluates the risk in serverless functions in your AWS environments. CloudGuard scans and analyzes your functions and their dependent libraries for vulnerabilities, IAM permissions that are not necessary, and sensitive information such as passwords and keys. It then calculates a **Posture** score, based on the number, nature, and severity of vulnerabilities found, and generates alerts for each vulnerability, that show the specific issues and, in many cases, the actions necessary to remedy them.

To activate Proact Serverless risk assessment, your AWS environment must be onboarded to CloudGuard (see *"Unified Onboarding of AWS Environments" on page 54*). Enable Serverless protection (see *"Enabling Serverless Protection" on page 527*) if you skipped this step during the onboarding procedure.

Benefits

- Identify overly permissive IAM roles used by serverless functions
- Identify 3rd-party libraries for vulnerabilities
- Identify hard-coded credentials, secrets, and other sensitive information in serverless code
- Identify functions that are not used

Continuous Scanning and Analysis

CloudGuard Proact scans the functions in your cloud accounts when they are onboarded to CloudGuard. In addition, CloudGuard rescans functions when they are changed to provide a continuous and up-to-date risk assessment.

Posture Explorer

The Posture Explorer is a graphical view of the security posture of a serverless function, based on an analysis of the function.

Legend:

Item	Description
1	The serverless function
2	The cloud service types that can trigger the function
3	The service types that the function has permission to access

Scan in CI/CD

In addition, you can scan your functions in the CI/CD pipeline, before they are deployed to your cloud account, with the *"Serverless CI/CD Plugin" on page 563*. This runs as part of your CI/CD toolchain, scans for the same risks, and presents the results in the CI/CD tool. In addition, you can configure it to block the deployment of builds containing specific risks.

Finding Types

The table below lists the scan-finding types.

Types

Finding type	Description
Permissive Role	This Lambda function uses a role that has redundant permissions, which are not required by the function. Permissions that are not necessary increase the function attack surface, which can be leveraged by attackers to find sensitive data and possibly cause a takeover of all resources.
Vulnerable Dependency	This Lambda function uses a library that has known vulnerabilities.
Credentials Usage	This Lambda function has credentials hard-coded as part of the Lambda code or of the environment variables. Setting credentials hard-coded can cause a leak of credentials. Attackers can use this to find sensitive data and take over all the resources.
Unused Resource	This resource has not been in use for a while, which means that a Lambda function was inactive for more than 90 days. Serverless scanners use AWS CloudWatch metrics to verify the Lambda usage. Keeping unused resources in your account can increase the attack surface of your account.
Versions	This Lambda function has more versions than the maximum additional version limit. Some of them may not be in use.
Excessive Timeout	This function has a large timeout configured that is not necessary. Attackers can leverage the long execution of the function to cause more damage in case of a vulnerability in the function.

Finding type	Description
Obsolete Runtime	This Lambda function uses an obsolete runtime version. AWS support for this Lambda runtime has ended. The Lambda function no longer applies security patches or other updates to the runtime. The Lambda function does not block invocations of functions that use deprecated runtime versions. Function invocations continue indefinitely after the runtime version reaches the end of support. But Check Point strongly recommends that you migrate functions to a supported runtime version.

Actions

Enabling Serverless Protection

With serverless protection enabled:

- CloudGuard continuously scans your serverless functions for vulnerabilities and risks - **Serverless Risk Assessment**
- You can apply runtime protection to your functions when they are invoked - **AWS Serverless Function Runtime Protection**

Your AWS account must already be onboarded to CloudGuard. See *"Unified Onboarding of AWS Environments" on page 54* for details on how to do this.

To enable CloudGuard protection on your serverless functions, you must grant permission to CloudGuard to access these assets in your accounts. In addition, the permissions granted to CloudGuard in the account onboarding procedure. In the procedure described below, you use an AWS CloudFormation (CFT) stack, which you run in your account. To learn more about the CFT resources deployed on your account, see *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287*.

To enable Serverless Protection:

1. Navigate to **Assets** > **Environments**.
2. Click **Enable Serverless protection** for an AWS account from the list.
3. Follow the instructions on the wizard page that opens. Click **Create Cross-Account Role**.

   The prompt suggests you sign in to your AWS account and then it redirects you to the CloudFormation page. You can see a CFT stack that grants CloudGuard a cross-account role in your AWS account.

4. In the AWS console, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
5. Click **Create stack**.

   CloudFormation starts to create the stack.
6. In the AWS console, click the **Template** tab to view details for the permissions that CloudGuard obtains when the stack is created. After you create the stack, more permissions are granted to CloudGuard, and CloudGuard completes the procedure of enabling protection.

   When complete, the serverless functions appear in the CloudGuard portal in the protected assets list of the environment, on the **Protected Assets** page, as well as on the **Workload Protection** > **Serverless Functions** page.

Viewing Scan Results

You can see the Posture score for your functions and view findings generated by risks discovered in your functions.

To see scan results:

1. Navigate to **Workload Protection** > **Serverless Functions**. This page shows all the functions in your environments with enabled serverless protection. You can filter the list to show specific functions.
2. Select the function from the list. Select the **General** tab. You can see the details for this function. A Posture score shows the results of the scan.
3. Select the **Posture Findings** tab to see the findings generated from the scan with the **Serverless** source. The findings table is the same as the *"Aggregated Events" on page 120* table filtered for the specific Lambda.

Runtime Protection

Runtime Protection is available for AWS environments and Kubernetes clusters.

CloudGuard monitors AWS Lambda at runtime, checks their inputs and runtime behavior, and generates notifications for suspicious behavior. Additionally, you can apply Runtime Protection to AWS functions at the CI/CD stage on their deployment in your environment.

For Kubernetes clusters, Runtime Protection monitors the kernel system calls done by workload containers. Or you can configure CloudGuard to block unwanted, malicious, or anomalous activity that it discovers.

AWS Serverless Function Runtime Protection

You can apply CloudGuard protection to serverless functions at runtime. This protects functions from malicious inputs or attacks, while it monitors the function behavior for anomalous behavior and acts as a workload firewall for inputs from malicious sources. It does not change the source code of the function and has minimal impact on the function's runtime performance.

Before you apply Runtime Protection to the serverless functions, you must onboard the AWS environment to Serverless Protection, see *"Enabling Serverless Protection" on page 533* below.

With Serverless Function Runtime Protection, you can:

- Apply a workload firewall on inputs to the function, which analyzes the input payloads for malicious attack patterns.
- Detect and, optionally, prevent anomalous runtime behavior.
- Use a function-specific profile of actual function behavior, to build an allowlist of normative activity (baseline).
- Detect or prevent attacks based on anomalous behavior.
- Get visibility into what the application code is doing, including monitoring process launching, network activity, and API calls.
- Reduce the effect on function runtime performance.

How it Works

You have two options to apply Runtime Protection in your account:

- For each serverless function individually or at the account level:
  - To detect issues with **Auto Protect**
  - To detect and prevent issues with **Auto Protect** and **Block on detect**
- To AWS functions at the CI/CD stage on their deployment in your environment, with the CI/CD Plugin.

When you apply protection, CloudGuard adds a small module to your function that is loaded in runtime, along with the function. This module monitors your function, while it adds some runtime overhead. It is also fully transparent - all reporting is done with the function logs, so you can review the metadata it collects.

Runtime protection dynamically inspects different points in the flow of functions, with mechanisms such as pattern matching, flow analysis, "denylisting" and "allowlisting", and applies policies such as reporting and blocking in response to suspicious activity.

Best Practice - Check Point recommends to enable Serverless Runtime Protection gradually, in several stages.

1. Enable Runtime Protection on your QA (staging, sandbox) environment.
   - Run your automation tools and make sure everything works properly.
   - Let Runtime Protection work for a while, based on your needs, to create an allowlist.
   - Make sure the allowlist is accurate, there is no false positive.
2. Enable Runtime Protection on your production environment with the Detect mode. Make sure everything works properly.
3. Enable Runtime Protection on your production environment with the Detect and Prevent mode.

Allowlist

When you enable runtime protection for a serverless function, CloudGuard uses machine learning techniques to profile the function in runtime and to create an "Allowlist" of its normative (permitted) actions. This includes:

- Processes
- Files and local storage accessed
- API functions (some calculated by a code analysis, some by runtime monitoring, which may include cases of code injection)
- External hosts accessed
- Network addresses communicating with the function

The **Runtime Protection** tab shows the Allowlist for a function. You can manually add activities to the Allowlist when you configure exclusions (*"Creating Serverless Functions Exclusions" on page 539*) or remove actions when you configure rules (*"Creating Serverless Functions Rules" on page 540*).

Rules and Exclusions

You can manually add or remove actions to the Allowlist. In this procedure, you can adjust the Allowlist that is automatically created when the function is profiled.

Rules remove actions from the allowlist, and exclusions add actions.

The **Rules & Exclusions** tab shows the rules and exclusions configured for a function.

Note - These rules and exclusions apply to the serverless function only. The rules and exclusions you configure on the *"Cloud Security Posture Management (CSPM)" on page 302* level apply to all functions and the environment.

Events

CloudGuard creates an event notification when it detects anomalous runtime behavior or malicious inputs. You can see these notifications for each function individually or on the *"Events" on page 119* page, together with notifications from other functions and other CloudGuard sources.

The table below lists the runtime alert messages created by Serverless Runtime Protection.

Alert Types

Alert Name	Description
Command Injection Attack	In Command Injection, the attacker extends the default functionality of the application, which runs system commands, without the necessity of injecting code. If successfully exploited, the effect could include loss of confidentiality, loss of integrity, loss of availability, and/or loss of accountability.
Code Injection Attack	Someone tried to run arbitrary code in the function. A successful attack can compromise the application's confidentiality, integrity, availability, or loss of accountability.
CSV Injection Attack	A successful attack can hijack the user's computer by exploiting vulnerabilities in spreadsheet software, such as CVE-2014-3524. Or it can hijack the user's computer by exploiting the tendency to ignore security warnings in spreadsheets downloaded from their website or exfiltrate contents from the spreadsheet, or other open spreadsheets.
Local File Inclusion Attack	Someone tried to access a file that the function did not initially intend to access. A successful attack usually results in sensitive data leakage. In addition, it can cause code execution through a file that contains attacker-controlled data, compromising your application confidentiality, integrity, or availability.
NoSql Injection Attack	Someone tried to change the intended logic of the query. A successful NoSQL injection exploit can have an effect on confidentiality, availability, and integrity by reading and changing database data.
SQL Injection Attack	Someone tried to change the initial SQL query of the function. A successful SQL injection exploit can read sensitive data from the database, change database data (Insert/Update/Delete), run administration operations on the database (such as shutdown the DBMS), recover the content of a given file that exists on the DBMS file system and in some cases issue commands to the operating system.

Alert Name	Description
XSS Injection Attack	Someone tried to run malicious JavaScript code. A successful XSS attack can access any cookies, session tokens, or other sensitive information retained by the browser and used with that site.
XML External Entity Attack	Someone tried to access a file or leak its data through a malicious XML payload. A successful attack usually results in sensitive data leakage, communication in the VPC, and more.
File Event	File access detected by CloudGuard as not approved.
Process Event	A production process detected by CloudGuard as not approved.
API Event	AWS API call detected by CloudGuard as not approved.
Host Event	Host call detected by CloudGuard as not approved.
IP Event	IP call detected by CloudGuard as not approved.
HTTP Event	HTTP call detected by CloudGuard as not approved.

Actions

Enabling Serverless Protection

With serverless protection enabled:

- CloudGuard continuously scans your serverless functions for vulnerabilities and risks - **Serverless Risk Assessment**
- You can apply runtime protection to your functions when they are invoked - **AWS Serverless Function Runtime Protection**

Your AWS account must already be onboarded to CloudGuard. See *"Unified Onboarding of AWS Environments" on page 54* for details on how to do this.

To enable CloudGuard protection on your serverless functions, you must grant permission to CloudGuard to access these assets in your accounts. In addition, the permissions granted to CloudGuard in the account onboarding procedure. In the procedure described below, you use an AWS CloudFormation (CFT) stack, which you run in your account. To learn more about the CFT resources deployed on your account, see *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287*.

To enable Serverless Protection:

1. Navigate to **Assets** > **Environments**.
2. Click **Enable Serverless protection** for an AWS account from the list.
3. Follow the instructions on the wizard page that opens. Click **Create Cross-Account Role**.

   The prompt suggests you sign in to your AWS account and then it redirects you to the CloudFormation page. You can see a CFT stack that grants CloudGuard a cross-account role in your AWS account.
4. In the AWS console, select the option **I acknowledge that AWS CloudFormation might create IAM resources with custom names**.
5. Click **Create stack**.

   CloudFormation starts to create the stack.
6. In the AWS console, click the **Template** tab to view details for the permissions that CloudGuard obtains when the stack is created. After you create the stack, more permissions are granted to CloudGuard, and CloudGuard completes the procedure of enabling protection.

   When complete, the serverless functions appear in the CloudGuard portal in the protected assets list of the environment, on the **Protected Assets** page, as well as on the **Workload Protection** > **Serverless Functions** page.

Enabling Runtime Protection in the CI/CD Pipeline

To enable Runtime Protection in the CI/CD pipeline, before functions are deployed to your cloud account, see *"Serverless CI/CD Plugin" on page 563*. The pipeline plugin checks the security policy to calculate which functions are configured to have the runtime protection module inserted into them. The insertion is transparent to function behavior and execution.

Enabling Runtime Protection at the Account Level

You can enable Runtime Protection at the account level, so CloudGuard automatically adds the protection layer to all existing and new deployments of the functions found in the account.

To enable Account Auto Protect:

1. Navigate to **Assets** > **Environments**.
2. Select the environment to enable Auto Protect.
3. Select the **Serverless** tab.

4. Below **Account Auto Protect**, in **Runtime Protection activation**, click the toggle button. The **Enable Auto Protect Mode** window opens.
5. Select the applicable version of the FSP.
6. Click **Enable**.

To apply Runtime Protection to selected serverless functions, see *"Serverless Functions" on page 536*.

Viewing AWS IAM Permissions

On the **IAM Policies** tab of the **Protected Assets** page, you can see the results of the function code scan. It can find different problems of different severity (critical, high, medium, or low). Click the **Suggested Policy** link to see the suggested policy remediation that you can copy to your clipboard.

Viewing Serverless Runtime Findings

Navigate to **Assets** > **Protected Assets**, select a function and go to its **Events** tab. The **Events** (**All**) page shows the table of findings from all CloudGuard features. Go to the **Threat & Security** page to show the table of security findings detected by Runtime protection. (For more information, see *"All Events" on page 120*).

Viewing Serverless Runtime Vulnerabilities and SBOM

Navigate to **Assets** > **Protected Assets** and select a Lambda function. To view its vulnerabilities, navigate to the **Vulnerabilities** tab. To view the SBOM, navigate to the **SBOM** tab.

Disabling Serverless Protection

When you disable Serverless Protection for your environment, you must delete the Cross-Account role created on your account. This procedure removes the stack created with the CloudFormation Template.

1. Navigate to **Assets** > **Environments**.
2. Select the environment to disable serverless protection.
3. Select the **Serverless** tab.
4. Scroll down to **Serverless Disabling**, click **Disable Serverless** and then click **Delete Cross Account Role**.

   CloudGuard redirects you to the cloud account on AWS, where you can delete the CFT stack.

Serverless Functions

To see your serverless functions, you must onboard the environment that contains these functions to CloudGuard. See *"Onboarding AWS Environments" on page 144* to onboard your environment.

When you enable Serverless Protection on your cloud environments, you can see all the functions that exist in these environments and their protection status on **Workload Protection >Serverless > Serverless Functions**.

Use the *"Filter and Search" on page 863* toolbar to select parameters to filter out the serverless functions with Runtime Protection enabled (**Protected**) or disabled, **Auto Protect** (Detect) enabled, or **Protection** mode set to **Detect** or **Prevent**.

The Serverless Functions page allows you to see the protection status of all functions:

- **Runtime** - Runtime language or framework.
- **Runtime Protection** - Shows **Protected** when Serverless Protection is enabled and the Cross Account Stack is updated.
- **Auto Protect** - Shows Auto Protect when enabled.
- **Protection Mode** - Shows **Detect** or **Prevent** (Block on detect) protection mode.
- **Learning** - Shows the progress of the profile learning to build the Allowlist.
- **FSP Version** - Shows the current FSP version.

Click the function name to see more details about its status, permissions, and posture findings.

Feature Status

The **General** tab of the onboarded serverless function shows its feature status. To learn more about each feature, read its tooltip information.

Feature Card Name	Feature	Status
Configuration Scanning	Posture Management is enabled after account onboarding to CloudGuard. Compliance engine scans the serverless function configuration	▪ **Active** - Always enabled
Vulnerability Scanning	Serverless Protection (Proact) scans serverless functions for known vulnerabilities and embedded secrets	▪ **Disabled** ▪ **Active** - You enabled Serverless Protection
IAM Hardening	Serverless Protection (Proact) does Deep Code Flow Analysis for application hardening and least privilege access	▪ **Disabled** ▪ **Active** - You enabled Serverless Protection

Feature Card Name	Feature	Status
Workload Firewall	FSP (Runtime Protection) validates workload runtime input	▪ **Disabled** ▪ **Active** - You enabled Runtime Protection
Behavioral Prevention	Behavioral Intrusion Prevention (Runtime Protection) learns specific workload behavior profile to detect and prevent anomalous behavior	▪ **Disabled** ▪ **Learning is progress** - You enabled Runtime Protection, which now builds the behavioral profile ▪ **Active** - You enabled Runtime Protection, and the behavior profile (Allowlist) is built

Actions

Enabling Runtime Protection on a Serverless Function

To enable Runtime Protection:

1. Navigate to **Workload Protection** > **Serverless > Serverless Functions**.
2. Select the function from the list (you can filter the list to narrow your search) and select the **General** tab.
3. Set **Auto Protect** to **ON**. CloudGuard starts to profile the function to build an Allowlist of normative activity. The results of profiling show below in the **Runtime Protection** tab.
4. In addition, you can see the version of the FSP (Function Self Protection) that runs on your function. If the arrow appears in the left upper corner, you can upgrade the installed FSP to a higher version. Click the arrow and follow instructions. Click the menu ⋮ button to change the version. You can select one of the available versions.

Enabling Runtime Protection for all Serverless Functions

You can enable Runtime Protection for several or all selected serverless functions.

1. Navigate to **Workload Protection** > **Serverless** > **Serverless Functions**.
2. Select one or more filtering criteria.

3. Click the **Select All Visible** check box in the table header to select all the shown functions.
4. On the toolbar, click **Auto Protect**. The **Modify Auto Protect Mode** window opens.
    a. Select to **Enable** the Auto Protect mode.
    b. Select the latest or other version of the FSP.
    c. Click **Apply**.

    CloudGuard applies protection to all selected serverless functions.

Enabling Runtime Protection on an Account Level

You can enable Runtime Protection on an Account level. This applies protection to all existing serverless functions, as well as the future functions on this account.

1. Navigate to **Workload Protection** > **Serverless | Serverless Functions**.
2. Select one or more filtering criteria.
3. Click the **Select All Visible** check box in the table header to select all the shown functions.
4. On the toolbar, click **Auto Protect**. The **Modify Auto Protect Mode** window opens.
    a. Select to **Enable** the Auto Protect mode.
    b. Select the latest or another version of the FSP.
    c. Click **Apply**.

    CloudGuard applies protection to all selected serverless functions.

Blocking Actions on Detect

When runtime protection is enabled, CloudGuard creates a runtime alert when the behavior of a function deviates from the allowlisted behavior detected during the profiling, or when inputs are received from known malicious sites.

In addition, you can configure CloudGuard to block the action, such as an attempt to access a file that is not on the allowlist, or the input.

To block actions for a function:

1. Navigate to **Workload Protection** > **Serverless Functions**.
2. Select the function from the list (you can filter the list to narrow your search) and then select the **General** tab.
3. Set **Block on detect** to **ON**.
4. In the **Defense Mode** window, click **Enable Blocking**.

To block actions for all functions:

1. Navigate to **Workload Protection** > **Serverless | Serverless Functions**.
2. Select one or more filtering criteria.
3. On the table header, click the **Select All Visible** checkbox.
4. On the toolbar, click **Protection Mode**. The **Modify Detect Mode** window opens.
    a. Select **Prevent and detect** to block the actions.
    b. Click **Apply**.

Creating Serverless Functions Exclusions

You can manually not include specific criteria not to block the activity of a function (or to cause an event if the function is not configured to block on detection). This allows you to adjust the criteria that Serverless Runtime Protection uses to monitor the runtime activities of the function. These exclusions stay until you change or remove them and apply each time the function is invoked.

To create an exclusion:

1. Navigate to **Workload Protection** > **Serverless** | **Serverless Functions**.
2. Select the function from the list (you can filter the list to narrow your search).
3. Select the **Rules & Exclusions** tab.
4. Expand the **Exclusions** section and click **Create New Exclusion**.
5. Enter a **Name** for the exclusion (as it shows in the list of exclusions in this tab).
6. Select a **Target**. The Target is the type of action that is not included in the monitoring, such as a procedure, a file, or a specific host. Select one target for the exclusion from the list:
    - AWS API
    - Network
    - Process
    - File
    - Input

    To select more than one target for an exclusion, create a different exclusion for each target.

7. Enter the **Pattern** for the exclusion. The pattern is a text string or list of strings, each on a different line. Actions that match the pattern are added to the function allowlist and excluded from the generation of an event notification or from the blocked actions. The pattern corresponds to the target type. For example, for an IP target, the pattern is one or more IP addresses.

   Note - The *File* pattern match differs from the *Host* or *Process* pattern match. To match a `*.TXT` file in `/tmp` and all its subdirectories, use this pattern: `/tmp/**/*.txt`
   If you want to match the `*.TXT` file in only `/tmp/logs`, without subdirectories, use this pattern: `/tmp/logs/*.txt`
   For more examples, see: https://help.sumologic.com/03Send-Data/Sources/04Reference-Information-for-Sources/Using-Wildcards-in-Paths

8. Select the **Scope** for the exclusion. This indicates if the exclusion applies to a specific function, a group of functions, or all functions in the environment (account).

   Note - The note after the scope *Apply only on protected functions (functions with FSP)* means that you can apply the scope when the *"Enabling Runtime Protection on a Serverless Function" on page 537*, that is, the Auto-Protect button is ON.

9. Click **Create**.

Creating Serverless Functions Rules

You can manually add specific criteria to block the activity of a function (or cause an alert if the function is not configured to block on detection). This allows you to adjust the criteria that Serverless Runtime Protection uses when monitoring the runtime activities of the function. These rules stay in place until they are changed or removed, and are applied each time the function is invoked.

To create a rule:

1. Navigate to **Workload Protection** > **Serverless Functions**.
2. Select the function from the list (you can filter the list to narrow your search).
3. Select the **Rules & Exclusions** tab.
4. Expand the **Rules** section and then click **Create New Rule**.
5. Enter a **Name** for the rule (as it shows in the list of rules in this tab).
6. Select a **Target**. The Target is the type of action, such as a procedure, a file, or a specific host. Select one target for the rule. To select more than one target for a rule, create a different rule for each target.

7. Enter the **Pattern** for the rule. The pattern is a text string, or list of strings, each on a different line. Actions that match the pattern are removed from the function allowlist and generate an event notification, or they are blocked. The pattern corresponds to the target type. For example, for an IP target, the pattern is one or more IP addresses.

   Note - The *File* pattern match differs from the *Host* or *Process* pattern match.
   If it is necessary to match a `*.TXT` file in `/tmp` and all its subdirectories, use this pattern: `/tmp/**/*.txt`
   To match the `*.TXT` file in only `/tmp/logs`, without subdirectories, use this pattern: `/tmp/logs/*.txt`
   For more examples, see: https://help.sumologic.com/03Send-Data/Sources/04Reference-Information-for-Sources/Using-Wildcards-in-Paths

8. Select the **Scope** for the rule. This indicates if the rule applies to a specific function, a group of functions, or all functions in the account.

   Note - The note after the scope *Apply only on protected functions (functions with FSP)* means that you can apply the scope when the *"Enabling Runtime Protection on a Serverless Function" on page 537*, that is, the Auto-Protect button is ON.

9. Click **Create**.

More Links

- *"Serverless Risk Assessment" on page 525*
- *"AWS Serverless Function Runtime Protection" on page 530*
- *"AWS Resources and Permissions for Serverless Runtime Protection" on page 287*

AWS Runtime Protection Implementation

This document describes how CloudGuard Serverless Runtime Protection is implemented. It explains what permissions and resources are used by CloudGuard and the effects they have on the runtime performance of serverless functions.

In addition, it explains how the Serverless Runtime Protection is added to a serverless function as an AWS Layer, and what it does at runtime.

Serverless Implementation

CloudGuard Serverless runtime protection is implemented as an AWS layer that is added to each function that is to be protected in runtime. This can be done in the build stage with the Serverless CI/CD Tools, or after the functions are deployed to an AWS account (that is onboarded to CloudGuard). This is done with the CloudGuard portal or API.

The layer communicates with the CloudGuard backend in runtime through these procedures:

- Use a dedicated *S3 bucket* to receive information from the backend. This S3 bucket is created when the Serverless Protection is enabled for an account (after the account is initially onboarded to CloudGuard). This is done through the CloudGuard portal. The bucket serves all functions in the account (each function has a different folder for communication from the backend to the layer).
- Use the *function log group*, to send information to the backend. The layer writes log messages to the log group, in the same way, and to the same group, as the function itself. The layer then parses the log group for its messages and forwards them to the CloudGuard backend.

Serverless Runtime Protection Module

Serverless Runtime Protection is added to each protected function as an AWS layer.

The layer does these activities:

- Monitors and blocks (if runtime protection policy prescribes so) function activities.
- Reads on the CloudGuard backend the runtime protection policy from a folder in an S3 bucket created in the user's AWS account.
- Writes information on the function log group as log strings, which is then forwarded to the CloudGuard backend.

Permissions

The Serverless Runtime Protection uses the same IAM permissions as the function to which it is added.

When the AWS account is enabled for serverless protection in CloudGuard, cross-account IAM permissions are added to all functions to let the module connect with the S3 bucket and the log group.

Instrumentation of the Serverless Runtime Protection Module

Instrumentation is the procedure of adding the Serverless Runtime Protection as a layer in the function. Instrumentation is done on each function individually. Runtime protection is enabled on each function individually with the CloudGuard portal or API.

In the instrumentation procedure, these activities are done:

1. The function configuration is changed to include the Serverless Runtime Protection (this change depends on the language of the function).
2. The Serverless Runtime Protection is copied from the CloudGuard AWS account and added to the function.

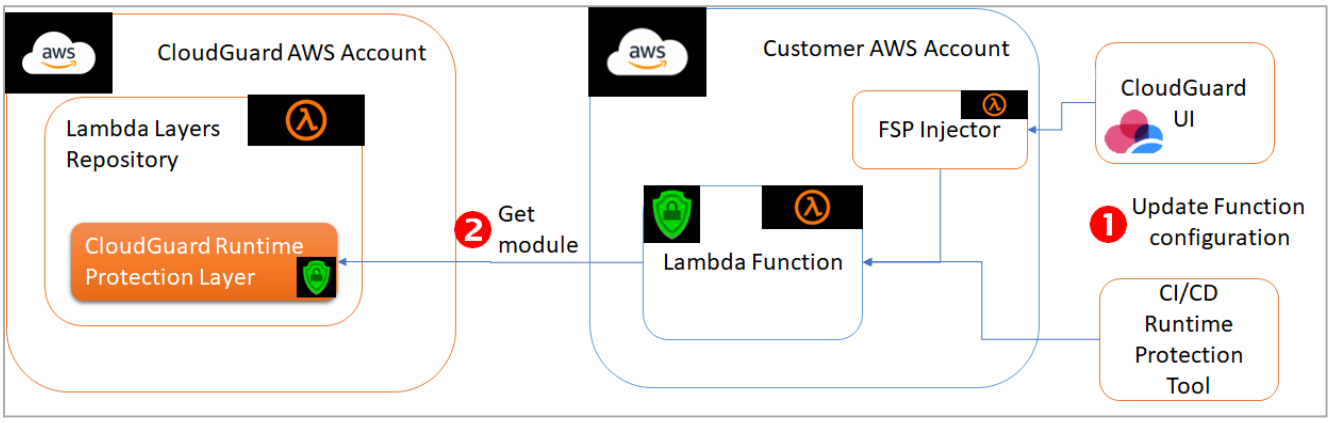

Function Languages and Layer Sizes

These languages and frameworks are supported for runtime serverless protection:

- Python
- Node.js
- Java
- .NET Core

The Serverless Runtime Protection layer has different sizes.

Runtime	Layer Name	ZIP Layer Size (MB)	Extracted Layer Size (MB)
.NET Core	cloudguard-fsp-csharp-layer	1.8	6.6
Node	cloudguard-fsp-nodejs-layer	4.1	14.6
Python	cloudguard-fsp-python-layer	3.2	10.8

Runtime	Layer Name	ZIP Layer Size (MB)	Extracted Layer Size (MB)
Java	cloudguard-fsp-java-layer	4.3	11

Using Serverless Runtime Protection

When you complete the instrumentation stage, the Serverless Runtime Protection can be added to the function at these times:

- In the CI/CD pipeline, with the Serverless CI/CD Runtime Protection Tool
- When the function is deployed in the cloud account with the Enable Runtime Protection function from the CloudGuard portal (Serverless assets page) or API

These two methods invoke a Serverless Runtime Protection injector function, which is deployed in the customer AWS account when Serverless protection is enabled.

Runtime Activity

The Serverless Runtime Protection is run each time the function is invoked.

The module operates in two modes:

- Profiling function activity
- Monitoring/enforcing runtime protection

These modes are described below.

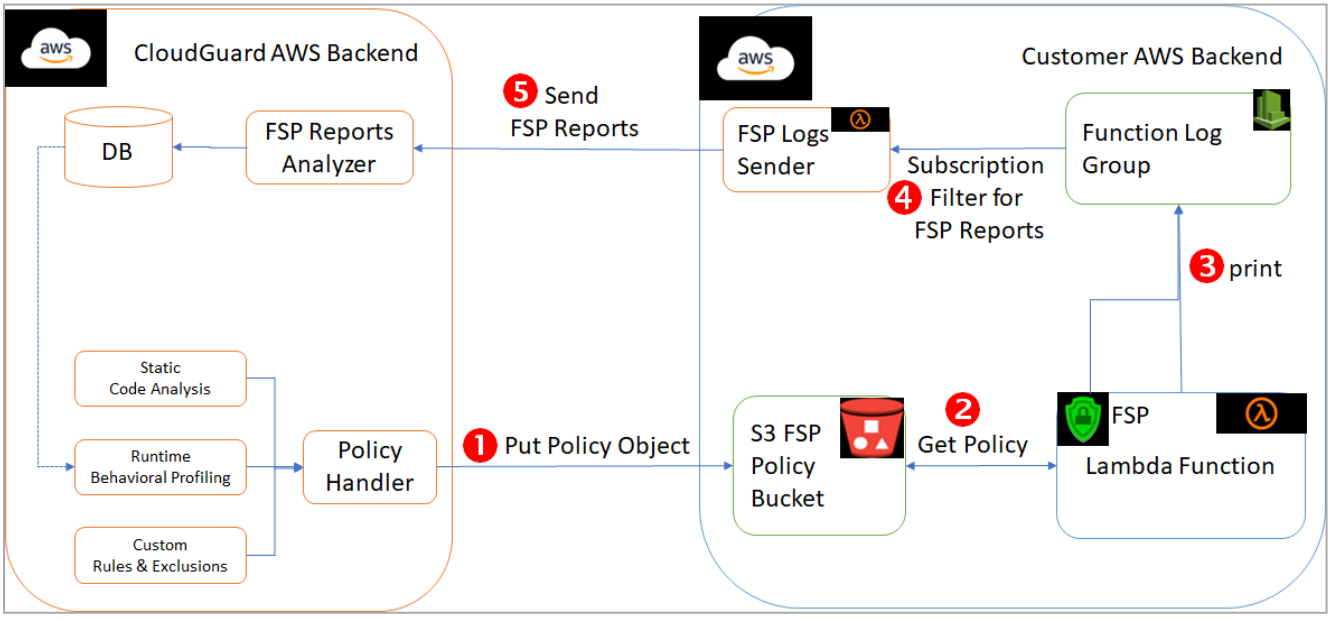

Profiling (allowlist)

The Serverless Runtime Protection monitors all function activities and writes this information as log group messages **[3]**, which are then forwarded to the backend **[5]**. The backend analyzes this information and creates an allowlist that represents the usual, baseline activity for the function (the user can view this allowlist on the CloudGuard portal and change it by adding exclusions).

This profiling continues until the backend writes an allowlist to the S3 bucket **[1]** and instructs the Serverless Runtime Protection to start enforcing it **[2]**. The profiling period is typically when the function is first invoked after being deployed, or after it is changed.

Monitoring and Runtime Protection Enforcement

When instructed by the backend (with the use of the S3 bucket) **[2]**, the Serverless Runtime Protection starts to enforce the allowlist.

It monitors activities done by the function and checks for actions not on the allowlist. It reports these events as log group messages **[3]** (which then become CloudGuard Alerts). It can be configured to block these actions. This is calculated by user settings when runtime protection is applied to the function.

The module monitors these actions of the function in runtime:

- Inputs
- File access
- Host connections
- API actions
- Processes that the function launches

Runtime Protection Policy (Allowlist)

The module receives information from the CloudGuard backend that shows which of the above actions are usual (allowlist) actions. The module monitors the actions of the function

Users can change the allowlist by creating exclusions, which add actions to the allowlist.

Connection to CloudGuard Backend

The module communicates with the CloudGuard backend to receive profile information and to send log messages for profiled or monitored activities.

For communications from the backend, an S3 bucket is used (with a dedicated folder for each function)

For communications to the backend, the Serverless Runtime Protection sends log messages to the function log group.

Runtime Performance

Serverless Runtime Protection has a minimal effect on the runtime performance of the function. The table below summarizes the effect of different actions taken by the module in terms of time and memory.

The latency and memory overheads for the actions represent the average time for the performance of one action, for example, one file access. FSP Init represents the overhead that the module introduces on each execution regardless of the action done by the Lambda function. The overheads of File, Process, and AWS API include the FSP Init overhead.

The table below shows the average reading captured for over 1000 invocations of a Python 3.8 function with and without the Runtime Protection module.

Action	Latency (ms)	Memory Overhead (MB)
Cold Start	1	11
FSP Init	0.22	0.05
I/O	0.38	0.39
File access	0.33	0.53
Process launch	5.27	0.19
Network connection	0.05	0.01
AWS API	0.57	0.28

Data Handling

Data Collection

CloudGuard collects data about AWS serverless functions at two points in the function's use:

- *When the function is scanned, before deployment* - CloudGuard collects data about the function when it is built before it is deployed to the cloud account.
- *At runtime* - The Serverless Runtime Protection collects information about actions the function does in runtime.

Function Scan Data

CloudGuard scans serverless functions in your AWS accounts (that are onboarded to CloudGuard and which have Serverless protection enabled). This occurs when you apply to them serverless protection, with the CloudGuard web interface or API, or when somebody changes the function code.

The scan uses the AWS *GetFunctionConfiguration* API method to get information about the serverless function. This method returns information about the function, but not the function source code, which is not collected by CloudGuard.

Collected Information

The function scan gets this information about the function:

- Inventory of functions
- Environment variables used by the functions
- Runtime environments used by the functions
- IAM roles used by the functions
- resource-based policies applied to the functions

Code Scan

CloudGuard scans the code on serverless functions in your AWS accounts that are onboarded to CloudGuard, and have Serverless protection enabled. This occurs when the function is deployed or when you apply to it serverless protection, with the CloudGuard web interface or API, or when somebody changes the function code.

CloudGuard does the code scan in the user's AWS account with functions deployed by CloudGuard in the user account as part of the procedure of enabling Serverless protection. The function source code is not exposed outside of the user's AWS account. The functions that do the scan send the results of the scan to the CloudGuard backend. This information is used to prepare a risk assessment for the function (for example, to show where there are too many IAM permissions given to the function).

Collected Information

The code scan checks the function source code for these events:

- API calls used by the function
- Dependency list of libraries and other modules
- Hard-coded credentials, such as passwords and keys

CloudGuard uses this information to prepare the risk assessment for the function. The scan does not send to CloudGuard actual user data, such as API payloads, values of hard-coded passwords, or keys.

Serverless Runtime Protection Data

At runtime, the Serverless Runtime Protection parses the function log group and extracts information about the actions that the function does as follows:

- Inputs
- File access
- Host connections
- API actions
- Processes that the function launches

It sends the extracted information to the CloudGuard engine by HTTPS. In addition, the module shows if it blocks the action or only detects it. You determine the preference in the CloudGuard web interface or API when it is necessary to apply Serverless Runtime Protection to the function.

Data in Motion

Information sent to the CloudGuard engine from the function (which the Serverless Runtime Protection does) is written to the function log group and, from there, fetched (by a function deployed in the user account) and sent to CloudGuard by HTTPS.

Data at Rest

Information that CloudGuard gets about serverless functions from function and code scans or runtime monitoring is stored in a database in the CloudGuard engine in the CloudGuard AWS account.

Information in this database is encrypted at rest.

Data Privacy

CloudGuard protects information collected about and from serverless functions based on Check Point Privacy Policy.

Kubernetes Runtime Protection

Kubernetes Runtime Protection allows CloudGuard to monitor containers in real time and prevent runtime threats that can compromise the user's live environment. The Runtime Protection mechanism combines several engines that monitor kernel system calls, file access, and network activity, and protects against possible security threats.

Prerequisites

Your Kubernetes cluster must already be onboarded to CloudGuard. See *"Onboarding Kubernetes Clusters" on page 188* for details on how to do this.

Runtime Protection Daemon requires the installation of Linux kernel headers on each node before the deployment.

To install the kernel headers, use the commands below:

Operating System	Supported Version	Command
Ubuntu / Debian		`apt-get install -y linux-headers-$(uname -r)`
CentOS / RHEL		`yum -y install kernel-devel-$(uname -r)`
▪ Amazon Linux ▪ Google COS ▪ AWS Bottlerocket (Signature, Profiling) ▪ RHCOS (OpenShift) ▪ AKS Ubuntu	Linux kernel v4.14 and higher	*No installation required*
▪ AWS Bottlerocket (File Reputation)	Linux kernel v1.11.1 and lower	*No installation required*

For Kubernetes Runtime Protection minimum requirements, see *"Supported Versions" on page 415*.

How it Works

Architecture

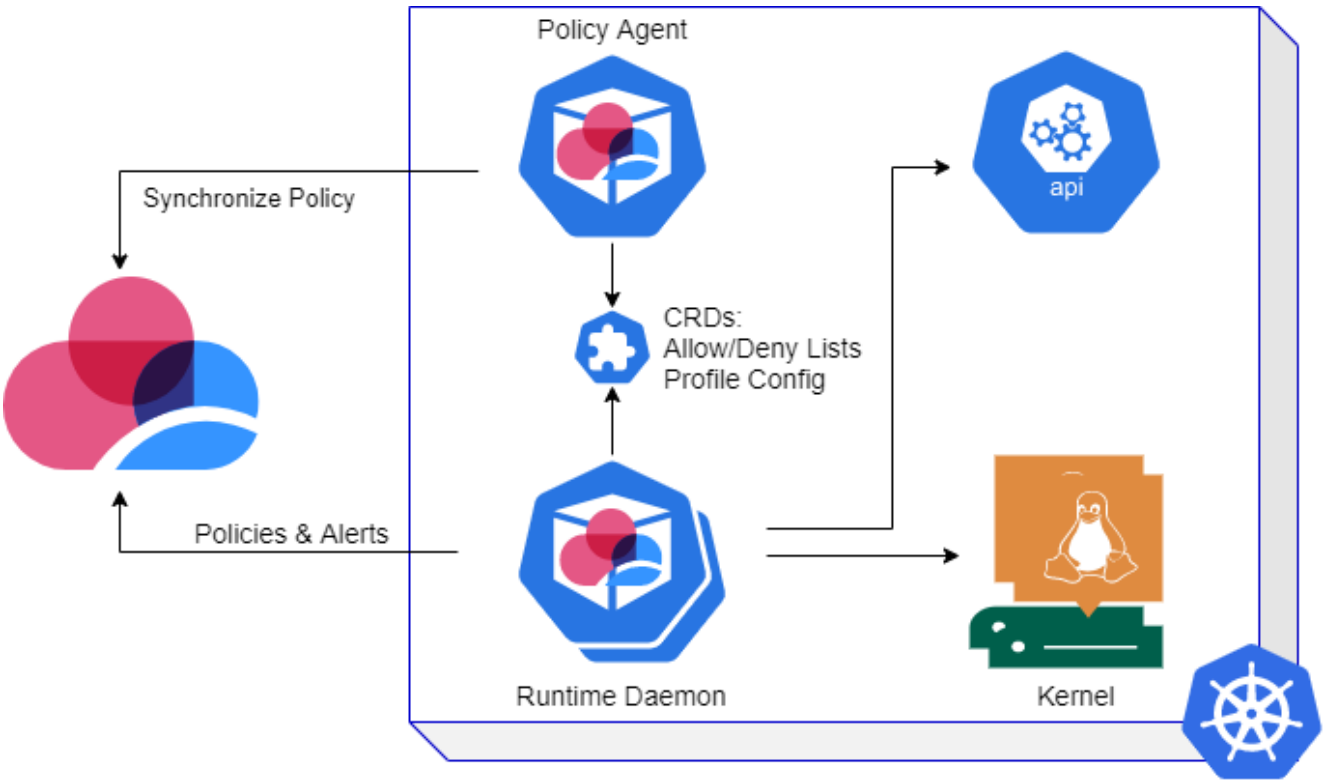

Resources

Runtime Protection includes these resources:

- **Runtime Protection Policy Agent** - A single-replica Deployment responsible to retrieve CloudGuard configuration and policies and keep them on the cluster.
- **Runtime Protection Daemon** - A DaemonSet running the Runtime Protection engines that interact with the underlying node to monitor the container activity.
- **Custom Resource Definitions (CRD)** - Kubernetes custom resources (CR) for Runtime Protection policies. The CRD is managed by the Runtime Protection Policy agent. The CR is only accessible to the Runtime Protection Policy agent (`modify` and `read` permissions) and the Runtime Protection Daemon agents (`read` permissions). Access to the CR, which include `read` permissions like `get` and `list`, should be restricted because it contains the cluster security configuration.
- **Inventory Agent** - A single-replica Deployment responsible to report inventory information on the cluster resources to CloudGuard.

Pod Groups

Runtime Protection refers to all pods of the same owner as members of a **Pod Group**. The Pod Group is the basic entity to which the Runtime Policy applies. The relationship between dependent resources appears in the *ownerReferences* property of all workload resources. A Pod without an owner is called an **Independent Pod**. An Independent Pod is considered its own Pod Group. In general, each entity in an environment can be considered a member of the Pod Group represented by the highest root entity above it.

All Runtime Protection rules, exclusions, and profiles that are managed on the Pod Group apply to all the pods in the group. For example, a new rule created in the context of a deployment Pod Group applies to all pods within this deployment.

Independent pods that you explicitly provision differ from the pods that a Kubernetes controller creates as part of the life cycle of another (parent) resource (a Pod Group).

Independent Pod Groups are:

- Kubernetes Deployments
- Kubernetes DaemonSets
- Kubernetes ReplicaSets, unless it is a dependent resource
- Kubernetes Pods, unless it is a dependent resource

Note - Pods derived from other controllers, such as Jobs, are considered independent.

Signatures, File Reputation, and Profiling

Kubernetes Runtime Protection has three primary engines:

- **Signatures** (Limited General Availability) - Compares the observed behavior of a workload with known signatures that potentially show malicious behavior. For example, the execution of processes related to crypto-mining software.
- **File Reputation** (Limited General Availability) - Detects in real time malicious access to executable files with Check Point Reputation Service.
- **Profiling** (Public Preview) - Detects anomalies in behavior compared to a baseline profile created during a profiling phase.

Signatures and File Reputation are enabled by default for all Pod Groups.

Signatures

The Signatures engine monitors the workloads' system calls and compares them to a predefined set of signatures corresponding to potentially malicious behavior. It also prevents possible threats and when it finds a malicious signature, it kills the main process of the container. The Check Point research team dynamically updates the signatures in CloudGuard to address emerging threats and vulnerabilities.

Examples signatures:

- Modification of `/etc/passwd` file
- Execution of a crypto miner binary

Note - For testing, it is possible to create a simulated signature event. For this, create an executable file named `eicar` and run the shell command:
`<path to eicar file> --cptest`

File Reputation

The File Reputation engine monitors real-time executables that run on the workload. With Check Point ThreatCloud solution, the engine detects and prevents the execution of malicious files.

Note - For testing, it is possible to create a simulated file reputation event. For this, execute the commands below with a Bash shell (for example, BusyBox) in one of the pods:

```
echo -n 'X5O!P%@AP[4\PZX54(P^)7CC)7}$EICAR-STANDARD-ANTIVIRUS-TEST-
FILE!$H+H*' > eicar
chmod +x eicar
./eicar
```

After successful simulation, the prompt shows syntax error messages, for example:

```
./eicar: line 1: syntax error near unexpected token `P^'
./eicar: line 1: syntax error: unexpected word (expecting ")")
```

Profiling

The Profiling engine works in two phases:

1. Profile Learning
2. Profile Enforcement

During the first phase, the Profiling engine learns the behavior patterns of the workloads that run in a given cluster and builds their profile. The profile (allowlist) is a policy that describes what this workload **can** do: which processes it **can** run and which domains it **can** access. When the learning is complete, the second phase begins. CloudGuard can detect violations and trigger alerts when one of the containers in a Pod Group performs an operation or accesses a network not observed during the profiling phase.

Profile Learning

For each node, the CloudGuard agent monitors the system calls that containers make in the node. When a new Pod Group is deployed in a cluster, CloudGuard monitors all the processes and network activity on all the Pod Groups until the profiling period ends. By default, the profiling period is 24 hours. It starts from the moment when the CloudGuard agent discovers a new Pod Group deployment or when Runtime Protection is enabled for existing Pod Groups. At the end of the period, the agent creates a security profile (baseline) of the Pod Group.

To see the learned profile:

1. Open a container with enabled Runtime Protection (see *"Enabling Runtime Protection on a Cluster" on page 557*).
2. In the **Protected Assets** tab, expand the **Kubernetes Deployment** group.
3. Find the deployment that contains *runtime-policy* in its name and click it.
4. The **Runtime Protection** tab contains the profile details in the **Processes** and **Network** sections. Each process and network learned at this stage has the **Runtime** tag in its line.

- **Processes** - CloudGuard designates each allowed process with a combination of the process path and the path of the Parent process (the one that initiated it). At the policy enforcement stage, CloudGuard does not allow a process if a different Parent process initiates it.
- **Network** - A combination of the allowed network domain and the parent process.

Notes:

- Profiling can fail if the number of observed behavior events, processes, or hosts is above 100. In this case, CloudGuard cannot enforce protection and allows all activity (all processes or all networks) on the Pod Group.
- Profiling includes network activity on domains only. IP addresses are not supported.

During the learning process, the Runtime Protection tab shows the current profiling status:

Status	Description	Profiling Period
In Progress (%)	CloudGuard is learning the workload behavior. **End Time** indicates when the learning is completed.	Not finished
Finalizing	CloudGuard is creating the profile based on the collected data.	On hold

Status	Description	Profiling Period
Partial	CloudGuard finished the profiling period and enforced the profile but detected no startup event. The **Partial** profile can advance to the **In Progress** status again when CloudGuard detects a startup process.	On hold
Completed	CloudGuard has a full profile of the Pod group.	Finished
Disabled	Runtime Protection is disabled.	Not started

The profile contains information on the workload life cycle from its deployment to the stage of stable running. If CloudGuard starts the profile learning after the workload has been deployed, its startup information is missing from the profile, and CloudGuard cannot complete it during the profiling period. The profiling status appears as **Partial**. To consider a profile complete, two conditions are necessary:

- CloudGuard detects a pod startup event
- The profile period is completed.

A **Partial** profile is enforced and can become **Completed** once the startup event is detected. Until the profile is complete, CloudGuard cannot enforce Runtime Protection.

To complete the profiling process, select one of two options below:

- Wait until the workload life cycle spawns a new pod naturally.
- Delete the pod.

Important - If you decide to delete a pod, do it with caution. This may have undesirable side effects on specific workloads, although it triggers the creation of a new pod.

Profile Enforcement

When the profile is complete, the agent stops to collect events and starts to enforce protection. When one of the containers in the Pod Group does an operation not observed during the profiling phase, the engine detects a violation and triggers an alert.

CloudGuard associates a Pod Group profile with a version and identifies it as the combination of all the containers' images. When it detects a container update with a new image version (for example, with the same image but a new image tag), it resets the profile, and a new learning period begins. User-defined settings, such as signature exclusions, stay intact.

More Links

- *"Kubernetes Runtime Protection Rules and Exclusions" on page 556*
- *"Kubernetes Runtime Protection Troubleshooting" on page 561*
- AWS Bottlerocket

Kubernetes Runtime Protection Rules and Exclusions

These Runtime Protection rules and exclusions allow you to customize the CloudGuard security policy:

- **Exclusion** - Classify a Security Event as benign and ignore similar future events.
- **Deny rule** - Prevent malicious event recurrence by deleting (killing) the container that executes the operation.

Rules and Exclusions by Engines

You can define deny rules for Signatures, but not for File Reputation and Profiling. For File Reputation and Profiling engines, you can define only custom exclusions.

Signatures - Rules and Exclusions

You can create a rule or exclusion for signatures to enforce your security policy in CloudGuard. These actions are possible only after a security event has occurred and appeared in the Events table.

- To allow the event and set it as benign, add an exclusion. When CloudGuard detects this behavior again, it does not trigger a security alert.
- To block the operations related to the signature, add a rule. When CloudGuard detects this activity again, the Signatures engine kills the container that executes the operation.
- To automatically block all signatures on a cluster, set a toggle button on the cluster's Runtime Protection tab - For more information, see Creating Rules.

  The blocked event has a special icon and the **Blocked** indication in the event details.

File Reputation - Rules and Exclusions

You can create an exclusion based on the manual classification of an executable file as benign. This allows you to override a false identification of the file as malicious.

You can create a cluster deny rule for all File Reputation events.

- To automatically block all File Reputation violations on a cluster, set a toggle button on the cluster's Runtime Protection tab - For more information, see Creating Rules.

  The blocked event has a special icon and the **Blocked** indication in the event details.

Profiling - Exclusions

You can manually allow (add an exclusion) specific processes and networks in the profile. This helps you fine-tune the profile if it falsely identifies some behavior as unwanted, malicious, or anomalous, while in practice it is benign (false-positive detection). For example, a monthly maintenance process that occurs after the profile finalization is flagged as anomalous.

Profile exclusions contain processes (commands) and networks that you explicitly allow for the applicable scope. Each process and network added to the profile with exclusion has the **Exclusion** tag in its line.

For profiles, you cannot define rules, that is, forbid a process or network.

Security Events Deduplication

Alerts deduplication mechanism allows CloudGuard reduce the clutter caused by repeated alerts.

When the Runtime Protection engine detects an alert that repeats frequently over a short period, it reduces the number of reported alerts. The engine only reports a sample of these repeated alerts.

Actions

Enabling Runtime Protection on a Cluster

Runtime Protection is enabled by default for all Pod Groups.

If you have disabled the functionality or did not select it when onboarding your cluster to CloudGuard, follow the steps below to enable Runtime Protection.

1. In the CloudGuard portal, navigate to **Workload Protection** > **Containers Assets** | **Environments**.
2. Select the environment from the list (you can filter the list to narrow your search) and open the **Blades** tab.
3. Set **Runtime Protection** to **ON**.

   The confirmation window opens with instructions on how to install the agent on the cluster.
4. Install the agent on your cluster.
5. In the confirmation window, click **Yes** to confirm your action.
6. Optionally, set **Behavioral Profiling** to **ON**. This is a Public Preview feature disabled by default.

It takes CloudGuard several minutes to enable Runtime Protection on your environment.

Creating Rules

You can create rules only for triggered security events, that is, from the recorded alerts. After you apply the rule, CloudGuard kills the executing container when the same malicious activity occurs again.

You cannot create rules for Profiling.

To create a rule:

1. Navigate to **Events** > **Threat & Security Events**.
2. Find an event with the **Containers Runtime Protection** source and **SignatureEvent** in the title and select it. If you need it, adjust the time frame filter to see all related events.
3. On the toolbar, click **Add Deny Rule**.

   The Deny Rule Confirmation window opens.
4. You can apply the rule to the selected pod group or to all pods in the cluster. Select the applicable option and click **Create**.

   CloudGuard adds the rule in the **Rules** section on the Runtime Protection Rules tab of the corresponding pod group.
5. For the cluster scope, create deny rules with the toggle buttons on the **Runtime Protection** tab of the cluster.
   a. Open the cluster page.
   b. Go to the **Runtime Protection** tab.
   c. In the **Rules** section:
      - Set the **Block All Malicious Signatures** toggle button to **ON** to automatically block all **Signatures** in a cluster:
      - Set the **Block All Malware** toggle button to **ON** to automatically block all **File Reputation** violations on a cluster:

Creating Exclusions

You can define and enforce exclusions in the scope of a specific workload, a group of workloads, or all the Pods in a specific cluster (environment).

Engine	Available Context
File Reputation	Create an exclusion from: ■ The menu of the finding or event in the **Events** > **Threat & Security Events** table, select a **Malicious File** event ■ The **Runtime Protection** tab associated with a cluster or a pod group
Signatures	Create an exclusion from the **Events** tab and from a specific security event.
Profiling	Create an exclusion from: ■ The menu of the finding or event in the **Events** > **Threat & Security Events** table ■ The **Runtime Protection** and **Runtime Protection Rules** tabs related to a Pod Group ■ The **Runtime Protection** tab associated with a cluster (environment)

To create an exclusion:

1. Select the context for the new exclusion and open the relevant entity.
2. In the Events table, click **Exclude**. On the entity tab, click **Create New Exclusion**.
3. In the **Create New Exclusion** window, enter the details:
    a. **Name** - The name for the exclusion that appears in the list of exclusions in this tab.
    b. **Target** - Type of action to exclude from monitoring: a process or a specific host. For Signatures, the Signature type is preselected.
    c. **Pattern** - Path for processes or domain name for hosts. For Signatures, the pattern is the preselected name of the Signature.
    d. **Process / Parent Process** - Path of the process that can run the excluded domain or process.
    e. **Scope** - Application to a specific Pod group, a list of Pod groups, or all the Pods in a certain cluster (environment).
4. Click **Create**.

Changing the Profile Learning Period

The default profile learning period is 24 hours. You can modify the profiling time based on your knowledge of the workload behavior. If you set the learning period to a value that is shorter than the completed period, the profile learning stops immediately.

To change the profiling period:

1. Open the **Runtime Protection** tab of the workload (Pod Group).
2. On the **Profiling status** bar, click the menu button ⋮ on the right and select **Settings**. The Pod Group Settings window opens.
3. With the up and down arrows, adjust the **Days**, **Hours**, and **Minutes** as necessary.
4. Click **Save**.

Note - CloudGuard saves the setting of the new profiling period for future versions of the Pod Group.

More Links

- *"Kubernetes Runtime Protection" on page 549*
- *"Kubernetes Runtime Protection Troubleshooting" on page 561*

Kubernetes Runtime Protection Troubleshooting

Verification of Installation

Follow these steps to verify the installation of the Runtime Protection agent:

1. In the CloudGuard portal, navigate to **Assets** and select the Kubernetes cluster of your interest.
2. Go to the **Blades** tab and make sure that all the agents under **Runtime Protection** have the **OK** state. This operation sometimes takes several minutes to update the information.
3. In the Kubernetes cluster, make sure that all agent resources are installed in the specified namespace and are in the **Ready** state. If your Internet connection quality is low, it sometimes takes several minutes to download all the images.
4. Make sure that the Runtime Protection Daemon runs on the correct number of nodes, based on the defined node selector and tolerations.
5. Check logs of all the agent components and make sure that there are no errors.
    - inventory-agent pod
    - runtime-policy pod
    - runtime-daemon pod (on each of the nodes):
        - *probe* container
        - *daemon* container

Agent Status Errors

The environment page of a Kubernetes cluster shows information about its agents' status.

The Runtime Protection agent status can show these error messages:

Error	Description
Signatures engine error	The Runtime Protection signature engine has failed to initialize
Profiling engine error	The Runtime Protection profiling engine has failed to initialize
File-Reputation engine error	The Runtime Protection file-reputation engine has failed to initialize
Container Runtime error	The Runtime Protection daemon has failed to communicate with the container runtime

Error	Description
Internal agent error	An undefined error occurred

Image Pulling Errors

- If you configured custom images, make sure that your nodes can access them.
- If the images are stored in a protected image registry (which is usual for the default EA images), make sure the credentials of the image registry specified during agent installation are correct.

Errors Log

If the logs contain errors related to access to api-cpx.dome9.com or a similar endpoint:

- Make sure that your nodes need the proxy to access Internet services, refer to *"Cluster Behind a Gateway" on page 199*.
- Make sure that you configured *clusterID* and *credentials* properly when you installed or upgraded the agents.

If logs contain errors related to access to the Kubernetes APIs, and a custom service account is configured, make sure it has a correct configuration.

Empty or Partial Profiles

Partial profiles do not contain processes launched during the Pod startup. To complete the process, see How to complete the profile learning.

Behavior on Profile Re-learning

Runtime Protection constructs a new behavior profile for a workload on each change of a container image related to the workload. The decision on the image change depends on the image reference string in the Pod template specification. This means that if, for example, it refers to a generic image tag, such as `:latest`, it does not trigger re-learning. This agrees with how Kubernetes identifies changes in Pod templates (it triggers a rollout of the new version).

Check Point does not recommend you to use the `:latest` image tag as it is considered insecure and does not necessarily cause profile re-learning.

More Links

- *"Kubernetes Runtime Protection" on page 549*
- *"Kubernetes Runtime Protection Rules and Exclusions" on page 556*
- *"Which CloudGuard endpoints do I have to allow on my network?" on page 917*

Serverless CI/CD Plugin

CloudGuard serverless protection lets you *shift left* security posture into the CI/CD pipeline. It allows you to configure the level of risk that prevents the deployment of serverless applications into your environments. You can download and integrate the CloudGuard Serverless CI/CD Plugin with many popular CI/CD tools, configure it to scan builds before their deployment, and limit deployment to environments on the severity level of vulnerabilities found.

When the CloudGuard serverless CI/CD plugin rejects a CI/CD deployment, it provides developers and DevOps engineers with clear guidance on how to remedy the detected risks. In addition, it provides developers with the ability to directly check their security posture, before they supply code into the pipeline.

The serverless CI/CD plugin supports Java, Python, Node, C#, and is designed to identify security risks spanning the serverless ecosystem (functions code, permissions, third-party libraries, and more).

The Plugin scans your code and configuration for the following:

- Identify overly permissive IAM roles used by serverless functions
- Identify 3rd-party libraries for vulnerabilities
- Identify hard-coded credentials, secrets, and other sensitive information in serverless code
- Identify functions that are not used

How it Works

The CloudGuard Serverless CI/CD Plugin Deep Code Flow Analysis analyzes your serverless function code to understand how it operates. During deployment code/byte-code is analyzed to understand what the code "does". Code is parsed into an abstract syntax tree (AST), and then the execution of the code is emulated before the code is run. This is a complex process that requires the processing of non-deterministic state changes, and it allows CloudGuard to create very accurate results.

Actions

You can configure the Plugin to enable Runtime Protection on functions before they are deployed to your cloud account. When they are onboarded to CloudGuard, they have Runtime Protection already enabled.

Downloading and Installing the Plugin

You must have a CloudGuard account to use the Plugin in your CI/CD tools.

1. Navigate to **Workload Protection** > **CI-CD Tool** and select one of the available methods:
   - **Serverless Plugin Integration** - allows installation of AWS FSP and AWS Proact.
   - **CloudFormation** - allows installation of AWS FSP and AWS Proact.
   - **CLI Plugin Integration** - allows installation of AWS FSP and AWS Proact.
2. Follow the on-screen instructions.

Cloud Detection and Response (CDR)

Cloud Detection and Response (previously known as *Intelligence*) allows you to visualize and analyze account activity and network traffic into and out of your cloud environment or container cluster. With this, you can, for example, identify traffic from unwanted or malicious sources, or misconfigurations that attackers can use to their advantage.

CloudGuard provides preconfigured queries for CDR, and you can create more custom queries with a graphical query builder based on the CloudGuard Governance Specification Language (GSL).

CDR combines cloud or Kubernetes assets and configuration information with real-time monitoring data from network traffic logs account activity logs, and current threat intelligence feeds, IP reputation databases, and geolocation databases. This results in enriched logs and enhanced visualization. For example, sources of network traffic from other AWS elements are shown based on the type and malicious external sources are marked as such.

CDR can give you near real-time views of account activity and network traffic. You can also see and do an analysis of past logs. You can configure CDR to send you real-time alerts for specific events that may occur in your cloud environment and therefore enable you to respond quickly.

CDR features:

- Near real-time view of events
- Adjust queries for specific events and threat hunting
- Enriched contextual information from different log sources allows you to get a quicker and clearer understanding of events that occur in your cloud environment

Benefits

1. **Streamline Network Security Operations**: With CDR, you can do network operations such as:
    - Security architecture review based on real-time traffic analysis
    - Increase visibility into your traffic flow
    - Troubleshoot and identify misconfigurations that cause intrusions and policy violations
    - Identify unusual account activity
    - Detect malicious sources that are sending traffic to your network assets

2. **Decrease meantime for threat detection**: On average, it takes about 200 days for incident responders to detect a breach. With CDR, you can identify and zoom in on a suspected asset and understand the full context from a configuration and traffic activity perspective, thereby reducing your mean time to detect threats.
3. **Detect Privilege Escalation / Credential Compromise**: CloudGuard has the full context of your account activity and the types of assets in your environment. With CDR, you can create lists of asset types that shouldn't be instantiated. If someone obtains unapproved privileges to launch an expensive EC2 instance that is perhaps used for crypto-mining operations or to steal API keys, and is at this time misused, CDR can detect such unapproved IAM changes or specific EC2-type traffic and immediately provide detailed alerts.
4. **Expedite and help in Compliance Validation**: With the Explorer, you can see a replay of traffic that can be used to prove that your cloud environment is adhering to different compliance standards (Control effectiveness).
5. **Detect unusual or abnormal use** of your cloud resources, network activity, logins, etc. For example, detect activity from geographic locations not permitted, suspicious port use, unusual logins, or authentication attempts.

CDR Dashboard

The CDR Dashboard shows information about security events organized by event type, rule, and platform. Click on elements in the dashboard to see more information about them.

To open the CDR Dashboard:

From the left toolbar, click **CDR** > **CDR Dashboard**.

To filter the CDR Dashboard:

A filter applies to the entire CDR Dashboard.

1. In the top left area of the dashboard, click the filter button ( ).
2. Apply filters.
3. In the upper right area of the dashboard, select a time period for the filter.

CDR License

As part of the CloudGuard CNAPP license, CloudGuard customers get the account activity component of CDR. This basic CDR functionality provides 12 GB of logs with retention for one month for each billable asset. This means that when you onboard your environment to CDR, it can analyze 12 GB of logs, no matter how much time it takes, and keep them for no more than one month. When you finish the quota, CDR stops to process your data.

In addition, you can purchase a license for **CloudGuard CDR Pro**, Check Point's Cloud Security Threat Defense Analytics platform. CDR Pro can help you detect and mitigate threats to users' cloud environments, analyze activity, and leverage UEBA algorithms to fend off cloud attacks. Same to basic CDR, the license for CDR Pro is based on consumption. All processed data is counted against the quota set by your license. The license covers data ingestion while there is capacity left.

Some assets can create a tremendous number of network traffic logs, which you may not want to be analyzed by CDR. You have to examine these numbers and, in general, the architecture of your network before you onboard your environments to CDR. To prevent the cases when your quota is consumed too quickly or spent on irrelevant data, you have to control which logs to onboard to CDR.

To manage the network traffic logs that your AWS assets send to CDR:

1. Enable Flow Logs only on the applicable VPC, subnet, or ENI.
2. Configure that the Flow Logs that you want to send to CDR are published to a specific S3 bucket which you then onboard.

When the license quota is near its end, you receive a warning message that the number of ingested logs is at 80%, 90%, or 100% of its capacity. When it reaches 100%, the logs ingestion stops until a new license is purchased. The ingested logs are stored and accessible during the full retention period stipulated by your license.

CDR Connectivity

Based on your Data Center location (region), CDR must have connectivity to this endpoint:

Region	Dome9 Portal	Infinity Portal
United States	magellan.dome9.com	magellan.us1.cgn.portal.checkpoint.com
Europe	webserver.logic.eu1.dome9.com	webserver.logic.eu1.cgn.portal.checkpoint.com
Australia	webserver.logic.ap2.dome9.com	webserver.logic.ap2.cgn.portal.checkpoint.com
Canada	webserver.logic.cace1.dome9.com	webserver.logic.cace1.cgn.portal.checkpoint.com
India	webserver.logic.ap3.dome9.com	webserver.logic.ap3.cgn.portal.checkpoint.com

Region	Dome9 Portal	Infinity Portal
Singapore (for Dome9 accounts only)	webserver.logic.ap1.dome9.com	

More Links

- *"Intelligence Onboarding and Offboarding" on page 570*
- *"Configuring CloudGuard Exclusions" on page 80*
- *"Remediation" on page 632*
- *"CloudGuard Permissions for Intelligence" on page 621*
- *"Kubernetes Intelligence" on page 617*

Getting Started with Intelligence Policy

An Intelligence policy has a ruleset (containing event definitions), one or more environments on which the events are applied, and a notification indicating where findings must be sent.

To set up an Intelligence policy:

1. Navigate to the **Policies** page in the **CDR** > **Threat Monitoring** menu.
2. Click **Add Policy** on the right.
3. Select a platform on which the policy applies and click **Next**.
4. Select one or more environments to which the policy applies. CloudGuard shows only those environments onboarded to Intelligence. Click **Next**.
5. For the initial Intelligence configuration, use the configured CloudGuard-managed rulesets. From the list, select one or more rulesets for the policy and click **Next**.
6. To add a new Notification, click **Add Notification**.
7. In the Create New Notification window, enter the notification name and, optionally, a description. For this initial policy, you can use the default settings. Make sure that the **Alert console** is selected. This option allows you to see all findings on the **Events > Threat & Security Events** page.
8. Click **Save**.
9. Select the Notification for the association.
10. Click **Save**.

Your policy appears on the Policies page.

More Links

- *"Intelligence Onboarding and Offboarding" on page 570*
- *"Configuring CloudGuard Policies" on page 78*
- *"Intelligence Queries" on page 630*
- *"Intelligence Security Events" on page 634*
- *"Notifications" on page 852*

Intelligence Onboarding and Offboarding

Onboarding

To use Intelligence for your assets, you have to onboard environments to Intelligence. This allows CloudGuard to use applicable logs, for example, CloudTrail or Flow Logs.

Note - CloudGuard **Intelligence Pro** license provides **account and traffic activity** based on platform logs, such as AWS CloudTrail and Flow Logs. The basic **Intelligence** function provides **only account activity** based on activity logs, such as CloudTrail. For more information about the license, see *"CDR License" on page 566*.

- For AWS environments, see *"AWS Intelligence" on page 571*
- For Azure subscriptions, see *"Azure Intelligence" on page 600*
- For Kubernetes clusters, see *"Kubernetes Intelligence" on page 617*
- For GCP projects, see *"GCP Intelligence" on page 609*

Offboarding

Intelligence offboarding removes CloudGuard permissions to access logs and stops logs sending to the Intelligence engine.

AWS Intelligence

Onboarding

There are multiple ways to onboard your AWS environment to Intelligence:

- *"Onboarding AWS Environments to Intelligence" on page 572*
- *"Onboarding AWS Environments to Intelligence with API" on page 590*
- *"Custom Onboarding of AWS Environments to Intelligence" on page 595*

Offboarding

There are multiple ways to remove Intelligence from your AWS environment:

- *"Removing Intelligence from AWS Environments" on page 588*
- *"Removing Intelligence from AWS Environments with API" on page 593*
- *"Manual Removing of Intelligence from AWS Environments" on page 599*

Onboarding AWS Environments to Intelligence

This topic describes how to onboard an AWS environment to Intelligence with an automated onboarding experience. For a legacy procedure with manual onboarding, see *"Custom Onboarding of AWS Environments to Intelligence" on page 595*.

Your AWS environment has to be onboarded to CloudGuard before you can onboard it to Intelligence. If your account is not yet onboarded, follow the instructions in *"Unified Onboarding of AWS Environments" on page 54*.

How It Works

Intelligence uses *"VPC Flow Logs" on page 382* and CloudTrail logs from your AWS account. These logs have to be stored in an Amazon S3 bucket. Intelligence analyzes only the new logs that CloudGuard starts to receive after you onboarded the account to Intelligence. The S3 bucket sends PutObject event notifications to an SNS topic, and the topic is connected through an SNS Subscription to the Intelligence SQS-queue endpoint. When CloudGuard receives this event notification, it runs the S3 GetObject on the path specified in the notification to retrieve the log.

During the onboarding process, CloudGuard establishes the connection to receive notifications regarding new log files and obtains permissions to run the S3 GetObject on one or more log buckets. For easier handling of the required configuration, CloudGuard can create a CloudFormation Template (CFT) to run in your AWS environment.

S3 bucket for each account

The architecture diagram below illustrates how to onboard an S3 bucket for each account.

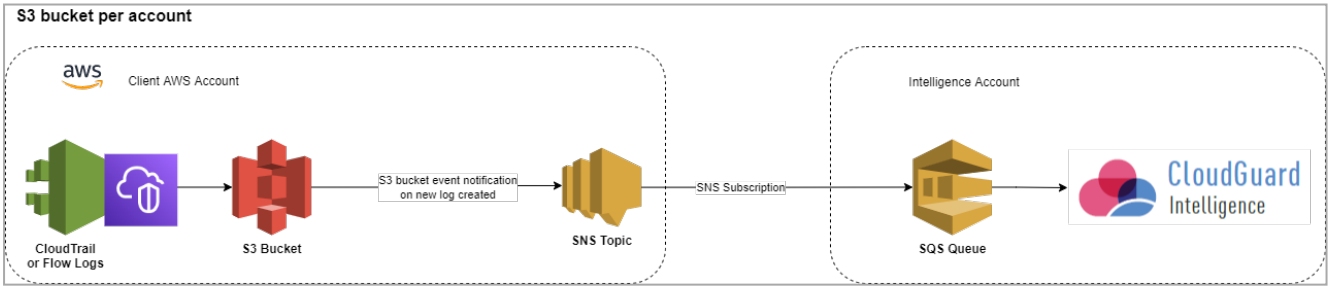

Centralized S3 bucket

A centralized S3 bucket is an Amazon S3 bucket that stores logs for multiple AWS accounts. Use the onboarded centralized S3 bucket to easily monitor and run log analysis of all your AWS accounts that send logs to this S3 bucket.

The onboarding initiates from the AWS account that has the centralized S3 bucket already configured in its environment. CloudGuard then retrieves a list of AWS accounts that send Flow Logs or CloudTrail logs to the centralized S3 bucket. You can select the applicable AWS account(s) to onboard from this list. To onboard all AWS accounts to Intelligence, use *"Automatic Onboarding" on page 582*.

The onboarding of an organizational unit is supported only when a centralized S3 bucket is configured in your AWS organization structure.

The architecture diagram below illustrates how to onboard several AWS accounts to a centralized S3 bucket.

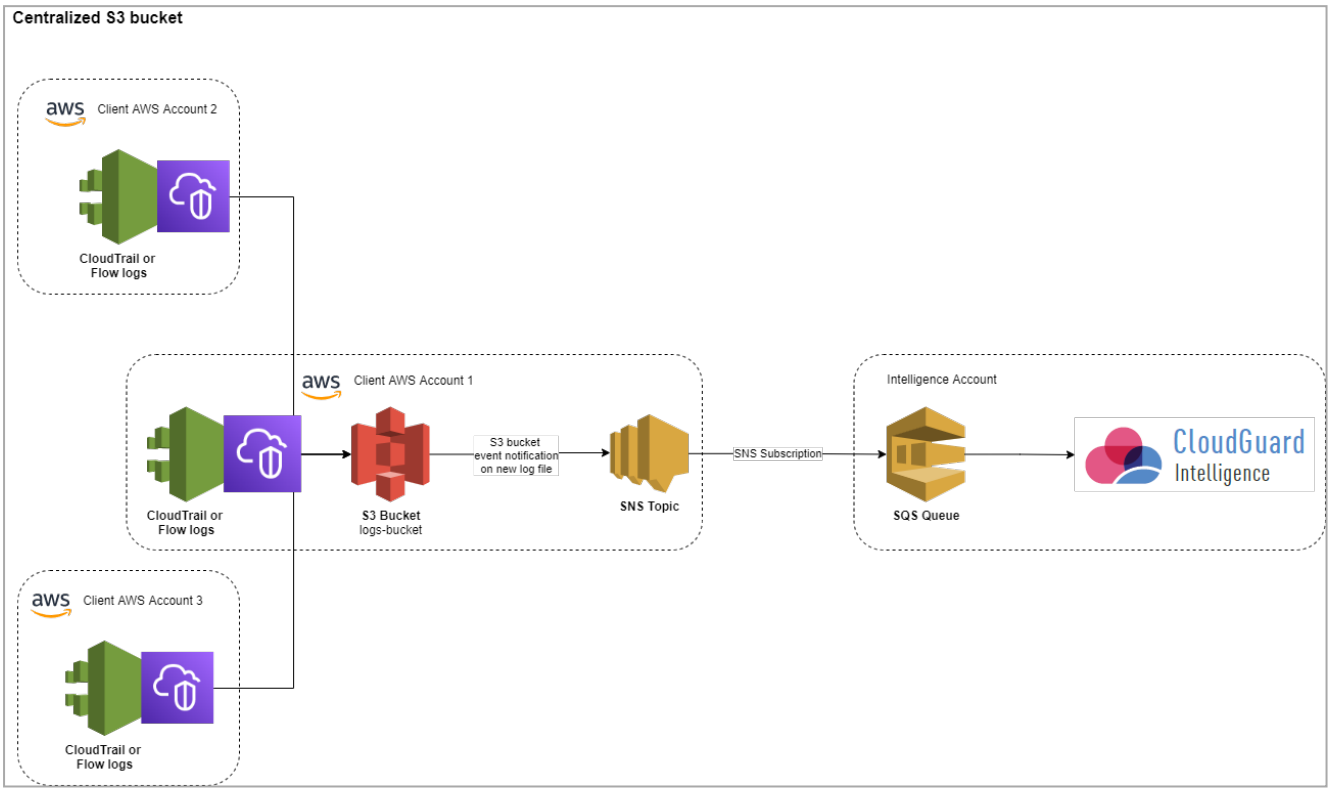

Known Limitations

- AWS allows you to set only one event notification on a specific prefix with a specific event type. You cannot onboard an S3 bucket if the bucket has the event notifications set, and one of the notifications meets two conditions below:
  - The notification has an empty prefix filter set (that is, all of the bucket) or an explicit AWS log-prefix set.
  - The event type is *PutObject* or all object-created events.
- For Intelligence, you can onboard an S3 bucket through one SNS topic only.
- Intelligence cannot analyze the involved IAM policies: S3 bucket policy, existing SNS topic policy, and policies of the CloudGuard IAM trust role. Therefore, permission-related issues can occur when you create the onboarding stack.

For these and other CloudGuard limitations, see *"Known Limitations" on page 924*.

Onboarding to Account Activity or Traffic Activity

Note - You must onboard Flow Logs and CloudTrail separately for each account.

Follow these steps in CloudGuard to enable Account Activity with CloudTrail:

1. In CloudGuard, open the **Environments** page from the **Assets** menu.
2. Select the AWS environment to be onboarded to Intelligence.
3. In the account row and the **Account Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   As an alternative, you can click and open the environment page. From the top right, select **Add Intelligence** > **CloudTrail**.
4. Follow the on-screen instructions to complete the wizard.

Follow these steps in CloudGuard to enable Traffic Activity with Flow Logs:

1. In CloudGuard, open the **Environments** page from the **Assets** menu.
2. Select the AWS environment to be onboarded to Intelligence.
3. In the account row and the **Traffic Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   Or, you can click and open the environment page. From the top right, select **Add Intelligence** > **Flow Logs**.
4. Follow the on-screen instructions to complete the wizard.

Wizard Stages

1. **Welcome** - Read carefully the onboarding prerequisites and make sure that the AWS environments to be onboarded meet all the required conditions.
2. **S3 Buckets**
   a. Select one or more S3 buckets to be onboarded.
   b. Set **Auto Onboard** to **ON** to let CloudGuard:
      - detect all onboarded AWS accounts that send logs of a certain type to this centralized S3 bucket
      - automatically onboard these accounts to Intelligence

   The toggle button is enabled only for new buckets that you select to onboard. To update the onboarded bucket(s) mode, go to the **Intelligence** tab on the environment page. For more information, see *"Automatic Onboarding" on page 582*.

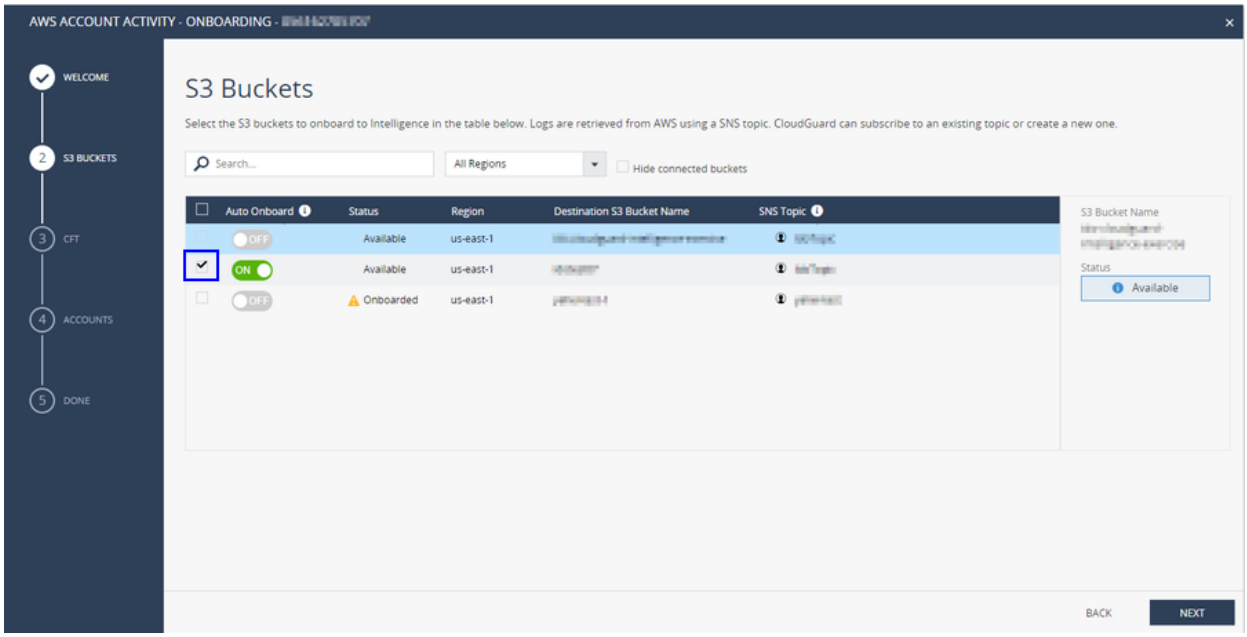

You can see the status details for each S3 bucket on the right pane. The details help you troubleshoot issues that prevent onboarding of the S3 bucket. For more information on possible issues, see *"Errors, Warnings, and Troubleshooting" on page 582*.

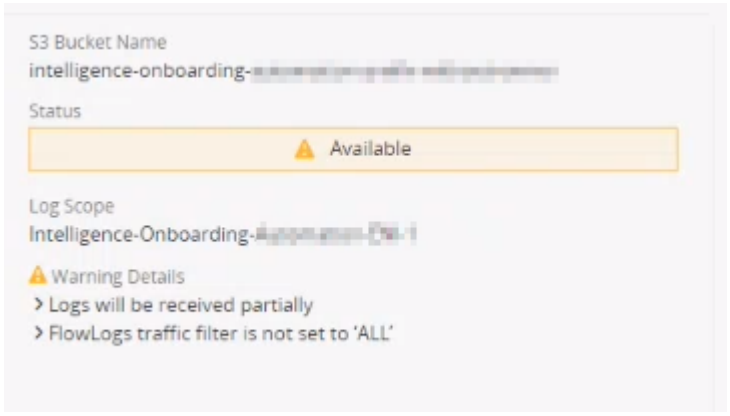

3. **CloudFormation Template** - Create a stack from the provided CFT. For more information, see *"CFT Resources and Permissions" on page 579*.

   During the process, CloudGuard shows the onboarding status on the bottom part of the page. Click **Check Now** to see the current status of the configuration. For more information, see *"Status Check" on page 581*.

   CloudGuard skips this step if your AWS account already has all the required resources.

4. **Accounts** - For centralized S3 buckets, select the AWS accounts that you want to onboard and see their logs.

   CloudGuard skips this step if only one AWS account sends logs to this S3 bucket.

5. **Done** - Make sure the onboarding is successful.

AWS Onboarding Permissions

When you launch a CloudFormation stack, the stack gets its permissions through two primary processes.

1. It inherits the permissions of the user who creates the stack. The user has one of these credentials:
   - IAM user
   - Federated user (SSO/SAML)
2. It receives an IAM Role assigned directly to the stack.

   To assign a dedicated role to the onboarding stack, the role must allow for the AWS CloudFormation service to assume it.

For successful onboarding, it is necessary to give the stack all the necessary permissions.

Some of the permissions include a *Delete* step. It is applicable when it is necessary to create a dedicated IAM role for the stack, rather than use the user's permissions. As a result, this role is used for the deletion of the stack.

Permissions

Four primary conditions that have an effect on your customized CloudFormation template:

1. A bucket that is already preconnected to an existing SNS topic with an Event Notification.

   This is the most basic, in terms of the resources created, template. It requires these permissions:

   ```
   "iam:DetachRolePolicy",
   "iam:GetRolePolicy",
   "iam:PutRolePolicy",
   "sns:Subscribe",
   "sns:Unsubscribe"
   ```

2. A bucket will connect, as part of the current template, to an existing SNS topic. In that case, it is necessary to add permissions* to create an AWS Lambda function that adds

the applicable event notification.

```
"ec2:DescribeNetworkInterfaces",
"iam:AttachRolePolicy",
"iam:CreatePolicy",
"iam:CreateRole",
"iam:GetPolicy",
"iam:GetRole",
"iam:GetRolePolicy",
"iam:DeletePolicy",
"iam:DeleteRole",
"iam:DeleteRolePolicy",
"iam:DetachRolePolicy",
"iam:ListPolicyVersions",
"iam:PassRole",
"iam:PutRolePolicy",
"lambda:CreateFunction",
"lambda:DeleteFunction",
"lambda:GetFunction",
"lambda:GetFunctionCodeSigningConfig",
"lambda:GetFunctionConcurrency",
"lambda:InvokeFunction",
"lambda:PutFunctionConcurrency",
"logs:CreateLogGroup",
"logs:DeleteLogGroup"
```

*Note - In addition, the permissions from point 1 are necessary.

3. A bucket connected, as part of the current template, to a new SNS topic (that is created in the current template). More permissions* are necessary for the creation of a SNS topic.

```
"sns:CreateTopic",
"sns:GetTopicAttributes",
"sns:ListSubscriptionsByTopic",
"sns:SetTopicAttributes",
"sns:TagResource"
```

*Note - In addition, the permissions from points 1 and 2 are necessary.

4. If you select buckets from multiple regions as part of one process, stack sets are used that necessitate more permissions*:

```
"cloudformation:CreateStackInstances",
"cloudformation:CreateStackSet",
"cloudformation:DeleteStackInstances",
"cloudformation:DeleteStackSet",
"cloudformation:DescribeStackSet",
"cloudformation:DescribeStackSetOperation",
"cloudformation:GetTemplateSummary",
"iam:CreateRole",
"iam:GetRole",
"iam:GetRolePolicy",
"iam:DeleteRole",
"iam:DeleteRolePolicy",
```

```
"iam:DetachRolePolicy",

"iam:PassRole",

"iam:PutRolePolicy"
```

*Note - In addition, the permissions from points 1 and 2 or from points 1 and 3 are necessary.

CFT Resources and Permissions

Common Permissions

CloudGuard generates a customized CloudFormation Template based on your bucket selection. Based on the setup in your account, the template can include these resources:

- **IAM inline policy** added to the CloudGuard trust role.

  The policy grants these permissions:

  - `s3:GetObject` - Only on the onboarded log buckets
  - `sns:Subscribe` - Only on the topics connected to the onboarded log buckets
  - `sns:Unsubscribe`
  - `kms:Decrypt` - Only if the onboarded log bucket is KMS-encrypted or the stored CloudTrail log files are encrypted
- **SNS Topic** - If the log bucket is not pre-connected to SNS topics
- **SNS Topic Policy** - The policy allows all of the buckets in the AWS account to publish event notifications to the topic. The policy allows subscription to the Intelligence SQS queue endpoint.
- **SNS Subscription** to the Intelligence SQS queue endpoint.
- **Lambda function** - It is used to create a correct event notification on the S3 bucket if there is no such notification.
- **CloudWatch Log Group** - To store logs from Lambda.
- **Lambda Execution IAM Role** and **Lambda Execution IAM Policy** - To grant the lambda the following permissions:

  - `logs:CreateLogStream`, `logs:PutLogEvents` - Allow logging in CloudWatch, standard lambda permissions.

- `s3:GetBucketNotification`, `s3:PutBucketNotification` – Allow reading the existing notifications on the bucket and appending the correct one. This permission applies only to the bucket onboarded in the process.

Resources for Multiple Buckets

If you select multiple buckets from different regions as part of one onboarding process, then is necessary for CloudGuard to have more resources. These resources are necessary because CloudFormation does not allow interaction with multiple regions from one stack.

- **Stack Sets** – A stack set for each selected region.
- **Stack** – A stack for each region.
- **Two IAM roles** that CloudGuard role does not have permission to assume.
  - **IAM Role (Admin Custom Role)** - The role that can be assumed by CloudFormation and can assume the execution role below.
  - **IAM Role (Execution Custom Role)** – The role with the exact permissions that the CFT needs to complete.

IAM Custom Role

If necessary, configure a custom IAM role to reduce the scope of your administrative role and allow CloudFormation to use this role to create, change, or delete resources in the stack.

This custom role can be assigned through:

- A regular IAM user
- A federated user
- An IAM role

You can select this role in the **Permissions** section when you configure stack options.

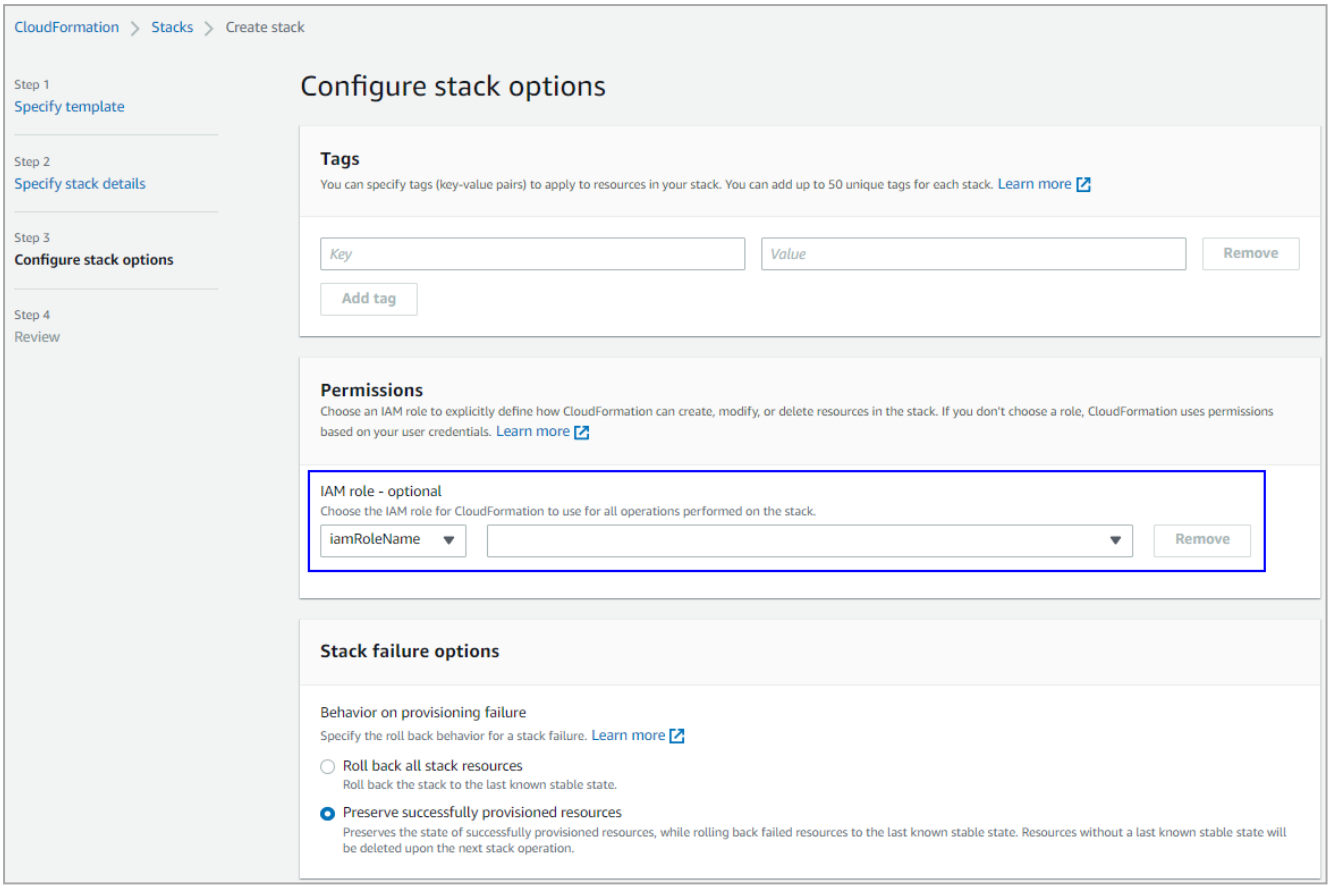

Important - To make sure that the stack creation is successful:

- Make sure that the IAM user (or the federated user, or the IAM role) that creates the stack has all the permissions for the involved resources listed above.
- After you have generated a template, **do not** change the account configuration until the template reaches a steady (not *in progress*) state.

Status Check

You can see in real time the status of part of the configurations added as part of the CFT.

This process verifies that:

- All required S3 bucket event notifications were added to the correct SNS topics, with the correct prefix filter and event type.
- All connected SNS topics are at this time subscribed to the correct Intelligence SQS endpoint based on the onboarding log type.

This process does not make sure that the permissions are added successfully to the CloudGuard trust role. And it does not make sure that the stack reached a **Create_Complete** state.

Possible results:

- **Complete** - The checked configurations appended successfully. You can continue with the onboarding process.
- **Not Complete** - Some parts of the configuration are missing. Possibly, the stack creation did not complete. Wait until the stack reaches a stable state and retry the call.

Automatic Onboarding

CloudGuard can detect all onboarded AWS accounts that send logs of a certain type to the centralized S3 bucket and automatically onboard them to Intelligence.

- For accounts that send CloudTrail logs to the centralized S3 bucket, CloudGuard can automatically onboard them to Account Activity.
- For accounts that send Flow Logs to the centralized S3 bucket, CloudGuard can automatically onboard them to Traffic Activity.

If you do not select this option, you have to onboard each account to Intelligence manually.

To update the Auto Onboard status:

1. On the **Assets** > **Environments** page, click to open an environment onboarded to Intelligence.
2. Go to the **Intelligence** tab.
3. In the **Auto Onboard** column, click the toggle button to change the status of automatic onboarding.

Note - Automatic Onboarding is applicable to a certain log type on the entire bucket. On the environment page, Intelligence tab, you can sometimes see more than one row of a bucket for different log scopes (VPCs). However, when you enable Automatic Onboarding on one row, it is enabled on all other rows of the same bucket and the same log type.

Errors, Warnings, and Troubleshooting

When you select from available S3 buckets for onboarding, you can see the bucket status on the right pane. The bucket status can be:

- **Available** - You can continue to the next step and onboard the S3 bucket.
- **Onboarded** - This S3 bucket is already onboarded. For centralized S3 buckets, you can continue and onboard more accounts with the bucket.
- **Cannot onboard** - You cannot onboard this S3 bucket. Resolve the related issue or select a different S3 bucket.

  Note - Some AWS account configurations do not allow S3 bucket onboarding, for example, when the S3 bucket is already connected to a non-SNS endpoint.

Errors prevent your S3 bucket from onboarding; with warnings, you can continue the process.

Best Practice - Attempt to resolve issues from a status warning before you continue to onboard the S3 bucket. Although CloudGuard allows you to move on to the next step, this can cause Intelligence not to do a full analysis of your account.

Warning and Error Messages

Warning / Error	Details	Corrective Action
Warning	Flow Logs traffic filter is not set to 'ALL'	Add new Flow Logs with **ALL** traffic selected
Warning	Flow Logs custom log format is insufficient	Add new Flow Logs with all default attributes present.
Warning	Couldn't get S3 bucket encryption configuration	For the `AccessDenied` error, check the S3 bucket policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the `S3:GetEncryptionConfiguration` action on the bucket resource.
Warning	S3 Bucket wasn't listed properly	For the `AccessDenied` error, check the S3 bucket policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the `S3:listobjects` action on the bucket resource.
Warning	Couldn't resolve bucket's kms key	For the `AccessDenied` error, check the S3 bucket policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the `kms:listaliases` action.
Warning	SQS Subscription to the onboarded SNS topic not found	-
Warning	Onboarded S3 bucket's SNS topic not found	-
Warning	CloudTrail is not multi-region	Change a trail that applies to one Region to apply to all Regions, see AWS documentation

Warning / Error	Details	Corrective Action
Warning	Logs will be received partially	Each S3 bucket can be onboarded through a topic, hence CloudGuard receives only part of the logs. Select the topic or configure the separation of the S3 bucket prefixes.
Warning	Bucket is in another account and requires action	The S3 bucket is in a different account. Select the account on the **Environments** page and onboard to Intelligence from the account.
Warning	"Access Denied" error received from AWS	After the warning is corrected (see *"AccessDenied Error" on page 621*), you must select the S3 bucket in the onboarding wizard and onboard it again. It is not necessary to offboard the bucket.
Warning	The configured event notification destination on the S3 bucket does not exist.	Remove the S3 bucket event notifications from the invalid destinations and reload the page. The invalid SNS topic's ARN that CloudGuard detected shows in the warning on the CloudGuard portal.
Warning	AWS `AuthorizationError` received for SNS topic encryption	Check the SNS topic policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the SNS:GetTopicAttributes action on the topic resource.
Error	Couldn't get S3 bucket notifications	For the `AccessDenied` error, check the S3 bucket policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the `s3:getbucketnotification` action on the bucket resource.
Error	Couldn't get S3 bucket location	For the `AccessDenied` error, check the S3 bucket policy and the CloudGuard trust role policies. Make sure the trust role is allowed to do the `s3:getbucketlocation` action on the bucket resource.

Warning / Error	Details	Corrective Action
Error	Couldn't resolve S3 bucket's account	■ Make sure the bucket is in the AWS account that is onboarded to CloudGuard ■ Review the CloudGuard trust role and the bucket policy to make sure they allow all required actions
Error	Encrypted SNS topics connected	Intelligence cannot onboard buckets connected through an encrypted SNS topic. Remove encryption if you want to use a specific topic.
Error	The S3 bucket has a direct connection that is not supported anymore.	Remove the S3 bucket event notification to the Intelligence endpoint and reload the page.
Error	We cannot connect to the existing bucket's event notification due to an AWS limitation.	There is an AWS limitation that cannot define different event notifications if there are overlapping prefixes for the same event type. Solution - See below.

Troubleshooting

Failed Connection to Existing Bucket's Event

This error shows because AWS allows the event notification to have only one endpoint of the 'put' type. It cannot define more than one event notification if there are overlapping prefixes for the same event type. To solve the issue, use the fact that an SNS topic can have multiple subscriptions.

If you already have a destination, create an SNS topic to send the data to the CloudGuard SQS and the existing destination, which can be an SNS, Lambda, or SQS.

For this, build the structure on the AWS side and use the option of manual onboarding of Intelligence in CloudGuard. This allows you to select the SNS topic manually.

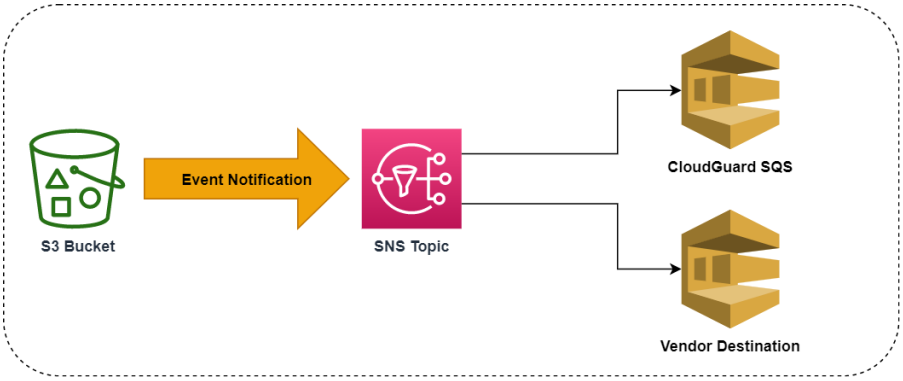

For more information, see AWS documentation.

S3 Bucket of Logs Does Not Appear in the List

The issue of not seeing an S3 bucket of logs in the onboarding wizard can occur in both CloudTrail and Flow Logs, though it is more common with Flow Logs. This is because CloudTrail is often enabled across the entire organization, whereas Flow Logs can be missing due to a lack of network setup in the environment.

In a typical AWS environment, especially in organizations, there is often a dedicated logging environment with a centralized S3 bucket storing logs for the entire organization. Sometimes, this logging environment does not have Flow Logs enabled because there might not be any network setup or logs to record.

To resolve the issue:

1. Verify that Flow Logs/CloudTrail are configured in the AWS environment where the S3 bucket is located.
2. If Flow Logs/CloudTrail are not configured in the account of the bucket, consider these options:
   - Create Flow Logs/CloudTrail in the AWS environment of the bucket, which can be removed after onboarding if not needed.
   - Proceed with custom (manual) onboarding using the link provided in the top banner of the onboarding wizard. For more information, see *"Custom Onboarding of AWS Environments to Intelligence" on page 595*
   - Use API-based onboarding as detailed in *"Onboarding AWS Environments to Intelligence with API" on page 590*.

Logs Do Not Appear on the CloudGuard Portal

In some cases, you do not see the CDR logs on the CloudGuard portal because an encrypted resource (S3 bucket or CloudTrail) is not located at the same account as its encryption key.

To solve the issue:

Add permissions for the CloudGuard role to allow to do the `kms:Decrypt` action under the KMS policy of your encryption key.

- For **Organization Onboarding** - Add permissions for the role existing in the account with the S3 bucket.
- For **CDR Onboarding with CloudFormation Template or Manual Onboarding** - Add permissions for the role existing in your current account.

If you cannot successfully finish the onboarding process:

- Contact *Check Point Support Center*
- After several hours attempt again. Sometimes, maintenance or network issues on AWS can cause interference with onboarding.
- Onboard with the legacy procedure in *"Custom Onboarding of AWS Environments to Intelligence" on page 595*

Onboarding Verification

When the onboarding is complete, make sure that the new logs of the onboarded AWS account start to appear in the CloudGuard portal in **Events** > **Cloud Logs** > **Account Activity** or **Network Traffic**. This can take less than 30 minutes.

More Links

- To onboard your AWS environments to Intelligence with API, see *"Onboarding AWS Environments to Intelligence with API" on page 590*.
- For manual onboarding, see *"Custom Onboarding of AWS Environments to Intelligence" on page 595*.
- To remove Intelligence from your AWS environments, see *"Removing Intelligence from AWS Environments" on page 588*.

Removing Intelligence from AWS Environments

This topic describes how to remove Intelligence from your AWS environment with a new onboarding experience. For a legacy procedure with manual removing, see *"Manual Removing of Intelligence from AWS Environments" on page 599*

You can remove Intelligence from your AWS environments. As a result, CloudGuard stops to receive all Account activity and Traffic activity (CloudTrail and Flow Logs) from your environment.

- For environments onboarded before with Standard Onboarding, the process removes the S3 event notification to the Intelligence SNS topic endpoint from all S3 buckets on this account.
- For environments onboarded with CFT or before with Custom Onboarding, the process removes the subscriptions to the Intelligence SQS queue endpoint.

Best Practice - First, in your AWS account, remove the stack created during the Intelligence onboarding. Second, remove Intelligence in the CloudGuard portal.

On your AWS account

To remove the onboarding stack on the AWS account:

1. Open the AWS console and find the stack created during the Intelligence onboarding process.
2. On the top menu, click **Delete**.

Stack deletion can fail if the stack resources were changed outside the stack context. For example, because of a manual change from the AWS console.

For troubleshooting issues, use the stack drift. From the menu, select **Stack actions** > **Detect drift**.

If the result that you get is not **IN_SYNC**, then correct the drift status before you delete the stack.

Notes:

- The stack removal does not delete the SNS topics and their policies if you use them with other resources.
- If you started offboarding from your CloudGuard account before the stack deletion, the subscription can appear as missing in the drift check. This does not prevent you from successful offboarding.

On your CloudGuard account

To remove Intelligence in CloudGuard:

1. Navigate to **Assets > Environments** and find your AWS environment with **Filter** and **Search** fields.
2. Click the environment to enter it.
3. In the top right menu, click **Remove Intelligence**.

   A verification window opens.
4. Click **Remove** in the verification window.
5. CloudGuard notifies you of the successful removal of Intelligence from your environment.

No more account logs are sent to Intelligence, and you cannot see the existing logs on the CloudGuard portal.

Note - CloudGuard stores the existing logs until the end of your retention period. If you onboard your account to Intelligence again during this retention period, you can see the logs for the period before the offboarding.

Onboarding AWS Environments to Intelligence with API

You can onboard one or more AWS environments to Intelligence with the CloudGuard REST API. For onboarding an AWS environment with the CloudGuard portal, see *"Onboarding AWS Environments to Intelligence" on page 572*.

Prerequisites

Before onboarding your AWS environments with API, make sure that you have prepared:

- CloudGuard account information: API key and secret for your account.
- AWS account information:
  - ID of your AWS account (Account Number)
  - Name of the bucket that stores Flow Logs or CloudTrail logs
  - ARN of the SNS topic that receives event notifications from the bucket
  - For a centralized S3 bucket: IDs of the other AWS environments that send log files to the centralized bucket
- AWS account setup that includes:
  - SNS topic
  - Event notification that the S3 bucket can send to the SNS topic
  - Permissions given to CloudGuard Intelligence

Creating an SNS Topic

Your AWS account requires an SNS topic that is ready to receive notifications from the logs bucket. The topic must have a policy that allows `SNS:Publish` on the logs bucket resource.

To create an SNS Topic:

1. Open the **Amazon SNS** console and navigate to **Topics**.
2. Click **Create topic**.
3. In the **Create topic** section, enter a name and description for the topic.
4. In the **Access policy** section, set who can publish and subscribe to the topic to **Everyone**.

   Best Practice - Check Point recommends to limit the publishing and subscription policy when the onboarding is done.
   For more information on SNS Access Policy, see AWS documentation.
5. Click **Create topic**.

Connecting a bucket to an SNS topic

You can connect an SNS topic and an S3 bucket with an event notification. The notification has to include the event type **Put** in the Object creation events (`s3:ObjectCreated:Put`).

Make sure the prefix filter of the notification is sufficient and includes all desired logs.

To attach the SNS topic to the S3 bucket with the event notification:

1. In the AWS Console, navigate to **Amazon S3** > **Buckets**.
2. In the buckets list, open the applicable bucket and go to the **Properties** tab.
3. Scroll to the **Event notifications** section and click **Create event notification**.
4. Below **General configurations**, enter the details of your setup,, and for **Event types** select **Object creation - Put**.

   Note - AWS does not allow overlapping prefixes for the same event type.

5. For **Destination**, select the **SNS topic** and specify it by selection or with an ARN.
6. Click **Save changes**.

Granting Permissions to CloudGuard Intelligence

Add the permissions below to the CloudGuard trust role that you created during onboarding.

1. To recover the trust role ARN, open the environment page, click **Edit Credentials** and copy the **Role ARN** value.
2. Add these permissions:
   - Action: `s3:GetObject`

     Resource: ARNs of all onboarded and to be onboarded buckets
   - Action: `sns:Subscribe`, `sns:Unsubscribe`

     Resource: ARNs of all of the onboarded SNS topics and the topics about to be onboarded
   - Action: `kms:Decrypt`

     Resource: ARNs of the KMS keys encrypting the buckets or the CloudTrail trails (if available)

Request

```
POST /v2/view/magellan/magellan-custom-onboarding
```

```
{
"bucketName": " intelligence-onboarding-*****-*****-******* ",
"bucketAccountId": "5*****************9",
"topicArn": " arn:aws:sns:us-east-1:5***********9:********",
"cloudAccountIds": [  ],
"onboardingType": "Cloudtrail"
}
```

For API documentation and code examples, see API Reference.

Authorization

Basic Authorization: Use the API key and secret as username and password.

Parameters

- **bucketName** - Name of the S3 bucket that stores the Flow Logs or CloudTrail logs
- **bucketAccountId** – AWS account ID that contains the S3 bucket (must be onboarded to CloudGuard)
- **topicArn** – ARN of the SNS topic that receives event notifications from the S3 bucket
- **cloudAccountIds** – For a centralized S3 bucket, the cloud account IDs of the other AWS accounts that send log files to the centralized S3 bucket.
- **onboardingType** – CloudTrail or Flow Logs for Account Activity or Traffic Activity

Response

```
200 – OK
```

Onboarding Verification

When done, make sure that:

- The subscription is added to the SNS topic.
- The new logs of the onboarded AWS account start to appear in the CloudGuard portal in **Events** > **Account Activity** or **Network Traffic**. This can take less than 30 minutes.

Removing Intelligence from AWS Environments with API

Offboarding removes the SNS subscription between AWS and CloudGuard Intelligence. After the SNS subscription is removed Intelligence no longer receives data from CloudTrail or Flow Logs for that specific CloudGuard account.

Note - Offboarding does not delete the SNS topic, S3 Buckets, or S3 Event Notifications.

To offboard an AWS account from CloudGuard Intelligence, send the API request through an API platform, such as Postman, or use code to run the API request automatically.

Prerequisites

Important - If you did onboarding with a CloudFormation stack, you must first do the steps in "On your AWS account" in *"Removing Intelligence from AWS Environments" on page 588*

Before offboarding your AWS environments with API, make sure that you have prepared:

- CloudGuard account information: API key and secret for your Dome9 account, see *"V2 API" on page 840*.
- AWS account information: ID of your AWS account (Account Number).
- Make sure the AWS environment shows in your CloudGuard account. Go to **Assets** > **Environments**. If the selected environment is connected to CloudGuard Intelligence, then below **Network Traffic** or **Account Activity** a green check mark shows.

Notes:

- Before you do these steps, it is necessary to remove the onboarded stack on your AWS account.
- Offboarding from Intelligence is a full offboarding of all S3 buckets for network traffic and account activity.

To offboard manually with API

Request

```
POST /v2/view/magellan/disable-magellan-for-cloud-account
```

```
"cloudAccountId": "......"
"Vendor": "AWS"
```

For API documentation and code examples, see API Reference.

Authorization

Basic Authorization: Use the API key and secret as username and password.

Parameters

- **cloudAccountId** - AWS account ID that contains the S3 bucket, which must be onboarded to CloudGuard Intelligence.
- **vendor** - Name of the cloud provider.

Response

```
200 – OK
```

To offboard automatically with code through API

Check Point recommends this option for users with multiple CloudGuard Intelligence accounts.

Requests

- For API documentation and code examples, see API Reference.
- For base URL information, see *"REST API" on page 907*.

Authorization

Basic Authorization: Use the API key and secret as username and password.

Parameters

- **cloudAccountId** - AWS account ID that contains the S3 bucket, which must be onboarded to CloudGuard Integration Guide.
- **vendor** - Name of the cloud provider.

Response

```
200 – OK
```

Offboarding Verification

When the offboarding is complete, make sure the subscription is removed from your AWS SNS topic.

Custom Onboarding of AWS Environments to Intelligence

This topic describes how to onboard an AWS environment with the manual onboarding experience. For the automated unified onboarding process, see *"Onboarding AWS Environments to Intelligence" on page 572*.

Your AWS environment has to be onboarded to CloudGuard before you can onboard it to Intelligence. If your account is not onboarded, follow the instructions in *"Unified Onboarding of AWS Environments" on page 54*.

Intelligence uses *"VPC Flow Logs" on page 382* and CloudTrail logs from your AWS account. These logs have to be connected to an AWS S3 bucket.

In the onboarding process below, you add an IAM policy to your AWS environment.

You must do some of the onboarding steps processes in the AWS console and other steps in the CloudGuard portal to onboard information from the selected AWS accounts to Intelligence.

Note - You must onboard Flow Logs and CloudTrail separately for each account.

Custom Onboarding

During the Custom Onboarding process, CloudGuard receives permission to create a subscription to an SNS topic and retrieve logs from the S3 bucket that sends logs to this SNS topic. This mode usually applies to three primary use cases:

- You have multiple environments that send logs to one (centralized) S3 bucket. The AWS environment that has the centralized S3 bucket and includes logs from all other connected accounts is your **Root Account**.

  During the onboarding process, you can select to onboard some accounts that send logs to the centralized bucket. Afterward, to onboard one of the accounts, start the onboarding wizard from the Root Account's page and not on the page of the account to onboard.

- You use a non-default prefix to organize data in the S3 bucket that holds your logs.
- You need to send your logs to another third-party destination, for example, to a SIEM. For a specific prefix, AWS only supports Event Notification to one destination. You can send the logs to an SNS topic and send them through this procedure to different subscribers.

Custom Onboarding includes these steps:

- **Prerequisites** - Make sure you have all required components before you start.
- **Configuration** - Configure an SNS topic: use the existing topic or create a new one if you do not have it and attach it to the S3 bucket. Note that only one SNS topic for each bucket is allowed.
- **Buckets** - Select the centralized bucket that holds your logs and sends events to the SNS topic.
- **Accounts** - Select the cloud accounts logs that you want to onboard to Intelligence.

  Note - You can have some **Connected** accounts that send their logs to the centralized S3 bucket of the **Root Account**. On the **Accounts** page, you can select only those accounts that are relevant for onboarding to Intelligence.
- **IAM Policy** - Prepare the IAM policy for CloudGuard Intelligence.
- **Summary** - Review the components to be onboarded to Intelligence.

Known Limitations

- The centralized S3 bucket cannot send events to two SNS topics. One S3 bucket = one SNS topic.
- You cannot onboard an account to Intelligence if you use an encrypted SNS.

For these and other CloudGuard limitations, see *"Known Limitations" on page 924*.

Onboarding to Account Activity with CloudTrail

Follow these steps in CloudGuard to enable Account Activity with CloudTrail:

1. In CloudGuard, click the **Assets** menu and make sure the **Environments** page opens.
2. Select the AWS environment that you want to onboard to Intelligence. For the centralized bucket onboarding, this environment must be your root account.
3. In the environment row and the **Account Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   Alternatively, you can click and enter the environment page. On the top right menu, click **Add Intelligence** and select **CloudTrail**.
4. In the top banner of the onboarding wizard, click the link for **custom onboarding (manual)**.
5. Follow the on-screen instructions to complete the wizard.

Onboarding to Traffic Activity with Flow Logs

Follow these steps in CloudGuard to enable Traffic Activity with Flow Logs:

1. In CloudGuard, click the **Assets** menu and make sure the **Environments** page opens.
2. Select the AWS environment that you want to onboard to Intelligence. For the centralized bucket onboarding, this environment must be your root account.
3. In the account row and the **Traffic Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   As an alternative, you can click and enter the environment page. On the top right menu, click **Add Intelligence** and select **Flow Logs**.
4. In the top banner of the onboarding wizard, click the link for **custom onboarding (manual)**.
5. Follow the on-screen instructions to complete the wizard.

Troubleshooting Intelligence Onboarding

You completed all the steps in the Onboarding wizard, but no logs show in the CloudGuard portal.

Possible causes:

- **IAM Role Permissions**

  The `CloudGuard-Connect` IAM role requires additional actions to allow Intelligence to get objects.

  Make sure that you have the `CloudGuard-for-intelligence` policy attached:

  - S3 bucket policy statement has explicit permission to `GetObject`
  - S3 bucket policy statement has the permission to `PutBucketNotification`

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Dome9S3ForLogic",
            "Action": [
                "s3:GetObject",
                "s3:PutBucketNotification"
            ],
```

```
            "Effect": "Allow",
            "Resource": []
        }
    ]
}
```

- **SNS Topic Permissions**

  - Make sure that the policy allows everyone to publish and subscribe to this SNS topic.

    Note - Check Point recommends that you limit the publishing policy after you complete the onboarding process.

  - Make sure that the SNS topic is configured and set as an Event Notification destination. The Event Notification event type that CloudGuard Intelligence requires is `Put: s3:ObjectCreated:Put`. Set the destination to the applicable SNS topic.

  - Make sure that the SNS access policy allows the `CloudGuard-Connect` IAM role to subscribe to the SNS. For this, update the SNS access policy and the `CloudGuard-Connect` IAM role policy to allow `SNS:Subscribe` and `SNS:Unsubscribe`.

  - Make sure that the subscription is confirmed in the SNS topic. For this, on the **Subscription** tab, make sure that the subscription status is **Confirmed**. If it is not confirmed, click **Request confirmation** and wait for about ten minutes to get the confirmation from CloudGuard.

- **Encryption Issues**

  - Make sure that logs/S3 buckets are not encrypted with a custom CMK (customer master key). If they are encrypted, add `kms:Decrypt` permissions of the specific key to the CloudGuard role.

  - Make sure that the SNS topic is not encrypted with a CMK. CloudGuard cannot retrieve the logs that are sent to an encrypted SNS topic.

See also *"Troubleshooting" on page 585*.

If the problem continues, contact *Check Point Support Center*.

Manual Removing of Intelligence from AWS Environments

This topic describes how to remove Intelligence from your AWS environment with a legacy onboarding experience. For a new experience with a simple onboarding process, see *"Removing Intelligence from AWS Environments" on page 588*.

You can remove Intelligence from your AWS environments. As a result, CloudGuard stops to receive all Account activity and Traffic activity (CloudTrail and Flow Logs) from your environment.

- For environments with Standard Onboarding of Intelligence, the process removes the S3 Event Notification from all S3 buckets on this account.
- For environments with Custom Onboarding of Intelligence, the process removes the subscription to the SNS topic to which the S3 bucket sent logs.

To remove Intelligence:

1. Navigate to **Assets > Environments** and find your AWS environment with **Filter** and **Search** fields.
2. Click the environment to enter it.
3. In the top right menu, click **Remove Intelligence**.

   A verification window opens.
4. Click **Remove** in the verification window.
5. CloudGuard notifies you of the successful removal of Intelligence from your environment.

No more account logs are sent to Intelligence, and you cannot see the existing logs on the CloudGuard portal.

Note - CloudGuard stores the existing logs until the end of your retention period. If you onboard your account to Intelligence again during this retention period, you can see the logs for the period before the offboarding.

Azure Intelligence

Onboarding

"Onboarding Azure Subscriptions to Intelligence" on page 601

Offboarding

"Removing Intelligence from Azure Subscriptions" on page 608

Onboarding Azure Subscriptions to Intelligence

You can use Intelligence to do an analysis of network and other account activities in Azure subscriptions. For this, onboard the subscriptions to Intelligence. This process creates a connection between Intelligence, Azure activity logs, and Azure network traffic logs. You can do this after *"Onboarding Azure Subscriptions" on page 169* to CloudGuard.

How it Works

During the Intelligence onboarding, CloudGuard sends you an ARM template to execute. The ARM template makes onboarding simple, and it is not necessary to give CloudGuard special permissions for your subscriptions. Then CloudGuard creates a `SystemTopic` for each of the selected storages if this does not exist in your environment, and an `EventGridSubscription` for each of the selected log types. For example, if you select a storage account *MyAccount* and logs of the types *Activity Log* and *Audit Log*, the resources created in the ARM are one `SystemTopic` and two `EventGridSubscriptions`.

The ARM gives the CloudGuard App Registration these permissions in the scope of the selected storage and system topics:

- Microsoft.EventGrid/systemTopics/eventSubscriptions/write
- Microsoft.EventGrid/systemTopics/eventSubscriptions/delete
- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read

Important - For existing Azure subscriptions previously onboarded to Intelligence, the subscription storage key is removed and replaced with App Registration. This is done to allow least privilege access to the Storage Account for the Activity Logs.

Network Security Group

You can onboard each individual Azure account to Network Activity through Network Security Group (NSG) that generates traffic logs. The NSG is the unit on which you can configure the Flow Logs in Azure.

Centralized Storage Account

A centralized storage account stores logs for multiple Azure accounts. Use the onboarded centralized storage account to easily monitor and run log analysis of all your Azure subscriptions that send logs to it. For Traffic Activity, the centralized storage supports pulling NSG flow logs and VNet flow logs.

Prerequisites

1. **Storage Account Creation**

   Make sure you have a storage account. To create a storage account:

a. In the Azure portal, select **Storage Accounts** and click **Create** in the top menu.

b. Enter all required information. Make sure to select only necessary storage settings to reduce the subscription costs.

2. **Storage account kind is StorageV2**

   Make sure your storage account kind is **StorageV2** (general purpose v2).

   a. In the Azure portal, open your storage account.

   b. Navigate to **Settings** > **Configuration** and verify that the **Account kind** is **StorageV2 (general purpose v2)**.

   c. If the Account kind is Storage (general purpose v1), click **Upgrade**.

3. **For NSG Network Activity only: Flow Logs version 2**

   You have to select **Flow Logs version 2** format for NSG flow logs to onboard them.

   For this:

   a. In the Azure portal, open your NSG.

   b. Navigate to the **NSG flow logs** and click the name of your flow logs to open the settings.

   c. On the **Flow log settings** window, below **Flow logs Version**, make sure the **Version 2** is selected.

4. **For Flow Logs with the centralized storage**

   NSG Flow Logs:

   a. In the Azure portal, look for the **Network security groups** in the search bar.

   b. Select the NSG to be onboarded to Intelligence.

   c. From the left menu, select **Monitoring** > **NSG Flow Logs**.

   d. Select the appropriate NSG from the list.

   e. In the **Flow Logs Version** section, select **Version 2**.

   f. Make sure that your storage account kind is StorageV2 (general purpose v2).

      - Navigate to **Settings** > **Configuration** and verify that the **Account kind** is **StorageV2 (general purpose v2)**.
      - If the Account kind is Storage (general purpose v1), click **Upgrade**.

   g. Click **Save**.

h. Navigate to the **NSG flow logs** and click the name of your flow logs to open the settings.

i. On the **Flow log settings** window, below **Flow logs Version**, make sure the **Version 2** is selected.

VNet Flow Logs:

a. In the Azure portal, search for **Virtual networks** in the search bar.

b. Click the VNet to be onboarded to Intelligence.

c. From the left menu, select **Monitoring** > **Virtual network Flow Logs**.

d. Click **Create** and then select the appropriate VNet from the **Target resource** list.

e. In the storage account field, enter the name of the centralized storage.

f. Click **Review + create** and select **Create**.

Enabling Account Activity with Activity Logs

You can find Account Activity Logs in Azure in different log categories, such as:

- Azure Activity log
- Microsoft Entra ID - Sign-In logs
- Microsoft Entra ID - Audit logs
- Azure Storage Analytics logs (classic)

CloudGuard wizard allows you to onboard several types of logs to Intelligence. If you select to onboard only one type of logs, then afterward you can start the wizard again and onboard other types of logs.

To onboard one type of Azure Activity Logs to Intelligence:

1. Navigate to the **Assets** > **Environments** page.
2. Use the filter **Platform**: **Azure** or the search bar to display the Azure subscription that you want to onboard to Intelligence.
3. In the subscription row and the **Account Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   As an alternative, you can click and enter the Azure subscription page. In the top right menu, click **Add Intelligence** and select **Activity Logs**.
4. Follow the on-screen instructions to complete the wizard.

Enabling Traffic Activity with Flow Logs

1. Navigate to the **Assets** > **Environments** page.
2. Use the filter **Platform**: **Azure**or the search bar to display the Azure subscription that you want to onboard to Intelligence.
3. In the subscription row and the **Traffic Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   As an alternative, you can click and enter the Azure subscription page. In the top right menu, click **Add Intelligence** and select **Flow Logs**.
4. Follow the on-screen wizard to complete onboarding for Intelligence.

When you complete these steps, CloudGuard starts the onboarding process for Intelligence. It can take several minutes.

Wizard Stages

1. **Welcome** - Read carefully the onboarding prerequisites and make sure that the Azure subscriptions to be onboarded meet all the required conditions.
2. **Storage** or **Network Security Group (NSG)**
   - NSG
     a. Select an NSG to onboard.
     b. See the status of each NSG. If the NSG is not connected to a storage account, you can create a new storage account.

     For each network, selecting an NSG selects all NSGs connected to the same storage. CloudGuard onboards the entire storage account to Intelligence and receives network traffic logs from all NSGs that send logs to this storage.

- Centralized Storage Account

  a. Select **one**:

  - **Enter Manually -** Select this option to manually enter the Account ID of a storage account from Azure. Click **Add**, paste the Account ID, and then continue to the **ARM Template** step of the Wizard.
  - **Select from List** - Select this option to choose storage accounts from a list that CloudGuard generates. Continue until the end of this procedure.

  b. Select one or more storage accounts to onboard.

  c. See the status of each storage account. For more information about the storage account status, see *"Storage Status" on the next page*.

  d. For account activity, select the log types.

  e. Set **Auto Onboard** to **ON** to let CloudGuard:

  - detect all onboarded Azure subscriptions that send logs of a certain type to this centralized storage account
  - automatically onboard these subscriptions to Intelligence

  Important - For existing Azure subscriptions previously onboarded to Intelligence, the subscription storage key is removed and replaced with App Registration. This is done to allow least privilege access to the Storage Account for the Activity Logs.

  The toggle button is enabled only for new storages that you select to onboard. For more information, see *"Automatic Onboarding" on the next page*.

3. **ARM Template** - Based on your storage account configuration, CloudGuard generates a custom ARM template. Click the link to deploy it. Click **Check Now** to see the current status of the configuration.
4. **Subscriptions** (For centralized storage accounts only) - Select the Azure subscriptions that you want to onboard and see their logs. You can select only those subscriptions that are not connected yet. For more information about subscription status, see *"Subscription Status" on page 607*.
5. **Azure Network Firewall** - Allow Intelligence access to your Azure storage account. Click **Check Access** to see the access status. For more information on Firewall IPs, see FAQ.
6. **Summary** - Make sure the onboarding is successful.

   Afterward, you can see the traffic activity or account activity on the Logs page when you navigate to **Events** > **Cloud Logs** > **Network Traffic** or **Account Activity**.

Storage Status

When you select from the available storage accounts for onboarding, you can see their status as follows:

- **Not connected** - Includes logs that you want to onboard, and none of them are onboarded to Intelligence. For account activity, you can select the applicable logs and continue to the next step.
- **Connected** - The account is onboarded with all applicable logs. For account activity, you can see the types of onboarded logs in the **Log Types** column.
- **Partially Connected** - Only part of the logs are sent to Intelligence.
  - For account activity:
    - Not all available log types are sent to Intelligence.
    - The storage is onboarded with non-centralized configuration. This means that the event subscription sends logs from a specific subscription or NSG and not from all subscriptions.
  - For network activity (centralized onboarding only) - The account is onboarded to Intelligence with specific NSGs and only their logs are sent to Intelligence. It is necessary onboard the account again with centralized configuration, so that all network traffic logs (from all NSGs) are sent to Intelligence.

  Note - If the partially connected storage is configured with non-centralized configuration, during the onboarding step with ARM template, the storage is changed and considered centralized.

Note - Regardless of the type of Azure log that you select, CloudGuard retrieves the logs from Azure Storage Accounts.

Automatic Onboarding

CloudGuard can detect all onboarded Azure subscriptions that send logs of a certain type to the centralized storage account and automatically onboard them to Intelligence.

- For accounts that send various activity logs to the centralized storage account, CloudGuard can automatically onboard them to Account Activity.
- For accounts that send Flow Logs to the centralized storage account, CloudGuard can automatically onboard them to Traffic Activity.

If you do not select this option, you have to onboard each account to Intelligence manually.

To disable automatic onboarding with API, use the onboarding API (see *"Onboarding with API" on the next page*) and add to the request `isAutoDiscoveryEnabled=false`.

Subscription Status

- **Connected** - The subscription is already onboarded to Intelligence
- **Ready to be connected** - The subscription can be onboarded to Intelligence
- **Cannot be connected** - The subscription is not onboarded to CloudGuard, so it cannot be onboarded to Intelligence

Onboarding with API

You can use API to onboard Azure subscriptions to Intelligence.

For more information, see

- API for Account Activity Onboarding
- API for Traffic Activity Onboarding

Removing Intelligence from Azure Subscriptions

You can remove Intelligence from your Azure subscription. This process removes all subscription events created during onboarding from all Storage accounts. As a result, CloudGuard stops to receive all Account activity and Traffic activity (Activity Logs and Flow Logs) from your subscription.

To remove Intelligence:

1. Navigate to **Assets > Environments** and locate your Azure subscription with **Filter** and **Search** fields.
2. Click the environment to enter it.
3. In the top right menu, click **Remove Intelligence**.

   A verification window opens.
4. Click **Remove** in the verification window.
5. CloudGuard notifies you of successful removal of Intelligence from your environment.

No more logs are sent to Intelligence, and you cannot see the existing logs on the CloudGuard portal.

Notes:

- The Storage accounts that CloudGuard created during the Intelligence onboarding remain in your Azure account. You can delete them if you do not need the accounts.
- CloudGuard stores the existing logs till the end of your retention period. If you onboard your account to Intelligence again within this retention period, you can see the logs for the period prior to offboarding.

GCP Intelligence

Onboarding

"Onboarding GCP Projects to Intelligence" on page 610

Offboarding

"Removing Intelligence from GCP Projects" on page 616

Onboarding GCP Projects to Intelligence

You can use Intelligence to analyze the account activity of your Google Cloud Platform (GCP) project. For this, onboard the project to Intelligence. This process creates a connection between Intelligence and the GCP project. You can do this after *"Onboarding Google Cloud Platform Projects" on page 182* to CloudGuard.

Note - The onboarding process is done separately for Activity Logs and for Flow Logs.

Prerequisites for Onboarding

- For Intelligence Account Activity, enable Activity Logs on your GCP project.
- For Intelligence Traffic Activity, enable VPC Flow Logs on your GCP project.
- Make sure you have permissions to run the Google Cloud Shell.
- Consider one of the options for the onboarding scope:
  - **Standard** - To onboard one GCP project.
  - **Centralized** - To onboard multiple projects through the centralized Pub/Sub configuration.
- For the **Standard** scope, make sure that you have the **Owner** or **Editor** permissions.
- For the **Centralized** scope, make sure you have these permissions:
  - For a centralized project - **Editor**, **Pub/Sub Admin**, **Logging Admin** roles, or a custom role. See below *"Centralized Custom Role Permissions" on page 612*.
  - For other projects - **Logging Admin** role or a custom role. See below *"Custom Role Permissions" on page 612*.
- **If your project was created on or before April 8, 2021, expand here for more information**

  If your project was created on or before April 8, 2021, then you must grant the `roles/iam.serviceAccountTokenCreator` role to the Google-managed service account `service-{PROJECT_NUMBER}@gcp-sa-pubsub.iam.gserviceaccount.com` on the project. This allows Pub/Sub to create tokens.

  But if your project was created after this date, it is not necessary to grant this role because the service account has the `roles/pubsub.serviceAgent` role with the same permissions. For more information, see Google Cloud's Push Subscription documentation.

Creating a Custom Role

If you do not want to use more permissive built-in roles, you can create and use a custom role. This lets you minimize and control specific required permissions.

To configure a custom role:

1. Log in to the Google cloud console.
2. From the navigation menu, go to **IAM & Admin** > **Roles** and click **Create Role**.
3. Define mandatory permissions for the custom role, add specific permissions for the GCP services in use, and search & select permissions from the list below:

- **Centralized Custom Role Permissions**
  - iam.serviceAccounts.actAs
  - iam.serviceAccounts.get
  - iam.serviceAccounts.list
  - iam.serviceAccounts.create
  - iam.serviceAccounts.delete
  - pubsub.subscriptions.create
  - pubsub.subscriptions.delete
  - pubsub.subscriptions.get
  - pubsub.subscriptions.list
  - pubsub.topics.attachSubscription
  - pubsub.topics.create
  - pubsub.topics.delete
  - pubsub.topics.get
  - pubsub.topics.getIamPolicy
  - pubsub.topics.list
  - pubsub.topics.setIamPolicy
  - logging.sinks.create
  - logging.sinks.delete
  - logging.sinks.get
  - resourcemanager.projects.get
  - serviceusage.services.get
  - serviceusage.services.enable
- **Custom Role Permissions**
  - logging.sinks.create
  - logging.sinks.delete
  - logging.sinks.get
  - logging.sinks.list
  - resourcemanager.projects.get

4. Assign the custom role to the user. For this:
    a. From the navigation menu, go to **IAM & Admin** > **IAM** and click **Grant Access**.
    b. In **Add principals**, enter the email of the user assigned with the role.
    c. In **Assign roles**, search and select the custom role you created in step 3.
    d. Click **Save**.

Onboarding Account Activity to Intelligence with Activity Logs

1. Navigate to the **Assets** > **Environments** page.
2. Click **Add Filter** > **Platform** > **GCP** or use the search bar to show the project to onboard to Intelligence.
3. In the project row and the **Account Activity** column, click **Enable** to start the Intelligence onboarding wizard.

    As an alternative, you can click and enter the GCP page. In the top right menu, click **Add Intelligence** and select **Activity Logs**.
4. Follow the on-screen wizard instructions to complete onboarding for Intelligence.

    This involves:
    a. Selecting the onboarding scope: **Standard** or **Centralized**.
    b. **For the Centralized scope only:**
        i. Selecting a new or existing Pub/Sub.
        ii. Selecting one or more projects.
    c. Opening the Google Cloud Shell and authorizing it to run the script. The script is available on GitHub for review from the **Welcome** or **Shell Script** page of the onboarding wizard.
    d. Opening your Google Workspace.
    e. Copying a command from the CloudGuard wizard and pasting it in the Cloud Shell terminal. The command creates in the project necessary CloudGuard resources. When the process is done, the Cloud Shell terminal displays a confirmation message.
    f. Examining the deployment status with the **Check Now** button.
    g. When you see a message that the shell script process is finished, click **Onboard**.

5. After you get a message that the project is successfully onboarded, click **Finish** to close the wizard. CloudGuard suggests you onboard a new environment or add alerts for this environment. For more information on adding alerts, see *"Getting Started with Intelligence Policy" on page 569*.

When you complete these steps, CloudGuard starts the onboarding process for Intelligence. It can take several minutes. All GCP environments with onboarded Activity Logs appear in **Assets > Environments** with Account activity enabled.

Afterward, you can see the account activity on the Logs page, when you navigate to **Events** > **Cloud Logs** > **Account Activity**, select the project name, and click **Run**.

Onboarding Traffic Activity to Intelligence with VPC Flow Logs

1. Navigate to the **Assets** > **Environments** page.
2. Click **Add Filter** > **Platform** > **GCP** or use the search bar to show the project that you want to onboard to Intelligence.
3. In the project row and the **Traffic Activity** column, click **Enable** to start the Intelligence onboarding wizard.

   As an alternative, you can click and enter the GCP page. In the top right menu, click **Add Intelligence** and select **Flow Logs**.
4. Follow the on-screen wizard instructions to complete onboarding for Intelligence.

   This involves:

   a. Enabling VPC Flow Logs on your project.

   b. Selecting the onboarding scope: **Standard** or **Centralized**.

   c. **For the Centralized scope only:**

      i. Selecting a new or existing Pub/Sub.

      ii. Selecting one or more projects.

   d. Opening a repository in Google Cloud Shell and authorizing Google Cloud Shell to run the script. The script is available on GitHub for review from the **Welcome** or **Shell Script** page of the onboarding wizard.

   e. Copying a command from the CloudGuard wizard and pasting it in the Cloud Shell terminal. The command creates in the project necessary CloudGuard resources.

      When the process is done, the Cloud Shell terminal displays a confirmation message.

   f. Examining the deployment status with the **Check Now** button.

5. After you get a message that the project is successfully onboarded, click **Finish** to close the wizard. CloudGuard suggests you onboard a new environment or add alerts for this environment. For more information on adding alerts, see *"Getting Started with Intelligence Policy" on page 569*.

Error and Warning Messages

Warning / Error	Details	Corrective Actions
You do not currently have an active account selected.	You did not select the **Trust repo** when opened it in Cloud Shell.	1. Click the **Session information** icon in Cloud Shell and select **Return to default Cloud Shell**. 2. Close the Cloud Shell window and click this link to run again the script. The onboarding wizard provides the link in step 2 of the Create Resources page.
You do not appear to have access to project [project_ID] or it does not exist. Are you sure you wish to set property [core/project] to project_ID?	You copied an incorrect project_ID.	Follow the instructions in the onboarding wizard and copy the correct project_ID.

When you complete these steps, CloudGuard starts the onboarding process for Intelligence. It can take about 30 minutes. All GCP environments with onboarded Traffic Logs appear in **Assets > Environments** with Traffic activity enabled.

Afterword, you can see the traffic activity on the Logs page, when you navigate to **Events** > **Cloud Logs** > **Network Traffic**, select the project name, and click **Run**.

Removing Intelligence from GCP Projects

You can remove Intelligence from your GCP project. This process breaks the connection between CloudGuard and your GCP project logs.

As a result, CloudGuard stops to receive the account and network activity from your GCP project.

Prerequisites for Offboarding

Make sure you have the permissions corresponding to the type and scope of your GCP project onboarding.

For more information, see *"Prerequisites for Onboarding" on page 610*.

Offboarding

To remove Intelligence:

1. Navigate to **Assets > Environments** and find your GCP project with **Filter** and **Search** fields.
2. Click the project to enter it.
3. In the top right menu, click **Remove Intelligence**.

   The offboarding wizard opens.
4. Follow the on-screen instructions in the offboarding wizard to complete Intelligence offboarding.

   This involves:

   a. Selecting **Delete Resources** if it is necessary that CloudGuard remove all applicable resources it created in your project during onboarding. Otherwise, the resources stays in the project, and you continue to Step 5.

   b. Opening Google Cloud Shell.

   c. Running the script provided in the wizard.
5. Click **Next** to see the summary.
6. Click **Finish** to close the wizard.

No more logs are sent to Intelligence. You can see the existing logs on the CloudGuard portal until the end of the retention period based on your Intelligence license.

Kubernetes Intelligence

Onboarding

"Onboarding Kubernetes Clusters to Intelligence" on page 618

Offboarding

"Removing Intelligence from Kubernetes Clusters" on page 620

Onboarding Kubernetes Clusters to Intelligence

You can use Intelligence to do analysis on network activities in Kubernetes clusters. For this, onboard the clusters to Intelligence. This process creates a connection between Intelligence and Kubernetes network traffic logs. You can do this after *"Onboarding Kubernetes Clusters" on page 188* to CloudGuard.

Enabling Traffic Activity

1. Navigate to the **Assets** > **Environments** page.
2. Click **Add Filter** > **Platform** > **Kubernetes** or use the search bar to see the Kubernetes cluster that you want to onboard to Intelligence.
3. In the cluster row and the **Traffic Activity** column, click **Enable**.

   As an alternative, click and enter the cluster page.
4. On the **Blades** tab, in the **Threat Intelligence** row, move the slider to **On**.
5. The Threat Intelligence window opens. It contains a command for the agent installation on your cluster. Copy the command and run it in your cluster with the correct parameters for strings in **< >**.
6. In the Threat Intelligence window, click **Yes**.

When you complete these steps, CloudGuard starts the onboarding process for Intelligence. It can take several minutes.

Afterward, you can see the traffic activity on the Logs page, when you navigate to **Events** > **Network Traffic**, select the cluster name, and click **Run**.

Troubleshooting Kubernetes Intelligence

To verify the correct installation of the Kubernetes Intelligence agents:

1. In the CloudGuard portal, go to **Assets** > **Environments** and open the required cluster page.
2. On the Blades tab, make sure that all agents below **Threat Intelligence** have the **OK** status.
3. In the Kubernetes cluster, make sure that the agent resources below are installed in the specified namespace, and have the **Running** state.
   a. For Flow Logs DaemonSet, make sure that it is running on the correct number of nodes, based on the defined node selector, tolerations, and more.
   b. For the Inventory agent, make sure that one replica is running.

More Links

"Intelligence for Kubernetes Containers" on page 645

Removing Intelligence from Kubernetes Clusters

You can remove Intelligence from your Kubernetes cluster. This process removes all applicable CloudGuard agents deployed on your cluster during onboarding. As a result, CloudGuard stops to receive all Traffic activity (Flow Logs) from your cluster.

To remove Intelligence:

1. Navigate to **Assets > Environments** and find your Kubernetes cluster with **Filter** and **Search** fields.
2. Click the cluster to enter it.
3. On the Blades tab, in the Threat Intelligence row, move the slider to **Off**.
4. The Threat Intelligence window opens. It contains a command to uninstall the agent from your cluster. Copy the command and run it in your cluster with correct parameters for strings in **< >**.
5. In the Threat Intelligence window, click **Yes**.
6. CloudGuard notifies you of the successful removal of Intelligence from your cluster.

No more logs are sent to Intelligence, and you cannot see the existing logs on the CloudGuard portal.

Note: CloudGuard stores the existing logs until the end of your retention period. If you onboard your cluster to Intelligence again during this retention period, you can see the logs for the period before the offboarding.

CloudGuard Permissions for Intelligence

In some cases when Intelligence cannot retrieve logs from your account, you receive a notification email from CloudGuard. It usually has the account number and the error type. Based on these details, you can understand the issue and troubleshoot it.

After you restore the permissions, CloudGuard can continue the log collection on your account, see *"Reactivating Intelligence" on page 627*.

Access to AWS

AccessDenied Error

A notification email informs you that CloudGuard cannot get logs from your S3 bucket <storage-name> because of the `AccessDenied` error. This error means that you removed some permission that allowed the collection of your logs.

The most usual cause of *denied access* is changes made in the IAM Policy. To verify your IAM Policy, make sure that you have the correct *CloudGuard-for-intelligence* policy attached:

Reason 1:

The Intelligence policy connected to the CloudGuard trusted IAM role was changed or deleted by one of the users in the AWS account.

Solution for Reason 1

It is necessary to restore the removed permissions.

To restore the removed permissions:

1. In the AWS console, open the CloudGuard trusted IAM role.

   To find the trusted IAM role, open the CloudGuard portal. After the cloud account is opened, the ARN role shows in the wizard in **edit permissions**.

2. Add a new inline/managed policy with the necessary permissions based on your S3 buckets.

   `s3:GetObject` - Allow this procedure for all S3 buckets onboarded to Intelligence. You must include the bucket ARN and the bucket ARN with an '/*' suffix to show that permissions apply to all objects.

```
{
    "Sid": "IntelligenceRequiredBucketPermissions",
```

```
    "Effect": "Allow",
    "Action": [
        "s3:GetObject"
    ],
    "Resource": [
        "arn:aws:s3:::cloudguard-intelligence-example-s3-
bucket",
        "arn:aws:s3:::cloudguard-intelligence-example-s3-
bucket/*",
    ]
}
```

`kms:Decrypt` - Allow this procedure on all the KMS keys the ARN used to encrypt the S3 bucket files or the CloudTrail logs directly.

```
{
    "Sid":
"IntelligenceRequiredBucketOrCloudTrailDecryptPermissions"
    "Action": "kms:Decrypt",
    "Effect": "Allow",
    "Resource": [
        "arn:aws:kms:_____:1**********2:key/********-****-
****-****-************"
    ]
}
```

`sns:Subscribe` and `sns:Unsubscribe` - These permissions are not directly required to retrieve the log files from the S3 bucket. Rather, Intelligence applies these permissions in its use enforcement process. Allow these procedures for all of the SNS topics used to send the S3 bucket event notifications to the Intelligence SQS queue.

```
{
    "Sid":
"IntelligenceUsageEnforcementRequiredSnsPermissions",
    "Action": [
        "sns:Subscribe",
        "sns:Unsubscribe"
    ],
    "Effect": "Allow",
    "Resource": [
```

```
            "arn:aws:sns:_____:1**********2:______"
        ]
}
```

Important - Do not change the inline IAM policy added as part of the CloudFormation Intelligence onboarding stack. These changes could be overwritten by future onboardings, but append a new inline/managed IAM policy.

Reason 1.2:

A different cause is a change to the CloudGuard trust policy.

The trust policy of the CloudGuard trusted role was changed, which causes permission issues in CloudGuard. For example, if the trust policy does not grant permissions to the right CloudGuard AWS account or the External ID required in the condition does not align with the one CloudGuard sends. To see the External ID from CloudGuard, in **Cloud Account** go to **Edit Credentials**.

Solution for Reason 1.2

It is necessary to correct the CloudGuard role trust policy.

To correct the CloudGuard role trust policy:

1. Give `sts:assumerole` permission to the CloudGuard AWS Account ID.
2. Set the External ID condition to the right value.

   The trust policy shows in the AWS Console below the **Trust relationship** section for the specific IAM Role.

To find the trust role and External ID value that CloudGuard uses:

1. Open the environment page in the CloudGuard Portal.
2. Select the applicable AWS account and click **Edit Credentials**.

   The value shows below the Role ARN and External Id sections.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
```

```
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam:: 1**********2:root"
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "StringEquals": {
                    "sts:ExternalId":
"************************"
                }
            }
        }
    ]
}
```

Reason 2:

The S3 bucket resource-based policy was changed to include an explicit denial that overrides the intelligence permission to do `S3:GetObject`.

Solution for 2

Exclude the CloudGuard trusted IAM role from the explicit deny.

To see the resource-based policy for the S3 bucket:

1. Go to the bucket view.
2. Select the **Permission** tab in the AWS Console.

Reason 3:

Encryption was added to the S3 bucket or directly to the CloudTrail.

Solution for 3

Add a new inline/managed policy to the CloudGuard trusted IAM role with the `kms:Decrypt` permissions with the correct KMS key resource.

To find the KMS KEY ARN:

1. Go to related CloudTrail and get the ARN, which is below the Log file SSE-KMS encryption.
2. Go to **bucket** > **Properties** > **Default Encryption**.

Add a new inline/ managed policy with the `kms:decrypt` resource to the CloudGuard trusted IAM role.

```
{
    "Sid":
"IntelligenceRequiredBucketOrCloudTrailDecryptPermissions"
    "Action": "kms:Decrypt",
    "Effect": "Allow",
    "Resource": [
        "arn:aws:kms:_____:1**********2:key/********-****-****-
****-************"
    ]
}
```

Important - Do not change the inline IAM policy added as part of the CloudFormation Intelligence onboarding stack. These changes could be overwritten by future onboardings. To prevent this, append a new inline/managed IAM policy.

AmazonSecurityTokenServiceException.AccessDenied Error

This error means that you removed the CloudGuard AWS account from your trusted entities.

To add the account to the trusted entities:

1. Log in to the AWS console (aws.amazon.com).
2. Select **Services** and select the **IAM service**.
3. Click **Roles** and search for the Role created for CloudGuard (usually, *CloudGuard-Connect*).
4. To verify the External ID on the Role, click the **Trust relationships** tab.
5. Make sure that the External ID is the same as given on the CloudGuard portal.

   Note - The External ID must not be empty.

6. If the External ID is empty or it is necessary to change it, click **Edit trust relationship** and correct it as required.

For more information on troubleshooting the AWS account permissions, see AWS documentation at https://aws.amazon.com/premiumsupport/knowledge-center and look for:

- S3 troubleshoot 403
- S3 access denied bucket policy
- S3 accidentally denied access

Access to Azure

A notification email informs you that CloudGuard cannot get logs from your Azure Storage account, with one of these errors: *AuthorizationFailure* or *PublicAccessNotPermitted*.

The authorization error can occur when you change or remove an existing access key (storage account key) generated as part of the CloudGuard onboarding process.

Fixing Authorization Issues

Method 1

Set the *AllowBlobPublicAccess* property to **Enabled**.

Method 2

1. In your Azure storage account, navigate to **Security + networking** > **Networking**.
2. On the **Firewalls and virtual networks** page, in the **Public network access**, see which option is selected:
   - If **Disabled** is selected, change it to **Enabled from selected virtual networks and IP addresses** option. Make sure you have a proper IP rule that allows Intelligence collectors to collect log data from the storage account.
   - If **Enabled from selected virtual networks and IP addresses** is already selected, make sure you have a proper IP rule that allows Intelligence collectors to collect log data from the storage account.

   Note - If you change the **Enabled from all networks** option to **Enabled from selected virtual networks and IP addresses**, do it with caution as it impacts all access to the storage account.
3. In the **Firewall** section, enter an IP address based on your Data Center location. For the full list of the IP addresses, refer to FAQ.
4. Click **Save**.

Method 3

CloudGuard supports retrieving logs using App Registration credentials. Sometimes, storage accounts onboarded with App Registration credentials can have issues with access to the CloudGuard portal.

To correct the issues, follow the instructions in *"Invalid Credentials or Missing Permissions" on page 176*.

For more details on public read access, go to Azure Storage documentation, search for **configure anonymous read access**, and follow the instructions in the article.

Reactivating Intelligence

After you restore the permissions, do one of these:

- Onboard your environment to Intelligence again and follow the steps as in *"Onboarding AWS Environments to Intelligence" on page 572* or *"Onboarding Azure Subscriptions to Intelligence" on page 601*. It is not necessary to do an offboarding process before this.
- Contact Check Point with a request to reactivate your log collection.

Intelligence Filtering

You can filter the information shown in the Intelligence graph and table views of log data, to refine the query and remove clutter from the view. This helps you identify and concentrate on events of interest in the log data.

You can apply filters in different places in these views.

Quick filters on graph data

When you have selected or created a query in Intelligence (for any view), then you can filter the displayed graph view. These quick filters remove events from the displayed graph.

The quick filter box is below the query box.

The filter options change based on the specific graph (Traffic or Event activity).

You can select more than one filter option. In addition, you can clear filters that have been selected. The selected filters are applied (with a logical AND) to the underlying query data shown in the graph. The initial query statement is unchanged.

Click **Start** to re-run the query with the filter selection.

Filters on specific entities

You can filter graphs or log table views for specific entities, based on the type and value (for example, an instance with a specific name, or a specific IP address). These filters add phrases to the initial query statement.

In Graph views

In the Traffic or Event Activity graph views, you can select entities from the detail pane (on the right) and add them to the query statement.

To create a filter from a graph view

1. Click on an entity in a graph (for Traffic or Event activities), to show detail for it (in the detail pane, on the right).
2. Click on an entity in the detail or Statistics pane.
3. Select the logic to connect it to the query (AND, OR, NOT). The phrase is added to the query statement.

In Table views

In table views of log data (opened by clicking **OPEN LOGS** from the graph views), you can add query phrases for specific entities to the query statement.

To create a filter from a table view

1. Click on a record in a table (for either Traffic or Event activities) to see its details.
2. Click on an entity in the detail pane.
3. Select the logic to connect it to the query (AND, OR, NOT). The phrase is added to the query statement.

Intelligence Queries

Build Queries in Intelligence to hunt out specific threats from log files.

Queries

Intelligence uses sophisticated queries to filter the information from cloud logs, to search for information or events of interest. These queries are built with the *"Governance Specification Language (GSL)" on page 326*, equivalent to queries for *"Running an Assessment" on page 304*. You can use these queries 'out-of-the-box' to quickly visualize traffic on your cloud environments. For example,

- *Inbound traffic* - Shows all inbound traffic
- *Rejected traffic* - Shows all rejected traffic to or from your VPC
- *Malicious accepted traffic* - Shows traffic that was accepted by your network, that originated from malicious IP addresses (as determined by threat intelligence sources)

In addition, you can configure custom queries, to filter for specific information not covered by built-in queries.

Build Custom Queries

To create custom queries for Intelligence Traffic and Activity Explorer views, use a graphic query builder, or enter the query directly as text.

The examples below illustrate how to create queries with these methods.

Example 1: Create a Query with Graphic Query Builder

This example creates a query in the Traffic Explorer view.

Rules are built up in the **Rule GSL** box, based on entities and operators that appear below the box. The set of entities and operators that are shown varies incrementally based on the context of the query as you develop it.

1. In the query box at the top of the page, click **open editor**. This opens the GSL editor. The rule is built in the **Rule GSL** box, on the left. You build the rule incrementally. At each stage, the entities that you can select are shown below the box (based on the context of the rule as it is being built).

   On the right is a dictionary of all the entities and properties that you can select, and the data type for each (use this when you create a rule with Free Text).

2. Select the cloud provider.

3. Select the source (**vpcfl** or **cloudtrail**, for AWS). This is the source of the log information. For AWS accounts, **vpcfl** logs are used for network querie and **cloudtrail** logs for account activity queries.
4. Next, select a condition (where). This is the only option at this stage. After this, you can select the left parenthesis, to open a clause or a property (of the source entity).
5. Select a property from those shown (status/protocol/action/src etc.). In this example, select *src.* You can then select more properties, to qualify the *src* property.
6. Select a different property to qualify *src*. In this example, select *address*, giving *src.address.*
7. Select an operator (*=, like, regexMatch*) and an argument. In this example, select the function *isPublicCIDR()*, for which an operator. is not necessary. This gives the query *vpcfl where src.address isPublicCIDR().*
8. Click **OK** to close the editor. The query is placed in the query box, ready to run. Select an account, and a time frame for the query, and then click **Run** to run the query. The results show all traffic that originates from a public IP address. The results appear in the Network Log Explorer view.

Example 2: Create a Query by Entering Text Directly

You can enter the text for a query directly in the Free Text box. To create the same query as in the earlier example:

1. In the query box at the top of the page, click **open editor**. This opens the GSL editor.
2. Enter the text of your query in the text box. For example, enter *vpcfl where src.geolocation.countryname='China' and action='ACCEPT' or protocol isPrivateCIDR() or packets isEmpty()*
3. Click **OK**, to close the editor and go back to the Traffic or Events activity page. Click **Run** to run the query.

Remediation

CloudBots automatically correct compliance issues discovered in your cloud environments by CloudGuard compliance checks. You can configure your CloudGuard account to use CloudGuardCloudBots. For more information about CloudBots, see *"CloudBots" on page 317*.

Prerequisites:

To use CloudBots, you must launch the CloudBot stack on your cloud and CloudGuard accounts. For detailed instructions, see the CloudBots platform documentation.

Configuring Remediation for Intelligence

Step 1: Onboarding CloudBots CFT

This runs the Lambda of the CloudBots on your cloud account to create CloudBots Lambda and SNS topic.

Note - For the ARN key value, in the AWS console, go to **CloudFormation** > **Stack**. Select **Outputs**. The key is `InputTopicARN` and the value is the ARN that shows.

Step 2: Adding a Remediation Rule

You can add remediation for a specific rule in a ruleset or all rules in a ruleset. You limit remediation to specific environments, entities, or environments and entities.

To add a remediation rule:

1. Navigate to **CDR** > **Threat Monitoring** > **Remediation**.
2. Click **Create New Remediation**, in the top right.
3. Select the rules for which the remediation applies, from the given options. You can combine the options, so the remediation applies to the combination of all the selected options.
    a. **Ruleset** (mandatory)
    b. A specific **Rule** in the ruleset (optional, if missing, all rules are implied)
    c. Select an **Environment** that applies the remediation to rules in the selected ruleset only when the ruleset is applied to the selected environments.
    d. A specific **Entity**, by its entity ID (optional, if missing, all entities are implied); this selects all rules involving the selected entities

4. Select the CloudBot, from the list. If the CloudBot does not show, select **Custom**, and then add the name of the CloudBot, along with the runtime arguments. The CloudBot must be deployed in the selected environment, in the same folder as the other bots.
5. Add a comment (mandatory) and click **Save**.

To delete a remediation ruleset:

1. Navigate to **CDR** > **Threat Monitoring** > **Remediation**.
2. Select one or more remediations to delete and click **Delete Selected**.

Step 3: Attaching a Policy to the Remediation Rule

It is necessary to add a CloudGuard policy to the configured CloudBot. In addition, in this step, you create a notification with the SNS from Step 1.

1. Navigate to **CDR** > **Threat Monitoring** > **Polices** > **Add Policy**.
2. Select a cloud platform and click **Next**.
3. Select the environment and click **Next**.
4. From the **Ruleset** menu, select a ruleset and click **Next**.
5. In **Notifications**, select **Add Notification**.

   It is necessary to create a specific notification for the CloudBot remediation stack in which you must enter the Topic ARN configured in the AWS console. For more information about SNS notifications, see *"Getting Started with Intelligence Policy" on page 569*.
6. Click **Save**.

Intelligence Security Events

You can configure Intelligence to trigger an alert when specific events occur in your cloud or cluster network. You or other recipients receive this alert as an email or as a different type of notification, so that you can respond to the event almost immediately.

To receive alerts, you must set up a policy. The policy includes a ruleset with specific Intelligence alert definitions, which are applied to selected cloud environments (VPCs) or Kubernetes clusters. With this policy, you associate a notification that specifies where you want to receive the alerts. Intelligence includes a number of preconfigured, CloudGuard-managed rulesets and policies.

In the Intelligence menu, you can set up your rulesets and policies.

Benefits

- Automatic and continuous monitoring of your cloud environments and clusters based on queries configured for your enterprise needs
- Automatic generation of near real-time alerts based on specific events and thresholds, issued to user-configured notification targets
- Built-in rulesets that cover many of the same enterprise needs, to apply to your environments and Kubernetes clusters out-of-the-box

Malicious IP Classification

For Intelligence rules that identify malicious IPs, CloudGuard uses the Check Point's ThreatCloud technology. The table below explains the meaning of each IP category.

Class	Description
Unclassified	The service could not classify the IP. There is not sufficient data about this resource.
Adware	The IP domains operate in the gray areas of the law, collecting private data on the users, and show unwanted content or a website that contains sub-application to download.
Volatile	The IP domains contain malicious software, for example, hacking websites.
Benign	Legitimate IP that is not malicious.
CnC Server	Command and control of malware.
Compromised Server	Legitimate IP that was hacked and operates a malicious function.

Class	Description
Phishing	The IP domains attempts to get sensitive information such as usernames, passwords, and credit card details (and sometimes, indirectly, money), frequently for malicious reasons, by masquerading as a trustworthy entity in electronic communication.
Infection Source	The IP domains can infect their visitors with malware.
Web Hosting	The IP domains allow you to rent out space for websites to have your business in.
File Hosting	The IP domains allow you to rent out space for storage to have your business in.
Parked	The IP domains permanently do not have content. Possibly, they contain advertising content on pages that have been registered but do not (at this time) have initial content.
Scanner	The IP is a known Internet scanner.
Anonymizer	The IP is a known Tor (The Onion Router) anonymity proxy server.
Cryptominer	The IP domains are used for crypto mining.
Spam	The IP domains are used for spam.
Compromised Host	The victim's IP.

Actions

Adding a Rule to a Ruleset

1. Navigate to the **Rulesets** page in the **CDR** > **Threat Monitoring** menu.
2. Select the ruleset to which the rule is added (or create a new one based on the steps below).
3. Click **New Rule** in the top right.
4. Enter a name and description for the rule.
5. Optionally, enter remediation text for the rule, indicating what steps can be taken to resolve the issue indicated by the rule. Afterward, the Event page shows the text.
6. Select the severity of the rule.

7. Enter a definition for the finding. This consists of the following details:
    a. For AWS and Azure rules, the source of the alert for traffic or activity logs.
    b. The GSL statement for the alert.
    c. The entity on which the event occurred (source or destination). This is applicable for events from VPC Flow Logs only.
8. Click **Save**.

Creating a Policy

An Intelligence policy has a ruleset (containing event definitions), one or more environments on which the events are applied, and *"Notifications" on page 852* indicating where findings must be sent.

1. Navigate to the **Policies** page in the **CDR** > **Threat Monitoring** menu.
2. Click **Add Policy** on the right.
3. Select a platform on which the policy applies and click **Next**.
4. Select one or more environments on which the policy applies and click **Next**.
5. Select one or more rulesets for the policy from the list and click **Next**.
6. Select one or more notifications from the list.
7. Click **Save**.

Creating a Ruleset

An Intelligence ruleset is a set of rule definitions. The rules inquire about specific events in VPC or CloudTrail logs, based on a *"Governance Specification Language (GSL)" on page 326* definition.

Intelligence includes some built-in rulesets. You can include them in policies and apply them to your environment.

In addition, you can create customized rulesets for your specific needs.

1. Navigate to the **Rulesets** page in the **CDR** > **Threat Monitoring** menu.
2. Click **Add Ruleset** on the right.
3. Enter a name and description for the ruleset and select the platform.
4. Click **Create**.

Viewing Events

You can see the events generated by Intelligence policies on the **Threat & Security** page of the **Events** menu. This page shows the events only if you set **Include in the CloudGuard Events page** option in the notification attached to the policy.

See *"All Events" on page 120* for more details about the Events page.

Intelligence Entities

Intelligence uses information from CloudTrail and VPC Flow Logs (for AWS) and enriches it with more details, such as geolocation, malicious IP sources/destinations, etc.

The entities for these are described below.

Some common terms are defined here:

- **identity** - Initiator of an action
- **event** - The event or action captured in the log record(s)
- **issuer** - The procedure by which the identity issued a token, for example, Access Key, Role, or the AWS Console
- **malicious info class** - The class of malicious information, for example, 'Compromised Host', 'Anonymizer', 'Phishing'
- **request** - The request, for example using an API, to do an action
- **user agent** - The agent through which the request was made, such as the AWS Management Console, an AWS service, the AWS SDKs, or the AWS CLI.
- **target** - Target entity of an action

CloudTrail (Account Events)

Properties

Property	Description	Type	Required
event_error_message	error message text	STRING	No
event_id	event id	STRING	No
event_name	event name	STRING	No
event_status	status for event (for example, 'success')	STRING	No
event_time	time the event occurred	TIMESTAMP	Yes
event_type	event type	STRING	No
identity_account_id	identity (initiator of the action) account id	STRING	No
identity_assetid	target asset id	STRING	No

Property	Description	Type	Required
identity_id	identity id	STRING	No
identity_image	identity image	STRING	No
identity_name	identity name	STRING	No
identity_region	cloud region for identity	STRING	No
identity_tags	tags for identity	STRING	No
identity_type	identity type	STRING	No
identity_useragent	identity user agent	STRING	No
identity_vpc	identity VPC	STRING	No
issuer_id	issuer id	STRING	No
issuer_mfa	MFA applied to issuer	BOOLEAN	No
issuer_name	name of issuer	STRING	No
issuer_region	cloud region for issuer	STRING	No
issuer_sts_token	AWS STS token for issuer	STRING	No
issuer_token	issuer access token	STRING	No
issuer_type	issuer (method used to gain access) type	STRING	No
request_parameters	additional parameters of the request, for the action	STRING	No
src_address	source IP address	STRING	No
src_address_ geolocation_ countrycode	source country code	STRING	No
src_address_ geolocation_ countryname	source country	STRING	No
src_address_ip	source IP address	STRING	No

Property	Description	Type	Required
src_address_ maliciousinfo_class	source malicious information class	STRING	No
src_address_ maliciousinfo_ malwarefamily	source malicious information family	STRING	No
src_address_ maliciousinfo_owner	source malicious information owner	STRING	No
src_assetid	source asset id	STRING	No
src_image	source image	STRING	No
src_ismalicious	source is malicious	BOOLEAN	No
src_name	source name	STRING	No
src_region	source cloud region	STRING	No
src_type	source type	STRING	No
src_vpc	source VPC	STRING	No
target_account_id	account id for target	STRING	No
target_arn	ARN for target entity	STRING	No
target_assetid	target asset id	STRING	No
target_id	id for target	STRING	No
target_image	target image	STRING	No
target_name	target name	STRING	No
target_region	cloud region for target	STRING	No
target_tags	tags for target	STRING	No
target_type	target entity type	STRING	No
target_vpc	target VPC	STRING	No

VPC Flow Logs (Traffic Events)

Properties

Property	Description	Type	Required
account	account	STRING	No
action	action	STRING	No
availabilityzone	AWS availability zone	STRING	No
bytes	# of bytes	INTEGER	No
direction	direction of communication	STRING	No
dst_address	destination IP address	STRING	No
dst_asset_assetid	destination asset id	STRING	No
dst_asset_ availabilityzone	AWS availability zone for destination	STRING	No
dst_asset_description	destination asset description	STRING	No
dst_asset_ groupbysgsid	destination asset grouped by SG id	STRING	No
dst_asset_image	destination asset image	STRING	No
dst_asset_ispublic	destination asset is public	BOOLEAN	No
dst_asset_name	destination asset name	STRING	No
dst_asset_nics_id	destination asset NIC id	STRING	No
dst_asset_nics_ privateipaddress	destination private IP address	STRING	No
dst_asset_nics_ publicdnsname	destination public DNS name	STRING	No
dst_asset_nics_ publicipaddress	public IP address of destination asset	STRING	No
dst_asset_nics_sgs	destination asset Security Groups	STRING	No

Property	Description	Type	Required
dst_asset_nics_ subnet_subnetid	destination asset subnet id	STRING	No
dst_asset_region	destination asset region	STRING	No
dst_asset_subtype	destination asset subtype	STRING	No
dst_asset_tags	destination asset tags	STRING	No
dst_asset_type	destination asset type	STRING	No
dst_asset_vpc	destination asset VPC	STRING	No
dst_geolocation_ countrycode	destination country code	STRING	No
dst_geolocation_ countryname	destination country	STRING	No
dst_ismalicious	destination is malicious	BOOLEAN	No
dst_maliciousinfo_ class	destination malicious information class	STRING	No
dst_maliciousinfo_ malwarefamily	destination malicious information malware family	STRING	No
dst_maliciousinfo_ owner	destination malicious information owner	STRING	No
dst_port	destination port	INTEGER	No
eni	elastic network interface	STRING	No
event_date	event date	INTEGER	No
packets	# packets	INTEGER	No
protocol	protocol	INTEGER	No
region	cloud region	STRING	No
src_address	source address	STRING	No
src_asset_assetid	source asset it	STRING	No

Property	Description	Type	Required
src_asset_ availabilityzone	source asset cloud availability zone	STRING	No
src_asset_description	source asset description	STRING	No
src_asset_image	source asset image	STRING	No
src_asset_ispublic	source asset is public	BOOLEAN	No
src_asset_name	source asset name	STRING	No
src_asset_nics_id	source asset NIC id	STRING	No
src_asset_nics_ privateipaddress	source asset private IP address	STRING	No
src_asset_nics_ publicdnsname	source public DNS name	STRING	No
src_asset_nics_ publicipaddress	source asset public IP address	STRING	No
src_asset_nics_sgs	source asset Security Groups	STRING	No
src_asset_nics_ subnet_subnetid	source asset subnet id	STRING	No
src_asset_region	source asset region	STRING	No
src_asset_subtype	source asset subtype	STRING	No
src_asset_tags	source asset tags	STRING	No
src_asset_type	source asset type	STRING	No
src_asset_vpc	source asset VPC	STRING	No
src_geolocation_ countrycode	source country code	STRING	No
src_geolocation_ countryname	source country	STRING	No
src_ismalicious	source is malicious	BOOLEAN	No

Property	Description	Type	Required
src_maliciousinfo_class	source malicious information class	STRING	No
src_maliciousinfo_malwarefamily	source malicious information malware family	STRING	No
src_maliciousinfo_owner	source malicious information owner	STRING	No
src_port	source port	INTEGER	No
starttime	start time	TIMESTAMP	Yes
status	status	STRING	No
stream_owner	stream owner	STRING	No
vpc	the VPC for the CloudTrail	STRING	No

Intelligence for Kubernetes Containers

To use Threat Intelligence on your Kubernetes cluster, you have to onboard it to Intelligence. For more information, see *"Onboarding Kubernetes Clusters to Intelligence" on page 618*.

With Kubernetes Intelligence, you can:

- Visualize and analyze north-south and east-west network traffic for your Kubernetes cluster
- Identify communications with malicious addresses
- Monitor cross-namespace communication
- Identify port scanning

CloudGuard provides a preconfigured Intelligence ruleset and custom queries created with a graphical GSL-based query builder.

Supported Versions

Name	Version
Kubernetes	v1.16 and higher
OS	Linux kernel v4.1 and higher

Architecture

Kubernetes Intelligence includes these components:

- **Inventory agent** - A single-replica Kubernetes Deployment responsible to report inventory information on cluster resources to CloudGuard.
- **Flow Logs DaemonSet** - A DaemonSet of agents that do this:
  - Interact with the underlying cluster node to monitor IP traffic between the virtual network interfaces in the cluster
  - Upload crafted logs to CloudGuard for analysis

Note - Check Point distributes agents as Helm Chart (see https://github.com/CheckPointSW/charts) and associated Docker images (see a private container registry - quay.io/checkpoint).

Rulesets and Policy

CloudGuard shows alerts for security events found in the Intelligence logs as part of the **Threat & Security Events** table on the Events page. To see events related to your cluster, it is necessary to configure a Kubernetes Ruleset or use the preconfigured CloudGuard-managed **Kubernetes CloudGuard Best Practices** ruleset. Then you set up a Policy that associates the ruleset with one or more Kubernetes clusters and assigns a notification.

Kubernetes Intelligence rulesets are equivalent to other Intelligence Rulesets (see *"Intelligence Security Events" on page 634*).

Actions

Use instructions in Intelligence Security Events for a Kubernetes cluster where an environment is mentioned.

Note - Kubernetes Intelligence does not support Audit Logs. It uses network traffic Flow Logs and Kubernetes assets data.

Known Limitations

For a full list of known limitations, see Known Limitations

More Links

"Onboarding Kubernetes Clusters" on page 188

"Kubernetes Containers" on page 415

Code Security

To assess and scan your code at the earlier stages of development, you can use Check Point Code Security which brings CloudGuard security abilities to detect and prevent risk in cloud deployments into the CI/CD pipeline. Code Security provides one interface for various CI/CD security steps.

Code Security capabilities:

- Scans your Infrastructure as Code (IaC) templates for risks
- Checks your software for known vulnerabilities
- Scans your collaboration and productivity tools

Your development team and/or DevSecOps team can use Code Security to scan:

- **Git repos**, which semantics Code Security understands and as such is especially useful to scan source code.
- **Software libraries**, such as Node.js npm packages, and Java Jars.
- **Software deliverables**, such as Android apps pulled from Google Play.
- **Production artifacts**, such as logs and storage.
- **Containers**, such as Docker containers locally built or pulled from Docker Hub.

Code Security focuses on scanning the text irrespective of programming language. It does this either directly or with the Code Security Toolchain which deals with unpacking and reading complex input sources. In addition, you can use Code Security in other different scenarios, for example:

- Add Code Security to CI as a build step
- Code Security scans your codebase with an ever-growing array of detectors
- If matches are translated to findings (sensitive data, credentials, or other risks), Code Security can fail the build (or not) and pinpoints the problem
- Provide a full report for tracing the finding to the source file and position for easy risk assessment and mitigation

You can also use SPEQL (proprietary mini-detector query language) to enrich Code Security built-in detectors with your security policies. When you configure new detectors, you check them in to the code repository. For more information, see *"Building Detectors" on page 758*.

Secure by Design

It minimizes the surface area of attack by:

- Doing the right thing by default and having security-by-design. For example, Code Security never shows you secrets it found, it will just lead you there.
- Code Security never communicates your private data with the outside world, ever
- Code Security never stores, indexes, or offloads sensitive file contents to another place - in or out of your data centers (not even temporary files) - it does all the scanning in real time in-memory. This is why performance in Code Security is by design, to support security by design.
- Code Security is built with a safe, compiled, borrow-checked, and a programming language that has proven itself to excel, and is widely adopted in the security domain: Rust. No scripting languages, no toy languages, no compromise.

Input

An input source is any folder that hosts files.

File systems:

- **Git repositories** - pre-commit or on CI
- **Home folders** (~) - protect entire desktops
- **Document folders** - scanning legal documents
- **Cloud storage folders** - scanning files stored on Dropbox, Google Drive, and so on

And artifacts, such as:

- **Container file systems** - as a container security solution
- **Android apps** - as app static analysis engine
- **Websites** - as a remote scraping and monitoring security solution
- **Npm modules** - as a build/production verification layer after pushing a new module to production

Detectors

A detector is a way to formulate a security best practice, detection of secrets, and a logical way to do security by design.

Code Security comes with a premium, high-grade array of detectors and it also allows and encourages you and your teams to add detectors to it.

By letting you write your own custom detectors, we want to empower your team to build various policies such as:

- Identify and sanitize private customer IDs and data
- Special home-brew secrets for internal systems
- Employ policies such as "a test-only credit card with a certain pattern should only exist in the `/examples` folder"

For more information, see *"Building Detectors" on page 758*.

Zero Configuration

Code Security integrates with your existing environment with zero configuration (without any setup from your side):

- **CI** - Such as TravisCI, CircleCI, and others
- **Logging provider** - Such as Elastic
- **Alert provider** - Such as Sentry and PagerDuty
- **Automation Infrastructure** - By provision of a raw JSON stream of events

Platforms

Code Security is a single binary (around 12 MB). It has no external dependencies, no OS dependencies, and no networking requirements.

It supports a binary for every major platform and has no problem to add new ones.

Binaries

Code Security can run on:

- Linux (GNU + musl)
- macOS
- Windows
- BSD

You can have Code Security run on every hardware, on-premises or cloud.

Getting Started with Code Security

Check Point Code Security uses the engine named **spectral** to scan your environments. It does not require that you onboard your environment to CloudGuard.

To get started with Code Security:

1. Log in to CloudGuard.
2. From the menu, select **Code Security**. CloudGuard suggests you set up your Continuous Integration (CI).
3. Click **Setup your CI**. The on-screen wizard opens on the **Sources** page (see *"Sources" on page 664*).
4. Follow the instructions in the wizard to configure the integration.

To install Code Security:

Run the command based on the Data Center location of your CloudGuard account and the operating system:

United States (US)

For	Run this command
Homebrew on Mac	`brew tap spectralops/tap && brew install spectral`

Scoop on Windows	On Infinity Portal:
	`scoop install https://spectral-us.checkpoint.com/latest/scoop/spectral.json`
	On Dome9 Portal:
	`scoop install https://spectral-us.dome9.com/latest/scoop/spectral.json`

Mac and Linux	On Infinity Portal:
	`curl -L https://spectral-us.checkpoint.com/latest | sh`
	On Dome9 Portal:
	`curl -L https://spectral-us.dome9.com//latest | sh`

For	Run this command

Windows PowerShell	On Infinity Portal:
	`iwr https://spectral-us.checkpoint.com/latest/ps1 -useb | iex`
	On Dome9 Portal:
	`iwr https://spectral-us.dome9.com/latest/ps1 -useb | iex`

Europe (EU)

For	Run this command
Homebrew on Mac	`brew tap spectralops/tap && brew install spectral`

Scoop on Windows	On Infinity Portal:
	`scoop install https://spectral-eu.checkpoint.com/latest/scoop/spectral.json`
	On Dome9 Portal:
	`scoop install https://spectral.eu1.dome9.com/latest/scoop/spectral.json`

Mac and Linux	On Infinity Portal: `curl -L https://spectral-eu.checkpoint.com/latest | sh` On Dome9 Portal: `curl -L https://spectral.eu1.dome9.com/latest | sh`

Windows PowerShell	On Infinity Portal:
	`iwr https://spectral-eu.checkpoint.com/latest/ps1 -useb | iex`
	On Dome9 Portal:
	`iwr https://spectral.eu1.dome9.com/latest/ps1 -useb | iex`

Note - You can use the DSN to get the Code Security enterprise offering, for example: `curl -L https://spectral-us.dome9.com/latest/sh?dsn=<YOUR_DSN> | sh`

To scan a directory:

For a sample spectral-test directory, run:

```
$ mkdir spectral-test && cd spectral-test
$ $HOME/.spectral/spectral scan
✓  no matches found
  scanned 0 bytes and 0 files in 2ms
```

Best Practice - Do not run `curl|sh` without inspecting the install script (shell script).

To create a dummy secret:

Run:

```
$ echo AKIAIOSFODNN7EXAMPLX > foo.txt
$ $HOME/.spectral/spectral run
/Users/superhero/spectral-test/foo.txt
        0:20 Error        Visible AWS Key CLD001
```

It shows the file, the location (`0:20`), severity (`Error`), description and detector code (`CLD001`).

Code Security does not dump the actual secret or key to the console, or anywhere. If you want to view it, add a `SPECTRAL_SHOW_MATCH=1` environment flag before running.

```
$ $HOME/.spectral/spectral run
...
[your-project] SVC006 - Exposed PubNub Secret on Client Side App
- res/values/strings.xml
```

To perform a sample scan:

You can use the `codesec-goat` file to perform a sample scan. Unzip the `codesec-goat` file and run:

```
$ cd codesec-goat
$ $HOME/.spectral/spectral scan
```

Dashboard

The Dashboard provides an overview of all your organization's assets.

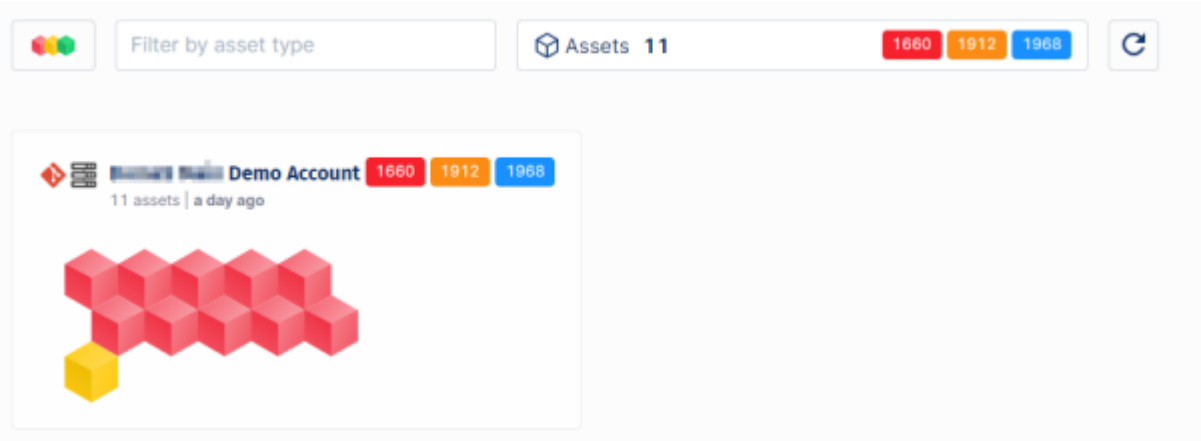

Each cube inside the cards represents an asset, like a repository or a container.

Code Security provides four different statuses for your assets:

- **Red** - An issue of type error was detected
- **Yellow** - An issue of type warning was detected
- **Blue** - An Issue of type Info was detected
- **Gray** - Code Security scanner did not scan the asset in the past seven days. This can be due to, for example, a user who did not run any build in the past week or deleted the spectral scanner from the CI. You must check with the Administrator or review your CI configuration.

The status filter in the upper-left corner can help you focus on assets of a specific status.

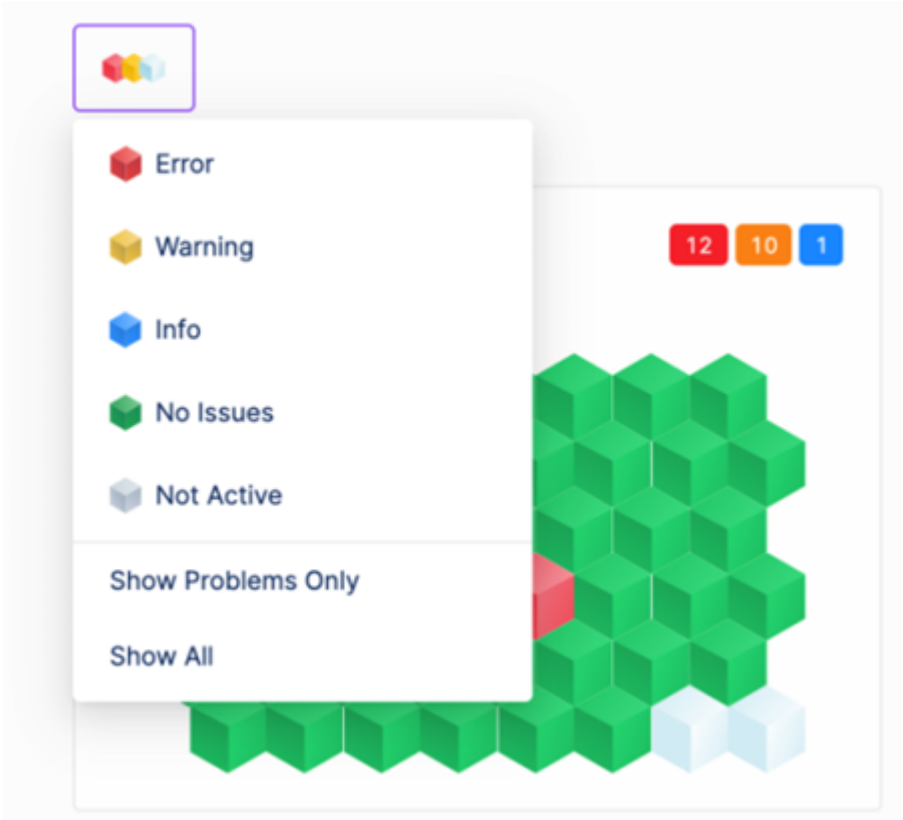

You can filter out any of the findings by clicking it. When you click one of the assets, the scan page of the asset opens.

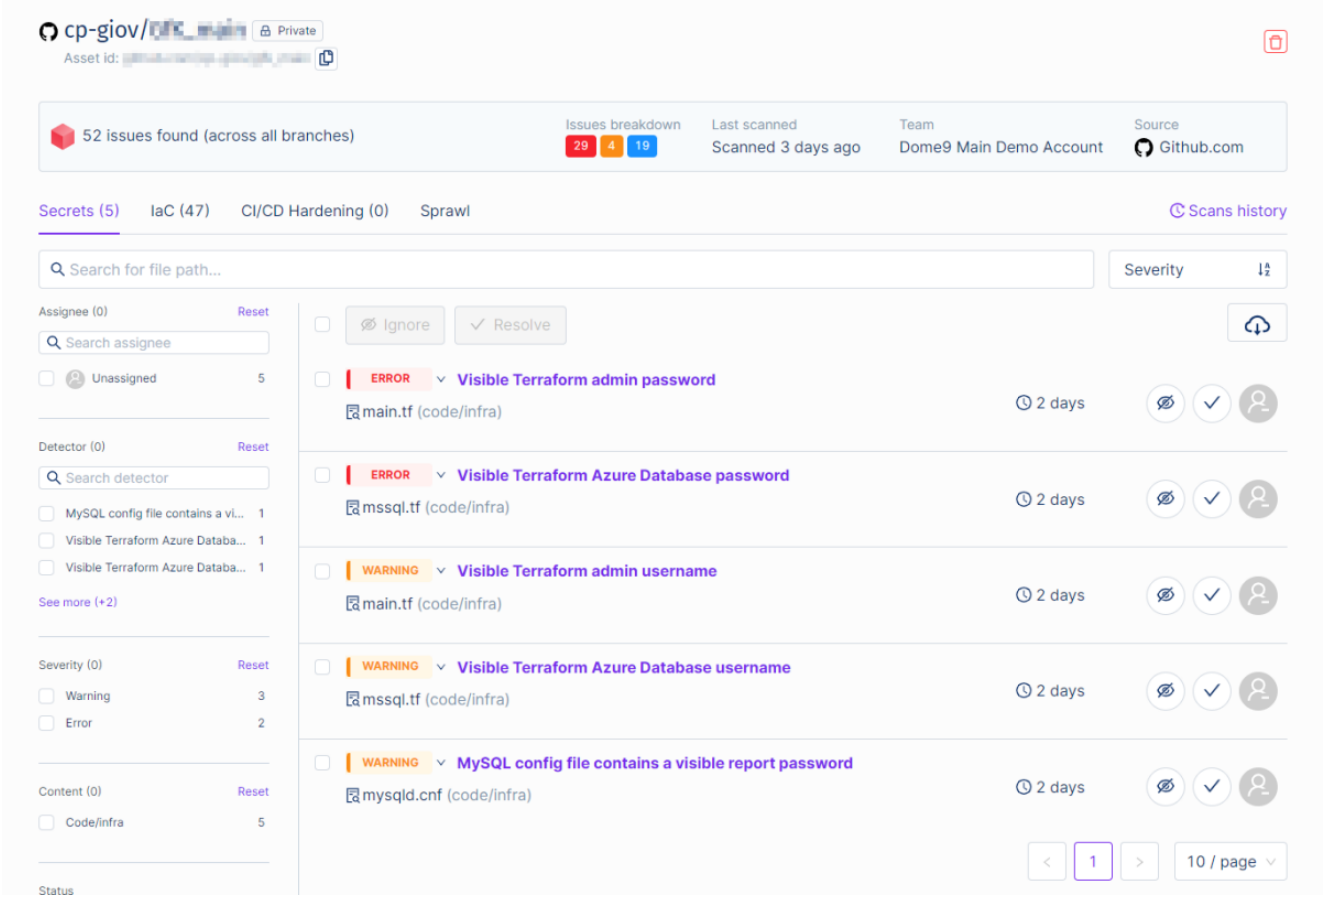

In the table, you can find the last scan results for a specific asset. For more details, see *"Assets" on page 659*.

Team and User Permissions

A Admin can assign roles to users and groups in **Settings** > **Organization**.

Role	Privileges
Owner	A superuser with full access to all Spectral features. Each organization can have only one Owner.
Admin	Has the same privileges as the Owner, except for modifying certain organization settings. You can grant Admin privileges to specific teams, or grant Admin privileges globally across the entire organization.
Member	Can view issues and take action on them, but only for assets they have permission to access.
Read Only	Has the same viewing permissions as a Member but can only modify their own account details, such as their personal API key or report subscriptions.

These are the privileges of the **Admin**, **Member**, and **Read Only** roles for specific features:

Feature	Admin	Member	Read Only
Assets scanning	Yes	Yes	Yes (only for unstaged scans)
Open tabs and do actions	Yes (can be restricted to specific teams)	Yes (only for their own team's assets)	View (only for their own team's assets)
Reports	Yes	Yes (only for their own team's assets)	View (only for their own team's assets)
Manage Personal Notifications Settings	Yes	Yes	Yes
Scan Configuration	Yes	View	View
Settings - Weekly/Daily reports	Yes	View	View
Generate API Keys	Team and Personal	Personal	Personal
Invite new users and manage pending invitations	Yes	No	No

Feature	Admin	Member	Read Only
Change user roles	Yes	No	No
Delete Asset	Yes	No	No
Integrations	Yes	No	No
Org Teams	Yes	No	No
Asset Mapping	Yes	No	No
Custom Rules	Yes	No	No
Remote Ignore	Yes	No	No

Teams and Asset Mapping

You can create a team and create a mapping between an asset and a team. That way, only the relevant users see the alerts and reports for the assets. For example, you can create a mapping for front-end code repositories only to front-end team members. An asset can be mapped to one team.

To create a team:

1. From the second toolbar from the left, click **Settings**.
2. From the **Settings** toolbar, click **Teams**.
3. Click **New Team**.

   The **Add team** sliding window opens.
4. Enter a Team name.
5. Add one or more webhook integrations for the team. For more information, see *"Code Security Integrations" on page 739*..
6. Click **Save**.

To create an asset mapping:

1. From the second toolbar from the left, click **Settings**.
2. From the **Settings** toolbar, click **Asset Mapping**.
3. Select one of these:
   - **Unmapped** - view assets that are not mapped to a team
   - **Mapped** - view assets that are currently assigned to a team

     Note - If you create a new mapping for a mapped object, the new mapping replaces the previous mapping.
4. **Optional -** In the filter assets bar, filter the assets by name.
5. Select one or more assets.
6. In the**Map to team** field, select or enter the name of a team.

Assets

Use the **Assets** page to overview, filter, and sort your assets. And, if applicable, see helpful information about IaC data in each asset.

If your asset has any IaC issues, you can find the **IaC** tag on that asset, and as top files with IaC issues (highest number of IaC issues).

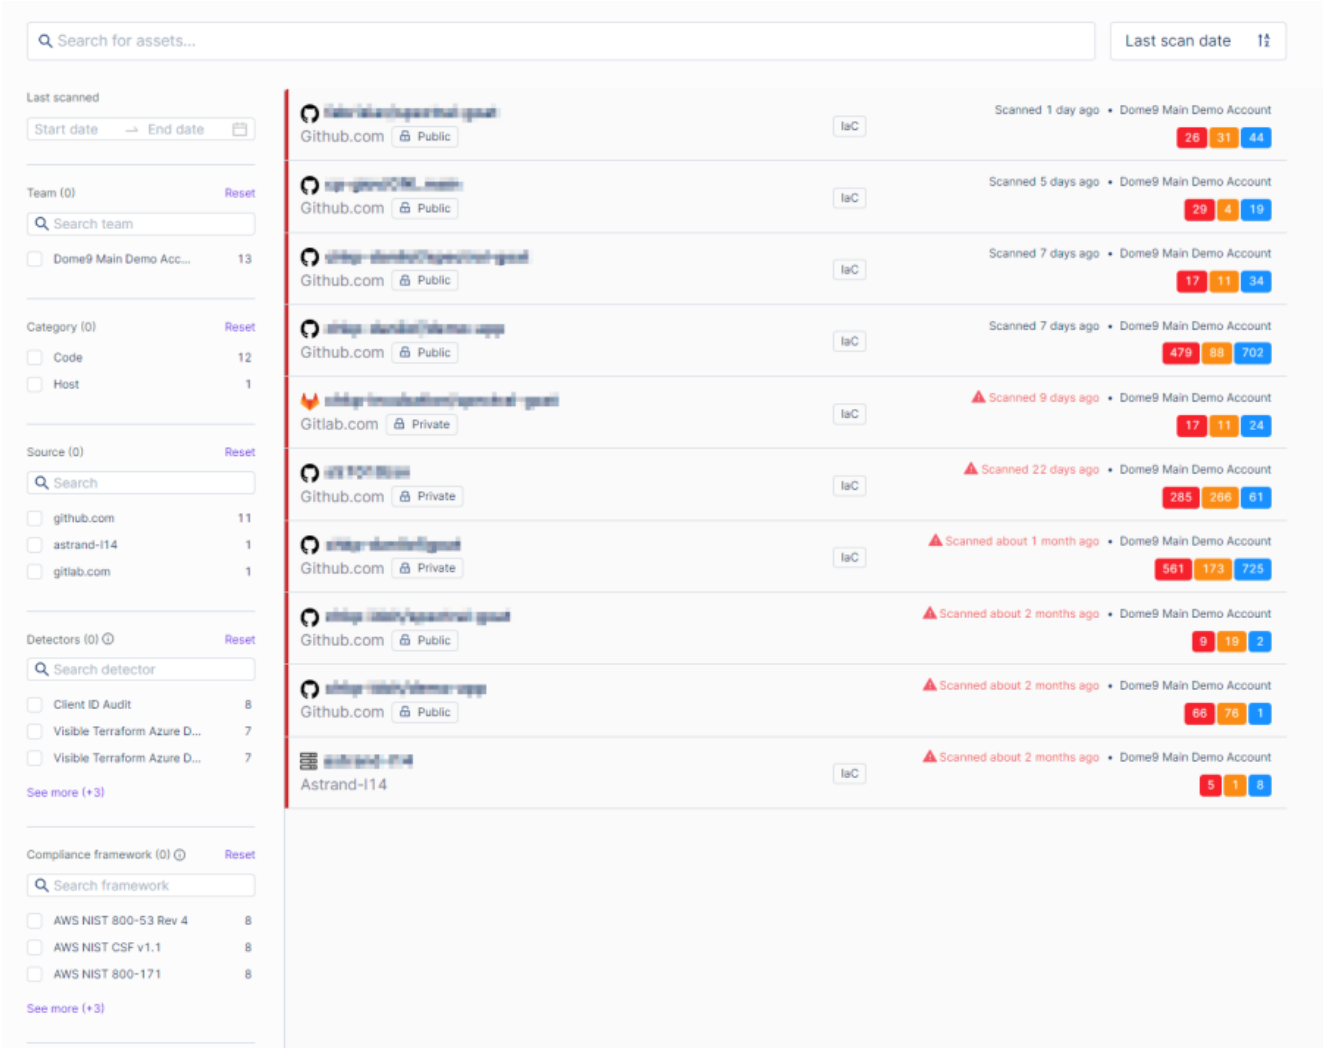

To see the asset page:

1. Click the name of an asset to expand it and see its findings.
2. Click **See all IaC issues** or **See exposed secrets** to open the asset page that provides details about a single asset.

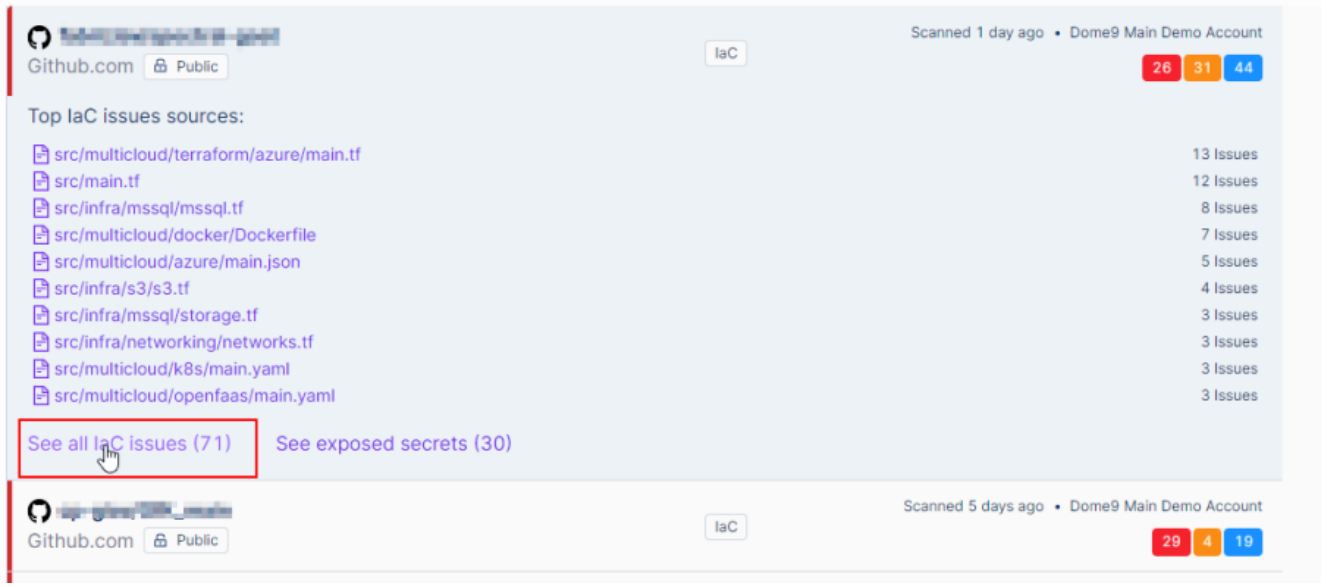

From the *"Dashboard" on page 653*, you can expand an asset to view issues. Click on an issue to see remediation instructions. Spectral classifies issues as code issues or infrastructure (**"infra"**) issues. For secrets issues, SpectralOps tests the validity of keys. A live key or token is labeled as valid.

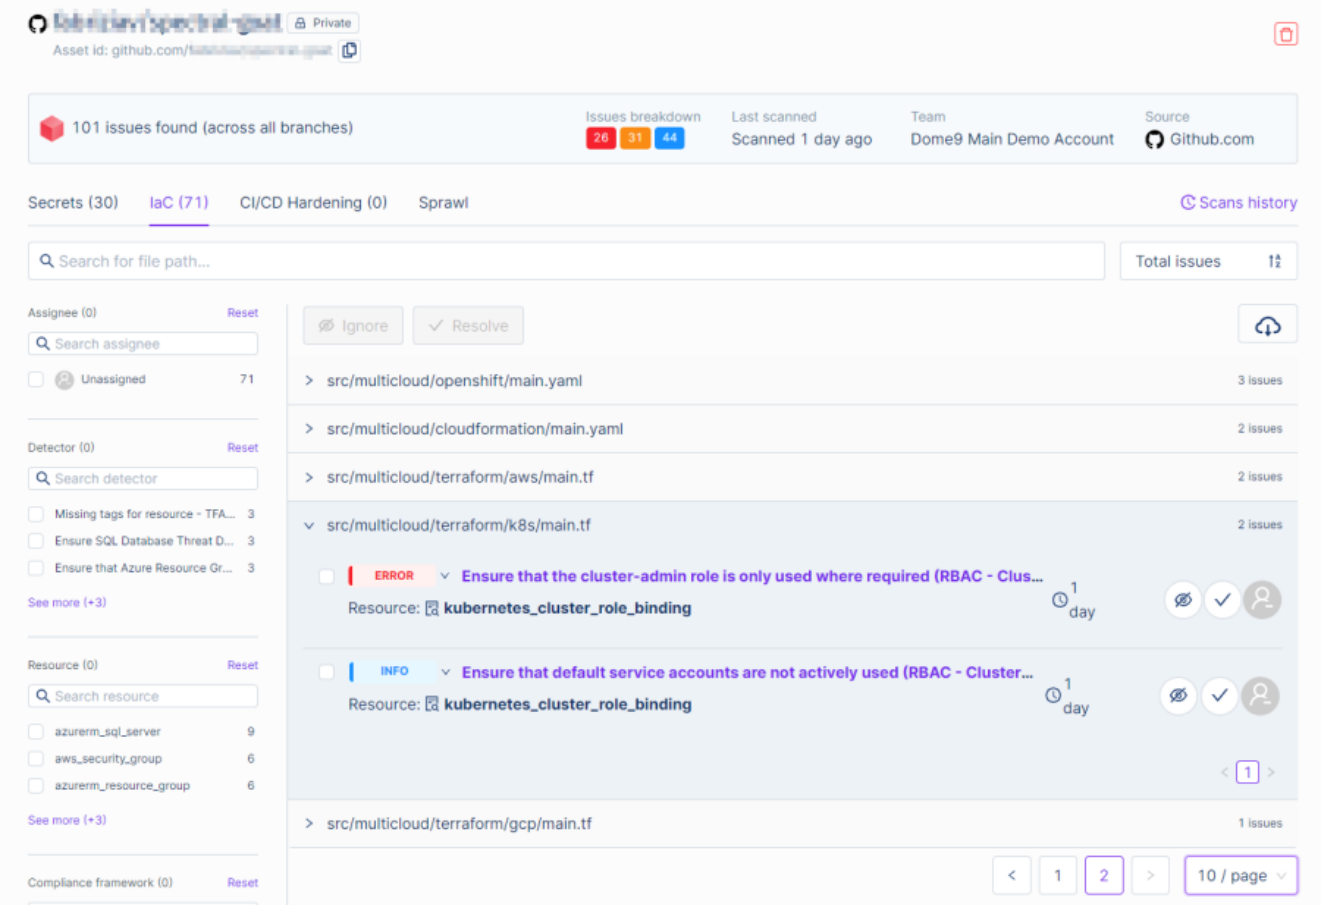

The **Assets** page includes these tabs:

- **Secrets** - Shows the exposed secrets that Code Security found.
- **IaC** - Shows the infrastructure-as-code issues that Code Security found in the asset grouped by file. Each IaC issue shows the IaC resource related to it.
- **CI/CD Hardening** - Shows issues in the repository settings, for example, when the *main* or *master* branch has no policy for a merge.
- **Sprawl** - Shows secrets that appear in multiple locations in an asset or across assets.

To view and change the severity of an issue

The severity of an issue appears to the left of the name of the issue. These are the severity levels:

Severity Level	Meaning	Suggested Action
Critical	An asset is compromised.	Resolve the issue immediately.
High	An asset is at high risk to become compromised.	Resolve the issue immediately.
Medium	A potential security risk exists for the asset.	Resolve the issue within a reasonable time period.
Low	There is no immediate risk, but the asset does not comply with best practices.	Take action when possible to implement best practices.
Informational	There is no immediate risk.	Use the information to improve your security posture.

To assign an issue to a user

1. To the right of the issue, click this button: .

   A list of users opens.
2. Search or scroll through the list to find the relevant user.
3. Click the name of the user.

To mark an issue as resolved

Click the check mark button to the right of the issue. If Spectral finds the issue again, it marks it as a regression. To unresolve an issue, click this button:  .

To export issues

1. From the top toolbar, open the relevant tab (for example: **Secrets**).
2. **Optional -** From the left toolbar, apply filters .
3. In the top right, click this button:.:
4. Select a format for the report:
   - **Export CSV**
   - **Export PDF**

   Your web browser downloads the report.

To create a ticket for an issue in Jira

1. Create the Jira integration. See *"Code Security Integrations" on page 739*
2. In the **Code Security Dashboard**, navigate to the issue.
3. Select the checkbox to the left of the issue.
4. From the toolbar above the table, expand **Actions** > click **Create Jira issue**.

   The **Create Jira Issue** window opens.
5. Fill the required fields.
6. Click **Create**.

After setting up your integration, select the issue and click the "Assign" button,

Select the Jira project in which you would like to open an issue, add a Summary of the problem and edit the description.

Once you create a Jira issue, a “Jira” label will be assigned to the spectral Issue so you can track the progress in Jira.

Scans

The Scans page presents all the different runs that were executed from other Code Security scanners in your organization.

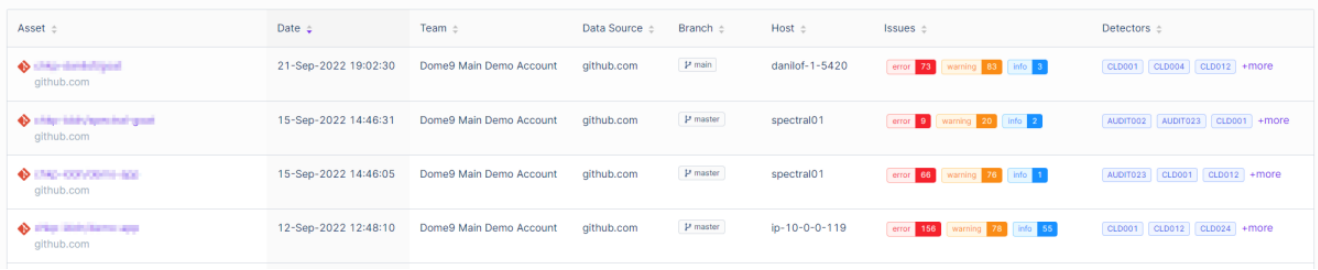

Asset	Date	Team	Data Source	Branch	Host	Issues	Detectors
github.com	21-Sep-2022 19:02:30	Dome9 Main Demo Account	github.com	main	danilof-1-5420	error 73 warning 83 info 3	CLD001 CLD004 CLD012 +more
github.com	15-Sep-2022 14:46:31	Dome9 Main Demo Account	github.com	master	spectral01	error 9 warning 20 info 2	AUDIT002 AUDIT023 CLD001 +more
github.com	15-Sep-2022 14:46:05	Dome9 Main Demo Account	github.com	master	spectral01	error 66 warning 76 info 1	AUDIT023 CLD001 CLD012 +more
github.com	12-Sep-2022 12:48:10	Dome9 Main Demo Account	github.com	master	ip-10-0-0-119	error 156 warning 78 info 55	CLD001 CLD012 CLD024 +more

You can filter the different scans based on the asset name, the asset source, its severity, and its date.

By default, Code Security presents the last 1000 scans.

Sources

From the **Sources** page, you can get instructions on how to download and configure your scanner.

There are three modes of secret detection by Code Security:

- **Developer mode** - a set of detectors that are recommended to run in CI/CD, with the highest precision (low rate of false positive secrets detected)
- **Security mode** - a set of detectors that provide higher precision at the expense of recall
- **Audit mode** - a set of detectors that have maximum recall (higher rate of false negative) but lower precision

  Best Practice - Check Point recommends to use the audit mode only in rare cases.

For more instructions on how to integrate Code Security, see *"Code Security CI/CD Integrations" on page 688*.

Securing Open Source Code

Spectral can scan for open source vulnerabilities in your repositories. Run:

```
spectral scan --engines oss
```

These programming languages and package managers are supported:

- C, C++ (conan)
- Dart (pubs)
- Dotnet (deps.json)
- Objective-C (cocoapods)
- Elixir (mix)
- Erlang (rebar3)
- Go (go.mod, Go binaries)
- Haskell (cabal, stack)
- Java (jar, ear, war, par, sar, nar, native-image)
- JavaScript (npm, yarn)
- Jenkins Plugins (jpi, hpi)
- Nix (outputs in /nix/store)
- PHP (composer)
- Python (wheel, egg, poetry, requirements.txt)
- Ruby (gem)
- Rust (cargo.lock)
- Swift (cocoapods)

Reports

The reports on this page relate to three categories:

- Code
- Host
- Productivity

The Code reports provide general statistics for your organization:

- Number of assets scanned
- Total number of scans till today
- Number of open issues in your assets (with a breakdown based on their severity)
- Organization trend
- Overall progress of your organization fixing issues in different assets

  It shows an aggregate result per day:

  - Green asset  has no findings.
  - Red asset  has at least one issue.
  - Gray asset  was not scanned lately, which makes it **inactive**.

The hotspot charts show:

- Top assets with the most issues.
- Top common issues.
- Top teams with most issues.

You can drill down from the chart to the **Asset** page and view the raw data.

Settings

Profile

The **Profile** screen includes settings that affect only your account. You can generate a personal authentication token (API key) for integrations (Admin only).

Organization

These settings allow you to edit the account-level DSN that your organization uses to authenticate with Code Security. You can set up:

- Team Key
- DSN (Data source name)

Scan Configuration

See *"Scan Configuration" on page 672*.

Source Code Management (SCM)

The Source Code Management settings page allows you to set up a custom domain of your on-premises instances. You can configure SCM for these domains by adding a DNS address for the domain:

- **GitHub Enterprise** instance
- **Gitlab Self-Hosted** instance
- **Bitbucket Server** instance

To configure a domain with a custom template, see *"Creating Custom Templates for Source Code Management" on page 670*.

Delete Asset

This page allows you to delete an asset.

Custom Rules

See *"Custom Rules" on page 668*.

Custom Rules

In addition to custom rules created locally for a specific repository, you can create custom rules at the organizational level. This allows security professionals to propagate custom rules to their organization (side by side with local scan configuration)

Step 1: Activate the Code Security Custom Rules Feature

1. In CloudGuard Code Security, navigate to **Settings** > **Custom Rules**.
2. In the confirmation window, click **Activate**.

Step 2: Create Custom Rules

1. Create a file structure. For this, run in the CLI:

   ```
   $HOME/.spectral/spectral custom-rules init
   ```

   The system creates a folder named *.spectral* in your current location. You can refer to the example rule file in *.spectral/rules* to get started.

2. We recommend that you sync the custom rules of your organization. To sync, run:

   ```
   $HOME/.spectral/spectral custom-rules get
   ```

3. Create a file per rule or a few rules in a single file. The name of the file must be in the format *custom_rule_*.yaml*. For more information on how to write indicators, see *"Building Detectors" on page 758*.

   Note - Make sure that the rules do not contain sensitive data.

4. **Optional -** For a centralized custom rules, add your own link to the playbook of your detector. To do this, add a `link` property to the rule. For example:

   ```
   link: https://guides.spectralopps.io/docs/cloud-keys
   ```

5. Make sure your rules do not expose sensitive data by targeting a secret directly or by targeting a secret using an obvious regex. You can run such commands locally in your environment, but we do not recommend adding them to rules.

6. Publish the rule. To publish, run:

   ```
   $HOME/.spectral/spectral custom-rules publish
   ```

The system scans the current folder using only your custom rules. Explore the findings of this scan to make sure your rules are working as expected. The system also scans the custom rules to identify sensitive data exposures. If the system detects issues in the custom rules, you are notified and you can cancel the publishing process.

Note - To disable the system scan of the custom rules, run:

```
$HOME/.spectral/spectral custom-rules publish --no-scan
```

7. Approve the custom rules. The system creates a new version of the custom rules.
8. The Administrator must approve or reject the new custom rules. If approved, the new custom rule is used by all the scans in the organization by default.

   Note - To exclude custom rules for a scan, run:

   ```
   $HOME/.spectral/spectral scan --exclude-tags custom-rules
   ```

Creating Custom Templates for Source Code Management

A Code Security account owner can create custom templates for Source Code Management (SCM) platforms that do not appear in the **Settings** > **SCM page**. To integrate a platform, create and add templates for these values:

- **Asset URI** - Code Security uses this to link to the line of code where the issue appears.
- **Blame URL** - Code Security uses this to show which contributors made specific changes to code

Step 1: Create custom templates for Blame URL and Asset URI

Use placeholders to create a template. When Code Security creates a URL to link directly to an issue, it replaces the placeholder with real data about the issue. Enter three curly braces before and after the name of the placeholder.

Supported Placeholders

Data	Placeholder for Data
URI of you repository	`{{{assetUri}}}`
Name of your repository (useful when the SCM has different host for the Git server)	`{{{assetName}}}`
Commit hash associated with the issue	`{{{commitSha}}}`
Branch associated with the issue	`{{{branch}}}`
Line of code where the issue begins	`{{{start}}}`
File path where the issue is located	`{{{path}}}`

Example Template for Asset URI

```
{{{assetUri}}}/blob/{{{branch}}}{{{path}}}#L{{{start}}}
```

Example Template for Blame URL

```
{{{assetUri}}}/blame/{{{branch}}}{{{path}}}#L{{{start}}}
```

Step 2: Use the custom templates to add an SCM platform

1. Go to **Settings** > **SCM**.
2. At the bottom of the page, use placeholders to create these templates:

- **Asset URI**
- **Blame URL**

3. Click **Save.**

   Warning - When you save the templates, Code Security updates all URLs that it already generated for the SCM platform. If you make templates incorrectly, you may cause unintended changes to data.

4. To test the connection between Code Security and the SCM platform, click the URL.

Known Limitations

- A custom SCM platform works only for issues that occurred after you created it. To use a custom SCM platform for older issues, remove the relevant assets from Code Security and then re-scan them.
- It is possible to save templates five times during a five-minute period.
- After you click **Save**, it is not possible to save a template while Code Security is creating a new custom SCM platform.

Scan Configuration

Use `spectral.yaml` to set the configuration for a specific asset if this is a file-based asset.

Code Security enables you to set the configuration on all assets of the same type.

To set the configuration for your account:

1. Navigate to **Code Security** > **Settings** > **Scan configuration**. Select the type of asset to which you apply the configuration.
2. Set the configuration in the input as YAML (in the same format and structure as in `spectral.yaml`).

Note - You can disable this configuration for an asset by adding the `--ignore-remote-config` flag while executing your scan. For example, `spectral scan --ignore-remote-config`.

Fallback

In the Fallback section, set the behavior of the scanner if misconfiguration or hardening settings affect the validity of the scan. Select **one** of these:

Setting	Behavior	Use Case
Strict	If the scan detects any misconfiguration, the scan stops and shows an error code. This is the default setting.	This mode is ideal for users who need a high level of confidence that their configurations are correct before proceeding with the scan.
Warn and proceed	The scan finishes and shows the fallback actions it took. If the scanner detects a misconfiguration, the scanner reverts the configuration to the default value and continues the scan. These are examples of misconfigurations ▪ Invalid `spectra.yaml` configuration file ▪ Invalid custom rules ▪ Invalid tags	This mode is beneficial for users who wish to continue scanning despite minor configuration issues and are comfortable with default values being applied for any misconfigured flags.

Hardening

Select actions to **Allow** or **Restrict** Spectral to do in assets for your organization:

Configuration	Source
OK	CLI, Local Configuration File
Exclude Tags	CLI, Local Configuration File
Exclude	CLI, Local Configuration File
Include	CLI, Local Configuration File
Ignores	Local Configuration File
Ignores Options	Local Configuration File, Inline Ignores
Fail on Error	CLI
Fail on Critical	CLI
Ignore Remote Config	CLI
Since	CLI
Max Size	CLI
Show Match	CLI
Send Local Ignores	Local Configuration File, Inline Ignores

Combining Configurations

- **match_ignores** - If a `spectral.yaml` file exists locally, its `match_ignores` section is merged with the asset type `match_ignores` section. This means that the list of ignores contain ignores configured locally in `spectral.yaml`, and also ignores defined per asset type.
- **projects** - If a `spectral.yaml` exists locally, the projects configuration of the asset type does not occur, and the local projects configuration is applied if it exists in `spectral.yaml`.

Asset Type Configuration Usage Indication

To know if an asset scan used asset type configuration, check if the scan banner named `remote_cfg` shows the value `Yes`.

Secrets Scanning

Use Secrets Scanning to avoid hardcoding and sharing secrets in your assets. Secrets Scanning user more than 2,500 built-in rules to scan for certificates, PEM files, API keys, passwords, and other sensitive information.

To do a Secrets Scan, run:

```
spectral scan
```

To see more results, run::

```
spectral scan --include-tags base,audit,audit3
```

Key Validation

When Spectral finds a token (for example: a valid GitHub token), it can test if the token is valid. Valid tokens in assets are a more serious security risk than invalid tokens.

Use the `--validate` flag to test the validity of tokens that Spectral finds in a scan. Run:

```
spectral scan --validate
```

Infrastructure as Code

Spectral Infrastucture as Code (IaC) scans your IaC files and notifies you about misconfigurations.IaC supports AWS, Google, Azure, Kubernetes and other platforms. Scan results show the misconfigured resource and give remediation suggestions.

Use the `--engines iac` flag to test the validity of tokens that Spectral finds in a scan: Run::

```
spectral scan --engines iac
```

Using Code Security CLI

Code Security is a CLI-driven toolchain.

To test Code Security

Run:

```
$HOME/.spectral/spectral run
```

To use Code Security to scan your environment

Run:

```
$HOME/.spectral/spectral scan
```

Commands

Command	Description
`run`	Run a scan interactively for exploring or auditing.
`scan`	Run a scan from your CI/CD pipeline.
`init`	Initialize configuration in your repo. This is how you customize ignores, detectors, and more. Once your run the `init` command, an hidden spectral folder is created and you can customize your configuration.
`fingerprint`	Encode a one-way fingerprint from a secret, for ignoring content.
`github`	Audit a GitHub organization, user, or repo.
`gitlab`	Audit a GitLab organization, group, user, or repo.
`history`	Run a Code Security git history scan. This flag enables you to scan the Git history and make sure there are no issues in the historical Git commits. **Note -** Issues found only show locally and are not included in the dashboard and UI (this is to prevent overhead).
`custom-rules`	Manage custom rules.

Environment Variables

Command	Description
`SPECTRAL_DSN=<DSN>`	Your private Code Security DSN, which connects to your account.
`SPECTRAL_SHOW_MATCH=1`	Show secrets in scan output (off by default).

Common Flags

Command	Description
`-t, --token`	Supply a token for GitHub, GitLab, or other for auditing.
`-h, --host`	Supply a git host (where relevant. For example, GitLab)
`-d, --dest`	Destination for git repos in case of auditing.
`-k, --kind`	Type of audit. For example, group, user, or org.
`--engines`	Engines to run in the current scan. Options: `secrets` (default), `iac`, oss.
`--include-tags base, audit`	Include additional ML-based detectors for full security coverage (more tags details: `spectral info --tags`).
`--include-tags iac`	Include IaC (Infrastructure as Code) security coverage (more tags details: `spectral info --tags`).
`---include CLD001, CLD002`	Scan only for specific detectors.
`--exclude CLD001, CLD002`	Exclude specific detectors from results. You can combine: `--include-tags base --exclude CLD002` to scan for a tag excluding specific detectors.
`--unstaged`	Scan with pre-commit and pre-receive hooks and send data to Code Security.
`-f, --fail-on-error`	Fail with non-zero exit code just on error severity matches.

Command	Description
`--fail-on-critical`	Fail with non-zero exit code only when Code Security detects High and Critical severity matches.
`scan`	Run a scan from your CI/CD pipeline.
`init`	Initialize configuration in your repo.
`fingerprint`	Encode a one-way fingerprint from a secret, for ignoring content.
`github`	Audit a GitHub organization, user, or repo.
`gitlab`	Audit a GitLab organization, group, user, or repo.
`gitlab -k all-groups`	Scan all GitLab groups.

Help

You can use the `--help` option for the main binary, or use a sub command, for example `$HOME/.spectral/spectral run --help` for the supported commands and flags.

```
spectral --help
Spectral Scan 1.8.37
Spectral Cyber Technologies Inc.

USAGE:
    spectral [FLAGS] [SUBCOMMAND]

FLAGS:
    -h, --help         Prints help information
        --nobanners    No help/free text banners. Make it easier to
parse output
    -V, --version      Prints version information

SUBCOMMANDS:
    config         Your local SPECTRAL_DSN config
    fingerprint    Fingerprint sensitive information for ignores
    github         Run a Spectral scan on a github organization, user,
or team. Alias: 'git'.
    gitlab         Run a Spectral scan on a Gitlab organization, user,
or team
    help           Prints this message or the help of the given
subcommand(s)
    history        Run a Spectral git history scan
    info           Spectral information
```

```
    init           Initialize Spectral configuration for a current
project. (Must be in the project root)
    local          Run a Spectral audit on local assets
    login          Log into your Spectral account
    logout
    logs           Run a Spectral logs scan
    register       Register for your own Spectral account
    run            Run a Spectral scan interactively
    s3             Run a Spectral AWS S3 scan
    scan           Run a Spectral scan in your CI pipeline
    version
```

Configuring Code Security

You can use Code Security without any configuration except in these cases:

- Special treatment of source root - In a given repo or folder, include specific folders, exclude others.
- Scan-time ignores - Ignore classes of files or detectors, or pieces of text at scan time (you can also perform ignores in your Code Security account).
- Detector inclusion or exclusion - For cases where you want to disable existing detectors, or enable experimental ones.
- Output formats - You can switch output formats to JSON, JUnit, and others, which can help outline your pipeline automation.
- Custom detectors - Code Security can load custom detectors that you build, and you can specify its location in the configuration.

You can use `spectral init` to generate a base configuration.

```
$ $HOME/.spectral/spectral init
            Initialized your spectral configuration in '.spectral/'.
```

This creates starter configuration layout in a special `.spectral` folder.

```
$ tree .spectral
                .spectral
                ├── rules
                │   ├── merchants.speql.yaml
                │   └── sample.yaml
                └── spectral.yaml

            1 directory, 3 files
```

Select this folder for source control.

Configuration File

The main Code Security configuration file is `spectral.yaml`. It configures these parameters:

- **Input sources** - Paths to scan.
- **Ignores** - Items to ignore and the stage to apply ignores.
- **Reporter outputs** -Reporter module to activate.

- **Detectors** - Detectors to include and/or exclude
- **Metadata** - Additional functionality to activate, such as masking, debug run and so on.

.spectral/spectral.yaml

```
# you can omit the reporter section entirely (you'll get a stylish
reporter by default)
                reporter:
                outputs:
                stylish: {}         # nice looking CLI reports
                # stylish: { html: "output.html" }  # produce HTML
reports
                # stylish: { csv: "output.csv" }  # produce CSV
reports

                # log:              # use a logger
                #   json: true      # enable JSON logging
                #   file: out.json  # put output in a file

                # junit: {}        # Great for integrating with CI
systems that understand JUnit XML (all of them, probably)
                # ignores: {}      # A reporter output that streams
results as ignores

                #
                # Ignoring Matches
                #
                # you can specify ignores for matches that you know
exist
                # and acknowledge them, but you don't care about them
for now.
                #
                # These are regex: rule_id, rule_name, match_text,
path.
                # To get a fingerprint, run `spectral fingerprint --
text YOUR-SECRET`
                #
                # match_ignores:
                #     ignores:
                #     - match_text: MYSQL_ROOT_PASSWORD
                #     - rule_id: <rule id>
                #       rule_name: <rule name>
                #       match_text: <rule id>
                #       path: <path>
                #       match_fingerprint:
79cdb7f2e0e4a96520304ff641f45f230be4f362a4a16c704730115a85fa545f

                projects:
```

```
                sample:
                project:
                name: sample
                input:
                - local: .
                name: sources
                # you can add a few more.
                # everything is relative to working directory (where
you run spectral from)
                #
                # - local: nteract/node_modules
                #  name: nteract
                rules:
                roots:
                - rules      # folder(s) relative to this file
                # cherry-pick rules for these roots
                # include:
                #    tags:
                #    - node
                #    ids: []
                # exclude:
                #    tags:
                #    - node
                #    ids: []

                # add as many more projects as you like:
                # all_pythons:
            # ...
```

Exclude Rules by Severity

You can use the include or exclude tags. For example, if you want to see just errors, you can use the tag error in the include, or use the exclude tags warning and info which gives the same semantic.

You can use the include or exclude tags also through CLI options. For example, `--include-tags error` or `--exclude-tags warning,info`.

Ignores

You can run a scan, and choose to ignore results for known issues or issues that you prefer address later.

There are two main ways to perform ignores:

- Glob ignores - Ignores the files completely.
- Match ignores - Ignores the actual matches by file name, content, rule, and more.

Glob Ignores

This is similar to `.gitignore` file ignoring an entire folder, a glob of a file structure or a specific file, regardless of any scan.

For example, if there is a large (in GBs) `tf-models` file and you have reasonable certainty there no security issues.

To ignore this file, add a special `.ignore` file to your repo, and set its content similar to a regular `.gitignore`.

```
tf-models/*
```

Note - If there is `.gitignore` in the root of the scan, all the glob's inside the `.gitignore` file is ignored.

Match Ignores

This ignores the specified text, such as test keys, demo keys, and more.

You can also ignore a specific rule and a specific set of files under the rule.

For example, to ignore all credit cards showing under a **test** folder, specify the PCI rule. Below the rule, specify a file glob, such as `tests/.*`.

To add ignores, edit the main spectral configuration file:

.spectral/spectral.yaml

```
match_ignores:
                ignores:
                - match_text: MYSQL_ROOT_PASSWORD
                - rule_id: <rule id, regex>
                rule_name: <rule name, regex>
                match_text: <rule id, regex>
                path: <path, regex>
            match_fingerprint:
b76fe610abe3bdaa92d4002dc0516dfa21c2dbf520373c6203469d0dee369888
```

Fingerprinting

When you want to ignore a secret, or a piece of confidential text without specify it, you can use a cryptographically secure digest fingerprinting. To fingerprint your piece of text, dd this fingerprint to your ignore rule:

```
$ spectral fingerprint --text AKIAXXXXXXXXXXXXXXX
            b76fe610abe3bdaa92d4002dc0516dfa21c2dbf520373c6203469d0dee
369888
```

Inline Code Ignores

To provide more flexibility, we support adding ignores in your source code (as code comments) using this pattern:

```
spectral:ignore-[file|next-line|line] [detector|fingerprint|text]
[COMMENT]
```

Modes

- **spectral:ignore-line** - Ignore match in the same line
- **spectral:ignore-next-line** - Ignore next line match
- **spectral:ignore-file** - Ignore in all file

Categories (supports multiple values, comma separated):

- **detector** - Ignore by specific detector id.
- **fingerprint** - Ignore by specific fingerprint.
- **text** - Ignore text prefix

Examples

- Ignore match in the same line by detector:

  ```
  AKIAXXXXXXXXXXXXXXXX // spectral:ignore-line detector:CLD001
  ```

- Ignore match by a fingerprint:

  To get the fingerprint, run this Bash command:

  ```
  $ $HOME/.spectral/spectral fingerprint --text
  AKIAXXXXXXXXXXXXXXXX
  ```

  To ignore the match, run:

  ```
  AKIAXXXXXXXXXXXXXXXX // spectral:ignore-line
  fingerprint:b76fe610abe3bdaa92d4002dc0516dfa21c2dbf520373c6203469
  d0dee369888
  ```

- Ignore match by text prefix:

```
AKIAXXXXXXXXXXXXXXX // spectral:ignore-line text:AKI
```

- Ignore multiple categories:

```
AKIAXXXXXXXXXXXXXXX // spectral:ignore-line detector:CLD001
text:AKI
```

- Describe the ignore action in a comment field:

```
AKIAXXXXXXXXXXXXXXX // spectral:ignore-line detector:CLD001
comment:'Testing token'
```

- Ignore multiple detectors in a file:

```
// spectral:ignore-file detector:CLD001,CLD002
```

Ignore Rules Categorization

Categorization is designed to enhance policy compliance within organizations by providing visibility into all ignores, whether defined in the Spectral configuration YAML file or inline in the code. This feature allows categorization of ignores to ensure they are tracked appropriately, aligning with the organization's policy compliance efforts.

YAML Configuration

Example YAML configuration of match ignores:

```
match_ignores:
                ignores:
                - rule_id: CLD001
                categorization:
                kind: wont-fix | false-positive
            comment: "Test secret"
```

Inline Configuration

Add this initial categorization for inline ignores:

```
// spectral:<ignore-kind> detector:<detector> categorization:<false-
positive | wont-fix> comment:<comment>

                example
            // spectral:ignore-next-line detector:CLD001
categorization:false-positive comment:'Test secret'
```

When Spectral detects an ignore with categorization, the issue appears in the SpectralOps dashboard. users cannot change the ignore category or comment Users cannot resolve the issue from the dashboard. The scanner manages these issues.

Hardening - Hide Local Ignores

The hardening feature allows organizations to enforce stricter controls on local ignores. For example, if an organization wants to track all issues ignored by users, they can use the hardening feature to restrict local ignores, ensuring that all ignores are categorized and visible in the dashboard. All restricted local ignores must have a categorization. Ignores without a categorization cause the scan to fail, or they drop from the scan according to the fallback mode.

Projects

You can scan each sub-folders in a folder separately. For example, monorepo folder has sub-folders, such as client, backend, and infra (Infrastructure as code).

A sample configuration for this project:

```
projects:
                infra:
                project:
                name: infra             # Name of the project
                input:
                - local: ./infra        # Path of sources to include
in project
                name: sources         # Name of the path
                rules:
                roots:                  # Rules paths on client
machine - you can add paths to custom rules you created (relative to
`spectral.yaml` location - `.spectral/spectral.yaml`)
                - rules                 # This is the default if not
set (`.spectral/rules`)
                include:
                ids: []             # Ids of rules you wish to include
in this project
                tags:               # Tags you wish to include in this
project
                - iac
```

```
                exclude:
                ids: []             # Ids of rules you wish to exclude
from this project
                tags:               # Tags you wish to exclude
from this project
                - audit
                client:
                project:
                name: client
                input:
                - local: ./client
                name: sources
                rules:
                include:
                ids: []
                tags:
                - audit
                - base
                exclude:
                ids: []
                tags:
                - iac
                server:
                project:
                name: server
                input:
                - local: ./server
                name: sources
                rules:
                include:
                ids: []
                tags:
                - audit
                - base
                exclude:
                ids: []
                tags:
            - iac
```

This triggers three scans, one scan per project defined in the configuration.

Properties

- `project` - (Mandatory)
  - `name` - (Mandatory)
- `input (Array)` - Path of sources to include in the project. Mandatory and each element should consist of:

- `local` - (Mandatory) Path of resources to include in the project.
- `name` - (Mandatory) Name for given path.

- `rules` - Configuration regarding rules to scan in this project.
  - `roots (Array)` - (Optional) Rules paths on client machine. You can add paths to custom rules you created (relative to `spectral.yaml` location - `.spectral/spectral.yaml`).
  - `include` - Configuration about excluding rules in this projects. Mandatory if `exclude` was not provided.
    - `tags (Array)` - (Mandatory) Collection of tags to include in this project (`iac, audit`).
    - `ids (Array)` - (Mandatory) Collection of rule ids to include in this project (`DB001, TFAWS110`).
  - `exclude` - Configuration about excluding rules in this projects. Mandatory if `include` was not provided.
    - `tags (Array)` - (Mandatory) Collection of tags to exclude in this project (`iac, audit`).
    - `ids (Array)` - (Mandatory) Collection of rule ids to exclude in this project (DB001, TFAWS110).

Configuration per Asset Type

With `spectral.yaml`, you can set configuration for a specific asset, and only for file-based assets. For more information, see *"Scan Configuration" on page 672*.

Code Security CI/CD Integrations

To integrate Code Security with your CI/CD, download the Code Security engine and run a scan in your test or build scripts.

Notes:

- Code Security team cryptographically signs its binaries. However, you can download the Code Security binary and store it in your own artifact store periodically, then, you can pull from your own store in your CI workflows.
- Use `$HOME/.spectral/spectral scan --ok` if you do not want to break builds. It is common to use this mode when you ramp up your discovery of issues or when you have a different kind of security workflow.

Code Security for GitHub Actions

Get your DSN from the Code Security platform and set up your `SPECTRAL_DSN` in the CI environment variables as a secret. For more information, click here.

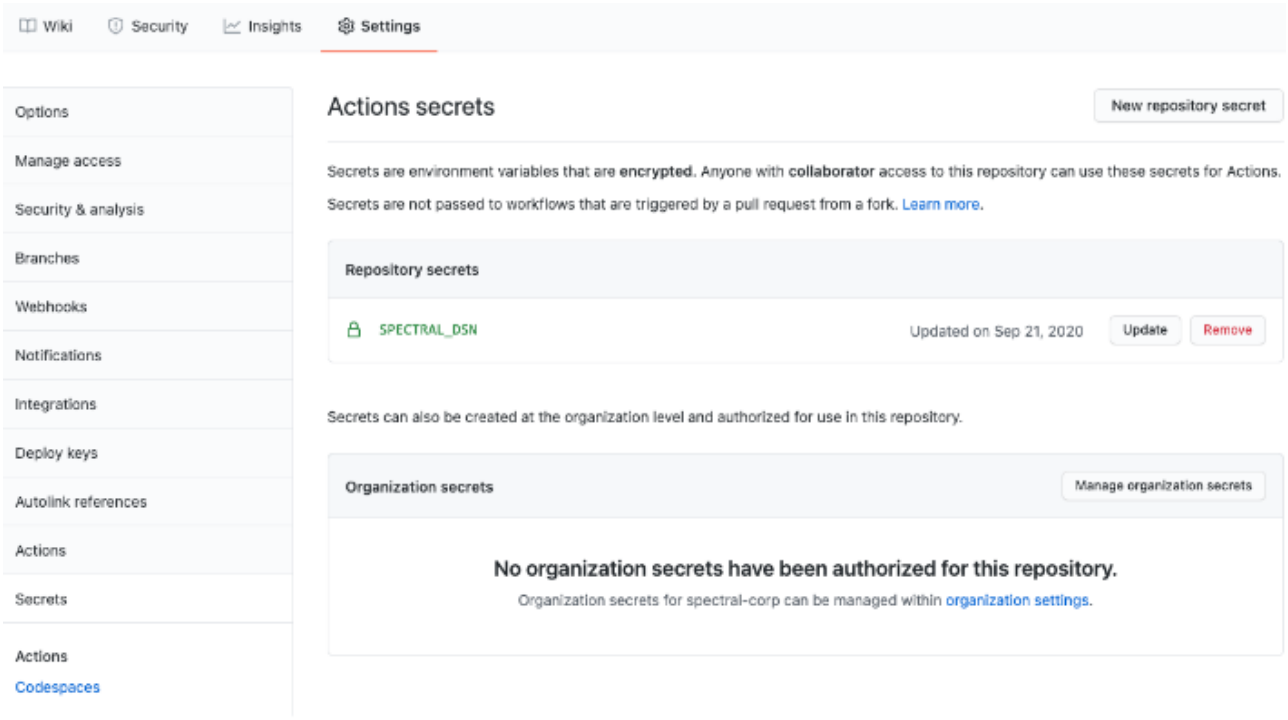

Example of the configuration:

.github/workflows/scan.yml

```
name: Main
on:
  push:
```

```
    branches: [ main ]

env:
  SPECTRAL_DSN: ${{ secrets.SPECTRAL_DSN }}

jobs:
  scan:
    runs-on: ubuntu-latest
    steps:
    - uses: actions/checkout@v2
    - name: Install and run Spectral CI
        uses: spectralops/spectral-github-action@v2
        with:
          spectral-dsn: ${{ env.SPECTRAL_DSN }}
          spectral-args: scan --ok
```

Example of a build integrated with Code Security:

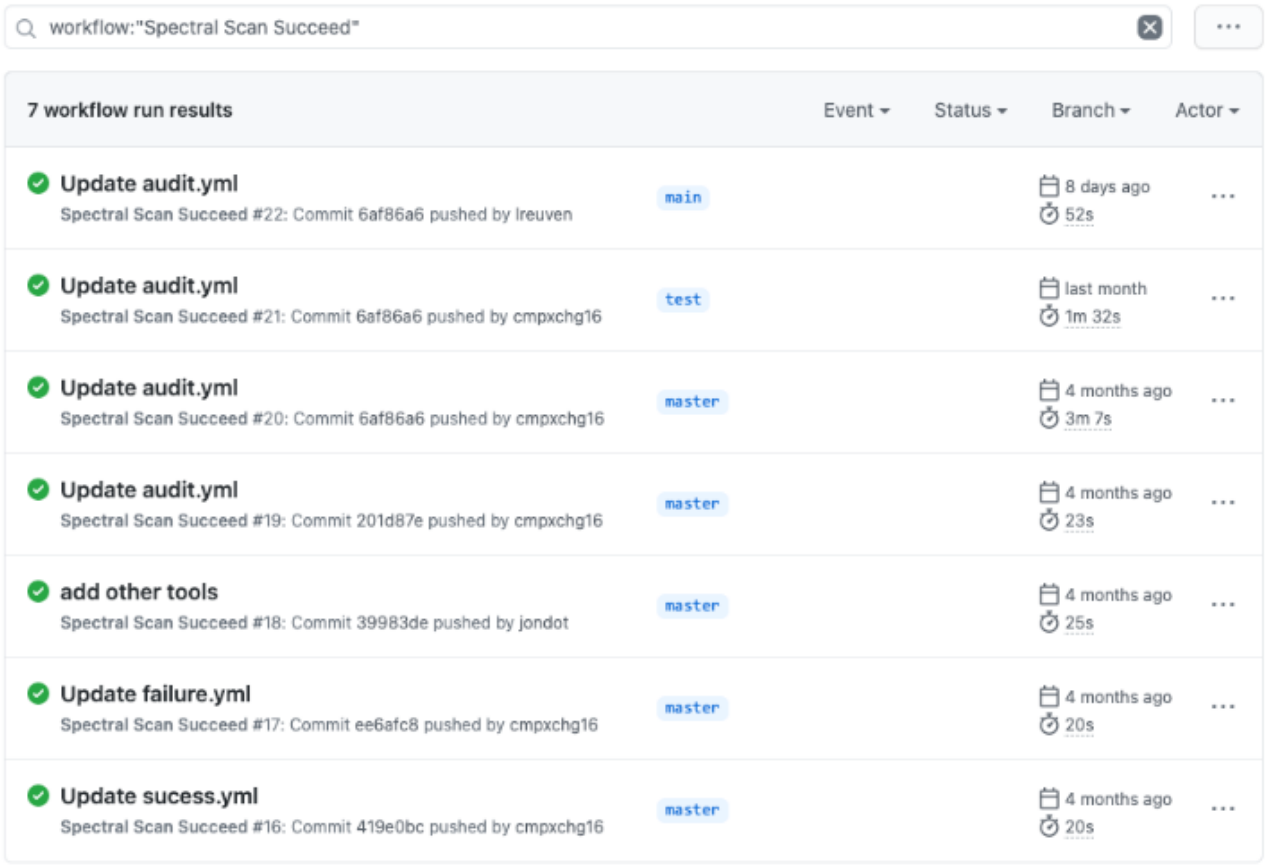

Code Security for GitLab CI/CD

Get your DSN from the Code Security platform and set up your DSN in the GitLab variables store as `SPECTRAL_DSN`.

Note - Make sure to select the Protected variable option.

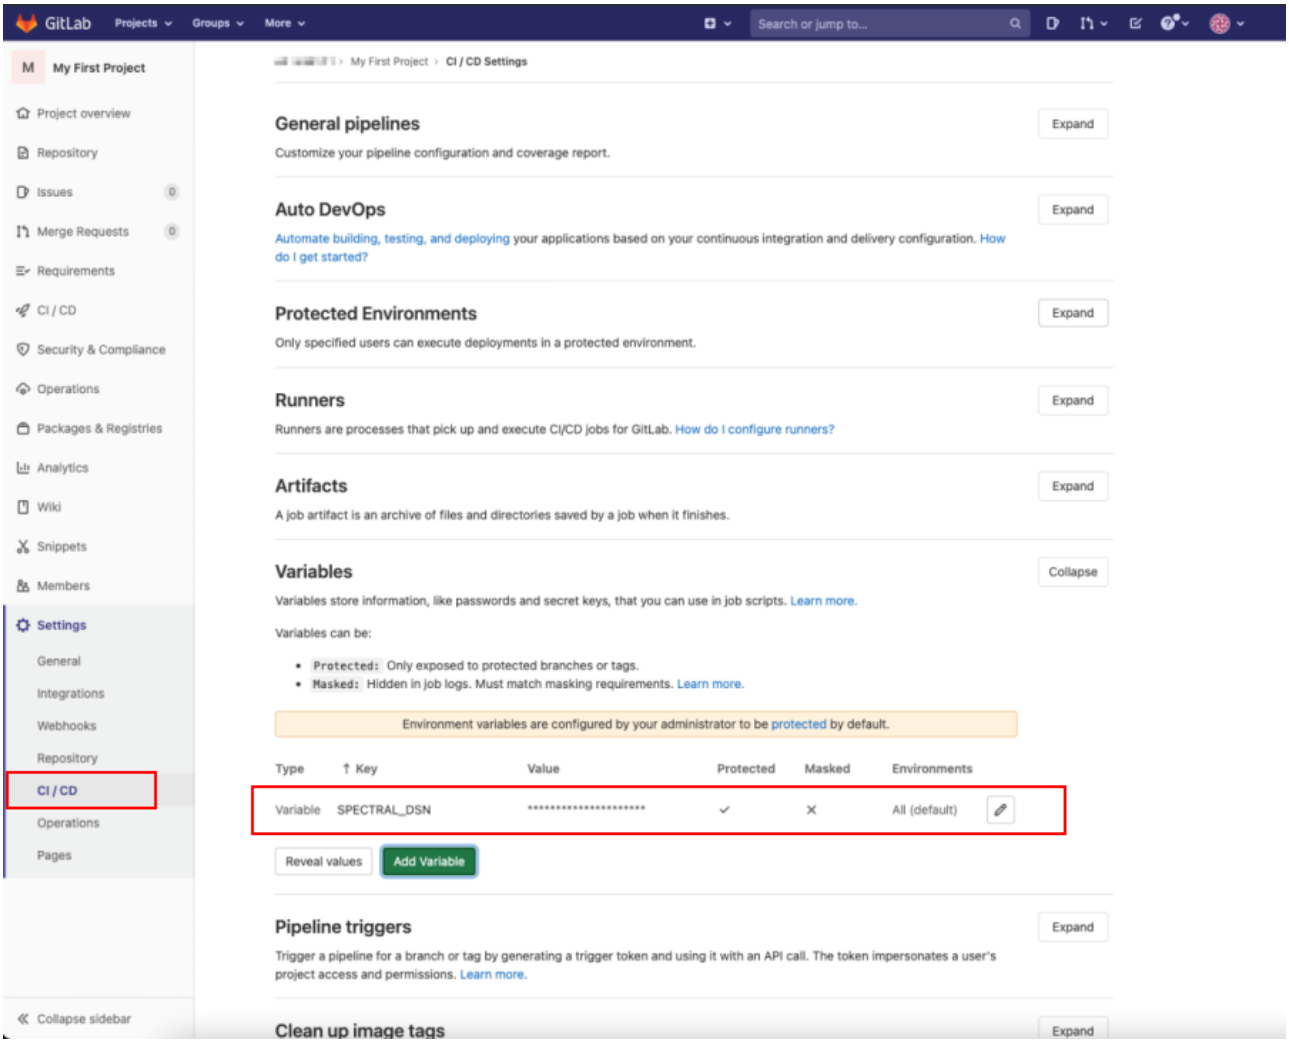

Example of configuration in the US region:

For Dome9 users

.gitlab-ci.yml

```
build-job:
  stage: build
  script:
    - curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
    - $HOME/.spectral/spectral scan --ok
```

For Infinity users

.gitlab-ci.yml

```
build-job:
  stage: build
  script:
    - curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
    - $HOME/.spectral/spectral scan --ok
```

Example of a build integrated with Code Security:

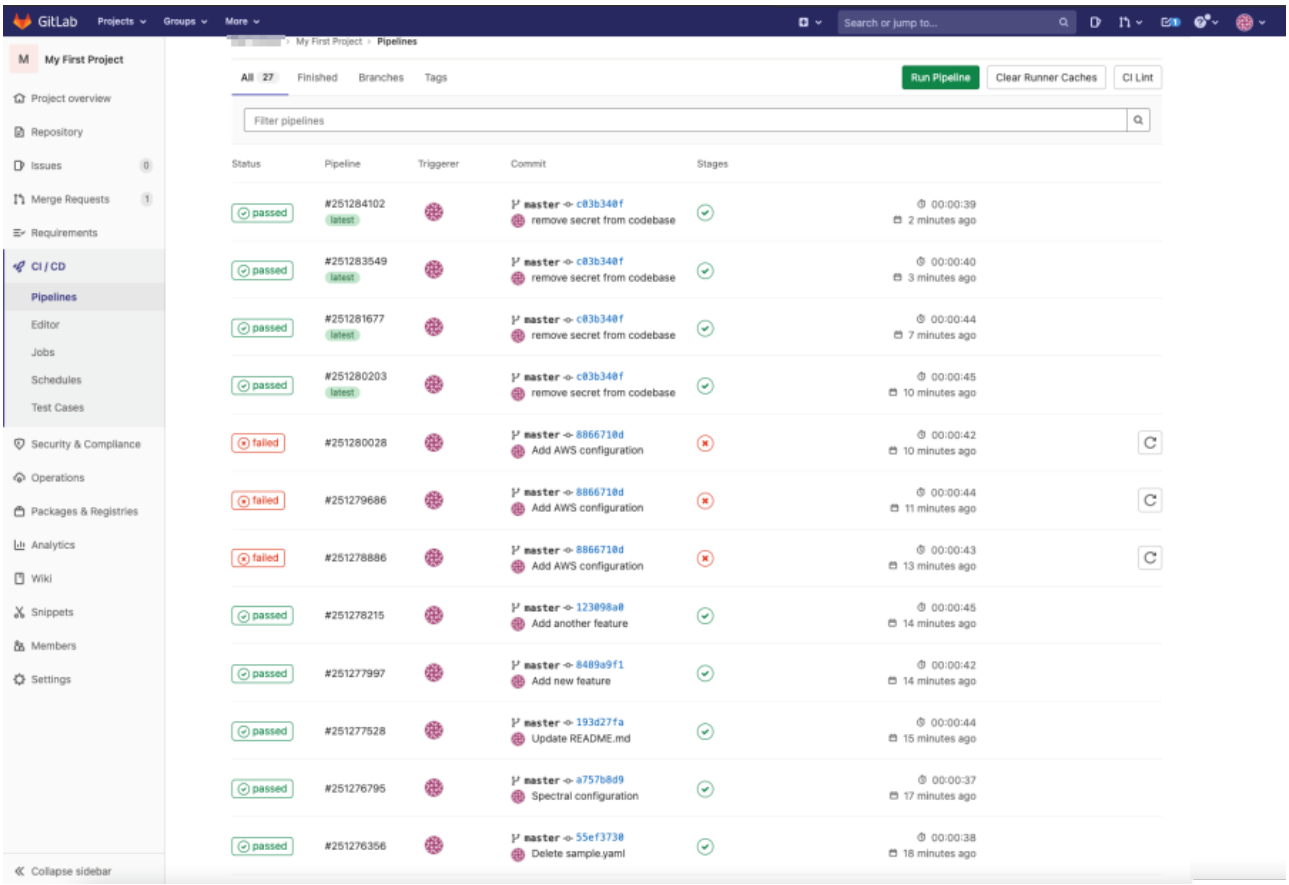

Code Security for Bitbucket Pipeline

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in variables and secrets.

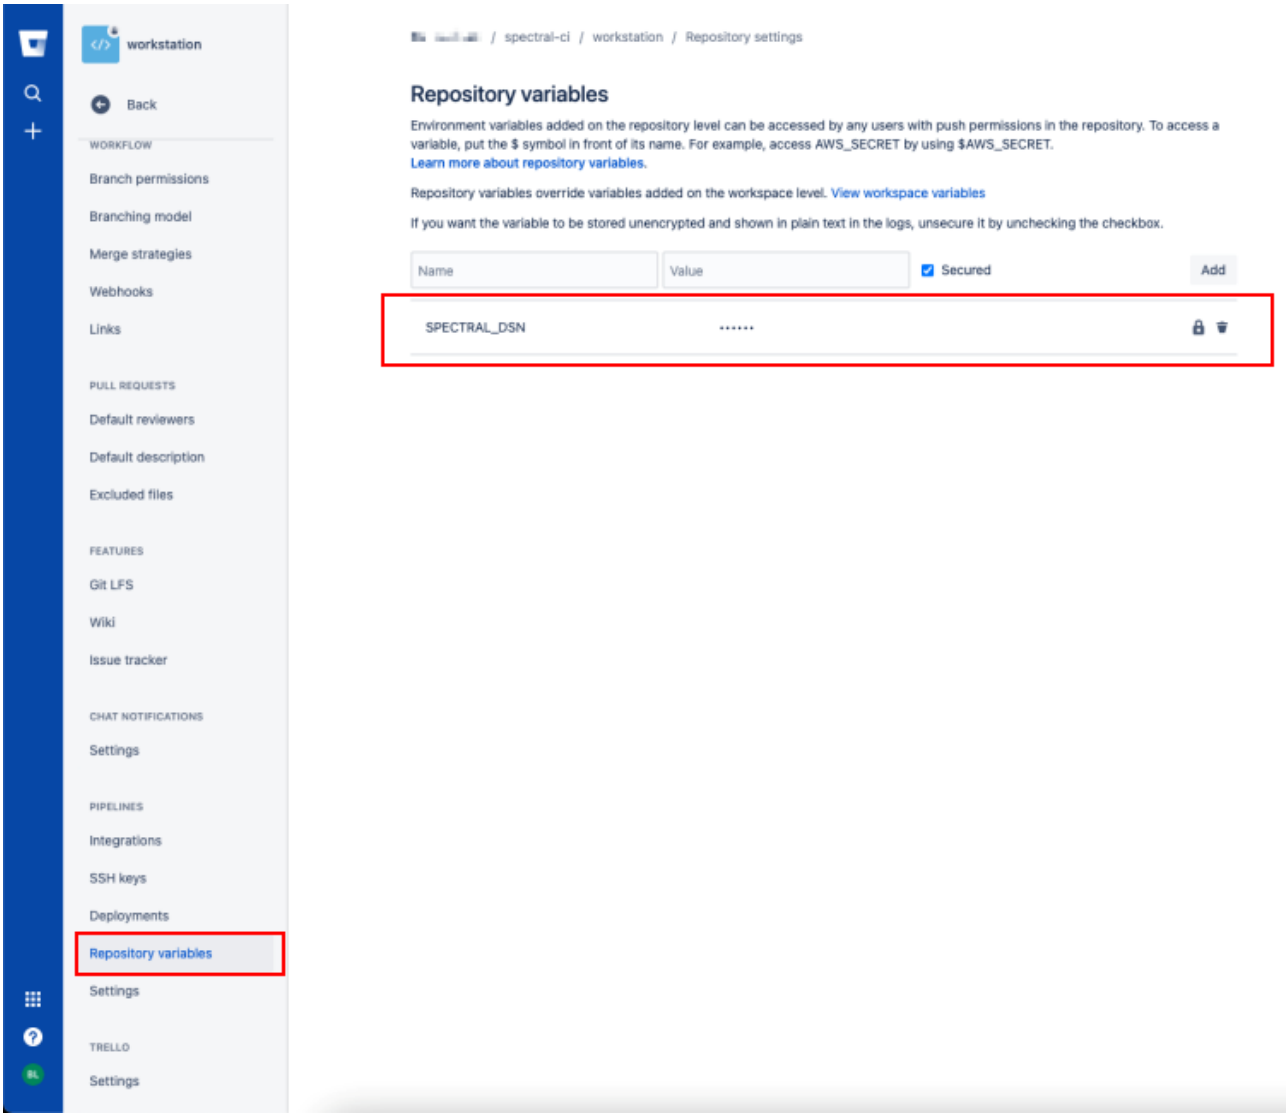

Example of configuration in the US region:

For Dome9 users

bitbucket-pipelines.yml

```
image: atlassian/default-image:2

pipelines:
  default:
    - parallel:
      - step:
          name: 'Install & run Spectral'
          script:
            - curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
            - $HOME/.spectral/spectral scan --ok
```

For Infinity users

bitbucket-pipelines.yml

```
image: atlassian/default-image:2

pipelines:
  default:
    - parallel:
      - step:
          name: 'Install & run Spectral'
          script:
            - curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
          - $HOME/.spectral/spectral scan --ok
```

Example of a build integrated with Code Security:

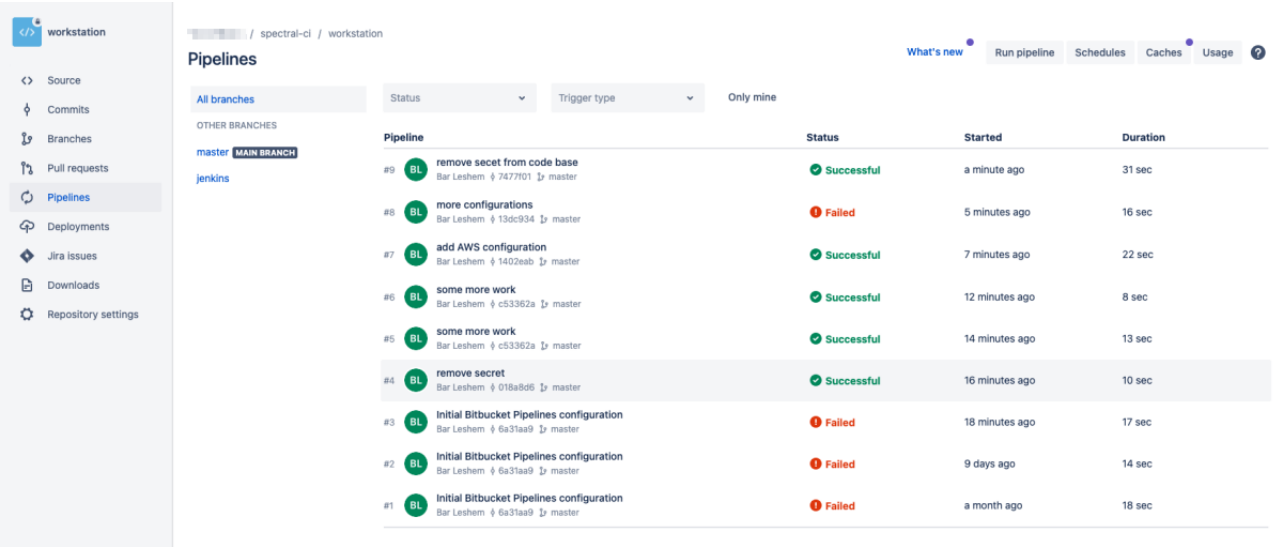

Code Security for Jenkins CI

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in Jenkins Credentials store.

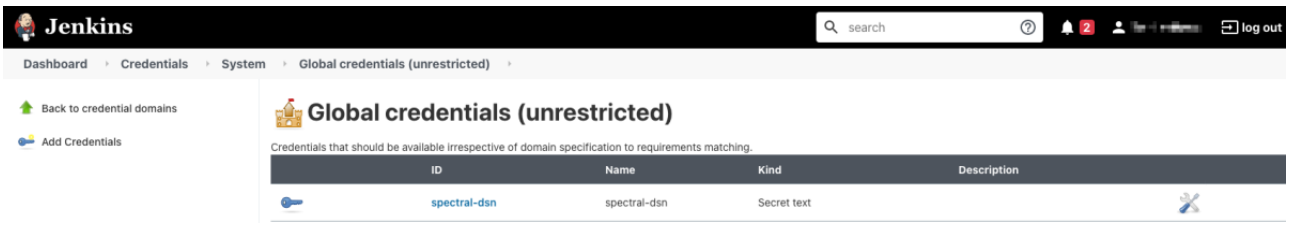

Example of configuration in the US region:

For Dome9 users

Jenkinsfile

```
pipeline {
  agent any
  environment {
    SPECTRAL_DSN = credentials('spectral-dsn')
  }
  stages {
    stage('install Spectral') {
      steps {
        sh "curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh"
      }
    }
    stage('scan for issues') {
      steps {
        sh "$HOME/.spectral/spectral scan --ok"
      }
    }
  }
}
```

For Infinity users

Jenkinsfile

```
pipeline {
  agent any
  environment {
    SPECTRAL_DSN = credentials('spectral-dsn')
  }
  stages {
    stage('install Spectral') {
      steps {
        sh "curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh"
      }
    }
    stage('scan for issues') {
      steps {
        sh "$HOME/.spectral/spectral scan --ok"
      }
    }
  }
}
```

Example of a build integrated with Code Security:

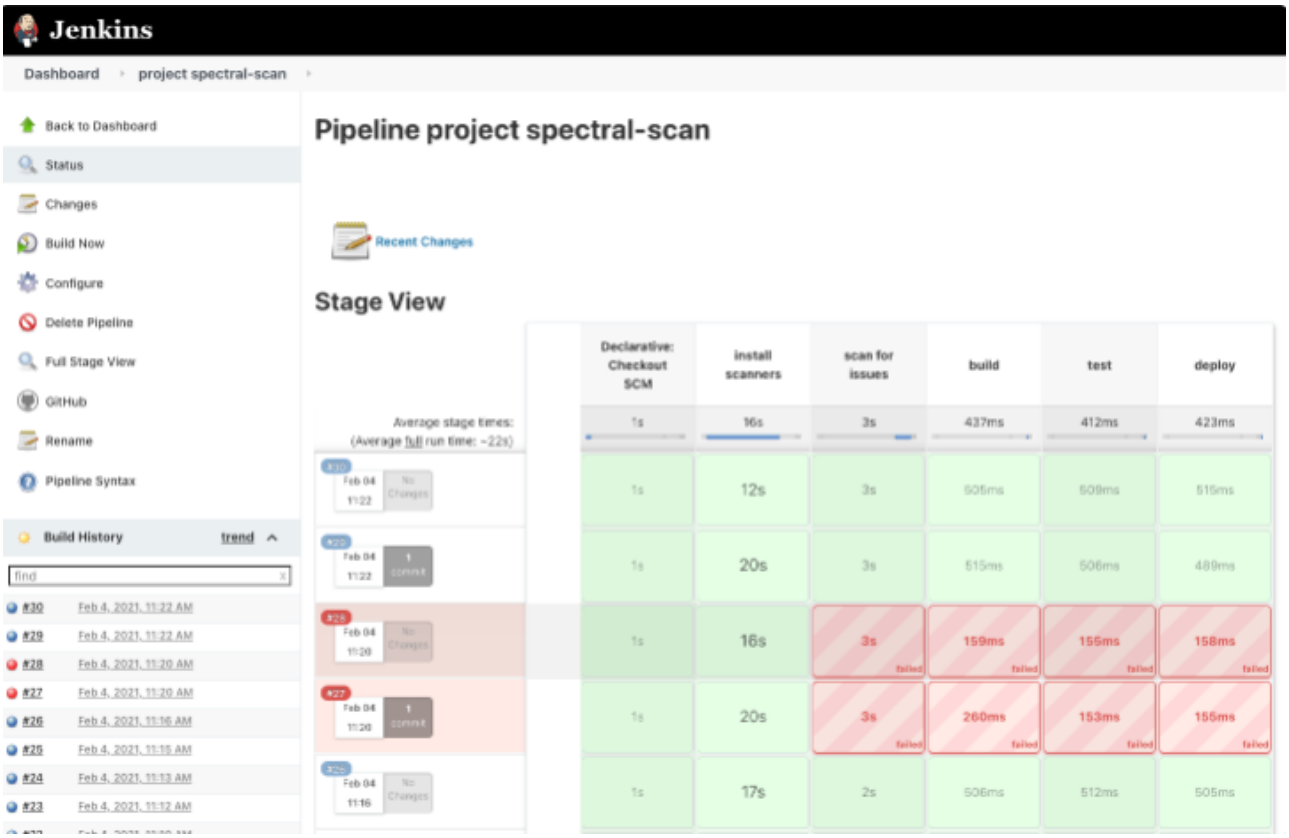

Code Security for CircleCI

Get your DSN from the Code Security platform and set up your DSN in CircleCI Secrets Store as `SPECTRAL_DSN`.

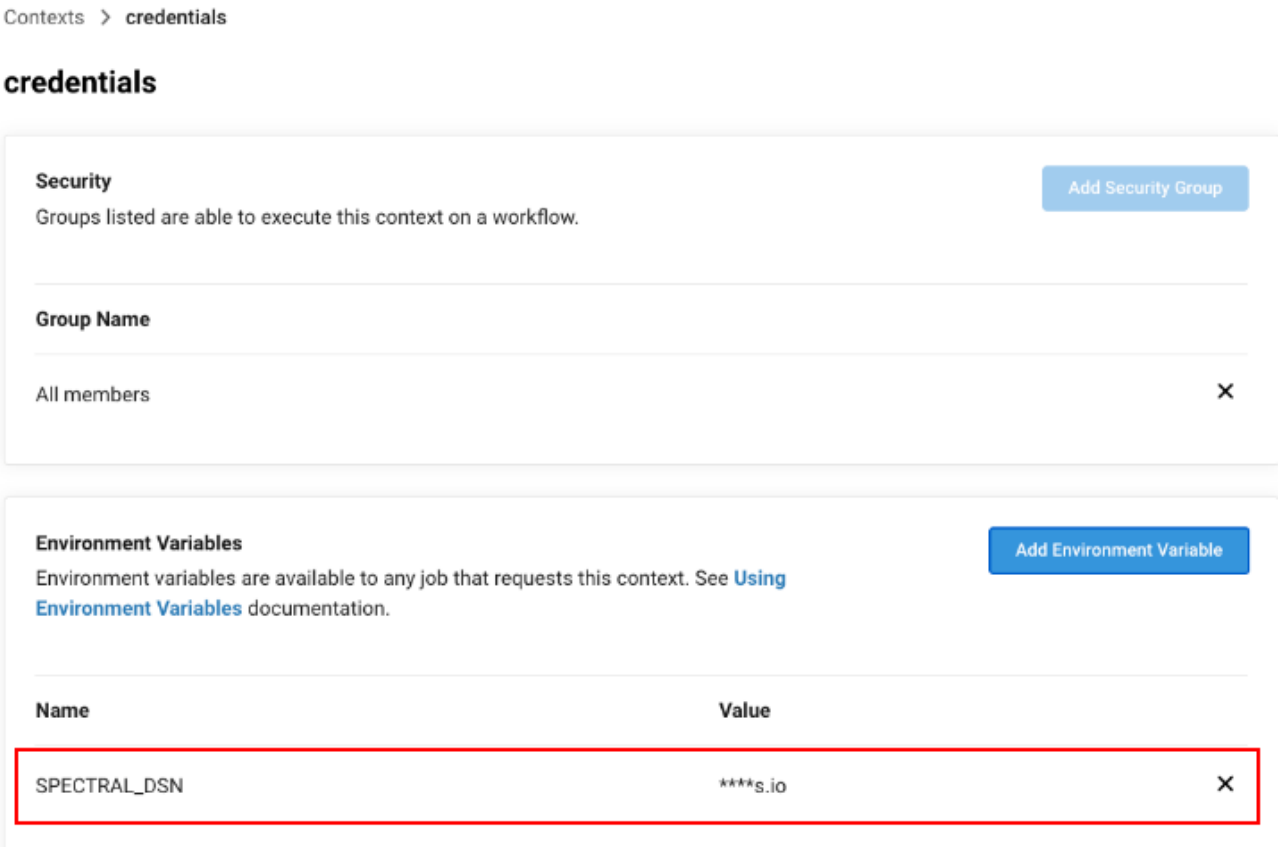

Example of configuration in the US region:

For Dome9 users

.circleci/config.yml

```
version: 2.1
workflows:
  test-env-vars:
    jobs:
      - build:
          context:
            - SPECTRAL_DSN
jobs:
  build:
    docker:
      - image: circleci/node:latest
    steps:
      - checkout
      - run: curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
      - run: $HOME/.spectral/spectral scan --ok
```

For Infinity users

.circleci/config.yml

```
version: 2.1
workflows:
  test-env-vars:
    jobs:
      - build:
          context:
            - SPECTRAL_DSN
jobs:
  build:
    docker:
      - image: circleci/node:latest
    steps:
      - checkout
      - run: curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
      - run: $HOME/.spectral/spectral scan --ok
```

Example of a build integrated with Code Security:

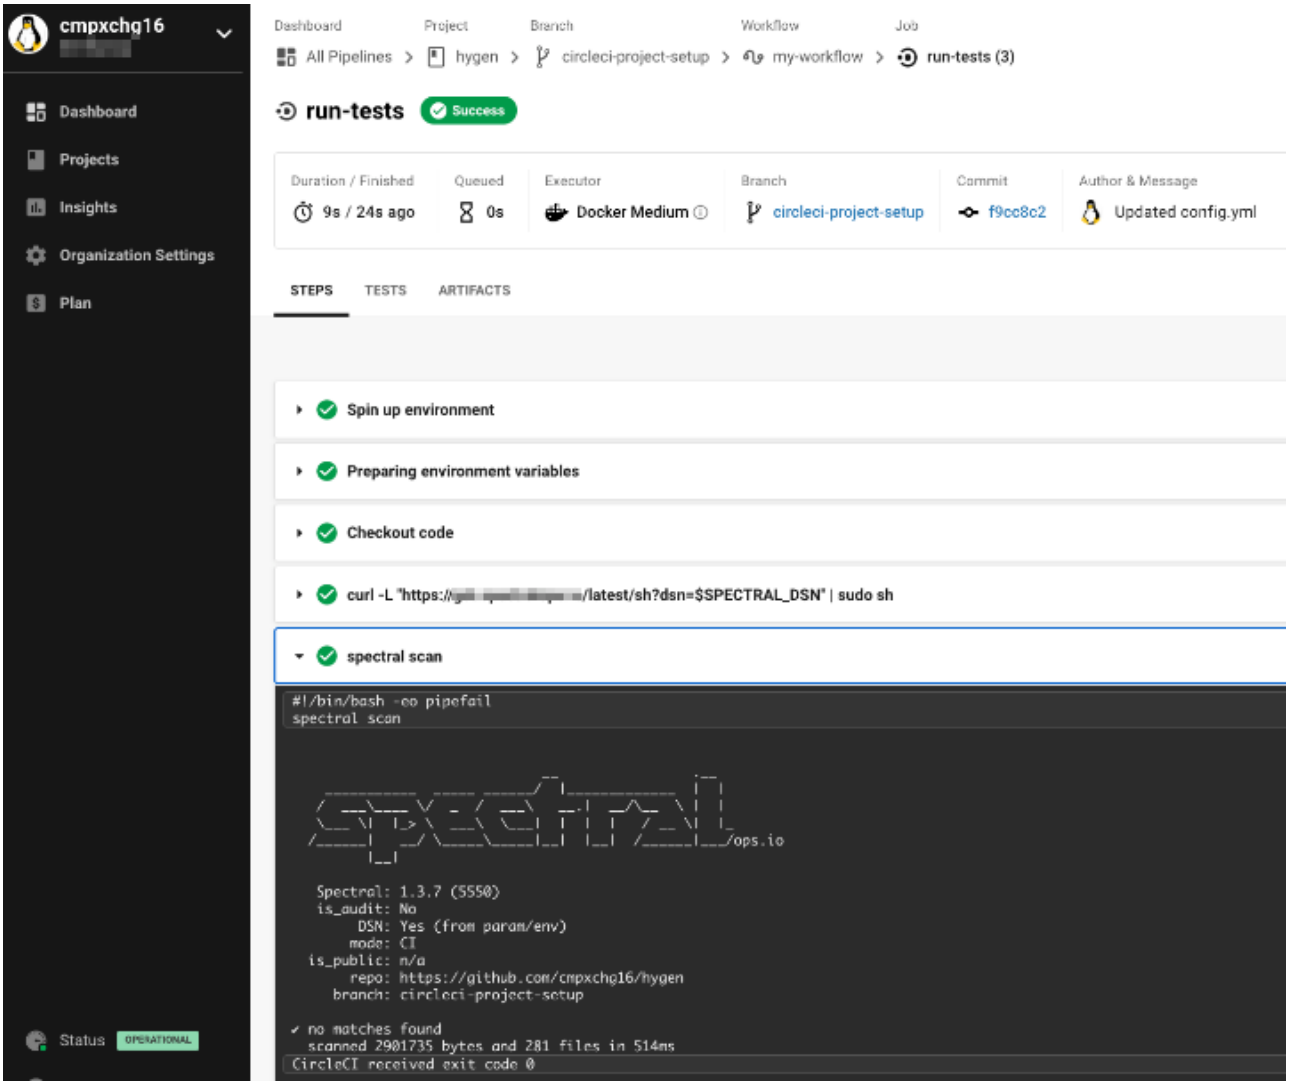

Code Security for Travis CI

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in the Travis Env Store.

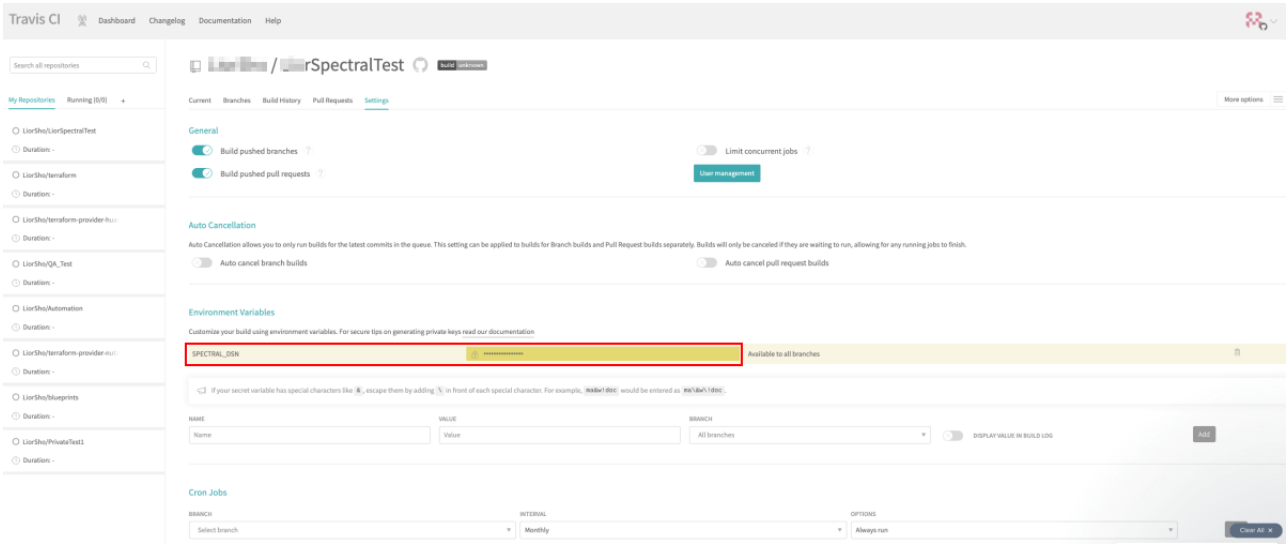

Example of configuration in the US region:

For Dome9 users

.travis.yml

```
language: ruby
install:
                          - curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
    - $HOME/.spectral/spectral --version
script:
    - $HOME/.spectral/spectral scan --ok
```

For Infinity users

.travis.yml

```
language: ruby
install:
                          - curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
    - $HOME/.spectral/spectral --version
script:
    - $HOME/.spectral/spectral scan --ok
```

Example of a build integrated with Code Security:

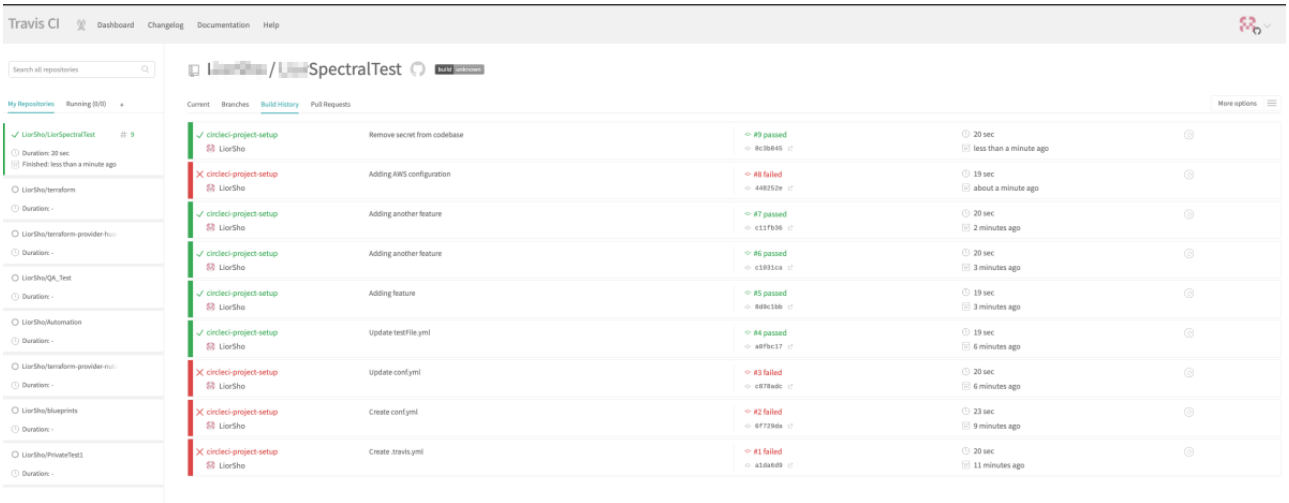

Code Security for AWS CodeBuild

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in Secrets Manager.

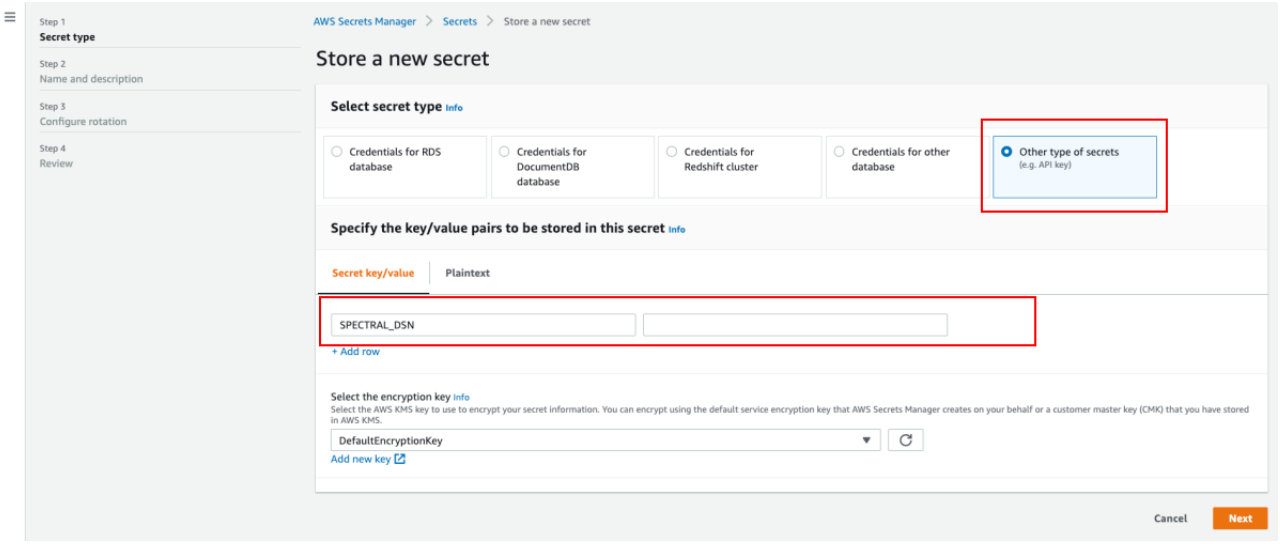

Example of configuration in the US region:

For Dome9 users

buildspec.yml

```
version: 0.2
env:
  secrets-manager:
    SPECTRAL_DSN: your-secrets-arn:SPECTRAL_DSN
phases:
  build:
    commands:
      - ...
  post_build:
    commands:
      - curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
      - $HOME/.spectral/spectral scan --ok
```

For Infinity users

buildspec.yml

```
version: 0.2
```

```
env:
  secrets-manager:
    SPECTRAL_DSN: your-secrets-arn:SPECTRAL_DSN
phases:
  build:
    commands:
      - ...
  post_build:
    commands:
      - curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$SPECTRAL_DSN" | sh
      - $HOME/.spectral/spectral scan --ok
```

Example of a build integrated with Code Security:

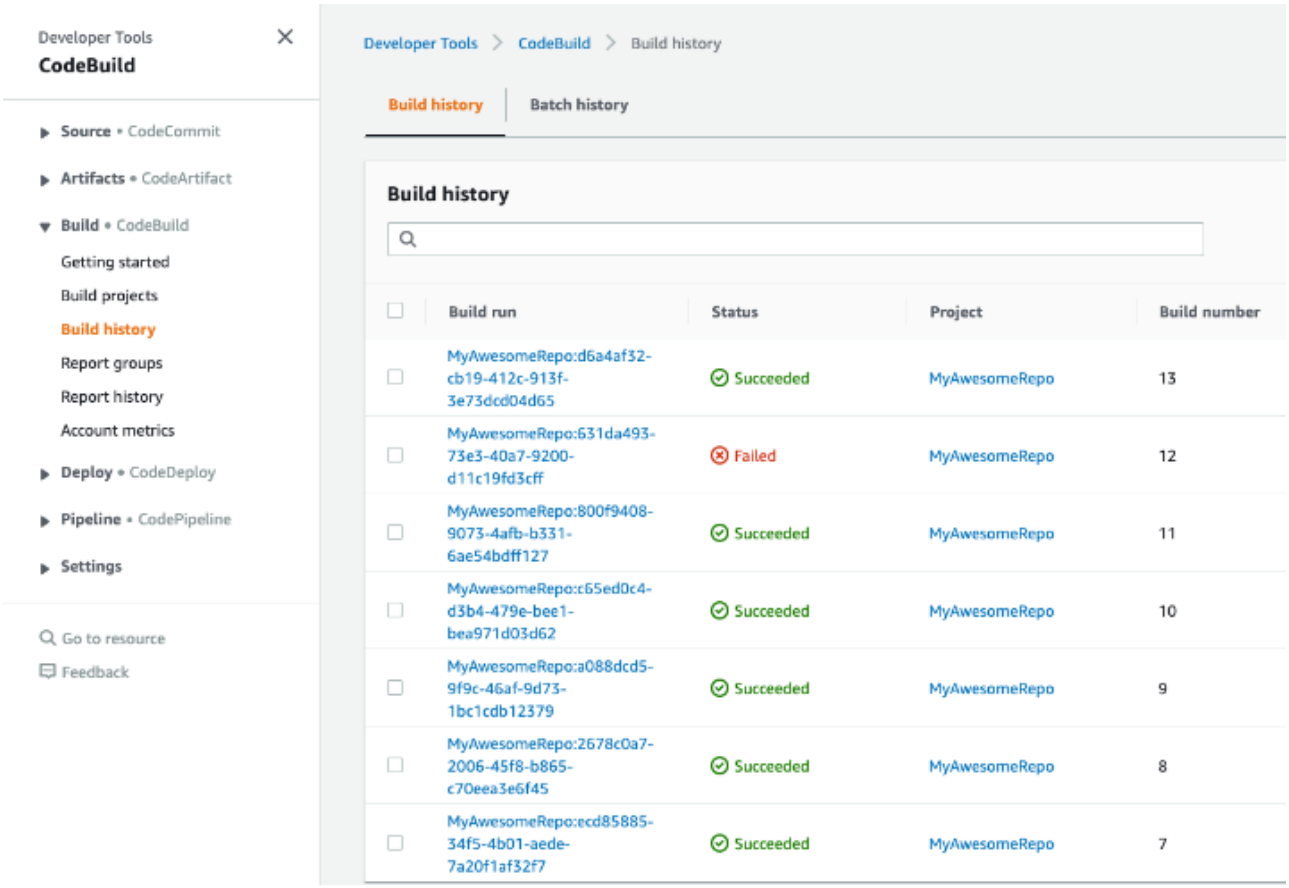

Code Security for Azure DevOps Pipeline

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in your secret variables.

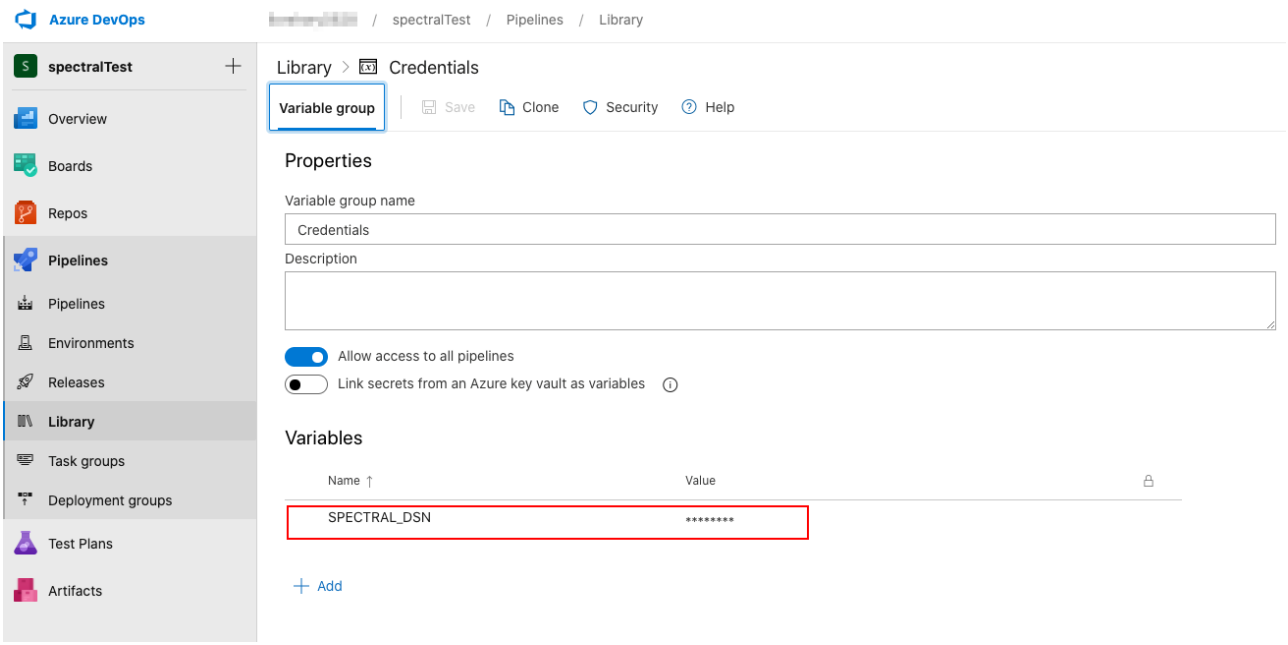

Example of configuration in the US region:

For CloudGuard Standalone (Dome9) users

build.yml

```
trigger:
- main
pool:
  vmImage: 'ubuntu-latest'
steps:
- task: CmdLine@2
  displayName: Checkout $(Build.SourceBranchName)
  inputs:
    - script: 'git checkout $(Build.SourceBranchName)'
- script: curl -L 'https://spectral-
us.dome9.com/latest/x/sh?dsn=$(SPECTRAL_DSN)' | sh
  displayName: 'Install Spectral'
- script: $HOME/.spectral/spectral scan --ok --dsn $(SPECTRAL_DSN)
  displayName: 'Spectral Scan'
```

For CloudGuard Infinity users

build.yml

```
trigger:
- main
```

```
pool:
  vmImage: 'ubuntu-latest'
steps:
- task: CmdLine@2
  displayName: Checkout $(Build.SourceBranchName)
  inputs:
    - script: 'git checkout $(Build.SourceBranchName)'
- script: curl -L 'https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$(SPECTRAL_DSN)' | sh
  displayName: 'Install Spectral'
- script: $HOME/.spectral/spectral scan --ok --dsn $(SPECTRAL_DSN)
  displayName: 'Spectral Scan'
```

Example of a build integrated with Code Security:

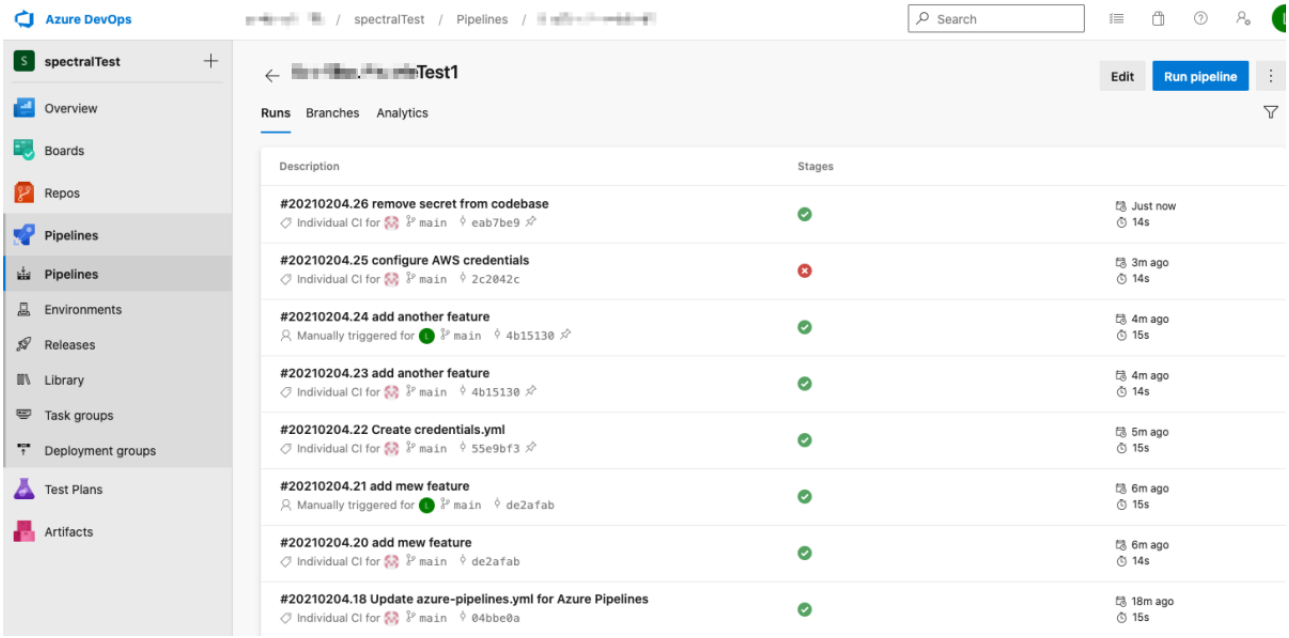

Code Security for Google Cloud Build

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in Secret Manager.

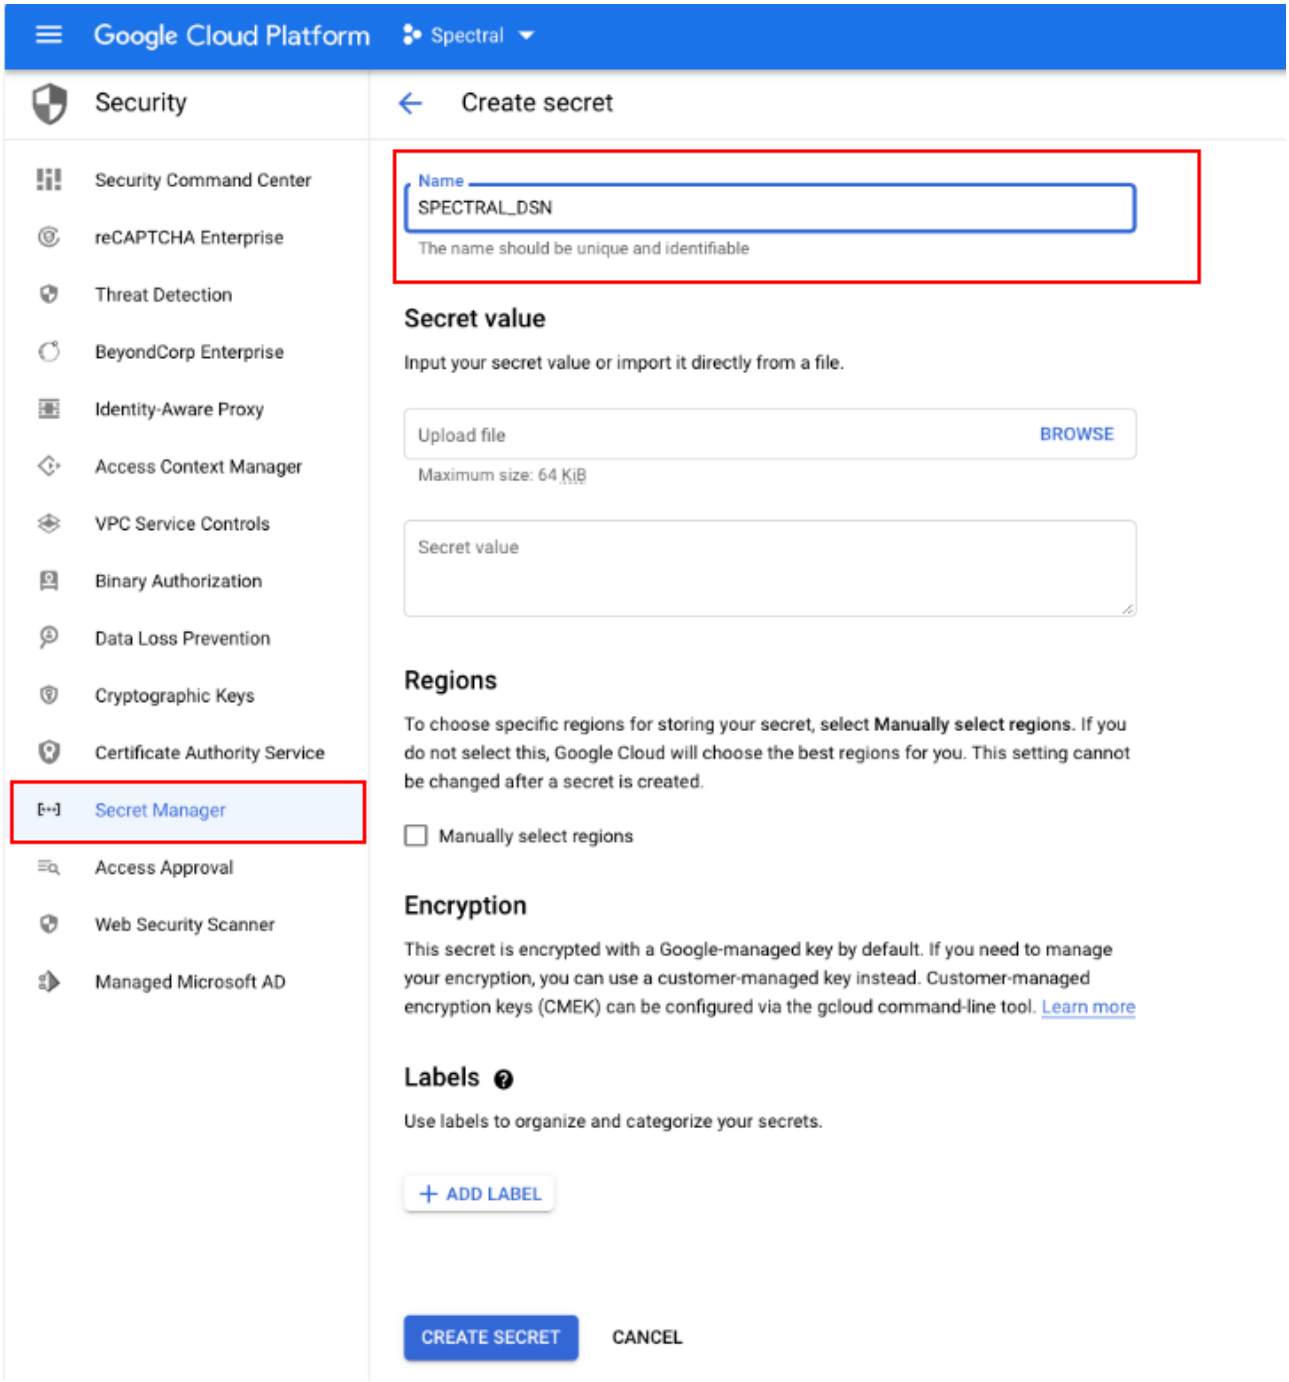

Example of configuration in the US region:

For Dome9 users

cloudbuild.yaml

```
steps:
  - name: gcr.io/cloud-builders/gcloud
```

```
  id: Spectral
  entrypoint: bash
  args:
    - -c
    - |
      curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$$SPECTRAL_DSN" | sh
      $$HOME/.spectral/spectral scan --ok

  secretEnv:
    - SPECTRAL_DSN

availableSecrets:
  secretManager:
    - versionName: projects/PROJECT_ID/secrets/SPECTRAL_
DSN/versions/latest
      env: SPECTRAL_DSN
```

For Infinity users

cloudbuild.yaml

```
steps:
  - name: gcr.io/cloud-builders/gcloud

  id: Spectral
  entrypoint: bash
  args:
    - -c
    - |
      curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$$SPECTRAL_DSN" | sh
      $$HOME/.spectral/spectral scan --ok

  secretEnv:
    - SPECTRAL_DSN

availableSecrets:
  secretManager:
    - versionName: projects/PROJECT_ID/secrets/SPECTRAL_
DSN/versions/latest
      env: SPECTRAL_DSN
```

Example of a build integrated with Code Security:

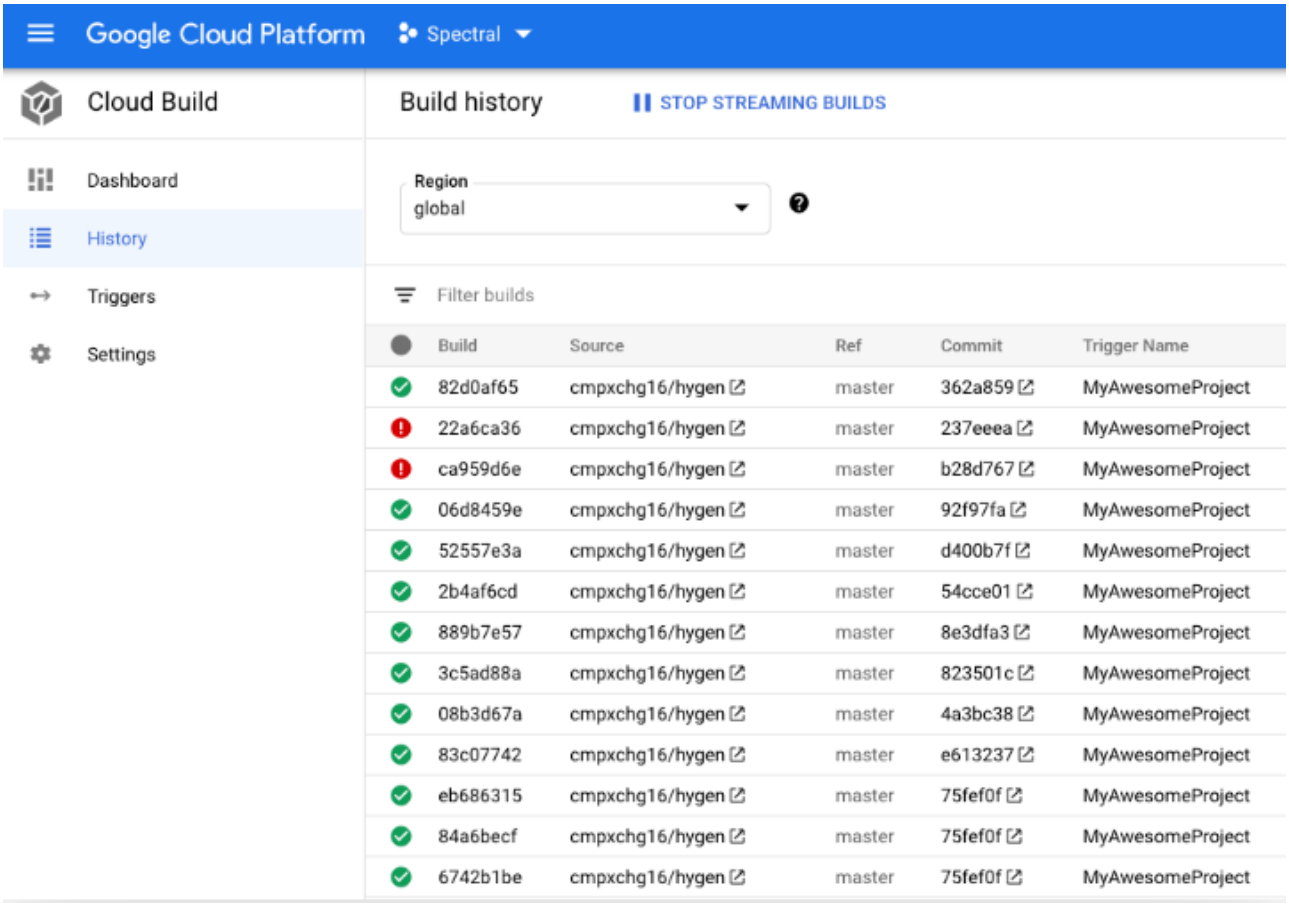

Code Security for JFrog Pipelines

Get your DSN from the Code Security platform and store it as `SPECTRAL_DSN` in the pipelines secrets.

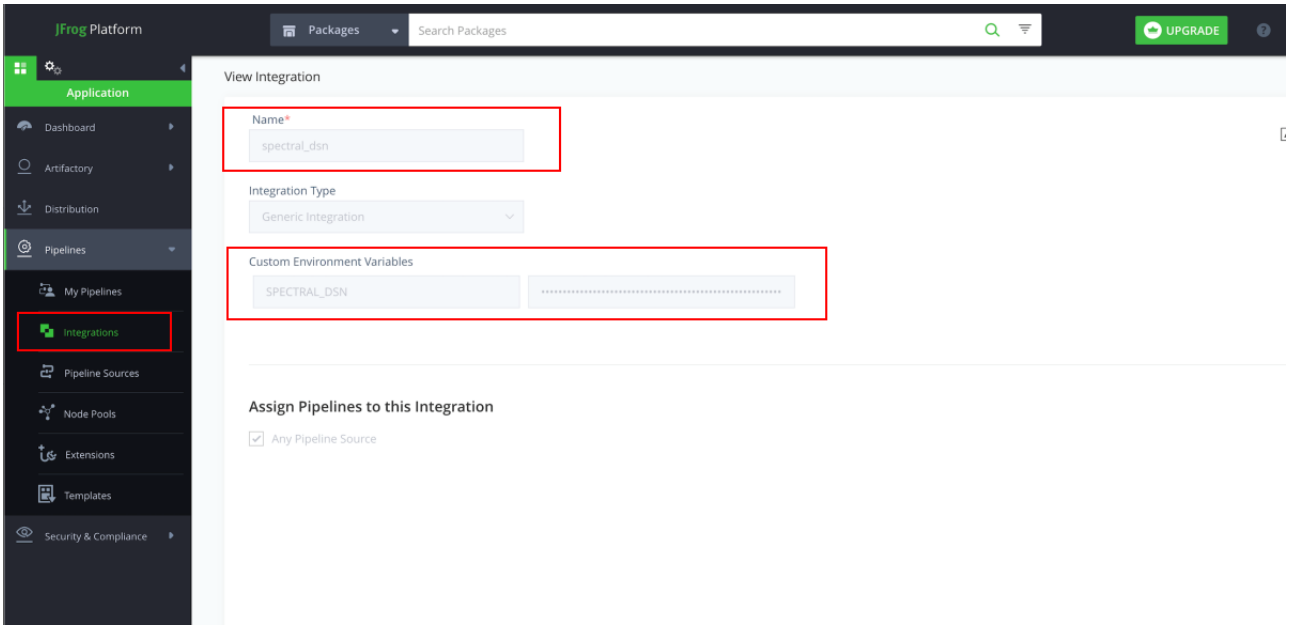

Example of configuration in the US region:

For Dome9 users

build.yaml

```
resources:
- name: myScannedRepo
  type: GitRepo
  configuration:
    # Your JFrog integration with Github
    gitProvider: "integration_name"
    # Github repository
    path: "org-name/repository-name"
    branches:
      include: main

pipelines:
- name: Spectral
  steps:
    - name: SpectralScan
      type: Bash
      configuration:
          integrations:
              - name: spectraldsn
          inputResources:
              - name: myScannedRepo
      execution:
        onExecute:
          - cd dependencyState/resources/myScannedRepo
          - curl -L "https://spectral-
us.dome9.com/latest/x/sh?dsn=$int_spectraldsn_SPECTRAL_DSN" | sh
          - $HOME/.spectral/spectral scan --ok --dsn $int_
spectraldsn_SPECTRAL_DSN
```

For Infinity users

build.yaml

```
resources:
- name: myScannedRepo
  type: GitRepo
  configuration:
    # Your JFrog integration with Github
    gitProvider: "integration_name"
    # Github repository
    path: "org-name/repository-name"
    branches:
      include: main
```

```
pipelines:
- name: Spectral
  steps:
    - name: SpectralScan
      type: Bash
      configuration:
          integrations:
              - name: spectraldsn
          inputResources:
              - name: myScannedRepo
      execution:
        onExecute:
          - cd dependencyState/resources/myScannedRepo
          - curl -L "https://spectral-
us.checkpoint.com/latest/x/sh?dsn=$int_spectraldsn_SPECTRAL_DSN" |
sh
          - $HOME/.spectral/spectral scan --ok --dsn $int_
spectraldsn_SPECTRAL_DSN
```

Example of a build integrated with Code Security:

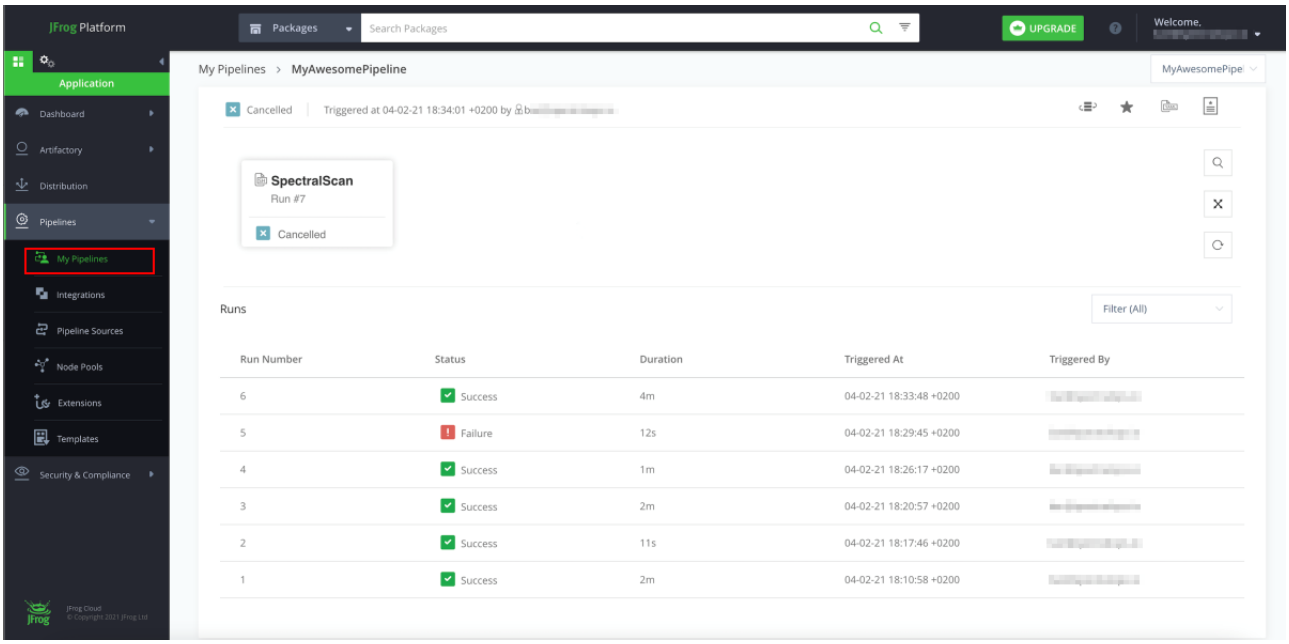

CI/CD Hardening

You can use Spectral to secure your CI/CD (continuous integration and continuous delivery/deployment) pipeline. A robust CI/CD pipeline helps to ensure that your software is delivered to your users in a timely and consistent manner, and for your developers to get immediate and fast feedback. By taking steps to harden your CI/CD pipeline with Spectral, you can help to reduce the risk of security breaches and improve the overall quality of your software while also keeping fast and efficient scans in check.

Spectral's CI/CD Hardening feature includes:

- Fast scans (keeping your pipeline fast)
- Full coverage of CI/CD steps, security rules, and guidelines. For example: locking a specific version for an Action on Github Action, other SLSA based (Supply-chain Levles for Software Artifacts) practices.
- Zero trust / fully airgapped scan: no additional permission requested and no data is sent out of your CI/CD pipeline

When you run Spectral in CI/CD hardening mode, Spectral securely fetches Github pipeline settings and security posture to your local computer. Spectral then merges this data with local Github pipeline settings and scans for issues. No traffic leaves your local computer.

Running Spectral in CI/CD Hardening Mode

To run Spectral in CI/CD hardening mode on a repository:

```
$HOME/.spectral/spectral discover github -k repo [YOUR_REPO]
```

To run Spectral in CI/CD hardening mode for a user:

```
$HOME/.spectral/spectral discover github -k user [YOUR_USER]
```

To run Spectral in CI/CD hardening mode for an organization:

```
$HOME/.spectral/spectral discover github -k org [YOUR_ORGANIZATION]
```

To run a Spectral CI/CD scan in your CI (continuous integration) or in the current project:

```
$HOME/.spectral/spectral discover github --kind repo .
```

Using Git Hooks

Adding a Simple Hook

Use the code to get started:

```
$ cd your-repo
$ echo '$HOME/.spectral/spectral scan --unstaged' > .git/hooks/pre-commit
$ chmod u+x .git/hooks/pre-commit
```

This requires Code Security installed globally for the user.

Each time you perform a commit, Code Security scans for issues.

To remove the hook, delete the file:

```
$ rm .git/hooks/pre-commit
```

Notes:

- As this repo is stored locally on your disk, Code Security cannot enforce the scan on the central server-hosted repo or on other computers.
- You can only have one running hook at a time.

Integrating Husky

You can integrate Code Security and Husky in two ways:

- Pick a hook manager for your own language.
- Pick a hook manager using for `Node.js`.

To integrate:

```
$ yarn init
$ yarn add --dev husky
$ npm set-script prepare "husky install"
$ npm run prepare
$ npx husky add .husky/pre-commit "$HOME/.spectral/spectral scan --unstaged"
$ git add .husky/pre-commit
$ git commit -m "Keep calm and commit"
```

Integrating other Hook Managers

If you prefer using a manager that is independent of a programming language, or if you prefer to use a manager specific to your programming language, refer to this website and select an applicable hooks manager.

GitHub Bot

Monitor, alert, and discover sensitive data in your code for each commit. Get instant feedback on any commit you push to your repository. Spectral Bot can be installed directly on organizations and user accounts and grant access to specific repositories. By default, only one instance the bot can be deployed to a region. To deploy more than one instance of GitHub bot to the same region, contact Customer Support.

Integration Environment Variables

Name	Required	Description	Valid Values	Default
`GITHUB_APP_ID`	Yes	GitHub app id		
`CHECK_POLICY`	Yes	If Spectral finds issues in a PR - how should we handle the PR check? The policies are based on the Spectral issue severity - critical, high, medium, low, and informational	■ `Fail on any issue` ■ `Fail on low and above` ■ `Fail on medium and above` ■ `Fail on high and above` ■ `Fail on critical only` ■ `Always Pass`	
`SPECTRAL_TAGS`	No	Include detectors by tag, separated by comma	`secrets / iac / oss`	
`SPECTRAL_ENGINES`	No	Engines list to run Spectral with, separated by commas		

Name	Required	Description	Valid Values	Default
`SPECTRAL_DSN`	No	Your Spectral DSN retrieved from SpectralOps (leave empty if you are using vault)		
`GITHUB_WEBHOOK_SECRET`	No	The GitHub app webhook secret, any strong secret would be fine (leave empty if you are using vault)		
`GITHUB_PRIVATE_KEY`	No	GitHub app private key base64 encoded (leave blank if stored in vault)		
`SECRETS_VAULT`	No	The vault you're storing your secrets in	aws_secrets_manager	
`VAULT_KEY_SPECTRAL_DSN`	No	The key in the vault where the Spectral DSN is stored. should be in the format of `Spectral_Dsn-*`		
`VAULT_KEY_GITHUB_WEBHOOK_SECRET`	No	The key in the vault where the GitHub webhook secret is stored. should be in the format of `Spectral_GithubBot_WebhookSecret-*`		

Name	Required	Description	Valid Values	Default
`VAULT_KEY_GITHUB_PRIVATE_KEY`	No	The key in the vault where the GitHub private key is stored. should be in the format of `Spectral_GithubBot_PrivateKey-*`		
`CUSTOM_COMMENT`	No	Add a custom text to the pull request comment and check summary - Markdown format		
`SHOULD_SKIP_INGEST`	No	If set to true - findings won't be sent to SpectralOps and won't be seen in your dashboard	`true` / `false`	`false`
`GITHUB_SHOULD_POST_REVIEW_COMMENTS`	No	Should review comments be posted on PR files where Spectral has identified issues	`true` / `false`	`false`
`GITHUB_SHOULD_SKIP_CHECK`	No	Disable GitHub check creation	`true` /`false`	`false`
`S3_BLACK_LIST_BUCKET_NAME`	No	Name of the bucket containing the blacklist file		
`S3_BLACK_LIST_OBJECT_KEY`	No	Blacklist file S3 object key	`true` /`false`	`false`

Name	Required	Description	Valid Values	Default
`STRICT_MODE`	No	If set to true - issues from changed files in the PR will fail the check even if the issues are in lines that hasn't been modified		
`HOME`	No	If the environment does not automatically set this variable, specify a path to an existing location that has write permissions. For example, in AWS Lambda, where this variable is not preset, assign it the value /tmp		
`GITHUB_SELF_HOSTED_DOMAIN`	No	Provide the domain if you're running a self-hosted Github (for example: *https://my-github-domain.com*).		
`CUSTOM_SSL_CERTIFICATE`	No	The SSL certificate content for on-prem Github. Use this variable if not using AWS Secrets Manager for the certificate.		
`VAULT_KEY_CUSTOM_SSL_CERTIFICATE`	No	The vault key to retrieve the custom SSL certificate.		`Spectral_custom_ssl_certificate`

Configuring Code Security Github Bot

1. Create a new GitHub app:
   - To install the app on an organizational account, use this link.
   - To install the app on a personal account, use this link.
2. Select on which depositories to install the app:
   - **All repositories** (applies to all current and future repositories owned by the resource owner)
   - **Only select repositories**

3. Use **one** of these methods to deploy the bot:

Method	How to Deploy

Cloud Formation	**Prerequisite**
	These permissions are required in your AWS account:
	`lambda:GetAccountSettings`
	`cloudformation:DescribeStacks`
	`iam:CreateRole`
	`iam:DeleteRole`
	`apigateway:POST`
	`logs:CreateLogGroup`
	`iam:PutRolePolicy`
	Procedure
	a. In AWS, launch the stack.
	b. In the settings of your new GitHubb app, generate a private key.
	c. Encode the private key in base64 without line breaks. Example OpenSSL command:
	`openssl base64 -A -in YOUR-PRIVATE-KEY.pem`
	d. In AWS > **GITHUB_PRIVATE_KEY** field, enter the private key that you created in the previous step.
	e. For Check Policy, select **one** of these:
	▪ **Fail on any issue**
	▪ **Fail on critical only**
	▪ **Fail on high and above**
	▪ **Fail on medium and above**
	▪ **Fail on low and above**
	▪ **Always pass**
	f. Set the value of the **STRICT_MODE** parameter. If this parameter is **true**, the status check runs on all issues. If this parameter is set to **false**, the status check runs only for new issues.

Terraform	Terraform - Deploy AWS resources using the Code Security **Terraform module**. Set the value of the `integration_type` parameter to `github`). **Note**: The Terraform deployment is supported starting from GitHub bot 2.x versions.
Docker	Follow the deployment instructions in the Code Security DockerHub repo.

4. Go to the **GitHub app settings** page.
5. In the **Webhook** section, enter these values:
   - **Webhook URL**
   - **Webhook secret**

   Note - In AWS Lambda, to generate the webhook URL, append `/api/github/event` to the ServiceEndpoint output from this stack: `https://<id>.execute-api.<region>.amazonaws.com/prod/api/github/event`

   Note - In Docker, to generate the webhook URL, append `/events/github` to the URL for your container: `https://<container-domain>/events/github`
6. In the **Webhook** section, set the **Webhook** to **active**.

Using a Vault

We recommend to store your secrets in a vault instead of directly in the configuration of your Lambda. Only AWS Secret Manager is supported.

To use AWS Secret Manager to store your secrets:

1. In Code Security, set the value of the `SECRETS_VAULT` environment variable to `aws_secrets_manager`.
2. In your vault, set these secrets:

Secret	Value
`GITHUB_PRIVATE_KEY`	`Spectral_GithubBot_PrivateKey`
`GITHUB_WEBHOOK_SECRET`	`Spectral_GithubBot_WebhookSecret`
`SPECTRAL_DSN`	`Spectral_Dsn`

Using Custom Vault Keys

Use these Code Security integration environment variables to create custom secret vault keys:

Note - If you use AWS Secrets Manager with a CloudFormation or Terraform deployment, the role you created for the Lambda functions gives permission to perform the action only for the default secrets.

Integration Environment Variable	Description
`VAULT_KEY_SPECTRAL_DSN`	Spectral DSN key name, in this format: `Spectral_Dsn-*`
`VAULT_KEY_GITHUB_WEBHOOK_SECRET`	GitHub app webhook secret, in this format: `Spectral_GithubBot_WebhookSecret-*`
`VAULT_KEY_GITHUB_PRIVATE_KEY`	Private key, in this format: `Spectral_GithubBot_PrivateKey-*`
`VAULT_KEY_CUSTOM_SSL_CERTIFICATE`	**Optional -** This is a key for a vault that contains a self-signed certificate. Default value: `Spectral_custom_ssl_certificate`

Advanced Configuration: Excluding Repositories

To exclude specific repositories from the scan, add them to an exclusion list of repo URLs.

1. Create a text file containing a list of full URLs of repositories to exclude (for example: *https://github.com/expressjs/express*). Put each URL on a new line in the file.
2. In Code Security, define these integration environment variables:

Integration Environment Variable	Description
`S3_BLACK_LIST_BUCKET_NAME`	The name of the bucket containing the exclusion list file

Integration Environment Variable	Description
`S3_BLACK_LIST_OBJECT_KEY`	The object key of the exclusion list file

Advanced Configuration: Configuring Multiple GitHub Apps with a Single Instance of Code Security Bot

To configure mutliple GitHub apps with a single instance of Code Security bot, configure a *multi_app.json* file in the root directory and add it to the frontend and backend Lambda functions.

Step 1: Edit the multi_app.json file

The GitHub object is a dictionary containing all of the GitHub Apps. For each app, you must set the *private_key*, which is a base64-encoded string of the GitHub App's private key. In the <app_id>

Procedure

1. In the AWS Lambda deployment, in the root directory, open the *multi_app.json* file.
2. In the *<app_id>* section, enter the GitHub App ID to configure a specific application.
3. **Optional -** Configure one or more of these optional parameters:

Parameter	Description
`check_failure_policy`	Determines the failure policy for scans from this specific GitHub App. If not set, it defaults to `CHECK_POLICY`. Refer to the table above for possible values.
`spectral_dsn`	Allows splitting scans using a different DSN for this GitHub App. It defaults to `SPECTRAL_DSN`.
`webhook_token`	Allows using a webhook secret for this GitHub App. It defaults to `GITHUB_WEBHOOK_SECRET`.
`secret_vault`	Enables storing secrets (webhook secret, app private key, spectral DSN) in a vault at the app level. When using a vault, you must specify all the keys for the webhook secret, GitHub App private key, and optional spectral DSN.

Parameter	Description
`self_hosted_domain`	Allows self hosted domain for this GitHub App. It defaults to `GITHUB_SELF_HOSTED_DOMAIN`.

Example multi_app.json Configuration File

```
{
  "github": {
    "<app_id>": {
      "private_key": "",
      "spectral_dsn": "",
      "webhook_token": "",
      "self_hosted_domain": "",
      "secret_vault": {
        "name": "aws_secrets_manager",
        "key_webhook_secret": "",
        "key_private_key": "",
        "key_spectral_dsn": "",
        "key_ssl_certificate": "",
      }
    }
  }
}
```

Step 2: Put the multi_app.json file in the AWS Lambda Deployment:

1. Download and unzip the Lambda code:
   a. Download the Lambda code as a *.zip* file.
   b. Unzip the downloaded file.
2. Add the *multi_app.json* configuration file you edited in Step 1 to the unzipped directory.
3. Re-zip the configuration file and upload it to the Lambda:
   a. Select all of the files in the unzipped directory.
   b. Zip the files.
   c. Upload the new *.zip* file to AWS Lambda. Make sure that the *multi_app.json* configuration file is in the *.zip* file.

Monitoring Code Security Github Bot with CloudWatch alarms

We recommend to monitor GitHub Bot errors using CloudWatch alarms in AWS Lambda. Follow this procedure in CloudWatch:

1. Click **Create Alarm**.
2. Click **Select Metric**.
3. Select the Errors metric for the Lambda.
4. Click **Select Metric**.
5. Set the statistic to **Sum** and select the required period to measure the threshold.
6. In the **Conditions** section:
    a. For **threshold**, select **static**.
    b. Select **Greater**.
    c. Set the threshold value to **1** or a greater number.
    d. Click **Next**.
7. Add the email addresses to receive the alarm.
8. Click **Next**.
9. Enter a name for the alarm.
10. Click **Next**.
11. Read the summary. If the configuration is correct, click **OK**.

Upgrading the Github Bot

The Code Security GitHub bot is versioned.

To upgrade the l GitHub bot with Lambda:

In Lambda, in the **Code** tab, upload a new *.zip* file of the Github bot. Upload the new version to the frontend Lambda and to the backend Lambda.

To upgrade the GitHub bot with Docker:

Deploy a new GitHub bot image from the docker hub.

Gitlab Bot

Monitor, alert, and detect sensitive data in your code for each push or merge request. Get instant feedback on any changes you make in a merge request. The Spectral Bot can be installed directly on organizations, user accounts, groups, and specific projects. It supports both system-level webhooks (GitLab Enterprise) and group-level webhooks.

By default, only single instance of the bot can be deployed to a region. To deploy several bot instances to the same region, contact Support.

Integration Environment Variables

To configure the Spectral Gitlab bot, set these variables in your environment:

Name	Required	Description
CHECK_ POLICY	Yes	Determines how the pipeline status should be handled if Spectral finds issues in a merge request. Options based on issue severity: "Fail on any issue", "Fail on low and above", "Fail on medium and above", "Fail on high and above", "Fail on critical only", "Always Pass"
GITLAB_SELF_ HOSTED_ DOMAIN	No	Provide the domain if you're running a self-hosted GitLab (e.g., https://my-gitlab-domain.com).
GITLAB_ WEBHOOK_ SECRET	No	A strong secret for securing the webhook. If you use a vault, leave this field empty.
GITLAB_ ACCESS_ TOKEN	Yes, if not using vault	Generate it in your GitLab profile under Access Tokens, with the "api" scope (leave empty if using a vault).
VAULT_KEY_ GITLAB_ ACCESS_ TOKEN	No	The vault key to retrieve the GitLab access token (e.g., Spectral_GitlabBot_AccesToken). Required if using a vault.
VAULT_KEY_ GITLAB_ WEBHOOK_ SECRET	No	The vault key to retrieve the GitLab webhook secret. Defaults to Spectral_GitlabBot_WebhookSecret.
VAULT_KEY_ SPECTRAL_ DSN	No	The key in the vault where the Spectral DSN is stored. should be in the format of Spectral_Dsn

Name	Required	Description
GITLAB_ SHOULD_SKIP_ PIPELINE	No	Specifies whether to skip pipeline creation on push events. Options: true, false. Default: false.
STRICT_MODE	No	If set to true, the check status is based on all issues found in the modified files, even if the issues are old.
SPECTRAL_ TAGS	No	Comma-separated list of tags to run Spectral with (e.g., base,iac,audit).
SPECTRAL_ ENGINES	No	Comma-separated list of engines to run Spectral with (e.g., secrets,iac,oss). Default: secrets.
SPECTRAL_ DSN	No	Your Spectral DSN from SpectralOps (leave empty if using a vault).
SECRETS_ VAULT	No	The vault where your secrets are stored. Currently supports aws_secrets_manager.
CUSTOM_ COMMENT	No	Custom text to add to the merge request summary (Markdown format).
SHOULD_SKIP_ INGEST	No	If set to true, findings won't be sent to SpectralOps and won't appear in your dashboard. Options: true, false. Default: false.
S3_BLACK_ LIST_BUCKET_ NAME	No	Name of the S3 bucket containing the blacklist file.
S3_BLACK_ LIST_OBJECT_ KEY	No	The S3 object key for the blacklist file.
CUSTOM_SSL_ CERTIFICATE	No	The SSL certificate content for on-prem GitLab. Use this variable if not using AWS Secrets Manager for the certificate.
VAULT_KEY_ CUSTOM_SSL_ CERTIFICATE	No	The vault key to retrieve the custom SSL certificate. Defaults to Spectral_custom_ssl_certificate.

Using a Vault

We recommend to store your secrets in a vault. To do this, set the SECRETS_VAULT environment variable to specify the type of vault you're using.

The AWS Secrets Manager is the only supported vault. Its environment variable is aws_secrets_manager vault.

Set these secrets in your vault:

- Spectral_GitlabBot_AccesToken (for GITLAB_ACCESS_TOKEN)
- Spectral_GitlabBot_WebhookSecret (for GITLAB_WEBHOOK_SECRET)
- Spectral_Dsn (for SPECTRAL_DSN)

If you deploy AWS Secrets Manager with CloudFormation or Terraform, make sure that the Lambda role can access the secrets using the secretsmanager:GetSecretValue action.

Deploy the Bot

Use CloudFormation, Terraform, or Docker to deploy the bot.

Cloud Formation deployment

1. Make sure the Lambda has these permissions in the AWS account:

   ```
   lambda:GetAccountSettings
   cloudformation:DescribeStacks
   iam:CreateRole
   iam:DeleteRole
   apigateway:POST
   logs:CreateLogGroup
   iam:PutRolePolicy
   ```

2. Launch the stack. See AWS documentation.
3. Generate a private key in the settings of your new GitHub app.
4. Encode the private key in base64 without line breaks. This is an example command in OpenSSL: `openssl base64 -A -in your-private-key.pem`
5. Enter the private key in the **GITHUB_PRIVATE_KEY** field.
6. For **Check Policy**, select one of these:
   - **Fail on any issue**
   - **Fail on critical only**
   - **Fail on high and above**
   - **Fail on medium and above**

- **Fail on low and above**
- **Always pass**

Note - If **STRICT_MODE** is set to **true** the status check includes all issues found in the modified files. If **STRICT_MODE** is not set to **true**, the status check includes only new issues.

Terraform deployment

Use the Spectral Terraform module to deploy AWS resources. Set the value of the `integration_type` parameter to `github`.

Deploy AWS resources using our Terraform module (set the integration_type parameter value to github).

Note -The Terraform deployment is supported starting from Github bot version 2.x.

Docker deployment

Follow the deployment instructions on our DockerHub repo for the GitHub scanner.

Configuring Multiple GitHub Apps with a Single Instance of Spectral Bot

Step 1: Deploy AWS Lambda

1. Download and unzip the Lambda code:
   a. Download the Lambda code as a *.zip* file.
   b. Unzip the downloaded file.
2. Add the *multi_app.json* file to the unzipped directory. Make sure that the frontend and backend lambda functions contain the file in their deployments. This is an example of the *multi_app.json* file:

```
{
  "github": {
    "<app_id>": {
      "private_key": ""
    }
  }
}
```

3. Zip the file and upload it to Lambda:

a. Select all of the files in the unzipped directory.

b. Zip the files again.

c. Upload the new *.zip* file to AWS Lambda.

Step 2: Configure the "github" object in the "multi_app.json" file

The github object is a dictionary containing all the GitHub Apps. Do these steps for each app:

1. Set the *private_key*. This is a bse64-encoded string of the GitHub App's private key.
2. Make sure the *<app_id>* contains the GitHup App ID to configure a specific application.
3. **Optional -** Set one or more of these optional parameters:
   - `check_failure_policy`: Determines the failure policy for scans from this specific GitHub App. The default value of this parameter is defined by the `CHECK_POLICY` Integration Environment Variable.
   - `spectral_dsn`: Allows splitting scans using a different DSN for this GitHub App. The default value of this parameter is defined by the `SPECTRAL_DSN` Integration Environment Variable.
   - `webhook_token`: Allows the use of a webhook secret for this GitHub App. The default value of this parameter is defined by the `GITHUB_WEBHOOK_SECRET` Integration Environment Variable.
   - `secret_vault`: Enables storing secrets (webhook secret, app private key, spectral DSN) in a vault at the app level. When using a vault, you must specify all the keys for the webhook secret, GitHub App private key, and optional spectral DSN.

Example multi_app.json Configuration

```
{
  "github": {
    "<app_id>": {
      "private_key": "",
      "spectral_dsn": "",
      "webhook_token": "",
      "secret_vault": {
        "name": "aws_secrets_manager",
        "key_webhook_secret": "",
        "key_private_key": "",
        "key_spectral_dsn": ""
      }
```

```
      }
    }
}
```

Storing Secrets in a Vault

We recommend to store your secrets in a vault. Currently, only AWS Secrets Manager is supported.

To store secrets in a vault

1. Set the value of the `Secrets_VAULT` Integration Environment Variable to `aws_secrets_manager`.
2. Set these secrets in your vault:
   - `Spectral_GithubBot_PrivateKey` (for `GITHUB_PRIVATE_KEY`)
   - `Spectral_GithubBot_WebhookSecret` (for `GITHUB_WEBHOOK_SECRET`)
   - `Spectral_Dsn` (for `SPECTRAL_DSN`)
3. Set these Integration Environment Variables to create custom secret vault keys:
   - `VAULT_KEY_SPECTRAL_DSN` - Spectral DSN key name.
     Format: *Spectral_Dsn-**
   - `VAULT_KEY_GITHUB_WEBHOOK_SECRET` - GitHub app webhook secret.
     Format: *Spectral_GithubBot_WebhookSecret-**
   - `VAULT_KEY_GITHUB_PRIVATE_KEY` - Private key.
     Format: *Spectral_GithubBot_PrivateKey-**

   Note - In a CloudFormation or Terraform deployment, the role that you create for the lambdas gives permission to perform the `secretsmanager:GetSecretValue` action only for those 3 secrets.

Exclude Repositories

To prevent the bot from scanning specific repositories you can supply the bot with the list of these "blacklisted" repo URLs.

To exclude repositories from the scan

1. Create a blacklist text file that contain a list of full URLs of repositorie to exclude(for example: *https://github.com/expressjs/express*). Start each URL on a new line.
2. Define these Integration Environment Variables:

- `S3_BLACK_LIST_BUCKET_NAM`E - the name of the bucket containing the blacklist file.
- `S3_BLACK_LIST_OBJECT_KEY` - the object key of the blacklist file

Complete the GitHub App Setup

1. Go to the GitHub app settings.
2. Enter the relevant **Webhook URL**:
   - For AWS Lambda: Append */api/github/event* to the ServiceEndpoint output from the stack.
     For example: *https://<id>.execute-api.<region>.amazonaws.com/prod/api/github/event*
   - For Docker:
3. Save the changes.

Monitoring

We recommend to monitor bot errors. If the bot is hosted by AWS lambda, you can configure CloudWatch alarms. For more information, see AWS documentation.

Gitlab Pipeline

You can run Code Security scans on a Gitlab pipeline to secure merge requests. To secure your merge requests without using pipelines, use the Code Security *"Gitlab Bot" on page 724*.

Basic Configuration

1. In the SPECTRAL_DSN file, set Integration Environment Variables to define the scan.
2. In your Gitlab environment, upload the *SPECTRAL_DSN* file to the Gitlab CI/CD Variables.

Best Practice - Verify the digest of a downloaded runnable file before you run it. You can use Code Security Preflight to verify the digest. Follow this link to see SHA digests of the binary, the download script and the gzip.

Code Security Integration Environment Variables for Gitlab Pipeline

Name	Required	Description
`GITLAB_TOKEN`	Yes	Generate it in your Gitlab profile -> Access Tokens, check the "api" scope (leave blank if you are using vault)
`SELF_HOSTED_GITLAB_DOMAIN`	Yes	If you are running a self-hosted Gitlab, enter the domain. For example: *https://my-gitlab-domain.com*
`SPECTRAL_DSN`	Yes	Your Spectral DSN retrieved from SpectralOps (leave blank if you are using vault)
`SPECTRAL_TAGS`	No	Tags list to run Spectral with, separated by commas (for example: base,iac,audit)
`SPECTRAL_ENGINES`	No	Engines list to run Spectral with, separated by commas (for example: secrets,iac,oss). Default is 'secrets'
`STRICT_MODE`	No	If set to true, check status is based on all issues found in the modified files (even if the issues are old)

Example of a Basic Configuration:

```
build-job:
  stage: build
  script:
    - curl -L "https://app.spectralops.io/latest/x/sh?dsn=$SPECTRAL_
DSN" | sh
```

```
    # This takes your SPECTRAL_DSN from the variables store in Gitlab
CI/CD
    - $HOME/.spectral/spectral scan --ok  --include-tags base,audit
```

Advanced Configuration: Gitlab Pipeline Scan of Changed Files

In an advanced configuration, you can limit the scan of the Gitlab Pipline on changed files. This feature is supported only on Standalone GitLab servers.

Prerequisites

- Only Standalone GitLab servers are supported
- You must use Docker in your Gitlab environment to limit the scan to changed files.

You can limit the scan to after a Merge Request event or to after a Standalone Push.

Setting	Description
Merge Request Events	When the job runs in a merge request context, Code Security scans only the files that were changed in this merge request. **Best Practice** - Change the repository configuration to allow merges only f the pipeline is successful.
Standalone Pushes	If the job runs without a merge request context, Code Security finds the diff (changed files) between these commits: ▪ the last commit ▪ the last commit before the push Then, Code Security scans only the diff between these commits.

To limit a scan of a Gitlab Pipeline to changed files:

In your Gitlab pipeline, define a job in the *.gitlab-ci.yml* configuration file to run a dedicated docker image named `checkpoint/spectral-gitlab-pipeline-scanner`

Example of a configuration in the *gitlab-ci.yml* file that limits the scan to after Merge Request Events

```
spectral-scan:
  stage: test # specify which stage the job should run
  allow_failure: true # should the job fail the whole pipline
  image: checkpoint/spectral-gitlab-pipeline-scanner:latest
```

```
  script:
    - /usr/src/app/scanner
  rules:
    - if: $CI_PIPELINE_SOURCE == 'merge_request_event'
```

Example of a configuration in the *gitlab-ci.yml* file that limits the scan to a direct push to the main branch (Standalone Push)

```
spectral-scan:
  stage: test # specify which stage the job should run
  allow_failure: true # should the job fail the whole pipline
  image: checkpoint/spectral-gitlab-pipeline-scanner:latest
  script:
    - /usr/src/app/scanner
  rules:
    - if: $CI_COMMIT_REF_NAME == $CI_DEFAULT_BRANCH
```

Gitlab Pre-Receive Hook

A Gitlab pre-receive hook prevents software developers from pushing commits to Gitlab if Code Security found security issues in the code.

Prerequisites

- The Gitlab instance must be installed on a Linux server.
- Code Security must be installed on the same Linux server as the Gitlab instance. See *"Getting Started with Code Security" on page 650*.

Configuration

Create a configuration file and install it as a server hook on the Gitlab Linux server.

1. On the Gitlab server, create a configuration file in **one** of these directories:
   - in your home directory: `.spectral/git-hook`

     for example: `/var/opt/gitlab/.spectral/git-hook/config.yml`
   - `/etc/spectral/git-hook`
2. Add parameters to the configuration file. The required parameters are `spectral_ds` and `gitlab_host`.

```
# file name for log, must be writable by "git" user hook_logfile:
/var/log/gitlab/hook.log # optional, but recommended # logging level.
For debugging purposes level 6 can be used log_level: 3 # optional,
default is 3 (show errors) spectral_dsn: https://spk-
******@get.spectralops.io # REQUIRED # Path to spectral scanner
executable. The default is "spectral" # which works if "spectral"
binary containing directory is listed in PATH environment variable. #
It is required to be installed into your gitlab instance. # For
installation instructions under your account at get.spectralops.io
press "add sources" , then "use CLI", # then "local repos". spectral_
binary: spectral # optional, default is "spectral" # Additional
arguments to "spectral history" command, for instance, to exclude some
checks. Use with caution! spectral_args: [] # optional, array of
strings # stop execution after specified number of seconds; child
process `spectral` will also be killed exec_timeout: 30 # optional,
seconds (int), default is 30 # if this string appears in any of the
commit messages being pushed, git hook will skip scanning; # use it
with caution when you are sure that commits trigger false positive and
ONLY false positive scan results; # likely, you may want to alter the
latest commit message with "amend" commit to skip scanning skip_scan_
keyword: ~ # string, default "skip-spectral-pre-receive" # gitlab
server host url gitlab_host: http://your-own-gitlab-server.com #
REQUIRED
```

3. Download the Gitlab pre receive hook.
4. Install the Gitlab pre-receive hook as a server hook.

Best Practice - Install the Gitlab pre-receive hook globally for all repositories.

Logging

Spectral pre-receive hook allows you to enable the hook logs from the configuration file when the pre-receive hook is triggered.

Before enabling the hook logging, make sure that GitLab logrotate is enabled or custom logrotate that you manage in the server. See the GitLab log guide for more details.

Spectral pre-receive does not delete or rotate the file that was written in your Gitlab machine.

To enable logging:

In the *configure.yaml* configuration file, change the value of the `hook_logfile` parameter to `/var/log/gitlab/gitlab-rails`.

Note - You can change the value of the `hook_logfile` parameter to a different path, if logrotate is managed for this path and the path exists in the *logrotate.conf* file. To get Gitlab the logrotate configuration, run: `$ cat /var/opt/gitlab/logrotate/logrotate.conf`

To change the number of logs Code Secruity writes to the log file:

In the *configure.yaml* configuraiton file, change the value of the `log_level` parameter (default value = `3`). To write more logs, increase the value. To write fewer logs, decrease the value.

To disable logging:

In the *configure.yaml* configuration file, clear the `log_level` field.

Bitbucket Pre-Receive Hook

A Bitbucket pre-receive hook prevents software developers from pushing commits to Gitlab if Code Security found security issues in the code. This feature is supported only for a Gitlab instance on a Linux server.

Prerequisites

- The Gitlab instance must be installed on a Linux server.
- Code Security must be installed on the same Linux server as the Gitlab instance. See *"Getting Started with Code Security" on page 650*.
- To load a custom Code Security configuration from the repository, you must have a Bitbucket service account with read-only personal access. If you do not have this type of Bitbucket service account, custom configurations from the repository do not apply.

Step 1: In Code Security, configure environment variables

Configure the Bitbucket pre-receive hook in Code Security by defining environment variables.

Environment variable	Description	Example value
`SPECTRAL_BIN`	Path to the Code Security binary.	`$HOME/.spectral/spectral`
`SPECTRAL_DSN`	DSN from Code Security (SpectralOps), in **Settings** > **Organization**.	`https://spu-xxxx@get.spectralops.io`
`SPECTRAL_COMMAND`	Code Security command line arguments.	`scan --include-tags base,iac`
`BITBUCKET_SERVER_BASE_URL`	The URL of the Bitbucket server.	`https://mydomain.com/bitbucket`

Environment variable	Description	Example value
`BITBUCKET_PAT`	The personal access token for Bitbucket (optional functionality to load *spectral.yaml* from repo).	`TmljZSB0cnkhIEJ1dCB0aGlzIGlzbid0IGEgcmVhbCBvbmUu`
`SPECTRAL_CHECK_POLICY`	The severity of findings allowed before the push is rejected. These values are supported: ■ `fail on any issue` (default value) ■ `fail on low and above` ■ `fail on medium and above` ■ `fail on high and above` ■ `fail on critical` ■ `always pass`	`fail on any issue`
`SPECTRAL_CUSTOM_REJECT_MESSAGE`	A custom text to append to the rejection message.	`Spectral has rejected your push`

Example configuration

```
SPECTRAL_DSN="https://spu-xxxx@get.spectralops.io"
SPECTRAL_CHECK_POLICY="fail on high and above"
BITBUCKET_SERVER_BASE_URL="https://mydomain.com/bitbucket"
BITBUCKET_PAT="TmljZSB0cnkhIEJ1dCB0aGlzIGlzbid0IGEgcmVhbCBvbmUu"
SPECTRAL_COMMAND="scan --include-tags base,iac"
SPECTRAL_BIN="~/.spectral/spectral"
```

Step 2: On the Bitbucket server, install the pre-receive hook

1. Download the Code Security plugin JAR file to your computer.
2. In the Bitbucket server UI, open the **Administration** settings.
3. Click **Manage apps** > **Upload app**.
4. Upload the Spectral plugin JAR file to the Bitbucket server.
5. Do one of these:
   - Enable the hook for the entire Bitbucket server:
     a. by going to the project -> Project settings -> Hooks and enabling "Spectral Pre Receive Hook"
   - You can also enable the hook for specific repository by going to the specific repository -> Repository settings -> Hooks and enabling "Spectral Pre Receive Hook"

Code Security Integrations

Code Security supports integrations with:

- Slack for communication
- Jira to open tickets and start your workflow based on the issue which was detected by Code Security
- Monday to create items and start your workflow based on the issue which was detected by Code Security
- Pager Duty for alerts management
- Custom webhooks
- Events Webhook integration

All the integration settings are global. You can also define a team-level integration except for Custom webhooks.

To open the Code Security Integrations menu:

1. From the left toolbar, click **Code Security**.
2. From the second toolbar from the left, click **Settings**.
3. From the Settings toolbar, click **Integrations**.

Slack Integration

1. Create a Slack app.
2. Enable incoming webhooks.
3. Copy your webhook URL, paste in the field

For more information, see the Slack documentation.

Jira Integration

You can use Spectral to scan your Jira environment. To connect your Jira account, you must have:

- Jira domain URL
- Jira admin/user email
- Jira API Token, you can view how to get the key here

For more information and configuration instructions, see *"Code Security Integration with Jira" on page 749*.

Confluence Integration

You can use Spectral to scan the content in a Confluence instance. The integration uses a Lambda function in AWS and a webhook in Confluence. For more information, see *"Code Security Integration with Confluence " on page 747*.

Monday Integration

Global Integration

To set the integrations globally, and connect your Monday account, you require the Monday API Token. To obtain the key, click here.

Upon successful integration, the Monday account domain is presented at the top of the Monday integration section.

Team level Integration

You can set a default Monday workspace for a team. By choosing a workspace for your teams, this workspace is prioritized in the workspaces list, while creating a Monday item.

Create an Item

You can create a Monday item, in a specific Monday board, from the issues table (Code/Log/Host/Productivity/Asset quick view). The workspaces list is sorted by:

If the asset of the selected issue is mapped to a team, and a Monday workspace is configured for this team. This workspace is on the top of the workspaces list.

If the user belongs to one or more teams, the workspaces that are configured for these teams is presented at the top of the list.

Additional Fields

You can get the same view in Monday as in Spectral by adding these fields:

- **Severity** > (severity, text/label)
- **asset** > (asset, text)
- **content** > (content, text)
- **firstSeen** > (first seen, date)

The additional fields are presented in the create item form - with no edit option.

Comparison of Events Webhook Integration and Custom Webhook Features

	Events Webhook Integration	Custom Webhook
Events	Notifies users of changes in the state of issues or assets. Covers a range of event types including updates to issues and assets (such as `issues_updated`, `issues_created`, and `assets_updated`). Notifies users when Spectral finishes the ingest process.	Notifies users only about new issues.
Payload	Contains only the resource IDs of the changed items	Contains detailed issue data, including information such as asset ID, detector ID, detector name, severity, and URI to view the issue in the source code.

Custom Webhook Integration

You can configure Code Security to send notifications through a Custom Webhook Integratoin.

An account administrator can configure custom webhook integrations in the **Settings** screen > **Integrations** tab.

How a Custom Webhook Integration Works

1. You enter the webhook URL.
2. Code Security generates a token for the integration. You cannot change this token.
3. Code Security uses the token to sign every notification using HMAC with SHA 256 encryption.
4. Code Security sends notifications to the webhook URL as a POST request (JSON enconded). The request must be 200 OK.

This is an example of the POST request that Spectral sends to the webhook URL. It is in an open API schema:

```
issues:
                type: array
                items:
                type: object
```

```
                properties:
                assetId:
                type: string
                title: Asset id
                description: The asset in which this issue has been
discovered
                detectorId:
                type: string
                title: Detector id
                description: The detector id
                detectorName:
                type: string
                title: Detector name
                description: The detector description
                severity:
                type: string
                title: Severity
                description: The severity of this issue (error /
warning / info)
                uri:
                type: string
                title: URI
            description: View the issue in your source code via this
uri
```

This is a sample of the content:

```
{
                issues: [
                {
                assetId: 'github.com/maggie/jelly-sandwich',
                detectorId: 'SENF026',
                detectorName: 'Ruby On Rails secret token
configuration file',
                severity: 'error',
                uri: 'https://github.com/maggie/jelly-
sandwich/blob/ae556dbaa9a8ce4e37a5fe9e95b7823eff79f379/config/initiali
zers/secret_token.rb#L7'
                },
                {
                assetId: 'github.com/maggie/jelly-sandwich',
                detectorId: 'CLD001',
                detectorName: 'Visible AWS Key',
                severity: 'error',
                uri: 'https://github.com/maggie/jelly-
sandwich/blob/ae556dbaa9a8ce4e37a5fe9e95b7823eff79f379/db/secret2#L4'
                },
                ]
```

```
            }
```

Verifying the Notification

The signature can be found in the `x-spectral-signature` header. An extra layer of security is added by using the timestamp (ISO format) to create the signature hash.

To sign the content, `[TIMESTAMP]_[SECRET-TOKEN]` is used.

To make sure the request has not replayed, you can verify the `x-spectral-signature` header in the request is in the last 5 seconds, for example. You can verify the signature using:

```
// NodeJS example
                import crypto from 'crypto'

                const verifySignature = (signature, signatureToken,
timestamp, payload) => {
                const calculatedSignature = crypto
                .createHmac('sha256', `${timestamp}_
${signatureToken}`)
                .update(payload)
                .digest('hex')

                return (signature === calculatedSignature)
            }
```

You can verify your implementation:

```
verifySignature
('108a388872e723c0b4be83a0b37abb5a55f9a89c17209c47bb11e8064ccd811d',
'myverysecrettoken', '0', 'bar') // returns true
```

Events Webhook Integration

The Events Webhook Integration feature allows users to receive notifications when the state of an issue or asset changes. This feature ensures that users are promptly informed about updates, facilitating better monitoring and management of their resources. To reduce the number of events, Spectral combines related events.

Event Payloads

This section describes the structure of event payloads for changes in issues or assets, and for ingest completion.

All events include the name of the event and the id or ids of the resource and the time of the event. To retrieve all the data associated with these resource IDs, see the Spectral API documentation.

Note - The webhook times out after 10 seconds. Make sure that your webhook endpoint can respond within 10 seconds.

Changes in Issue or Asset

When an issue or asset changes, Spectral creates an event with these fields:

```
{
                "name": String,
                "ids": [String],
                "time": YYYY-MM-DDTHH:mm:ss.SSSZ
            }
```

Field	Description
`name`	Shows the type of change events. Possible values: ■ `issues_updated` ■ `issues_deleted`
`ids`	Shows a list of ids that changed. Issue ids appear in `issues_updated`. Asset ids appear in `assets_deleted`.
`time`	Shows the time when the event occurred.

Ingest Completion

After Spectral finishes the ingest process, it creates an event with this payload:

```
{
                "name": String,
                "id": String,
                "issues_deleted_count": number,
                "issues_created_count": number,
                "issues_updated_count": number,
                "issues_resolved_count": number,
                "time": YYYY-MM-DDTHH:mm:ss.SSSZ
            }
```

Field	Description
`name`	The name of every Ingest Completion event is `asset_ingest_completed`.
`id`	Shows the asset id related to the ingest event.
`issues_deleted_count`	Shows the number of deleted issues in the ingest event.
`issues_created_count`	Shows the number of new issues in the ingest event.
`issues_updated_count`	Shows the count of updated unresolved issues in the ingest event.
`issues_resolved_count`	Shows the count of resolved issues in the ingest event.
`time`	Shows the time when the event occured.

Testing Your Webhook Integration

You can send a test notification to your webhook endpoint to simulate an event.

To test your webhook integration:

1. In the **Integrations** page, scroll to the section for the integration (for example: the **Custom Webhook Global Integration** section)
2. Click **Test**.

   A popup message shows you the test results.

Verifying a Signature for a Webhook Integration

The signature appears in the `x-spectral-signature` header.

Spectral uses a timestamp in ISO format to create the signature hash. Spectral uses this token to create the signature: `[TIMESTAMP]_[SECRET-TOKEN]`.

You can use the timestamp to make sure that Spectral does not send duplicate information.

This is an example of code to verify the signature:

```
// NodeJS example
                import crypto from 'crypto'
```

```
                const verifySignature = (signature, signatureToken,
timestamp, payload) => {
                const calculatedSignature = crypto
                .createHmac('sha256', `${timestamp}_
${signatureToken}`)
                .update(payload)
                .digest('hex')

                return (signature === calculatedSignature)
            }
```

This is an example of code to verify the implementation:

```
verifySignature
('108a388872e723c0b4be83a0b37abb5a55f9a89c17209c47bb11e8064ccd811d',
'myverysecrettoken', '0', 'bar') // returns true
```

Code Security Integration with Confluence

You can use Code Security to scan the content in a Confluence instance. The integration uses a Lambda function in AWS and a webhook in Confluence.

Prerequisites

- The Lambda function requires these permissions in AWS:

```
cloudformation:DescribeStacks
iam:CreateRole
iam:DeleteRole
apigateway:POST
logs:CreateLogGroup
iam:PutRolePolicy
```

- In your Confluence instance, you must install the webhooks manager extension.

To integrate Code Security with Confluence:

1. In AWS, launch this stack to deploy the Lambda function.
2. In your Confluence instance, add a new webhook in the webhook manager. Change `YOUR_ACCOUNT` to your instance domain in this url:

   https://YOUR_ACCOUNT.atlassian.net/wiki/plugins/servlet/ac/com.stiltsoft.confluence.cloud.webhooks/admin-webhooks-page)
3. Configure a webhook URL to point to your function endpoint. The function endpoint appears in the Lambda page in the AWS console.
4. Add the query string parameter `webhook_token`.
5. Add the query string parameter `webhook_token` to the same webhook token you put in the `CONFLUENCE_WEBHOOK_TOKEN` parameter in the Lambda function.
6. Add these event types to the webhook:
   - `attachment created`
   - `comment created`
   - `comment updated`
   - `page created`
   - `page updated`

- `content created`
- `content updated`

7. To test the integration, open a Jira issue with a fake secret (for example: `AKIA4HK52OLF2AAN9KWV`).

Example Configuration

Object	Value
Confluence event endpoint URL	*https://random123.execute-api.us-east-1.amazonaws.com/prod/api/confluence_event*
Token that you set in your `function env var`	`f4lmf4kl2ldoxxxxxx`
Webhook URL you configure in Confluence	`https://random123.execute-api.us-east-1.amazonaws.com/prod/api/confluence_event?webhook_token=f4lmf4kl2ldoxxxxxx`

Code Security Integration with Jira

You can use Spectral to scan the content of Jira issues, including summaries, descriptions, comments, and attachments. The integration uses a Lambda function in AWS and a webhook in Jira.

Integration Environment Variables

Variable	Required	Description
SPECTRAL_DSN	Yes	Your Spectral DSN retrieved from SpectralOps
JIRA_WEBHOOK_TOKEN	Yes	A token used to identify the sender, should be identical to the webhook token sent in the webhook_token param to the webhook endpoint
EMAIL	No	The email matching the Jira API token. If this is not provided, attachments will not be scanned
SPECTRAL_TAGS	No	Tags list to run Spectral with, separated by commas (eg base,iac,audit). Default is 'base'
REDACTED_MESSAGE	No	In case of active remediation - a custom message to replace findings
REMEDIATION_MODE	No	How to handle findings (Valid values: "Not active" / "Redact finding")
JIRA_API_TOKEN	No	A Jira api token to scan attachments as well. If this is not provided, attachments will not be scanned
JIRA_PROJECTS_BLACKLIST	No	A comma delimited list of project keys that you want to exclude from being scanned
JIRA_PROJECTS_WHITELIST	No	A comma delimited list of project keys that you want to scan. No other projects except these will be scanned

Configuration

Prerequisite

The Lambda function requires these permissions in AWS:

```
cloudformation:DescribeStacks
iam:CreateRole
iam:DeleteRole
apigateway:POST
logs:CreateLogGroup
iam:PutRolePolicy
```

To integrate Spectral with Jira:

1. Use one of these methods to deploy the Lambda function:
   - CloudFormation
   - Terraform
2. Copy the function gateway API URL.
3. In your Jira instance, add a new webhook in **System Settings** > **Webhooks to send events**. For example:

   https://YOUR_ORG_NAME.atlassian.net/plugins/servlet/webhooks
4. Mark these events for the webhook to send:
   - `issue->create+update`
   - `comment->create+update`
   - `attachment-> create`
5. Copy the Lambda URL from AWS and use it as the webook URL. Make sure to copy the full URL and a query string parameter for the webhook secret you entered when you installed the Lambda function. For example:

   https://random1.execute-api.us-east-1.amazonaws.com/prod/api/jira_event?webhook_token=[YOUR WEBHOOK SECRET]
6. To test the integration, open a Jira issue with a fake secret (for example: `AKIA4HK52OLF2AAN9KWV`).

Code Security Integration with Terraform Cloud

Protect your infrastructure by detecting potential issues in your Terraform configuration and plan before applying the changes to production. You can integrate Code Security with Terraform Cloud or Terraform Enterprise. This integration is based on an AWS lambda function which is being triggered by the Run Task at the relevant stage.

Terraform Cloud Integration Types

Code Security can be integrated with the Pre-plan stage or the Post-plan stage of the Terraform run:

Pre-plan stage

This stage takes place right before the plan stage.

In this stage,Code Security scans your Terraform configuration deployed in this run for misconfigurations.

Post-plan stage

This stage takes place between the plan and apply stages.

In this stage ,Code Security scans the generated plan of the current run for potential issues before applying the changes to your live infrastructure.

You can read more about run tasks in Terraform Cloud here and here.

Integration Environment Variables

Variable	Required	Description
`SPECTRAL_DSN`	Yes	Your Spectral DSN retrieved from SpectralOps
`CHECK_POLICY`	Yes	If Spectral finds issues - how should we handle the run? The policies are based on the Spectral issue severity - critical / high / medium / low / informational (Valid values: "Fail on any issue" / "Fail on low and above" / "Fail on medium and above" / "Fail on high and above" / "Fail on critical only" / "Always Pass")
`HMAC_KEY`	No	A key that will be used for securing your Run Task by validating the request payload signature, should be identical to the HMAC key you set will set in the Run Task
`TERRAFORM_USER_KEY`	No	User key created by Terraform (required for pre-plan run task) - can be created here

Configuration

Step 1: In AWS, create a Lambda Functiont)

1. Use **one** of these methods to create the required AWS resources:
   - Use this Terraform module (set the `integration_type` param value to `terraform`).
   - Launch the CloudFormation stack.
2. Add all the required environment variables, including `SPECTRAL_DSN`, `CHECK_POLICY` and `HMAC_KEY`.
3. **Optional -** If you plan to create a pre-plan Run Task:
   a. Create a user API key in Terraform
   b. Set the user API key in the `TERRAFORM_USER_KEY env` variable.
4. Copy the Gateway API URL and keep it in a safe place.

   Note - If you are using the Terraform module , use the `rest_api_url` output.

Step 2: In Terraform Cloud, create a Run Task

1. In Terraform Cloud, log in to your organization.
2. From the top menu, click **Settings**.
3. From the side menu, in the **Integrations** section, click **Run tasks**.
4. In the **Create a Run Task** section:
   a. Enter a name for the run task.
   b. In the **Endpoint URL** field, paste the **Gateway API URL** you copied from AWS.
   c. Enter a description for the run task.
   d. In the **HMAC key** field, past the value of the `HMAC_key` from Code Security.
5. Click **Create run task**.

Step 3: In Terraform Cloud, add the Run Task to your Workspace

1. Open the relevant Terraform Cloud Workspace.
2. Expand **Settings** > click **Run Tasks**.
3. In the **Configure Run Task** menu, set the **Enforcement Level** to **Mandatory**.
4. Click **Save**.

Step 4: In Terraform Cloud, test the Run Task

In Terraform Cloud, trigger a run to test the Run Task.

Detectors

Code Security comes with an industry's leading detector coverage with over 500 different detectors built-in, including machine learning based detectors. You can use Code Security's detector toolkit to build your own detectors for any purpose you find suitable.

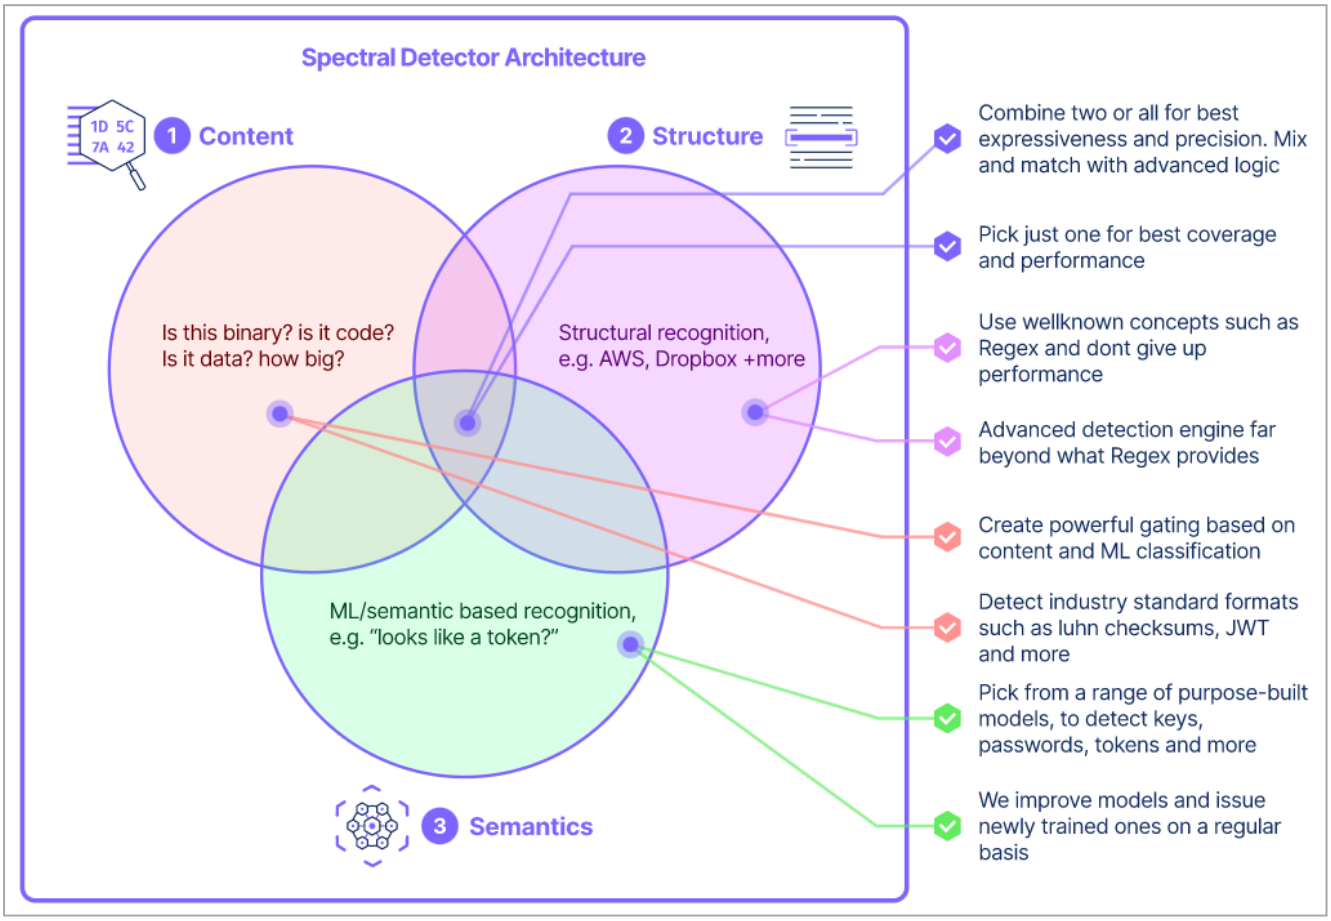

Creating a Detector Rule

You can create your own detector rule.

For example, if you want to create a detector rule for merchant IDs; `MKAT82jv72D2g92a`, `MKAT02maibnwMw37s`, `MKAT10dmwugmwpz9` and so on:

1. Run:

```
$ $HOME/.spectral/spectral init multi hello-rules --pattern "MKAT
[0-9a-zA-Z]+"
```

2. Code Security creates a detector rule:

```
rules:
- id: RUL001
  name: Your sample rule
  pattern_group:
    patterns:
    - pattern: "MKAT[0-9a-zA-Z]+"
      pattern_type: multi
```

3. Use the detector rule:

```
$ $HOME/.spectral/spectral run --just-ids RUL001 --nosend
```

If there are new merchant IDs; `KAT__310527195`, `KAT__310527110`, `KAT__310527100` and so on, you can update and use the existing detector rule as follows:

```
rules:
- id: RUL001
  name: Your sample rule
  pattern_group:
    patterns:
    - pattern: "MKAT[0-9a-zA-Z]+"
      pattern_type: multi
      pattern: "KAT__[0-9]+"
      pattern_type: multi
```

Hierarchical Expressions

Code Security runs a hierarchical matching engine, which means, a pattern - pattern expression within a pattern - pattern expression and so on.

For example, to add a condition if a KAT expression found and the file in which the expression is found must contain the word **production**:

```
rules:
- id: RUL001
  name: Your sample rule
  applies_to: # filter for the correct file pattern
  - ".*\\.py$"
  pattern_group:
    patterns:
    - pattern: "MKAT[0-9a-zA-Z]+"
      pattern_type: regex
    - pattern: "KAT__[0-9]+"
      pattern_type: regex
      pattern_group:   # the same pattern group construct repeats
        patterns:
            - pattern: ".*production.*"
              pattern_type: regex
```

Testers

Each pattern can have an assortment of testers that are sub detectors for different types of data.

For example, to ensure that the last part of MKAT validates as a machine generated secret:

```
rules:
- id: RUL001
  name: Your sample rule
  applies_to: # filter for the correct file pattern
  - ".*\\.py$"
  pattern_group:
    patterns:
    - pattern: "MKAT[0-9a-zA-Z]+"
      pattern_type: regex
    - pattern: "KAT__(.+)"
      pattern_type: regex
      test_token:
      - on: 1
        is: true
```

You can add any character in `KAT__(.+)` and Code Security uses machine learning based tester to find the right match.

Concepts

The Code Security rule engine contains basic constructs that in turn, have powerful knobs and switches. In combination, it makes for a robust query engine and static analyzer looking for security patterns.

Pattern

Patterns are the basic building blocks in Code Security. It is an expression you are searching in a text. Each pattern can optionally contain a pattern group, used as a verification or additional conditions on the pattern.

Pattern Group

A pattern group aggregates a bunch of patterns by a logic that you specify.

For example, for rules A, B, and C, you can create expressions like:

- A and B and C
- A or (B and C)
- not A or (B and C)

Filtering and Tagging

You can activate a rule only for certain file path patterns. In addition, you can add tag to rules so that you can filter the rules. For example:

```
rules:
- id: RUL001
  name: Your sample rule
  applies_to: # filter for the correct file pattern
  - ".*\\.py$"
  tags:
  - python
```

For more information on how to select categories of rules, see *"Configuring Code Security" on page 679*.

Building Detectors

A detector is a combination of one or more building blocks, such as `test_regex` (in the example below), and a workflow building block, such as `pattern_group` (in the example below), that consists of logic, patterns, machine learning and statistical tests.

```
rules:
  id: HELO001
  pattern_group:
    aggregate: or
    patterns:
    - pattern: "hello (.*)"
      test_regex:
      - on: 1
        pattern: universe|world
    - pattern: "namaste (.*)"
      test_regex:
      - on: 1
        pattern: bramhaand|vishv
```

Code Security optimizes:

- **Extreme performance** - Built with Rust's zero-overhead principles and low-level optimization.
- **Security** - Code Security is built with safe-only code and sandboxes detectors.
- **Productivity** - Code Security detectors are programming language agnostic there's no need to understand Java to scan Java; or any other programming language.
- **Rapid detector building** - No manual compilation, code and run. Code Security compiles and optimizes each detector automatically.
- **Declarative over programmatic** - Specify what you want to find, and Code Security finds it.
- **Automatic fingerprinting and tracking** - Code Security analyzes each finding and automatically create a secure irreversible and detectable fingerprint.

Detecting Sensitive Files or a Hardcoded JWT Secret in your Codebase

You can build a detector to find SSH-related sensitive files.

1. Run:

```
$ cd my-project
$ $HOME/.spectral/spectral init
```

2. Create a file in `.spectral/rules/my-rules.yaml`:

   - To detect a sensitive file:

```
rules:
  - id: SENS001
    applies_to:
      - "(?i).*_(rsa|dsa|ed25519|ecdsa)$"
    description: An SSH related sensitive file was found
    name: Sensitive SSH file
    recommendation_template: Please add sensitive files to
your .gitignore
    severity: error
    tags:
      - base
      - sensitive-files
    pattern_group:
      aggregate: or
      patterns:
        - pattern: "."
          match_on_path: true
          pattern_type: single
```

- To find a hardcoded JWT secret in your codebase:

```
rules:
  - id: SEC001
    description: A JWT (JSON Web Token) has been found to be
hardcoded
    name: Sensitive JWT (JSON Web Token)
    recommendation_template: Please remove the hardcoded
token, report it to SecOps for rotation, and fix with using
.env
    severity: error
    tags:
    - base
    - secrets
    pattern_group:
      aggregate: or
      patterns:
      - pattern: "=\\s+(.*)" # assignment
        pattern_type: multi
        test_jwt:
        - on: 1
          is: true
```

3. Run a scan:
   - For interactive sessions, use `$ $HOME/.spectral/spectral run --nosend`
   - To perform a scan in your CI, use`$ $HOME/.spectral/spectral scan`
4. To test the detector with a dummy file, run:

   `$ echo 'x' > id_rsa`
5. Run `$HOME/.spectral/spectral run` to view new detections.

Detector Item	Description
`rules:`	Indicates that it is a detector rules file.
`SENS001`	Detector ID that appears in the Code Security.
`applies_to`	Detect against the full file path. You also can use `applies_not_to` in combination.
`base`	Adds it to the base ruleset of Code Security. It is run by default with all other detectors.
`pattern_ group`	Sets a pattern group of a element with a logical OR relationship between elements. You can add more patterns.

Detector Item	Description
`match_on_path`	Rewires the engine to look at the file path as the tested content.
`pattern_type: single`	Attempts only once to find a match (no multiple results in same file here).
`pattern: "."`	Matches any character on the path.

Detecting an Actual Secret

You can ensure your codebase does not contain a well-known, organization-specific secret, such as:

- Development team credit card number
- Private network domains
- Internal server addresses
- Vendor and customer secrets

You can use fingerprinting to hide such sensitive information even in the detector rule.

1. Run:

```
$ cd my-project
$ $HOME/.spectral/spectral init
$ $HOME/.spectral/spectral fingerprint --text sekr3t
[fingerprint text]
```

2. Create a file in `.spectral/rules/my-rules.yaml`:

```
rules:
  - id: PRIV001
    description: A private organization secret is found hardcoded
in files
    name: Private secret
    recommendation_template: Please remove the hardcoded secret,
report it to SecOps for rotation, and fix with using .env
    severity: error
    tags:
    - base
    - secrets
    pattern_group:
      aggregate: or
      patterns:
      - pattern: "(:?key|token|secret|password|pwd|passwd)=(.*)"
# assignment
        pattern_type: multi
        test_fingerprints:
        - on: 1
          fp: [fingerprint text]
          is: true
```

3. Run `$ $HOME/.spectral/spectral run --nosend`

Detector Item	Description
`rules:`	Indicates that it is a detector rules file.
`SENS001`	Detector ID that appears in the Code Security.
`applies_to`	Detect against the full file path. You also can use `applies_not_to` in combination.
`base`	Adds it to the base ruleset of Code Security. It is run by default with all other detectors.

Detector Item	Description
`pattern_group`	Sets a pattern group of a element with a logical OR relationship between elements. You can add more patterns.
`match_on_path`	Rewires the engine to look at the file path as the tested content.
`pattern_type: multi`	Attempts multiples to find a match.
`pattern: "."`	Matches any character on the path.
`on: 1`	If `(:? .. )` is ignored, it scans in the first capture group.
`test_ fingerprints`	Scans for the fingerprints.
`is: true`	The fingerprint scan returned a match.

Logic Rules (OPA)

You can integrate Code Security with the Open Policy Agent (OPA) project such that it can use its logic facilities.

Rego in Code Security

A generic schema of a Code Security-compatible Rego policy:

```
package opa_rule

Policy[result] {
    < LOGIC >

    result := {
        "id": "CR-AF001",
        # search expression to detect line numbers when needed
        # "keysearch": "",
        "severity": "WARNING",
        "text": "Airflow encryption scheme should be more strict",
        # just a reference URL for playbooks, and additional
information
        "url": "https://example.com"
    }
}
```

Best practice for building detectors based on Rego:Principles:

- Every policy returns a result.
- Every result has a clear and strict schema, which Code Security requires to create a finding.
- Multiple policies can co-exist in the same rule, as long as each of those have a different ID.

Notes:

- `The package opa_rule` is mandatory.
- The `result :=` and its assigned document are mandatory.

Building a Custom Code Security Detector with OPA

Code Security supports two different match types:

- `opa_inline - pointing` to a rego rule inlined inside the yaml, in the pattern field. You must specify parser with `params.parser`.

- `opa_file - pointing` to a physical file on disk. pattern spec is `parser:namespace:file-relative-to-spectral.yaml`.

Example:

```
rules:
- id: AF001
  tags:
  - base # activate by default, part of the base package
  applies_to:
  - ".*\\.conf$"
  severity: info
  pattern_group:
    aggregate: or
    scope: text
    patterns:
    # remember: rego file must declare package 'opa_rule'
    # <parser>:<policy name>:<path to rego file>
    - pattern: "ini:airflow-main-configuration:./af.rego"
      pattern_type: opa_file
      test_regex_prematch:
      # this detects that the configuration file is *actually* Airflow
related
      - pattern: "airflow_"

- id: AF002
  tags:
  - base # activate by default, part of the base package
  applies_to:
  - ".*\\.conf$"
  severity: info
  pattern_group:
    aggregate: or
    scope: text
    patterns:
    # pattern is the actual rego policy, and we specify a parser
in 'params:' later below
    - pattern_type: opa_inline
      params:
        parser: ini
      test_regex_prematch:
      # this detects that the configuration file is *actually* Airflow
related
      - pattern: "airflow_"
      pattern: >
        package opa_rule

        Policy[result] {
          true # <some logic>
```

```
          result := {
            "id": "AF002",
            "severity": "WARNING",
            "text": "Airflow encryption seed value should not be
visible",
            "url": "https://example.com",
          }
        }
```

Rule Settings

Code Security's OPA rules are a combination of:

- Routing - the concept of detecting configuration types and pointing to the correct parser and policy bundle.
- Logic - Defined through Rego.
- Rule settings (`rule_settings`) - Overlay of settings to allow for user customization.

Rule setting example:

```
RuleSettings{
  shared: json_value
  rules: json_value
}
```

Corresponding yaml file example:

```
rule_settings:
  shared:
    whitelist_email: ".*"
  rules:
    CR010:
      min_committers: 3
```

Usage in rules is optional and depends on your requirement to customize settings for your rules. If you want to use rules, then we recommend this pragma to set up rule settings inside the Rego policy:

```
package opa_rule

_s := object.get(input, "__rule_settings__", {})
```

```
rules := object.get(_s, "rules", {})

shared := object.get(_s, "shared", {})

Policy[result] {
    rule_id := "ELA001"
    settings = object.get(rules, rule_id, {})
    port := object.get(settings, "port", 9200)

    not input.play.server.http.port = port
    result := {
        "id": rule_id,
        "severity": "WARNING",
        "text": "Elastic should always be on port 9200",
        "url": "https://example.com",
    }
}
```

To override the port value to a different value, use this rule settings in `spectral.yaml`:

```
rule_settings:
  rules:
    ELA001:
      port: 20200
```

Keysearch

Use keysearch for configuration that are unpredictable in its shape to find a line number.

No capture

```
"keysearch": "some-string-to-find"
```

Capture

Typically used with multi-line. Enable multiline with `(?s)` flag.

```
"keysearch": "(?s)foo.*bar(capture-this = .*?)andthat"
```

Codeprinting

Code Security can detect code copies, partial copies and fuzzy copies. If a piece of your sensitive code, configuration or any textual assets must be in a specific predefined place, Code Security can create a custom detector that looks for stray copies of it or partial (modified) copies of it.

You can use codeprinting to:

- Locate a configuration sprawl. A securely stored sensitive configuration file that individuals copy-paste between projects that are deemed unsafe.
- Trace a file found in mobile apps by mistake and is now delivered to many end-user devices as part of an APK build.
- Locate a complete codebase that is misplaced on a production server, a sandbox computer or other unauthorized devices.

Creating Effective Codeprints

Code Security helps you create codeprints securely and locally. Your code is never transmitted anywhere, and all codeprint hashes are produced with a local and secure hashing algorithm.

Codeprint hashes are a safe one-way hash converted into a textual string, which also hold a comparative trait which Code Security uses to measure code copies, partial copies, or fuzzy copies. You can store these in your custom detectors for detecting code copying.

Warning - Keep your codeprints private to your organization. While codeprints are not secrets, and cannot be reversed to the original text, virtually all hashes or one-way functions, such as MD5, SHA256, and others can be used to extract indirect knowledge about an organization.

Quick Start

Creating a codeprint essentially is creating a custom detector with your codeprints in it.

To create a codeprint, run:

```
$ $HOME/.spectral/spectral fingerprint --codeprint [FILE1] [FILE2]
...
```

Pick files that represent code, configuration, docs or other pieces of information that is unique to your organization or are deemed sensitive.

Code Security generates a detector for you and generates an output:

```
rules:
  - id: CPRT001
```

```
applies_to:
  - ".*$"
description: Detect code copies via secure codeprinting
name: Codeprint detector
severity: info
tags:
  - base
  - codeprints
pattern_group:
  patterns:
    - pattern: ".*"
      match_on_path: true
      pattern_type: single
      test_codeprints:
      - print: ".."
      - print: ".."
```

- `applies_to` - Use this to block any unwanted file for scanning.
- `match_on_path` - Rewires spectral to look at file paths and not content. You can use `pattern` to apply a secondary filtering rule (regex).
- `test_codeprints` - Actual codeprints.

You can copy or pipe to your own `spectral/rules/rules.yaml` file and store the file in a secure location.

Do's

- For each file Code Security scans, it matches against one of the codeprints in the list, so you can add more than one codeprint.
- If you have a sensitive file that you want to codeprint, you can create a detector just for that one.
- If you have a large codebase or assets you want to protect, try to identify the most unique-to-you files and create a codeprint for all of those.

Don'ts

- Use a very small file (smaller than 2 KB), because it might not contain enough data to be unique.
- Avoid using a public-domain, or a file that is not originally yours, such as a piece of open source code. It matches all the instances in the open source library, which can be used by a lot of codebases

Security

Codeprint is one-way. Code Security compresses and encrypts to avoid brute-googling of the simhashes.

Detector Engine

Query Structure

Detectors are composed of rules, or queries that are compiled into an efficient detector and are run with the Code Security engine against files.

Each query is a group of patterns, called a `pattern_group` and is hierarchical (a pattern group can contain more pattern groups and so on).

A pattern group is a collection of patterns with an *aggregate relation*.

```
    pattern_group:
      aggregate: or | and | append
      patterns:
      - pattern: "(:?key|token|secret|password|pwd|passwd)=(.*)" #
assignment
        pattern_type: single
      - pattern: "hello" # assignment
        pattern_type: multi
```

`pattern_group`

Aggregating patterns can:

- `and` - Bail out on the first mismatch.
- `or` - Try any of the matches.
- `append` - Try any of the matches and collect matches from those that matched.

`pattern`

Each pattern is of a *type* (`pattern_type`):

- `single` - Match once per file.
- `multi` - Match many times per file.

Both accept a performant, binary and text traversing regex.

Prematch Testers

A prematch tester is a test that runs before applying an in-depth matching and detection logic. As an example, it is more appropriate to use bail out detection for a small binary file with class documentation.

`test_content_prematch`

This meta-tester uses content classification and inference engine. It is a collection of testers that are useful when deciding if a certain file is worth getting a deep dive into.

By testing for content, you can:

- Filter for an unexpected binary file.
- Ensure a non-empty file goes through for further detection.
- Be able to run on *classes* of files, where Code Security has classified those by their content nature.

Content class	Example
`code/infra`	Ruby, Python
`data/infra`	SQL, JSON
`binary`	Binary files
`docs`	Markdown, Text
`tests`	Unit tests, other test code
`examples`	Example code, demo code and others
`vendor`	3rd party code sitting in `node_modules` and others
`files`	A general file class not fitting a single class.

Usage

```
pattern: ".*"
test_content_prematch:
    binary: false
    minlen: 20,
    maxlen: 2000,
    content_classes:
    - code/infra # our own classfication engine results
    # content_classes_not:
    # - Code # the inverse of content_class
    content_types:
    - Python # a programmingg language *name* (if you want
extension, there's ways for that too)
    content_types_not:
    # - Ruby # the inverse of content_types
```

Test positive

```
<a Python file, size at 20-2000 bytes>
```

Test negative

```
<an SQL file, or a small file, or a binary file, etc.>
```

`test_regex_prematch`

You can test for a specific pre-match structure before Code Security deep dives into further matching.

By testing for Regex prematch, you can:

- Make sure a certain file structure exists before applying further testing, such as variable assignments.
- Verify that a certain 'sentinel' word exists in a large file by applying a generic word lookup, before applying a more specific matching.

Usage

```
pattern: "pass:(.*)"
test_regex_prematch:
    - on: 0 # on full text
      pattern: "aws\\.amazon\\.com"
```

Test positive

```
<large documentation file>
Here is how to connect to our database
1. Log into AWS console (console.aws.amazon.com)
2. Use following details:
DB pass: shazam123
```

Test negative

```
<Big file, not containing any mention of AWS detail>
```

Content Testers

`test_fingerprints`

Code Security can create one-way fingerprints for you to use when you want to detect pieces of information you cannot reveal.

By using `test_fingerprints`, you can:

- Detect credit cards
- Find classified or private domains or hosts

First, you must generate your fingerprint. It is done locally on your machine using a secure and salted one-way hash:

```
$ $HOME/.spectral/spectral fingerprint --text <your private text>
< fingerprint >
```

Then, copy the resulting fingerprint.

Usage

```
pattern: "host=([a-zA-Z0-9_-.]+)"
test_fingerprints:
  - on: 1
    with: "<your fingerprint>"
    is: true
```

```
Note that by specifying the character class and narrowing it down,
we give some
useful information to attackers looking to bruteforce private
information. Always be mindful that your character classes and
secrets are wide enough.
```

Test positive

```
<private host>
```

Test negative

```
<any other text>
```

`test_from_env`

You can collect secrets from your ENV, rather than encode those as fingerprints and still search for them in your code. Code Security supports fetching those from your ENV, and relaying to the detector to use.

By using `test_from_env`, you can:

- Detect secrets that you already have in your environment (local machine or CI) without exposing them.
- Find secrets that you do not want to expose in a persistent way.

To test, make sure to export it first:

```
$ SOME_SECRET_VAR=shazam $HOME/.spectral/spectral scan --nosend
```

Usage

```
pattern: "host=(.*)"
test_from_env:
  - on: 1
    with: "SOME_SECRET_VAR"
    is: true
```

Test positive

```
shazam
```

Test negative

```
foobar
```

`test_luhn`

The Luhn algorithm is used for check-sum of a credit card and many forms of Social Security Number (SSN) numbers of the US, Canada and Israel.

By testing for Luhn, you can:

- Ensure a number is a valid credit card number.
- Verify that a given string match passes as a valid SSN, which helps identify fake from test strings.

Usage

```
pattern: "account=([0-9]+)"
```

```
test_luhn:
  - on: 1
    is: true
```

Test positive

```
79927398713
```

Test negative

```
79927398710
```

References

- Wikipedia

`test_number`

Available from: v1.4.2

Test for an generic representation of a number.

By testing for numbers, you can rule out a value that is supposed to be a password or a token.

Usage

```
pattern: "key=(.*)"
test_number:
  - on: 1
    is: false
```

Test positive

```
key=<random token>
```

Note that by returning `false` and `is: false`, `test_number` provides a positive outcome.

Test negative

```
key=0.1234
```

`test_base64`, `test_base64bin`

Verify that a text is a base64 encoded or binary encoded. Supports all common variants of encoding (URL safe and others).

By testing for base64, you can:

- Ensure that a match is base64 and fail fast in a sequence of tests when you are looking for a token.
- Validate that a string is base64 encoded given you suspect that it may contain sensitive information.

Usage

```
pattern: "account_encoded='([[:alnum:]/+]+[=]{0,2})'"
test_base64:
  - on: 1
    is: true
```

Test positive

```
account_encoded='eyAiYWNjb3VudCI6ICJzZWNyZXQtbnVtYmVyIiB9'
```

Test negative

```
account_encoded='replace_me'
```

The binary variant first decodes the base64 encoded string, and then tests whether it is binary or not:

```
pattern: "account_encoded='([0-9]+)'"
test_base64bin:
  - on: 1
    is: true
```

test_binary

As Code Security detectors are binary-aware, you can test for binary matches in any capturing expression.

By testing for binary data, you can flag and avoid matches that are false and contain no text.

Usage

```
pattern: "token=(.*)"
test_binary:
  - on: 1
    is: false
```

Test positive

```
<BINARY DATA>token=<BINARY_DATA>
```

Test negative

```
token=48SfRa4idxxUVyPAejafXxwjkreyj8MoJkjV
```

The binary variant first decodes the base64 encoded string, and then tests whether it is binary or not:

```
pattern: "account_encoded='([0-9]+)'"
test_base64bin:
  - on: 1
    is: true
```

test_maxlen, test_minlen

Test for content size, minimum or maximum.

By testing for content size, you can:

- Ensure to fail fast for very short strings or very large content, and skip the match.
- Validate that on top of the various structural captures that you have done, you end up with a reasonable sized match.

Usage

```
pattern: "account_encoded='(.*)'"
test_minlen:
  - on: 1
    score: 2
```

Test positive

```
account_encoded='eyAiYWNjb3VudCI6ICJzZWNyZXQtbnVtYmVyIiB9'
```

Test negative

```
account_encoded='XX'
```

In the same way, you can use `maxlen`:

```
pattern: "account_encoded='(.*)'"
test_maxlen:
  - on: 1
    score: 2000
```

Structural Testers

`test_jwt`

A JWT(JSON Web Token) test is an Internet proposed standard for creating data with optional signature and/or optional encryption, whose payload holds JSON that asserts claims, often used for service-to-service authentication.

By testing for JWT, you can:

- Make sure the key structure fits a standard JWT.
- Verify that a certain JWT is semantically valid (header is valid).

Usage

```
pattern: "token=(\\S+)"
test_jwt:
  - on: 1
    is: true
```

Test positive

```
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJsb2dnZWRJbkFzIjoiYWRtaW4iLCJ
pYXQiOjE0MjI3Nzk2Mzh9.gzSraSYS8EXBxLN_oWnFSRgCzcmJmMjLiuyu5CSpyHI
```

Test negative

```
bad_token
```

References

- JSON Web Token at Wikipedia

test_uri

A URI or URL parsing test. A given string is tested to be a valid URI.

By testing for URI, you can:

- Isolate URLs that are sensitive before applying further matching logic.
- Detect various kinds of authentication, such as Bearer, Basic and more, given a URL request structure (for example `curl`'ing URLs).

Usage

```
pattern: "curl\\s.*(http.*)"
test_uri:
  - on: 1
    is: true
```

Test positive

```
curl -L -o https://dev.acme.corp/secure/credentials.json -
H"Authorization: Bearer <token>"
```

Test negative

```
sh curl.sh arg1 arg2
```

test_tvar

Test for various template variables, common in configuration and IaC files.

By testing for template variables, you can filter for legitimate configuration that was built with proper template variables instead of hardcoded secrets.

Usage

```
pattern: "DB_PASS=(.*)"
test_tvar:
  - on: 1
    is: false
```

Test positive

Note - `is: false` is a positive outcome if the candidate does not contain a template variable.

```
DB_PASS=my-secret-password
```

Test negative

```
DB_PASS={{.Env.DBPass}}
```

`test_changeme`

Available from: v1.4.2

Test for various *changeme* values. Developer sometimes indicate a value to be replaced by various commonly-known idioms, such as `fixme` and `XXX`, which is also known as *changeme*.

By testing for *changeme*:

- You can filter for mock values, or "TODO: replace this" values.
- Use this in combination with other testers to create a powerful detector.

Usage

```
pattern: "DB_PASS=(.*)"
test_changeme:
  - on: 1
    is: false
```

Test positive

Note - `is: false` is a positive outcome if the candidate does not contain a *changeme* value.

```
DB_PASS="<real password>"
```

Test negative

```
DB_PASS="XXX"
```

`test_assignment`

Available from: v1.4.2

Test for an assignment structure.

By testing for assignment:

- You can set the scene for detectors which are only interested in one part of an assignment clause.
- Combine an expected assignment with another tester to create a more powerful detector.

Usage

```
pattern: "DB_PASS(.*)"
test_assignment:
  - on: 0 # on the complete expression
    is: true
test_token:
  - on: 1
    is: true
```

Test positive

```
DB_PASS=<random token>
```

Test negative

```
DB_PASS, foo, bar
```

`test_uuid`

Test if a given string is a UUID. It supports all UUID types and formats (with or without hyphens, and with or without a prefix).

By testing for UUID, you can ignore suspect strings that are randomly generated but in fact are IDs (database IDs or other).

Usage

```
pattern: "key=(.*)"
test_uuid:
  - on: 1
    is: false
```

Test positive Note `is: false` so a positive outcome is candidate NOT containing a UUID:

```
key=my-secret-key
```

Test negative

```
key=<UUID representing a DB table primary key>
```

`test_regex`, `test_regex_not`

A `test_regex` is a tester that can verify a structural form after a match is a found. You can verify the match further.

By using `test_regex`, you can:

- Apply a clearer set of validations, readable and maintainable.
- Split verification into stages to pronounce a specific use case:
  - Capture *something vague*. For example, `Bearer (.*)`).
  - Run a semantic tester. For example, `test_token` on the token part of the bearer.
  - Run a structural tester. For example, "it should look like a `curl` request" with `test_regex`.
- Apply verification that is beyond a Regex DFA capabilities. For example, a state machine with more aggressive but performant backtracking can first be achieved by running two separate ones and combining later.

As an array based tester, an `AND` relation is created between elements, and short-circuiting (failing fast) is applied.

- `test_regex` - all must apply, fail if one does not apply
- `test_regex_not` - all must *not* apply, fail if one applies

Usage

```
pattern: "token=(.+)"
test_regex:
    - on: 1
    pattern: "([0-9].*){2}" # the value include at least 2 numbers
    - on: 1
    pattern: "([a-zA-Z].*){2}" # the value include at least 2
letters
```

Test positive

```
token=48SfRa4idxxUVyPAejafXxwjkreyj8MoJkjV
```

Test negative

```
token=env.get('token')
```

Usage (`test_regex_not`)

```
pattern: "token=(\\S+)"
test_regex_not:
    - on: 1
    pattern: "[$][a-zA-Z0-9_-]+" # the value include valid template
variable.
    - on: 1
    pattern: "(?i)(exmaple|test|fake|1234|abcde|xxxx|foobar)" # the
value include some word or pattern that can tell that this is just a
token placeholder.
```

Test positive

```
token=48SfRa4idxxUVyPAejafXxwjkreyj8MoJkjV
```

Test negative

```
token=$my_token
token=48SfRa4idxxUVyPAejafXxwjkreyjEXMAPLE
token=testRa4idxxUVyPAejafXxwjkreyj8MoJkjV
token=XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
token=test1234abcdefoobarfake
```

Semantic Testers

`test_cword`

Test for the percentage of common words in a given string. Based on a unique and massive tech-related common words dictionary model.

By testing for common words. you can:

- Rule out non-machine generated keys.
- Validate that a given match passes as a machine generated secret.

Usage

```
pattern: "pass=(.*)"
test_cword:
  - on: 1
    from: 0.0 # defines a range of accepted percentage
    to: 0.2   # low percentage of common words (up to 20%)
```

Test positive

```
zx28821a{_)
```

Test negative

```
hello
```

`test_zx`

Test for password strength based on the popular zxcvbn library.

By testing for `zx` (abbreviated), you can:

- Detect strong passwords amongst fake.
- Apply existing policies for enforcing password strength.

Usage

```
pattern: "pass=(.*)"
test_zx:
  - on: 1
    score: 4.0 # same standard score scale (0-4) from zxcvbn
```

Test positive

```
zxHELLOyw{_)
```

Test negative

```
foobar
```

`test_pass`

Test for password strength (own model). Pick a threshold on a scale of 0-100.0. A password with strength `> 80` is considered *strong*.

By using `test_pass`, you can detect strong passwords amongst fake.

Usage

```
pattern: "pass=(.*)"
test_pass:
  - on: 1
    score: 80.0 # scale: 0-100
```

Test positive

```
zxHELLOyw{_)
```

Test negative

```
foobar
```

test_token

Test for tokens, keys, and machine-generated secrets (own model).

By using `test_pass`, you can:

- Detect real tokens, keys, and secrets.
- Verify that a machine generated token is secret by model attributes.

Usage

```
    - pattern: "token=(.*)"
      pattern_type: multi
      test_token:
      - on: 1
        score: 0.6  # True if the score is bigger then 0.6
                    # max is 1, min is 0
```

Test positive

```
token=48SfRa4idxxUVyPAejafXxwjkreyj8MoJkjV
```

Test negative

```
token=AnotherVariableOfClientData[0];
```

test_entropy

A normalized entropy test. We do not recommend this test as entropy is a metric not optimized for finding secrets and sensitive information. You can use entropy if you use legacy infrastructure and policies.

Usage

```
    - pattern: "token=(.*)"
      pattern_type: multi
      test_entropy:
      - on: 1
        score: 4.0  # True if the entropy of the value is bigger
then 4
                    # max is 5, min is 0
```

Test positive

```
token=48SfRa4idxxUVyPAejafXxwjkreyj8MoJkjV
```

Test negative

```
token=G6q5oRa4idxxxxxxxxxxxxxwjkreyj8MoJkjV
token=FooBarFooBarFooBarFooBarFooBarFooBar
token=asdfdsafdsfasdfadsfadsfdasfasdfsdafsdf
```

Testing Detector

To test, you can selectively include your new detectors by using `--just-ids` and/or `--just-tags`. With these you can use any of the common Code Security commands:

If you want to run your new rule on your entire Github org:

```
$HOME/.spectral/spectral github ... --just-ids PRV001
```

Alternatively, just to scan your current repo:

```
$HOME/.spectral/spectral run ... --just-tags acme-security
```

Submit the Detector for Review

Contact *Check Point Support Center* to review your detector. Ensure to redact sensitive information in the detector before your submit it. Check Point can help you build it and give you a free detector building session.

Output Formats

You can view Code Security output in various formats.

CLI

For CLI-based output:

.spectral/spectral.yaml

```
reporter:
    outputs:
        stylish: {}  # produce CLI based reports
```

HTML

For the output in the HTML format for infosec reviews, secops reviews or sending your periodical security email:

.spectral/spectral.yaml

```
reporter:
    outputs:
        stylish: { html: "output.html" }  # produce HTML reports
```

CSV

For the CSV output, the configuration is for each repository, which means that after you run `spectral init` the configuration file `.spectral/spectral.yaml` is added to that repository.

Note - You can have one configuration file for multiple repositories and point with the `-c flag` to the global config file's location.

For output in the CSV format:

.spectral/spectral.yaml

```
reporter:
    outputs:
        stylish: { csv: "output.csv" }  # produce CSV reports
```

JSON

For output in the JSON format:

.spectral/spectral.yaml

```
reporter:
    outputs:
        stylish: { json: "output.json" }  # produce JSON reports
```

Log

You can use the log format (with timestamps and logfmt fields) to parse or push to log indexing services, such as Elastic.

In addition, you can switch to JSON log.

.spectral/spectral.yaml

```
reporter:
    outputs:
        log: {}
        # log: { json: true }    # use JSON logging
```

Junit (junit-xml)

The Junit format is suitable for interop with CI/CD products that accept `junit-xml` test result format.

Code Security generates a failing test that you can inspect in your CI dashboard like any other test.

.spectral/spectral.yaml

```
reporter:
    outputs:
        junit: {}
```

The XML results are generated in `junit-out`. To use, for example with CircleCI, point your CI to this folder:

.circleci/config.yml

```
version: 2
jobs:
  build:
    docker:
      - image: circleci/node:latest
    steps:
      - checkout
      - run:  $HOME/.spectral/spectral run
      - store_test_results:
          path: ./junit-out/
```

Reports

Reports show summaries, trends, and insights based on data that CloudGuard collects. Cloud administrators and corporate executives can use reports to understand their cloud environments and prioritize top issues. These are the ways to generate reports:

- Schedule CloudGuard to send reports automatically as PDF email attachments.
- Send reports manually from CloudGuard as PDF email attachments.
- Download reports from CloudGuard with your web browser.

Creating a report

1. From the navigation toolbar, click **Reports**.
2. In the upper right, click **+ Create**.

   A new window opens.
3. In the **Name** field, enter a unique name for the report.
4. Select **one** report template:
   - **Compliance Executive Report**

     Shows high-level compliance findings, trends, and summarized data .
   - **Compliance Detailed Report - Detailed per account**

     Shows detailed compliance findings, trends, and summarized data per account.
   - **Compliance Detailed Report - Detailed per ruleset**

     Shows detailed compliance findings, trends, and summarized data per ruleset.
   - **Intelligence (CDR) Executive Report**

     Shows high-level Intelligence findings, trends, and summarized data.
   - **Intelligence (CDR) Detailed Report - Detailed per account**

     Shows detailed Intelligence findings, trends, and summarized data per account.
5. In the **Scope** section, select kinds of data to send in the report. The **Scope** section is different for different reports. If you do not select kinds of data, the report includes all relevant data.
   Example- For the **Intelligence (CDR) Executive Report**, you can specify which **CDR Log Types** to send.

6. Select a **Time Range** for which to send reports. The Time Range includes the current Day/Week/Month.
   Example - If today is Thursday and the **Time Range** is **3 Days**, the report includes data for Tuesday, Wednesday, and Thursday.
   a. Select a unit:
      - **Days**
      - **Weeks**
      - **Months**
   b. Select a number.
7. Select filters for the scope of the report. If you do not select a filter for a field, the report includes all relevant information for that field.
   a. Select one or more **Severity** levels:
      - **Critical**
      - **High**
      - **Medium**
      - **Low**
   b. Select one or more **Cloud Providers**.
   c. Select one or more **Organizational Units**.
   d. Select one or more **Environments**.

      Note - If the report template is **Detailed per account**, you cannot leave this field blank.

   e. Select one or more **Rulesets**.

      Note - If the report template is **Detailed per ruleset**, you cannot leave this field blank.
8. Enter one or more email addresses to which to send the report. After you enter an email address, press one of these: Enter, Spacebar, Comma (**,**).
9. **Optional -** To schedule CloudGuard to send the report automatically:
   a. Select **Schedule Settings**.
   b. Select a frequency, day, and time to send the report.

Note - If you do not schedule CloudGuard to send the report automatically, CloudGuard generates the report only when you download the report or send the report manually. See *"Downloading a report" below* and *"Sending a report manually" below*.

10. Click **Save**.

Editing a report

1. From the navigation toolbar, click **Reports**.
2. Select **one** report.
3. In the upper right, click **Edit**.

   A new window opens.
4. Edit the report.
5. Click **Save**.

Sending a report manually

1. From the navigation toolbar, click **Reports**.
2. Select one or more reports.
3. In the upper right, click **Run**.

   CloudGuard emails the report to the specified email addresses. It can take up to a few minutes for the report to arrive.

Downloading a report

1. From the navigation toolbar, click **Reports**.
2. Select **one** report.
3. In the upper right, click **Download**.

   Your web browser downloads the report.

Deleting a report

After you delete a report, the report does not appear in the table and CloudGuard does not send it in the future. Past reports continue to be visible as attachments in emails that CloudGuard sent.

1. From the navigation toolbar, click **Reports**.

   A new window opens.
2. Select one or more reports.

3. Click **Delete**.

   A new window opens to ask if you are sure that you want to delete the selected report (s).

4. Click **Delete**.

Integration Hub

The CloudGuard Integrations central hub enables seamless integration with internal and external third-party applications, APIs, and services. Integrations enhance CloudGuard functionality to provide a unified security view of your cloud environments.

On the **All Integrations** page, you can create, search, edit, and delete CloudGuard integrations.

On the **Configured Integrations** page, you can view, search, edit, and delete configured CloudGuard integrations.

Editing or deleting an Integration

1. From the left menu, click **Integration Hub**.
2. Click on the icon for the integration.

   The sliding window opens.
3. Edit the integration, or click the delete (trash can) icon to delete it.

   Note - You cannot delete an integration that is currently in use.
4. Click **Save**.

Integrations that Can Be Configured on the Integrations Page

Events and Logging

- **Splunk** is a data collection, monitoring, and analysis system. You can configure CloudGuard to send Posture findings to it, from where they can be seen, searched, and analyzed. See Configure Splunk as a Log system for CloudGuard
- **IBM QRadar** is an enterprise security information and event management (SIEM) system. It collects log data from an enterprise, its network devices, host assets and operating systems, applications, vulnerabilities, and user activities and behaviors. See *"Sending Findings to QRadar" on page 800*
- **Sumo Logic** is a cloud-native, real-time, unified logs, and metrics analytic platform. You can configure CloudGuard to send Posture findings to Sumo with an HTTP endpoint. See Configure Sumo Logic as a Log system for CloudGuard

Ticketing Systems

- **ServiceNow** is a SaaS incident response system. You can configure CloudGuard to send alerts to ServiceNow, with a custom application, available in the ServiceNow Store. See *"Sending Alerts to ServiceNow" on page 816* and Use CloudGuard as a ServiceNow Provider
- **Jira** is a platform that combines issue collection and agile project management capabilities. You can configure CloudGuard to send Posture findings to Jira with an HTTP endpoint. See *"Sending Reports to Jira" on page 814*
- **PagerDuty** is a SaaS-based incident response system. You can configure CloudGuard to send Posture findings to PagerDuty, from where they can be managed as incidents. See How to configure PagerDuty with CloudGuard

Collaborations and Messaging

- Generic Webhook
- **AWS SNS**, which stands for Simple Notification Service, is a cloud-based web service that sends messages. You can configure CloudGuard to send its system events to an SNS target and then configure SNS to forward these messages to different destinations (emails included). See *"Sending System Notifications to AWS SNS" on page 806*
- **Microsoft Teams** is a business communication platform developed by Microsoft that offers workspace chat and videoconferencing, file storage, and application integration. You can configure CloudGuard to send summaries of Posture findings to Teams through a Notification that connects to Teams with an HTTP webhook. See *"Integration with Microsoft Teams" on page 817*
- **Sentra** is a Data Security Posture Management (DSPM) platform that classifies cloud assets by data sensitivity. See *"Classifying Assets with Sentra " on page 830*.
- **Slack** is a SaaS-based collaboration and messaging tool. You can configure CloudGuard to send summaries of Compliance findings to Slack, with a Notification that connects to Slack with an HTTP webhook. See How to configure CloudGuard to send events to Slack
- Email

Vulnerability Security Scanner

- **Tenable.io** provides information about vulnerabilities in cloud environments. It can be configured to provide this information to CloudGuard, where it can be seen on the Events page. See *"Configuring Tenable.io as a Provider for CloudGuard" on page 828*.

Cloud Services

- **Microsoft Defender for Cloud** - Microsoft solution for cloud security posture management (CSPM) and cloud workload protection (CWP) that finds weak spots across your cloud configuration, helps strengthen the overall security posture of your environment, and can protect workloads across multi-cloud and hybrid environments from evolving threats. You can configure CloudGuard to send Compliance findings to the Defender. See *"Sending Findings to Azure Defender for Cloud" on page 826*
- **GCP Security Command Center** - The Google Cloud Security Command Center is a GCP service for security management and data risk assessment. It aggregates information security issues and risks on your GCP resources and gives centralized visibility and control of your cloud data and services. You can configure CloudGuard to send Compliance findings to the Command Center. See Configure CloudGuard as a source for the Google Cloud Security Command Center (CSCC)
- **AWS Security Hub** - AWS Security Hub provides you with a comprehensive view of your security state in AWS and helps you assess your AWS environment against security industry standards and best practices.Security Hub collects security data across AWS accounts, AWS services, and supported third-party products and helps you analyze your security trends and identify the highest priority security issues. See *"Configuring CloudGuard as an AWS Security Hub Provider" on page 821*.

System Audit

- **SNS Audit** - see *"Sending System Notifications to AWS SNS" on page 806*.

Other Integrations

Data Sensitivity (DSPM)

- **Amazon Macie** - see *"Data Sensitivity" on page 113*.
- **Azure PureView** - see *"Data Sensitivity" on page 113*.

Cloud Services

- **Amazon Inspector** is an automated vulnerability management service that continually scans AWS workloads for software vulnerabilities and unintended network exposure. See *"AWS Policies and Permissions" on page 279*.
- **Amazon GuardDuty** is an Amazon threat-detection service that continuously monitors logs for signs of malicious activity, infected hosts, and unauthorized behavior in your AWS account. See *"Integrating Amazon GuardDuty Findings with CloudGuard" on page 823*

- **GCP Eventarc** is a Google Cloud Platform service that allows you to asynchronously send events from other Google services, SaaS, and your apps. See *"Sending Findings to Eventarc" on page 810*
- **Microsoft Defender for Cloud** - Microsoft solution for cloud security posture management (CSPM) and cloud workload protection (CWP) that finds weak spots across your cloud configuration, helps strengthen the overall security posture of your environment, and can protect workloads across multi-cloud and hybrid environments from evolving threats. You can configure CloudGuard to send Compliance findings to the Defender. See *"Sending Findings to Azure Defender for Cloud" on page 826*
- **GCP Security Command Center** - The Google Cloud Security Command Center is a GCP service for security management and data risk assessment. It aggregates information security issues and risks on your GCP resources and gives centralized visibility and control of your cloud data and services. You can configure CloudGuard to send Compliance findings to the Command Center. See Configure CloudGuard as a source for the Google Cloud Security Command Center (CSCC)

Platforms (Cloud providers)

- **AWS** - see *"Unified Onboarding of AWS Environments" on page 54*.
- **Azure** - see *"Onboarding an Azure Subscription" on page 61*.
- **GCP** - see *"Onboarding a Google Cloud Platform (GCP) Project and Google Workspace" on page 183*.
- **Alibaba Cloud** - see *"Onboarding Alibaba Cloud Accounts" on page 75*.
- **OCI** - see *"Onboarding Oracle Cloud Infrastructure Environments" on page 186*.
- **Kubernetes** - see *"Onboarding Kubernetes Clusters" on page 188*.
- **Container Registry** - see *"Onboarding Container Registries" on page 204*.
- **ShiftLeft** - see *"Create a ShiftLeft Environment and Service Account" on page 469*.

Sending Findings to QRadar

IBM QRadar is an enterprise Security Information and Event Management (SIEM) system. It collects log data from an enterprise and its network devices, host assets and operating systems, applications, vulnerabilities, and user activities and behaviors.

Configuring QRadar

1. From IBM App Exchange, download and install the Dome9 QRadar application on a QRadar console or app host.

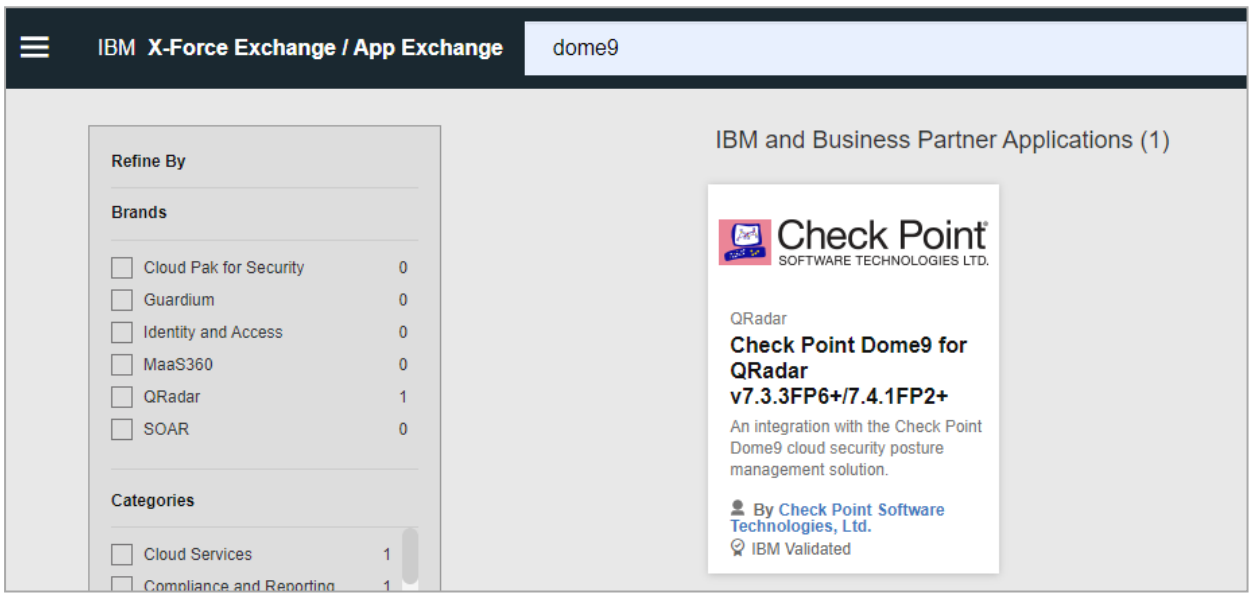

2. In the QRadar admin console, create a new QRadar role that only specifies access to the Dome9 application.

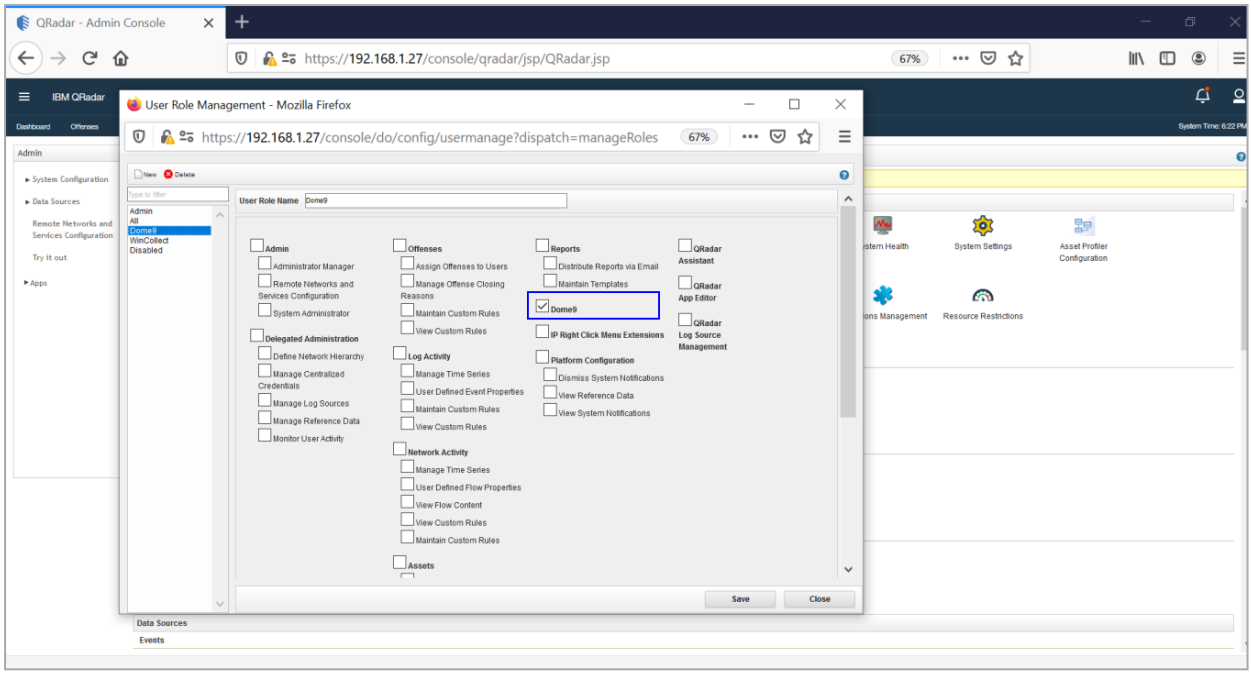

3. Create a new QRadar-authorized service that uses the role created in the previous step. Copy the Authentication Token for future use.

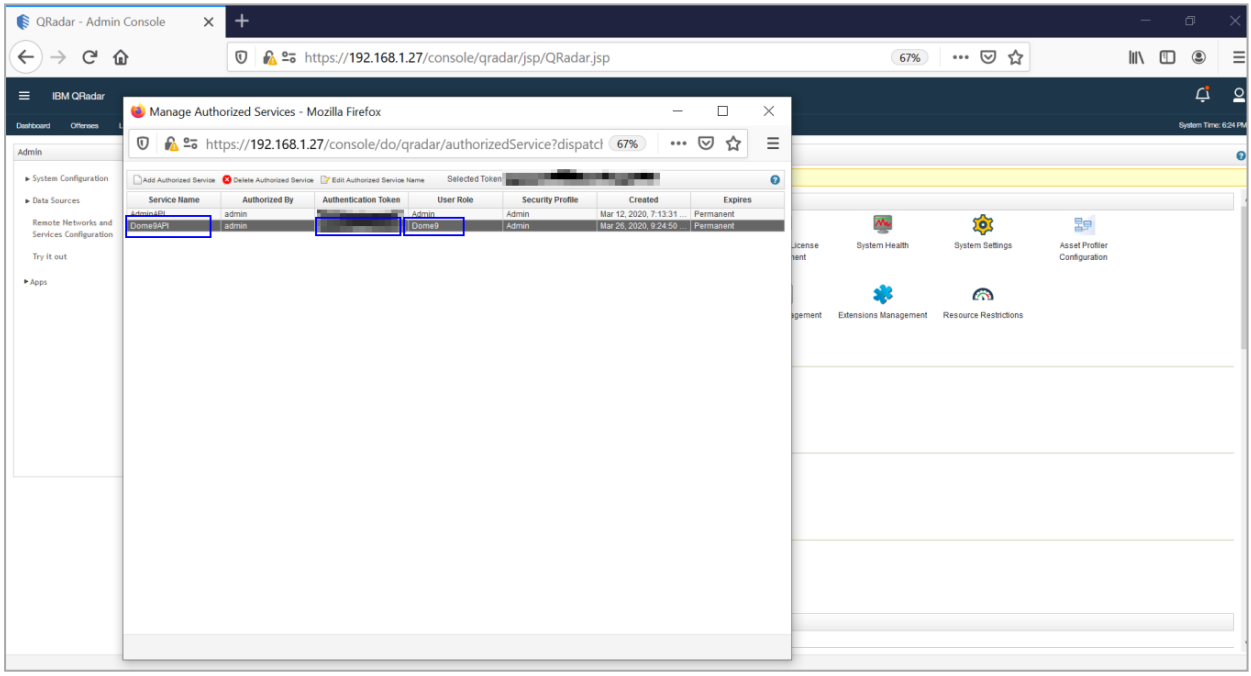

4. Below the **System Settings**, in the **Advanced** menu, set the **Max TCP Syslog Payload Length** value to **16,384**.

   If necessary, deploy the changes.

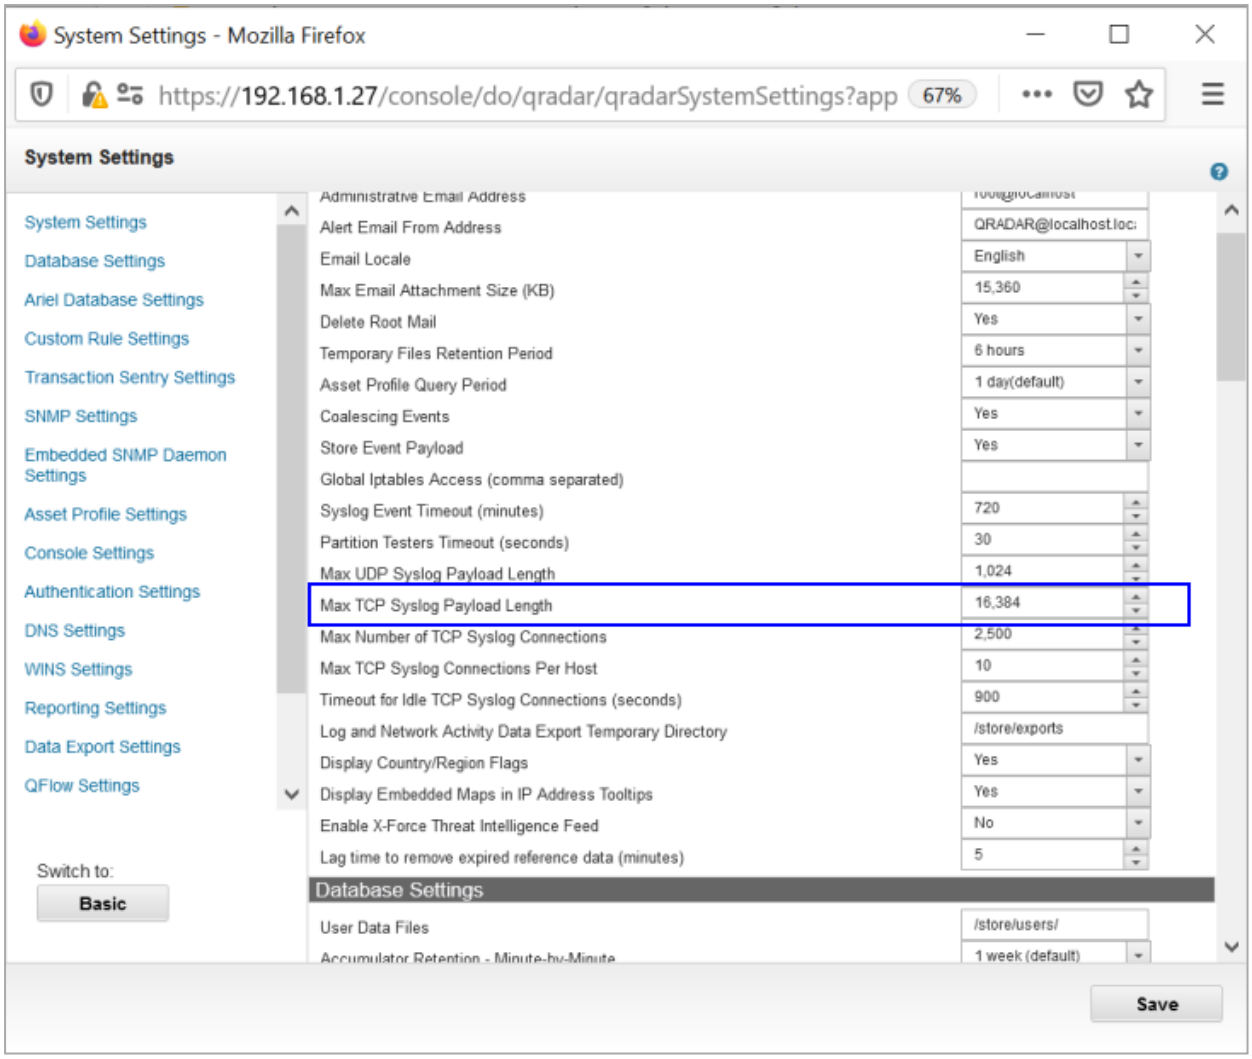

5. Create a new integration through the Dome9 Settings:
   - Copy the **Notifications HTTP Endpoint** value for future use.
   - (Optional) Provide the CloudGuard API credentials for the integration. With these credentials, you can acknowledge findings or create exclusions directly in QRadar.

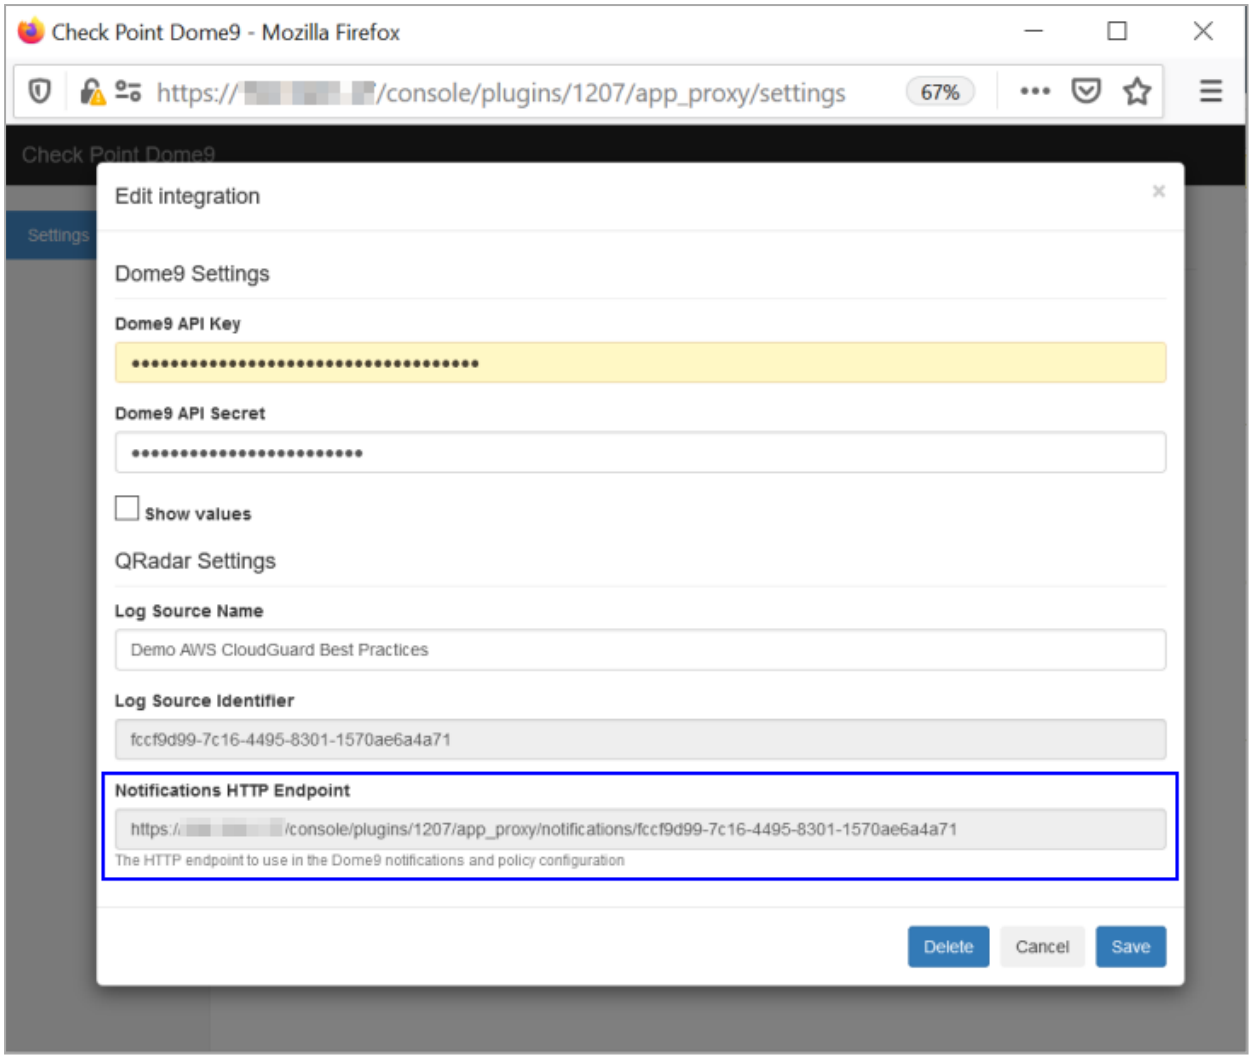

Configuring CloudGuard

1. In the CloudGuard portal, from the left menu, click **Integration Hub**
2. In the Events and Logging section, click **Qradar**.

   The **Qradar** sliding menu opens.
3. Create the integration.

Testing the Integration

1. In QRadar, make sure that Dome9 notifications show in the QRadar events database.

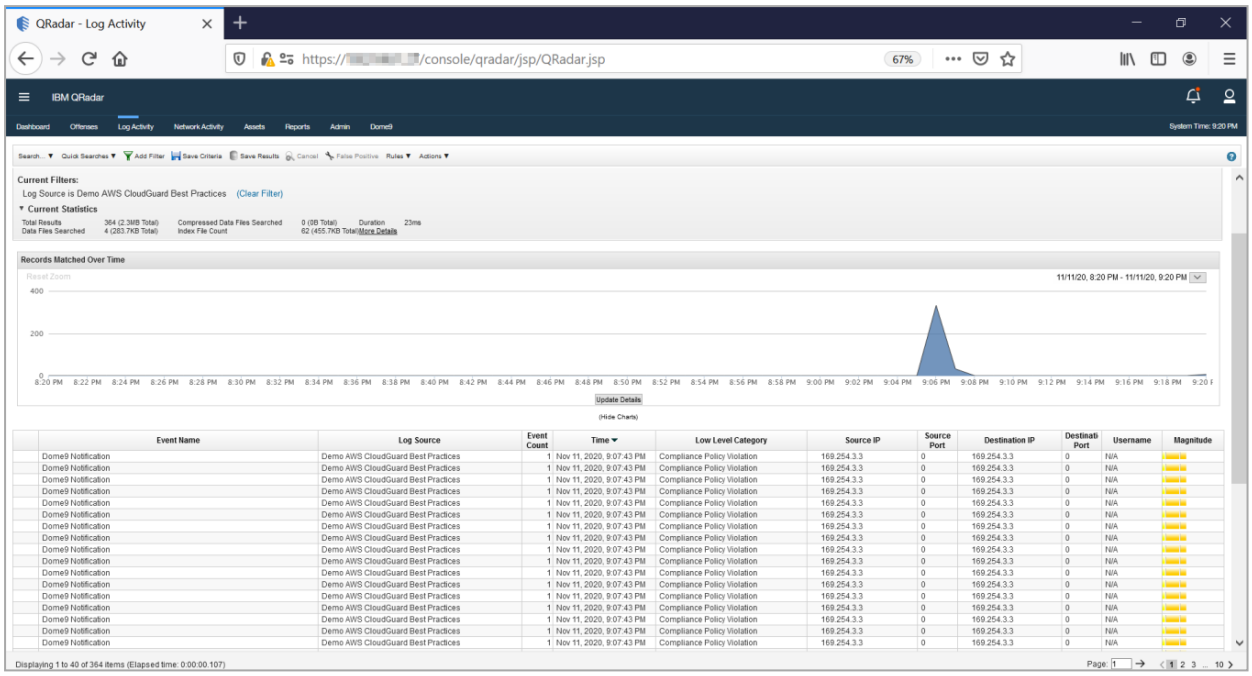

2. Make sure that custom properties are populated as expected in a sample event.

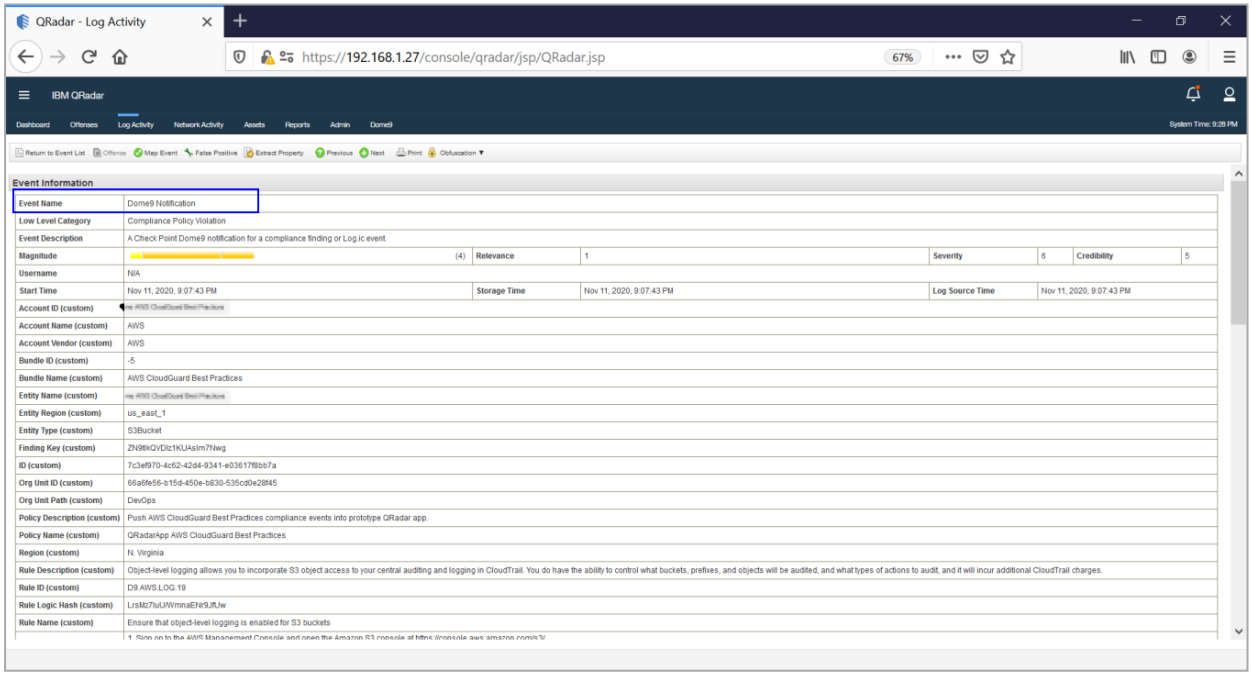

3. Browse events in the viewer in the Dome9 QRadar application.

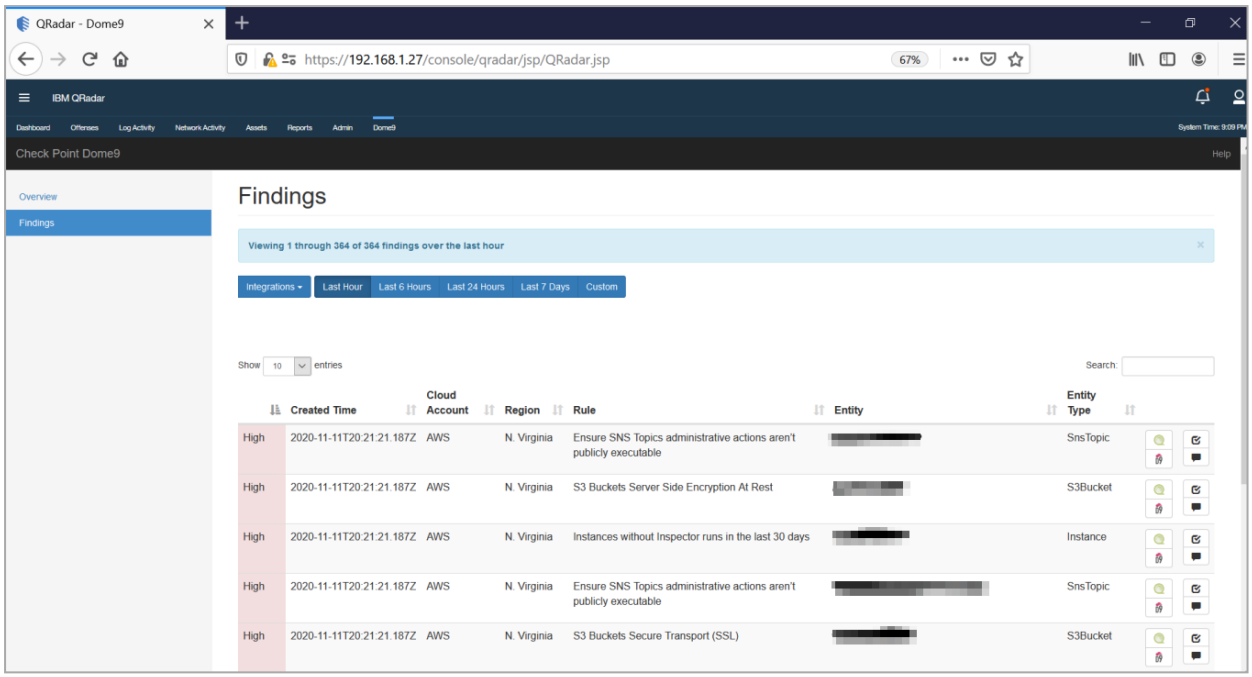

Sending System Notifications to AWS SNS

CloudGuard can send notifications on its system events and audit logs to your email or an SNS (Simple Notification Service) topic in your AWS account. If your SNS topic is configured with an SNS subscription, then you can configure CloudGuard to push notifications to your subscription's destination, such as Lambda, SMS, or email.

At AWS SNS, you receive notifications only for those events that you select in the settings for **Email Notifications**.

Connecting CloudGuard Events and AWS SNS

Step 1: In CloudGuard, enable SNS integration

1. In the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Cloud Services** section, click **SNS**.

   The **SNS** sliding menu opens.
3. Create the integration.

   Important - The environment number changes and is dependent on which data center your account belongs to.

Step 2: In AWS, create an SNS topic

1. In the AWS console, go to **Services** > **All services** > **Simple Notification Service**.
2. Select **Create topic**.
3. Below **Details**, select **Standard** and enter a name for the SNS topic, for example, *cloudguard-sns*.
4. Open the **Access policy** section and select:
   a. In **Choose method**, select **Basic**.
   b. In **Define who can publish messages to the topic**, select the option **Only the specified AWS accounts** and paste below the environment number copied in Step 1.
   c. In **Define who can subscribe to this topic**, select the option **Only the topic owner**.
5. Click **Create topic**.
6. In the **Details** window that opens, copy the **ARN**.

Step 3: Add an SNS subscription to the topic

After you create an AWS SNS topic, you must add subscriptions to integrate the information (notifications) to an endpoint. In AWS, navigate to the SNS page and select the SNS topic which you created.

1. In **Subscriptions**, click **Create subscription**.
2. Below **Protocol**, select a protocol, for example, **Email**.

   Enter the details for the endpoint that is to receive the subscription. For example, for an email, the endpoint is the email address.
3. Click **Create subscription**. The subscription status is set to *pending* until it is confirmed. For email subscriptions, when an email is sent to the endpoint address, it is necessary to confirm the email.

   Note - SNS is not sent to a subscription that is not confirmed.
4. Go back to the browser tab where CloudGuard is currently open.
5. Paste the topic ARN into the SNS configuration text box.
6. Click **Save**.

   In CloudGuard, SNS integration is **Enabled**.

Step 4: Create a new KMS key

1. In the AWS Management Console, go to **Services** > **Security, Identity, & Compliance** > **Key Management Service**.
2. Click **Create a key**.
3. On the **Configure key** page:
   a. In the **Key type** section, select **Symmetric**.
   b. In the **Key usage** section, select **Encrypt and decrypt**.
   c. If the KMS and CloudGuard are in different regions, in the **Advanced options** section select **Multi-Region key**.
   d. Click **Next**.
4. On the **Add labels** page:
   a. In the **Alias** section, enter an alias for the key.
   b. Click **Next**.

5. On the **Define key administrative permissions** page, configure administrative permissions for the key.
6. On the **Define key usage permissions** page:
    a. In the **Key users** section, select IAM users and rules to allow to use the key.
    b. In the **Other AWS accounts** section, click **Add another AWS account** and paste the environment number for CloudGuard that you copied in Step 1.
    c. Click **Next**.
7. On the **Review** page, click **Finish**.

    The **Customer managed keys** screen opens and the new key appears in the table.

Step 5: Associate the KMS Key with your SNS topic

1. In the AWS Management Console, go to **Services** > **Application Integration** > **Simple Notification Service** > **Topics**.
2. Click the SNS topic you created for CloudGuard.

    The topic page opens.
3. Click **Edit**.

    The **Edit topic** page opens.
4. In the **Encryption** section:
    a. Toggle the **Encryption** button to the "On" position.
    b. In the **AWS KMS key** field, select your KMS key.
5. Click **Save changes**.

Integration of Findings Notification

Configure the SNS topic to send single findings (not reports).

To configure SNS integration for single findings:

1. In the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Collaborations and Messaging** section, click **SNS**.

    The **SNS** sliding menu opens.
3. Create the integration.

More Links

- *"System Audit Logs" on page 134*
- *"Email Notifications" on page 860*

Sending Findings to Eventarc

Google Cloud Platform (GCP) Eventarc is a service that allows you to asynchronously deliver events from different event sources to different event consumers. When integrated with Eventarc, CloudGuard sends events to Eventarc, with the API interface.

In this integration, it is necessary to create an Eventarc Trigger that operates with other Google Cloud Platform components. Then you connect the Trigger with a Channel, which is a GCP resource in your project that represents the source of events from CloudGuard. Creating and activating a Channel serves as your explicit request to receive events from CloudGuard at Eventarc. Triggers filter and route events from a specific Channel.

The diagram below shows how to set a connection between CloudGuard and Eventarc:

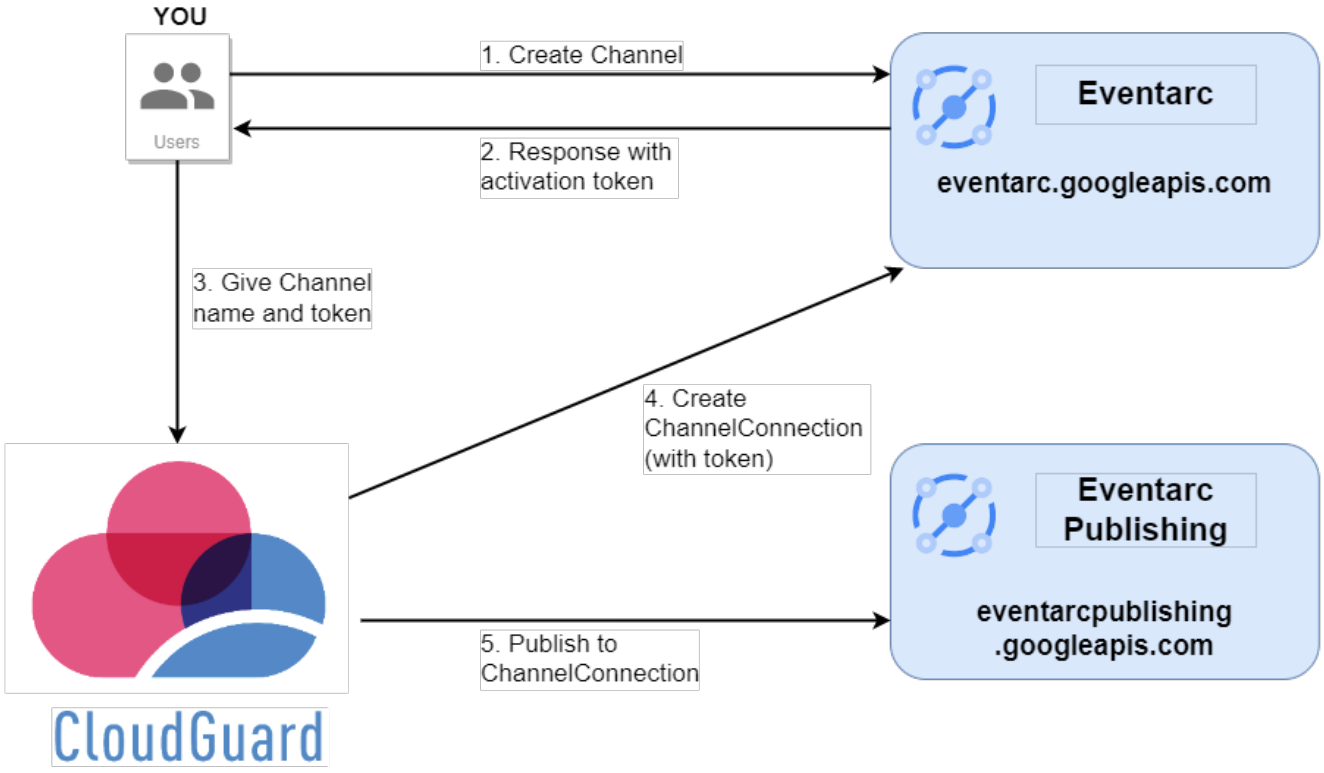

Steps 1 and 2 - Creating a Channel

Create a channel in the Google Cloud Platform. Eventarc responds with the channel name and activation token.

1. Log in to the Google console.
2. In the search bar, type **Eventarc** and select it from the list of Products & Pages to open the service.
3. In Eventarc, click **Create Trigger**.
4. In the Create Trigger window, set the required details:

a. **Trigger name** - Select a name for the trigger, for example, Events-from-CloudGuard.

b. **Event provider** - Start to enter **Check Point CloudGuard** and click to select it from the list.

c. **Channel** - Select an active channel or click **Create Channel** to create a new channel (Step 1 on the diagram). To learn how to create a channel, see Google documentation.

   After you create the channel, Eventarc shows this information that you need to save for the CloudGuard API call:

   - **Channel full name** - The channel name which includes the Google Cloud project ID and location of the resource. For example: `projects/your-project-123/locations/us-central1/channels/your-channel`
   - **Activation token** - The token is valid for 24 hours after the channel creation. After 24 hours, the channel becomes *Inactive*.

   When you click **Done**, the new channel appears from the list of available channels with the *Pending* status.

   Important - Make sure that the channel status changes to *Active* when the provider (CloudGuard) activates the channel with the channel name and activation token (Step 4 on the diagram).

d. **Event** - Select `cloudguard.v1.event`.

e. **Region** - Select one of the available regions.

   Note - The region selected for the channel must be the same region that you use in CloudGuard with this GCP project.

Steps 3 and 4 - Sending the Channel Name and Token to CloudGuard

In CloudGuard, use an API call to send to CloudGuard the channel name and token received from EventArc. With this information, CloudGuard creates a ChannelConnection (activates the channel) and sends you an identifier for the notification.

Request

```
POST
/v2/Compliance/ContinuousComplianceNotification/eventarcChannelConnection
```

```
{
  "channelFullName": "string",
```

```
  "activationToken": "string"
}
```

For API documentation and code examples, see API reference guide.

Authorization

Basic Authorization: Use the API key and secret as username and password respectively.

Parameters

- **channelFullName** - Full name of the channel created in the Google Cloud Platform
- **activationToken** – Activation token generated by Google Cloud Platform

Response

```
200 – OK
```

```
"string"
```

Use the response string in Step 5 for **channelConnectionId**.

Important - The channel status changes to *Active* when CloudGuard successfully activates the channel with the channel name and activation token.

Step 5 - Creating a Notification with Eventarc Target

Use an API call to create a notification, which allows CloudGuard to publish events to ChannelConnection.

Request

```
POST /v2/Compliance/ContinuousComplianceNotification
```

```
{
     "changeDetection": {
          "eventarcData": {
               "channelConnectionId": "string"
          },
          "eventarcIntegrationState": "Enabled"
     }
}
```

For API documentation and code examples, see the API reference guide.

Authorization

Basic Authorization: Use the API key and secret as username and password, respectively.

Parameters

- **"changeDetection": "eventarcIntegrationState":** "Disabled" - Set the status to "Enabled"
- **"eventarcData": "channelConnectionId"**: "string" - Use the string from the Step 4 response

Response

```
200 - OK
```

Validating the Integration

When you complete these steps, CloudGuard starts to send the events to Eventarc.

On Eventarc, you can see the graph of the trigger invocations when you navigate to the trigger's Details page.

Sending Reports to Jira

Jira is a platform that combines issue collection and agile project management capabilities. You can configure CloudGuard to send findings to Jira with an HTTP endpoint.

Step 1: Configure Jira

1. In Jira, create an account.
2. For this account, create a token for API requests.

Best Practice - Check Point recommends to create a new Jira account to securely use its credentials in the configuration below.

Step 2: In CloudGuard, create a Jira integration

1. In the CloudGuard portal, from the left menu, select **Integration Hub**.
2. In the **Ticketing Systems** section, click **Jira**.
3. In the **Jira** sliding menu, click **Add** to create a new integration.
4. Enter these details:
   - **Name** to identify the integration
   - **Endpoint URL** - your Jira domain URL
   - **Username** - email of your Jira account
   - **Password** - API token created in Step 1
5. Select the option **Ignore certificate validation** if you work with self-signed certificates. This state is typical only for development and integration environments and is not recommended for production environments.
6. Click **Save**.

Step 3: In CloudGuard, create a notification for the Jira integration

1. Create or edit a notification. For more information, see *"Notifications" on page 852*.
2. In the **Immediate Notification** section, select **HTTP endpoint notification per newly created finding**.
3. Select **Jira**.
4. From the list, select a Jira configuration.
5. Configure the Json Payload.

The section below the Jira configuration contains a JSON template that represents your ticket template in Jira. On the left side, use the expressions from the list on the right. These expressions are replaced with the alert data after evaluation of the cloud entity.

You can fill the built-in Jira fields with values when you add them to the JSON payload (in the API example, *Labels* takes an array of values). Custom Jira fields are not accessed with their names directly, but through the format “customfield_#####”. For more information, see the REST API example.

Important - You must configure:

- in the *project* object, the project name for the *key* parameter
- in the *issuetype* object, the ticket type name for the *name* parameter

(**WJIT** and **Bug** in the example below, respectively).

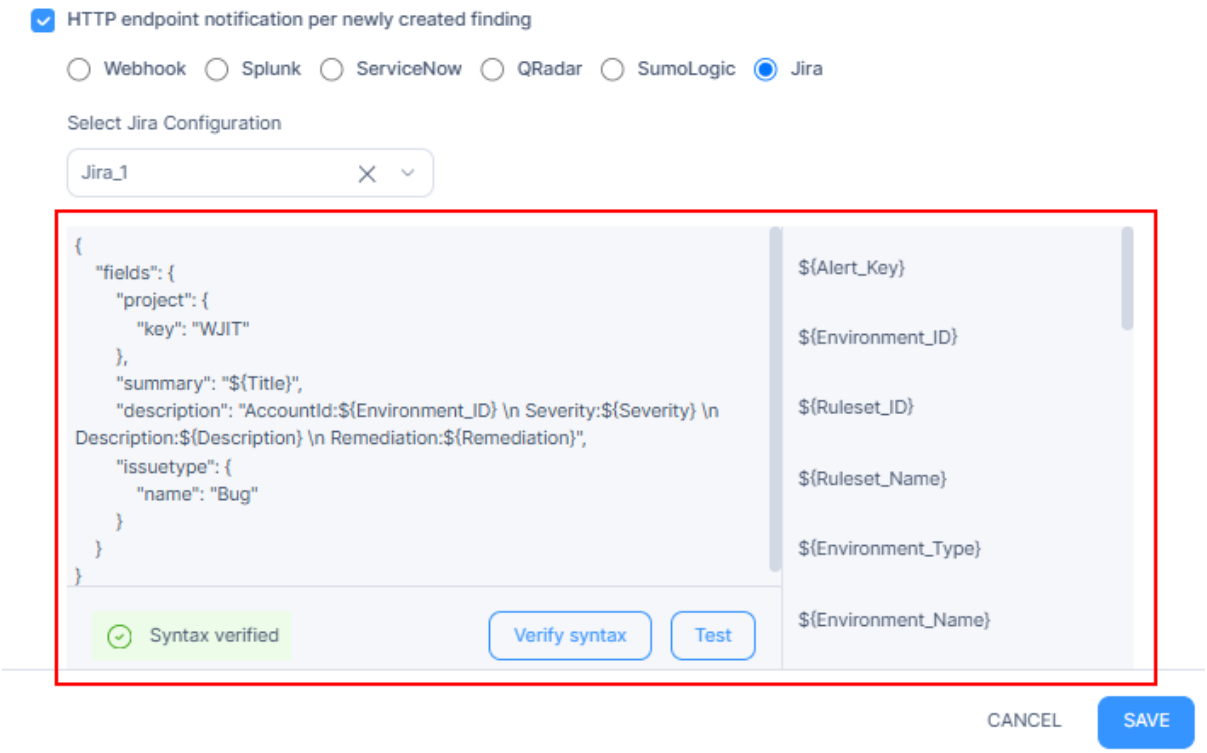

6. Click **Verify syntax** to verify the integration.
7. Click **Test** to test the integration.
8. Finish creating the notification.
9. Add the notification to a policy. For more information, see *"Configuring CloudGuard Policies" on page 78*.

Sending Alerts to ServiceNow

ServiceNow is a platform for managing tickets, incidents, and organizational flows. CloudGuard sends CSPM findings and container vulnerability data to ServiceNow, which records them as new ServiceNow incidents. These can be managed as tickets in ServiceNow and, when resolved, are cleared from CloudGuard in the next assessment.

This page describes the CloudGuard configuration only.

To integrate CloudGuard with ServiceNow:

1. In your ServiceNow account, install the CloudGuard CNAPP from the ServiceNow store and follow the instructions in the Installation Guide (click to download).
2. In the CloudGuard portal, from the left menu, click **Integration Hub**.
3. In the **Ticketing Systems** section, click **Service Now**.

   The **Service Now** sliding menu opens.
4. Create the configuration.

To send a report from CloudGuard to ServiceNow:

1. In CloudGuard, from the left menu go to **Settings** > **Configuration** > **Notifications** and click **Add**.
2. Enter the applicable options as described in *"Notifications" on page 852*.
3. In the **Immediate Notification** section, select **HTTP endpoint notification per newly created finding**.
4. Select **ServiceNow**.
5. Select the relevant ServiceNow configuration.
6. Click **Save**.

Integration with Microsoft Teams

You can configure CloudGuard to send notifications to Microsoft Teams. The integration uses a webhook that you create in Microsoft Teams. These CloudGuard features can send notifications to Microsoft Teams:

- Cloud Security Posture Management (CSPM)
- Cloud Detection and Response (CDR)
- *"Admission Control" on page 476*
- Runtime Protection
- *"Toxic Combinations" on page 92* - The Toxic Combinations feature sends its own notifications. After you create a Microsoft Teams integration, you can configure CloudGuard to send notifications to a Microsoft Teams channel when it detects a Toxic Combination. See *"Toxic Combinations" on page 92* and *"Action Hub" on page 93*.

Important - Microsoft will deprecate Office 365 Webhook Connectors on January 31st, 2025. To continue using an integration of Microsoft Teams with CloudGuard, you must create a new webhook workflow in Microsoft Teams. For more information, see Microsoft documentation.

Configuration

Step 1: In Microsoft Teams, create a webhook URL for CloudGuard

Note - If you are using Microsoft Teams Classic, click on **Apps** in the sidebar, search for Workflows, and then add Workflows. Click the **Post to a channel when a webhook request is received** workflow. Continue the instructions below in Step 1 > **5**.

1. Open Microsoft Teams.
2. For the relevant Microsoft Teams channel, click ··· and select **Workflows**.
3. In the **Notify a team** section, click **Post to a channel when a webhook request is received**.
4. Add a name for the new workflow.
5. Click **Next**.
6. From the **Microsoft Teams Team** list, select your team.
7. From the **Microsoft Teams Channel** list, select your channel.
8. Click **Add workflow**.
9. In the **Workflow added successfully** field, click  to copy the webhook URL. Keep this URL in a safe place.
10. Click **Done**.

Step 2: In CloudGuard, create a Microsoft Teams integration

1. From the left menu, select **Integration Hub**.
2. In the top right corner, select **All Integrations**.
3. In the **Collaborations and Messaging** section, click **Teams**.

   The **Teams** sliding window opens.
4. Click **Add**.
5. Enter a name for the integration.
6. In the **Teams webhook URL** field, paste the webhook URL you copied from Microsoft Teams in Step 1 > **9**.

Step 3: In CloudGuard, configure a notification to send to the Microsoft Teams integration

1. When you create or edit a notification, select one or more of these configurations::
   - **CSPM - Summary report to Teams channel**
   - **Send critical security events to Teams channels (CDR, Admission control and Runtime protection only)**
2. Add the notification to an applicable policy. For example, add the **CSPM - Summary report to Teams channel** notification to a CSPM policy.

For more information, see *"Notifications" on page 852*.

Step 4: Test the Integration

In CloudGuard, manually send the notification. In Microsoft Teams, check if the notification appears as expected. For more information, see *"Notifications" on page 852*.

Troubleshooting

If a test of the integration fails, follow this procedure to resolve a rare issue in Microsoft Teams

To troubleshoot the integration between CloudGuard and Microsoft Teams

1. Log in to your Microsoft Power Apps.
2. Click **Flows** and select the flow that you created for CloudGuard.
3. Click **Edit**.
4. Select **Send each adaptive card**.
5. From the **Select an output from previous steps** list, select **Post card in a chat or channel**.
6. From the **Post as** list, select **User**.
7. Click **Save**.
8. Test the integration between CloudGuard and Microsoft Teams.

Sending Security Events to Microsoft Sentinel

You can configure Microsoft Sentinel to pull security events from a CloudGuard service account and show them in Microsoft Sentinel.

Microsoft Sentinel Requirements

- You must have Security Administrator permissions.
- You must have Owner or Contributor role permissions in the Log Analytics workplace.

CloudGuard Requirement

- There must be one or more CloudGuard policies configured for Security Events. For more information, see *"Intelligence Security Events" on page 634*.

To configure Microsoft Sentinel to pull security events from CloudGuard:

Important - Keep Microsoft Sentinel and CloudGuard open during this entire procedure.

1. In CloudGuard, from the left menu, expand **Integration Hub** and click **Integrations**.
2. In the **Events and Logging** section, click **Microsoft Sentinel**.

   In CloudGuard, the **Microsoft Sentinel** configuration window opens.
3. Click **Add**.
4. Follow the steps shown in the **Microsoft Sentinel** configuration window to connect a CloudGuard service account to Microsoft Sentinel. For more information about CloudGuard service accounts, see *"Service Accounts" on page 843*.
5. Click **Save**.

Configuring CloudGuard as an AWS Security Hub Provider

For Continuous Posture assessments (only), configure CloudGuard to send alerts to the AWS Security Hub.

To receive CloudGuard notifications on the Security Hub, you must onboard your AWS account to CloudGuard. See *"Unified Onboarding of AWS Environments" on page 54*. If you have already onboarded your AWS account, continue with the instructions.

To configure an AWS IAM policy for CloudGuard:

1. In the AWS console, navigate to the **IAM** dashboard and select **Roles**.
2. Select the *CloudGuard-Connect* role.
3. In **Permissions**, click **Add permissions** > **Attach policies**.
4. On the **Add permissions** page, click **Create policy**.
5. In **Create policy**, select **JSON**.
6. In the editor, paste this policy:

```
{
                    "Version": "2012-10-17",
                    "Statement": [
                    {
                    "Sid": "VisualEditor0",
                    "Effect": "Allow",
                    "Action": [
                    "securityhub:UpdateFindings",
                    "securityhub:BatchImportFindings"
                    ],
                    "Resource": "*"
                    }
                    ]
                }
```

7. Optionally, add tags.
8. Enter the policy name and click **Create policy** to save it.

To subscribe to the CloudGuard Integration in the Security Hub:

1. In the AWS Security Hub, navigate to **Integrations**.
2. In the search bar, enter **Check Point: CloudGuard Posture Management card** and click **Accept**.
3. In the confirmation window, click **Accept finding**.

   The status of the integration changes to **Accepting findings**.

To configure CloudGuard to send notifications to AWS Security Hub:

1. In the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Cloud Services** section, click **AWS Security Hub**.

   The **AWS Security Hub** sliding menu opens.
3. Create the integration.

Configure Multiple AWS Accounts to One Security Hub

You can associate other AWS accounts to one (master) account, to see event notifications for all of them on the Security Hub dashboard of the master account. This is done on the AWS Security Hub console page.

To configure multiple AWS accounts to one Security Hub:

1. The corresponding accounts from which it is necessary to see CloudGuard events must be onboarded to CloudGuard (if they are not, follow the instructions here).
2. The corresponding accounts must be linked to the master account in AWS (in the Security Hub console).
3. Create a CloudGuard Continuous Posture Notification that directs findings to the master account in the AWS Security Hub. Afterward, apply this policy to each of the accounts, which include the master account (see Configure a Notification on CloudGuard above).

More Links

- *"Continuous Posture" on page 315*
- *"Unified Onboarding of AWS Environments" on page 54*

Integrating Amazon GuardDuty Findings with CloudGuard

Amazon GuardDuty is an Amazon threat-detection service that continuously monitors logs for signs of malicious activity, infected hosts, and unauthorized behavior in your AWS account. To further streamline your security operations, you can integrate Amazon GuardDuty with CloudGuard. This integration enables your security team to access all AWS findings from a single dashboard, which makes it easier to manage and prioritize alerts. CloudGuard can provide more security measures, such as threat intelligence and automated incident response, to help mitigate any detected threats.

Benefits

- Provide a single-security view of your AWS environment for threats and security events.
- Enrichment for findings.
- Improve workflow - Manage events such as Acknowledge, Comment, and Archive similar to your other CloudGuard findings, see *"Action Menu" on page 121*.

Prerequisites

- Onboard your AWS account to CloudGuard, see *"Onboarding AWS Environments" on page 144*.
- In your AWS account, configure GuardDuty to store log files on an S3 bucket.
- Configure GuardDuty to export findings to an S3 bucket and give the necessary permissions (KMS), see the AWS GuardDuty User Guide.

How it Works

When AWS logs a GuardDuty finding, GuardDuty forwards the event to a region-specific S3 bucket. CloudGuard's CFT SNS topic then forwards the findings to CloudGuard Events.

Based on your network configuration and security requirements, you can configure an S3 bucket for each AWS account or configure one centralized S3 bucket to manage multiple AWS accounts.

S3 bucket for each account:

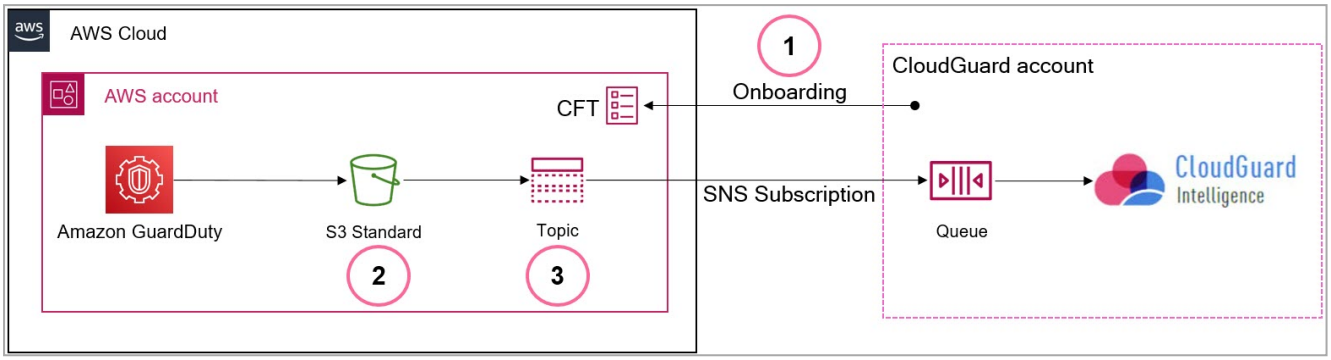

Item	Description
1	Use a CFT to onboard your AWS account to CloudGuard.
2	Create an S3 topic to send your GuardDuty findings to.
3	Configure an SNS Topic to send data from the S3 bucket to an SQS queue in CloudGuard.

Centralized S3 bucket:

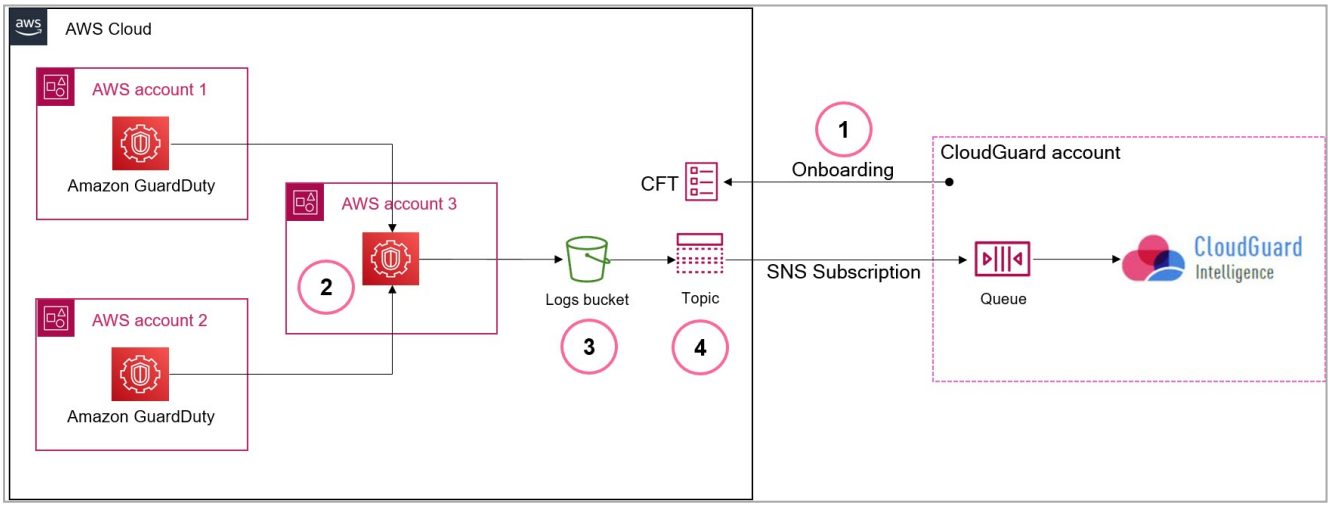

Item	Description
1	Use a CFT to onboard your AWS account to CloudGuard.
2	Configure GuardDuty to send its findings to a centralized GuardDuty.
3	Set up an S3 topic to send your GuardDuty findings.

Item	Description
4	Configure an SNS topic to send data from the S3 bucket to an SQS queue in CloudGuard.

Onboarding GuardDuty to CloudGuard

To onboard GuardDuty to Intelligence:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. In the table, below the **GuardDuty** column select **Enable GuardDuty**.

   Or,

   From the same table, select a specific environment. In the environment page that opens, select **Add GuardDuty**.
3. Follow the instructions in the onboarding wizard.
4. Click **Next**.
5. When the message "*Onboarding is completed successfully*" shows, click **Finish**.

To verify that GuardDuty is onboarded, in CloudGuard go to **Assets** > **Environments** and make sure that a checkmark shows below the GuardDuty column for the applicable account name.

To see GuardDuty findings, filter the Threat & Security Events table:

1. In the CloudGuard portal, navigate to **Events** > **Threat & Security Events**.
2. In the filter bar, click **Add Filter** and select **Source**.
3. Click **Source** and select **Amazon GuardDuty**.

   The event view shows all events filtered by GuardDuty as the source. The initial view of events takes approximately one hour from the actual onboarding.

To remove Amazon GuardDuty:

1. In the CloudGuard portal, navigate to **Assets** > **Environments**.
2. From the menu bar, select **Remove GuardDuty**.
3. In the window that opens, click **Remove**.

Sending Findings to Azure Defender for Cloud

You can configure CloudGuard to send findings on your Azure environments to theAzure Defender for Cloud. This allows you to see compliance issues for your Azure environments onboarded to CloudGuard on the Defender dashboard.

First, you must onboard your Azure account to CloudGuard. For more information, see *"Onboarding Azure Subscriptions" on page 169*. Second, set up a policy to assess the Azure subscription and include a notification to send findings to the Microsoft Defender for Cloud. In addition, you must configure your Azure subscription to accept findings from CloudGuard.

Step 1 - Configuring Azure Subscription Permissions

To receive CloudGuard findings, add more permissions to your subscription.

1. Log in to the Azure management portal.
2. Select the subscription you onboarded to CloudGuard.
3. Assign the **Security Admin** role to the application created during onboarding (*CloudGuard-Connect*).

Step 2 - Configuring CloudGuard to Send Notifications to Azure Defender for Cloud

1. IIn the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Cloud Services** section, click **Azure Defender For Cloud**.

   The **Azure Defender For Cloud** sliding menu opens.
3. Create the integration.
4. From the left menu, click **Settings** > **Configuration** > **Notifications**.
5. Click **Add**.
6. In the **Security Management Systems** section, select **Findings to Microsoft Defender for Cloud**.
7. Below **Select AzureDefender Configuration**, select the relevant integration.
8. Finish creating the notification.

Step 3 - Configuring CloudGuard Policy

From the left menu, go to **CSPM** > **Continuous Posture** and set up a Continuous Posture Policy with the Notification created in Step 2.

When CloudGuard runs an assessment for the environment, new findings in your Azure subscription are seen on the dashboard of Microsoft Defender for Cloud.

More Links:

- *"Continuous Posture" on page 315*
- How to set up Microsoft Defender for Cloud

Configuring Tenable.io as a Provider for CloudGuard

This section describes how to configure Tenable.io as a provider for CloudGuard. When configured, Tenable.io sends events to CloudGuard, which then shows CloudGuard's Events page.

Important - CloudGuard supports findings only for AWS EC2. Currently, there is no support for GCP or Azure.

Configuring Tenable.io to Send Events

Tenable.io integration allows findings in Tenable.io to be synced into CloudGuard, as long as the asset corresponding with the finding in Tenable.io exists in CloudGuard.

To send Tenable.io alerts to CloudGuard:

1. From your Tenable account, navigate to **Settings** > **Users**.
2. Create a new Tenable user, with the role **Administrator**.
3. Select the **New user**. The **New user** window opens.
4. Select **API keys** > click **Generate**.
5. Copy the **API Access Key** and **Secret Key**.
6. In CloudGuard, from the left menu go to **Settings** > **Configuration** > **Integrations**.
7. In the **Vulnerability Security Scanner** section, click **Tenable**.

   The **Tenable** sliding window opens.
8. Create the configuration.

Viewing Tenable.io Events

When your Tenable.io account is configured to send events to CloudGuard, the events show on the CloudGuard **Threat & Security Events** page. Only events for entities that are part of an environment that is onboarded to CloudGuard show.

To see Tenable.io events in CloudGuard:

1. In CloudGuard, navigate to **Events** > **Threat & Security Events**.
2. In the Filter, select **Source > Tenable.io**. If it does not show as an option, then it is not configured correctly. Make sure you did the configuration steps correctly.

The filtered list of events shows events from Tenable.io. To see more details, expand the event.

Building Rules and Queries Based on Tenable.io Findings

You can build CloudGuard Posture Management rules with conditions based on findings received from Tenable.io.

Example GSL rule that checks for instances in an environment for which external findings are sourced from Tenable.io.

```
Instance should not have externalFindings.findings with [
findingSource='Tenable.io']
```

Classifying Assets with Sentra

Sentra is a Data Security Posture Management (DSPM) platform that classifies cloud assets by data sensitivity. For example, Sentra identifies that a cloud asset contains financial information. You can forward classifications of cloud assets from Sentra to CloudGuard. The information from Sentra appears in the **Data Classification** and **Data Sensitivity** columns of the table in **Risk Management** > **Protected Assets**. For more information about data sensitivity classification, see *"Data Sensitivity" on page 113*.

To forward cloud asset classifications from Sentra to CloudGuard:

1. In the Sentra portal, create an API key with a **Viewer** role.
2. In the Sentra portal, copy the API key. Keep it in a safe place.
3. In the CloudGuard portal, from the left menu, click **Integration Hub**.
4. In the **Data Sensitivity (DSPM)** section, click **Sentra**.

   The **Sentra** window opens.
5. Click **Add**.
6. Enter a name for the integration.
7. Paste the API key you copied from Sentra.
8. Click **Save**.

To stop forwarding cloud asset classifications from Sentra to CloudGuard:

1. In the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Data Sensitivity (DSPM)** section, click **Sentra**.

   The **Sentra** window opens.
3. Click the trash can icon next to the name of the Sentra integration you want to delete.
4. Confirm in the confirmation window.

Classifying Assets with Cyera

Cyera is a Data Security Posture Management (DSPM) platform that classifies cloud assets by data sensitivity. For example, Cyera identifies that a cloud asset contains financial information. You can forward classifications of cloud assets from Cyera to CloudGuard. The information from Cyera appears in the **Data Classification** and **Data Sensitivity** columns of the table in **Risk Management** > **Protected Assets**. For more information about data sensitivity classification, see *"Data Sensitivity" on page 113*.

To forward cloud asset classifications from Cyera to CloudGuard

1. In the Cyera portal, create an API key with a **Viewer** role.
2. In the Cyera portal, copy the API key. Keep it in a safe place.
3. In the CloudGuard portal, from the left menu, click **Integration Hub**.
4. In the **Data Sensitivity (DSPM)** section, click **Cyera**.

   The **Cyera** window opens.
5. Click **Add**.
6. Enter a name for the integration.
7. Paste the API key you copied from Cyera.
8. Click **Save**.

To stop forwarding cloud asset classifications from Cyera to CloudGuard

1. In the CloudGuard portal, from the left menu, click **Integration Hub**.
2. In the **Data Sensitivity (DSPM)** section, click **Cyera**.

   The **Cyera** window opens.
3. Click the trash can icon next to the name of the Cyera integration you want to delete.
4. Confirm in the confirmation window.

Settings

Use Settings to change settings for your CloudGuard account, such as users and roles, and configure different notifications.

View and Change your CloudGuard Settings

Configure the following items for your CloudGuard account:

- Enable or disable **cross-account access**
- Set **credentials** for the Dome9 account: password, API key (for developers with the CloudGuard API), MFA for login
- Manage **mobile devices** connected to your Dome9 account with the CloudGuard mobile app
- Set **notifications** (by email) for different events and conditions
- Enable or disable **single sign-on** (SSO) to your Dome9 account (see *"Single Sign-On" on page 866*)
- Configure **Access Leases** - default lease time for connections
- Enable or disable the integration with AWS SNS (notifications) or Tenable.io

Account Info

See use statistics for your CloudGuard Account and your Account plan.

- Your CloudGuard **plan details** - The modules included. .
- Location of the **Data Center** (region) you are currently logged in to.
- Usage **statistics** (for instances) - The number of billable instances for each day (during 1, 3, or 6 months).
- Number of **users** on your Account.

License Activation

You can add to your CloudGuard Account more capabilities when you purchase more licenses. After the purchase, you must activate the new license.

Infinity Portal

You can manage the licenses on the **Services & Contracts** page of your Infinity Portal Account Settings (see Infinity Portal Administration Guide > Account Settings > Services & Contracts > Associated Accounts).

Note - You can associate your contract with only one Infinity Portal account. To transfer a contract between accounts, it is necessary to stop the association with the initial account.

Dome9 Portal

To activate the license for Dome9 accounts only:

1. Log in to the CloudGuard portal as the Account Owner. For more information on the Account Owner, see *" Users & Roles" on page 842*.
2. Navigate to **Settings** > **Account info** > **CloudGuard License Activation** and enter your User Center credentials.

   For more information on how to create a User Center account, see sk22716.
3. Click **Login**. You can see the list of available licenses in your Account name.
4. Select the desired license and click **Activate**.

   The message of successful license activation appears.

Billable Assets Calculation

This section shows your cloud assets on different cloud platforms that are billable by CloudGuard when they are onboarded to CloudGuard.

The CloudGuard CNAPP license is based on the number of utilized capabilities (Posture Management, Runtime Protection, IAM Safety) and the number of protected assets in your cloud environment.

The billable assets graph shows a number of billable assets in your CloudGuard Account over time. The red line is the limit defined in the license, and the blue line is the actual daily number of billable assets. You can download a summary report with this information as a CSV file.

Billing Reports

The details of your Account usage are available in the billing reports. CloudGuard provides two types of reports that you can download, if your CloudGuard permissions are set to *view all system resources* (**Auditor** Role) or greater.

Available reports:

- **Detailed Billable Assets Report** is a monthly billing report. It includes information on the Account billable assets, for AWS and Azure environments, and corresponds to your monthly invoice.
- **Total Billable Assets Report** is a summary report for a selected period (1, 3, or 6 months). It includes information on the total number of billable assets in the account, for all onboarded environments. The blue line on the graph represents the same information.

The detailed report you download from the portal presents information for only the month before. To download a billing report for a specific month and year, use the CloudGuard API.

Detailed Report information

The report presents these details in `CSV` format:

- **Platform** - AWS and Azure are supported.
- **Billing Entity** - An entity designated as a billing authority for this billable asset.
- **Billing Entity Name** - The name of the billing authority (paying entity).

- **Environment External ID / Resource group ID** - The External ID for AWS environments or Resource group ID for Azure subscriptions.
- **Environment Name / Resource group name** - The name of the AWS environment or Azure subscription the current asset corresponds with.
- **Asset Type** - The available types: Instance, Lambda Function, RDS, Virtual Machine, and more.
- **Billable Assets Count (Peak)** - The peak number of billable assets.
- **Billable Assets Count (Avg)** - The average number of billable assets.
- **Billable Assets Count (Avg - Normalized)** - The average number of assets after the application of a billing multiplier.
- **Billing Multiplier** - The conversion rate between the actual workload in your environment and billable assets calculated by CloudGuard (see below *"Calculation of Billable Assets" on the next page*).
- **Organizational Unit Name** - The organizational unit that contains the asset.
- **Active Environment** - The Boolean value of True for currently active environments or False for inactive (deleted) environments.

Note - In the CloudGuard portal, you can download the Billing Report for only one full month only, that is, the previous calendar month.

To download the Detailed Report from the CloudGuard portal:

1. In the CloudGuard portal, go to **Settings** > **Account Info**.
2. In the **Billable Assets** section, click **Export** on the upper right and select **Detailed Billable Assets Report**.
3. Save the file on your hard drive.

To download the Total Report from the CloudGuard portal:

1. In the CloudGuard portal, go to **Settings** > **Account Info**.
2. In the **Billable Assets** section, select the period for the graph and summary report.
3. On the top right, click **Export** and select **Total Billable Assets Report**.
4. Save the file on your hard drive.

To get the Detailed Report with REST API:

```
curl --location --request GET
'https://api.dome9.com/v2/account/billable-assets-
report?month=<requested_month_of_report>&year=<requested_year_of_
report>' \
                --header 'accept: application/json' \
                --header 'Authorization: Basic <<span class="Vars_CG_
Native.CloudGuard_product mc-variable">CloudGuard</span> API
key>:<CloudGuard API secret>'
```

For more information, see CloudGuard API Reference.

Known Errors

Error Message	Description
Internal Error	CloudGuard cannot get the AWS account number because of an error.
Customer managed permission error	CloudGuard does not have permissions to collect the account details
Subscription has been deleted	The account is no longer onboarded to CloudGuard

Calculation of Billable Assets

CloudGuard Cloud Native Application Protection Platform pricing is based on the number of utilized capabilities (Compliance/Runtime/IAM Safety) and the number of protected assets in the cloud environment.

CloudGuard Native Applications (CNAPP)

These CloudGuard CNAPP services are complimentary for all supported assets:

- Effective Risk Management (ERM)
- Cloud Infrastructure Entitlement Management (CIEM)
- Agentless Workload Posture (AWP)

A CNAPP license includes:

- 50 GB quote for CDR Pro
- 1 million HTTP requests for CloudGuard WAF
- 5 file emulations per month for API and Threat Emulation

CNAPP is provided with a built-in ability to analyze user account activity. 12 GB of log retention for 1 month is provided for each billable asset.

For licensing purposes, these tables show conversions from cloud capabilities / assets to CloudGuard billable assets:

CSPM

Workload	Billable asset
1 VM instance*	1 asset
1 database instance**	1 asset
60 Serverless Functions (Lambda)	1 asset
1 containers node	3 assets

The container posture includes container compliance, admission control and image assurance (runtime, registry and ShiftLeft). ShiftLeft scans are subject to fair use of up to 50 scans per month per node.

All nodes running containers on onboarded clusters are counted.

*AWS EC2 (not including Micro and Nano), Azure VM (not including 0 family (A0/D0)) , GCP VM (not including F1-micro), Alibaba ECS, OCI VM.

**AWS RDS, Azure SQL DB & SQL Servers, Ali RDS, OCI Autonomous DB

Runtime Protection

Workload	Billable asset
10M Serverless invocations (Lambda)	1 asset
1 container node	2 assets

All nodes running the containers runtime agent are counted.

For serverless runtime, the CP-CGD9-CNX/P-100-1Y SKU includes 200 million complimentary annual invocations and the CP-CGD9-CNX/P-25-1Y SKU includes 50 million complimentary annual invocations

Code Security

Workload	Billable asset
IAC posture scanning	complimentary
Source Code scanning	1 asset per developer
Productivity & Collaboration	1 asset per 2 developers

Infrastructure as code security is complementary with the Asset license.

Intelligence and Threat Hunting

Network Traffic Usage

This graph shows your usage of Traffic activity (Flow Logs) for different environments over time. It is available if you onboarded your environments to Traffic Activity with a CloudGuard **CDR Pro** license.

- **Network Traffic Usage** - Amount of Flow Logs from all cloud platforms sent to CloudGuard in percent of the purchased license quota.
- **Estimated End of Quota** - Estimation date based on the average of the last week's statistics.
- **Data Retention Period** - Period of logs retention in CloudGuard, in days.

On the graph, click the environment name to toggle its visibility or point to a date to see the usage distribution by environments.

The offboarded environments that were included in the overall usage calculation, but did not use the quote recently are shown as **Offboarded**.

Account Activity Usage

This graph shows your usage of Account activity (CloudTrail) for different environments over time. It is available if you onboarded your environments to Account Activity with a CloudGuard **CNAPP** license.

- **Account Activity Usage** - Amount of CloudTrail and similar logs from all cloud platforms sent to CloudGuard in percent of the purchased license quota.
- **Estimated End of Quota** - Estimation date based on the average of the last week's statistics.
- **Data Retention Period** - Period of logs retention in CloudGuard, in days.

On the graph, click the environment name to toggle its visibility or point to a date to see the usage distribution by environment.

The offboarded environments that were included in the overall usage calculation, but did not use the quote recently are shown as **Offboarded**.

Credentials

These credentials apply only to accounts on Dome9 portal.

Password

Click **Change Password** to change your main CloudGuard password.

V2 API

For API functionality with CloudGuard, it is necessary to have an API key. For this, log in to the portal as a user with the permissions to manage all system resources.

To create a V2 API key:

1. In the CloudGuard portal, navigate to **Settings** > **Account > Credentials**.
2. Below V2 API, click **Create API Key**.
3. CloudGuard creates a new API key. Copy the details from the window that opens and click **OK**.

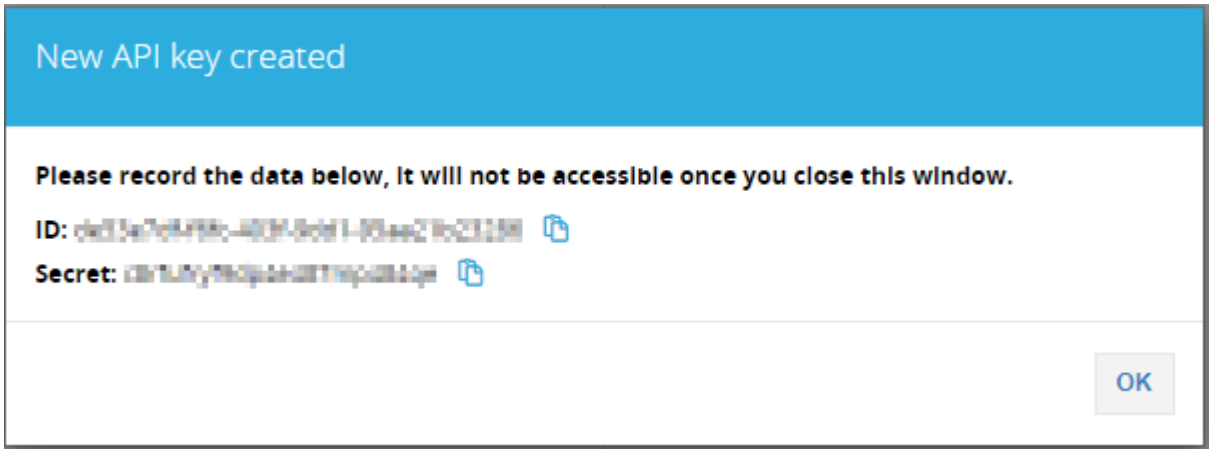

Note - Make sure to write down and record the ID and Secret that you receive. While you can see your ID later in the CloudGuard portal, you cannot retrieve the API Secret after you close this window.

CloudGuard Mobile Application

The CloudGuard mobile application allows you to configure access to a protected environment resource with Identity Safety and to set up a short-term dynamic access lease to your cloud resources. This application works only with the Dome9 portal on Apple iOS phones.

You can use the CloudGuard mobile app to create *"Dynamic Access Leasing" on page 376* or to open *"IAM Safety" on page 403*.

On an iOS phone, look for the CloudGuard app in the Apple App Store. After you install the app, you must pair it with your CloudGuard account.

To install and pair the mobile app:

Follow the steps below to install the CloudGuard app and pair it with your CloudGuard account. To use the app, you must have a screen lock protection method on your phone.

1. In the Apple App Store, find the CloudGuard app and install it.
2. In the CloudGuard portal on your computer, navigate to the **Settings** menu and select **Mobile Devices**.
3. In the on-screen instructions, click **GENERATE** to create a verification code.
4. On your mobile phone, start the CloudGuard app.
5. Enter the email address you use with your CloudGuard account and then enter the verification code.

Your mobile phone is paired with your CloudGuard account. It appears as **Active** from the list of **Devices** in the CloudGuard portal.

Users & Roles

Accounts created in the Dome9 portal and in the Infinity Portal handle users and roles slightly differently. Dome9 users are created in the Dome9 portal. In the Infinity Portal, users are created for the entire portal and then imported to the CloudGuard CNAPP integrated into it.

Users

Users interact with CloudGuard with:

- Web interface (through the portal)
- REST API

Infinity Portal

If you do not see the **Users** page in the **Settings** menu, the users on your CloudGuard account are fully managed by the Infinity Portal. For more information, see the Infinity Portal Administration Guide.

If you see the **Users** page in the **Settings** menu, then it is necessary to import users created in the Infinity Portal to CloudGuard. For more information, see *"Adding a New User in the Infinity Portal" on page 848*.

Dome9 Portal

The **Users** page under the **Settings** menu shows the users of the current CloudGuard account.

The user that creates the account is the **Account Owner**. This user manages CloudGuard Account-related issues, such as billing and subscription plan and has the privileges of a Super User. Only one Account Owner exists for each account. An Account Owner can assign a different user as the Account Owner. In this case, the previous Account Owner receives the role of Super User.

CloudGuard uniquely identifies a user with an email address. You cannot create more than one user for each email. If you need a user which is not bound to an email address, create a Service Account.

Caution - Make sure to delete unnecessary SSO users when they are deactivated or no longer need access to CloudGuard (see *"Deleting users" on page 850* for more information).

Service Accounts

You can create a **Service Account** to work with CloudGuard through the API. A service account interaction with CloudGuard using the web interface is not possible. You identify the service account with an API Key ID and API Key Secret. Unlike a regular user, this account is not bound to a specific email address. You can use the service account for administration, maintenance, and all other automation tasks, regardless of the person who does these tasks.

You can assign service accounts the same Roles as regular users. To create a service account, see *"Adding a New Service Account" on page 849*.

Roles

You can configure roles and assign them to users and service accounts. Then you assign permissions to a role. When you assign a role to a user, the permissions of the role are granted to the user, so it is not necessary to assign these permissions to the user explicitly.

In the Infinity Portal only, these e roles are synchronized with Specific Service Roles in your Infinity Portal account. You can assign the roles to users in the Infinity Portal. For more information, see Adding and Editing User Accounts.

You can configure any number of custom roles to include all the different types of users necessary for your CloudGuard account, each with the permissions applicable to it.

The preconfigured CloudGuard roles include:

- **Super User** for Dome9 or **Primary Admin / Admin** for Infinity Portal - Can access and manage all system resources, add new users, and change their privileges. There can be multiple Super Users in the system.
- **Auditor** for Dome9 or **User Admin / Read-Only** for Infinity Portal - Can see all system resources, but cannot create, change, or delete them.
- **Kubernetes Agent** - Internal role used by Kubernetes agents.

You cannot change or delete the preconfigured roles. You cannot delete a role that contains members.

Switch User Roles

In the Dome9 portal, use the menu on the top bar next to your username to select a different role in your CloudGuard account. The role must be configured and assigned to you.

Direct Permissions

You can grant direct permissions to users or roles to perform various actions in CloudGuard. Some permissions can be set separately or as part of other permissions. Some other permissions can only be granted collectively, such as View permissions given by inheritance. For example, the permission for managing **Policies** also grants permission to view **Rulesets** and **Notifications**.

To set direct permissions, select where to apply them (**Scope & Controls**, **Network Security**, or **Code Security**) and then drill down to set the required level of granularity. At each level, you can grant permissions to **View** or **Manage**.

To see permissions that you have already set, toggle the **Show Selected** button.

Scope & Controls

Scope

The All System Resources permission affects permissions to all resources in the system.

Scope	Resource Name	Includes	Impact
All System Resources	System configurations - Set only as part of the All System Resources	■ Accounts ■ Users and roles ■ Network security (can be set separately with the **Create Security Groups** permission) ■ Leases (can be set separately with the **Dynamic Access** permission) ■ Onboarding (can be set separately with the **Onboarding** permission)	
	CloudGuard resources - Can be set separately from the All System Resources	■ Notifications and integrations (can be set separately with the **Notifications** permission) ■ Rules & rulesets (can be set separately with the **Rules & Ruleset** permission) ■ Policies (can be set separately with the **Policies'** permission) ■ Alerts, exclusion, and remediation (can be set separately with the **Manage Alerts** permission)	Affects permissions to all of these resources: ■ Policy, Rules and Rulesets, Notifications & Integrations ■ Alerts, Exclusions, and Remediations
	Code Security resources	Select the permission level to which assign a Code Security role (Admin Access, Member Access, or Read-Only Access). If two permissions are assigned, the higher permission is granted.	

Scope	Resource Name	Includes	Impact
	All or specific assets - Can be set separately from the All System Resources.	■ Create / manage or view access to an Organizational Unit or any of the nested Organizational Units. ■ Select environments to give access to • all environments of a specific vendor • specific environments for a specific vendor	Affect the specified assets

Controls

Use these permissions to view all CloudGuard resources or manage them at the required granularity.

Controls	Resource Name	Description	Applicable Resources
All CloudGuard Resources	Alerts, exclusions, and remediations	Create and manage findings, exclusions, and remediations. Includes the **View** permission for Rules and Rulesets.	CloudGuard events, exclusions, remediations, rules, and rulesets
	Notifications and integrations	Create and manage *"Notifications" on page 852* for CloudGuard policies and integrations (see *"Integration Hub" on page 796*).	Notifications, integrations
	Policy	Create and manage CloudGuard policies: ■ Create a new policy, edit an existing policy, delete/unassociate a policy ■ Includes the **View** permissions for rulesets and notifications	CloudGuard policies, rulesets, rules, notifications
	Rules and Rulesets	Create and manage *"Rules and Rulesets" on page 309*	Rulesets, rules
Onboarding		Onboard and offboard environments from your CloudGuard account.	CloudGuard environments

Network Security

Dynamic Access

Use Dynamic Access Leases for secure access to your Security Groups (see *"Dynamic Access Leasing" on page 376*) for:

- AWS cloud accounts
- Organizational Units

Controls

Use controls to enable permissions to:

- Create security groups
- Create CloudGuard agents

Code Security

Access Level

Select the permission level to which assign a Code Security role:

- Read-Only Access
- Member Access
- Admin Access

If two permissions are assigned, the higher permission is granted.

Configurations

You can manage users, service accounts, and roles in the Users & Roles menu. In the users or roles table, click the menu in the first column to see and select available actions.

Adding a New User in the Infinity Portal

Super Users (Admins) and Account Owner (Primary Admin) can add new CloudGuard users to the account:

- Step 1 - Invite users to the Infinity Portal. To invite users to the Infinity Portal, refer to the instructions in the Infinity Portal Admin Guide and see Configuring Users > To add Users to the Infinity Portal account.
- Step 2 - Import the users to CloudGuard. This step is only applicable if you have imported users into the existing account before. In this case, you have the **Users** page in the **Settings** menu.

  If you create a new CloudGuard account, this step is not applicable.
- Step 3 - Assign roles or permissions to the users. For more, see *"Adding a Role" on page 850*.

To import users to CloudGuard:

1. Select the **Users** page in the **Settings** > **Roles** menu.
2. Click **Import User**.

   The new window opens.

3. Select the user from the list. You can select only a user who accepted your invitation to the Infinity Portal and authenticated with the email address and password.
4. Select a Role for the user. The permissions corresponding to the role are automatically granted to the user.
5. Click **Add**.

   The user appears in the users list with the assigned role.

Adding a New User in the Dome9 Portal

Super Users and Account Owners can add new CloudGuard users to the account.

1. Open the **Users** page in the **Settings** > **Users & Roles** menu.
2. Click **Add User**.
3. Enter details for the user. The user is identified by the email address. If the user signs in with Single Sign-On, see *"Single Sign-On" on page 866*.
4. Select Roles or Permissions for the user. The user receives the permissions corresponding with the role automatically, so no need to assign these explicitly in the Permissions section. If you do not assign a role, you must explicitly assign permissions to the user in this section. Users with direct permissions have the **Direct** tag in their Roles list.
5. Click **Close**. An email is sent to the new user, based on the email address entered for the user.

Adding a New Service Account

Super Users (Admins) and Account Owner (Primary Administrator) can add new CloudGuard Service Accounts.

1. From the **Settings** menu, select the **Users & Roles > Service Accounts** page.
2. Click **Add Account**.
3. In the **Add Service Account** dialog box, enter the account name.
4. Select a Role for the service account. You can select more than one role when you click each Role one by one.
5. Click **Add**. The **New Service Account Details** dialog box displays the API Key ID and API Key Secret values.
6. Click the **Copy** icon to copy the details of each value and save them for future use.
7. Click **Close**.

Adding a Role

You can configure roles with specific permissions and assign them to users and service accounts. The roles you configure are specific to your CloudGuard account. In the Infinity Portal, the roles are synchronized with Specific Service Roles in your Infinity Portal account. You can assign roles to the users in the Infinity Portal.

1. Open the **Roles** page in the **Settings** > **Users & Roles** menu.
2. Click **Add Role**.
3. Enter a name for the role and select permissions for it.
4. Optionally, select users and service accounts for the role. These users and accounts receive the permissions corresponding with the role.

Changing Roles or Permissions for a User

You can change details for a user or a service account, including their permissions.

1. Select the user or service account from the list.
2. On the menu bar, click **Edit** to make changes to the role(s) or permissions related to the user.
3. Click **Close**.

Connecting a user to SSO for Dome9 accounts

A Super User can configure a user to use Single Sign-On (SSO). To do this, first enable *"Single Sign-On" on page 866* for the account.

1. With a Super User account, log in to the CloudGuard portal and navigate to the **Users** page in the **Settings** > **Users & Roles** menu.
2. Select the user that is necessary to connect to SSO and click **Connect to SSO** on the menu bar.

Disconnect a user from SSO for Dome9 accounts

1. With Super User credentials, log in to the CloudGuard portal and navigate to the **Users** page in the **Settings** > **Users & Roles** menu.
2. Select the user that is necessary to disconnect from SSO and click **Disconnect from SSO** on the menu bar.

Deleting users

As a best practice, delete all unnecessary SSO and non-SSO users from the user list.

1. With a Super User account, log in to the CloudGuard portal and navigate to the **Users** page in the **Settings** > **Users & Roles** menu.
2. Select a user to delete and click **Delete** on the menu bar.

Disabling MFA for other users with Dome9 accounts

A Super User can disable MFA for other users.

To disable MFA for a different user:

1. In CloudGuard, navigate to **Settings > Users & Roles > Users** to see the list of all users in the CloudGuard account.
2. Select the user with enabled MFA.
3. On the menu bar, click **Disable MFA**.

   A confirmation window opens.
4. Click **OK**.

Setting a user as the Account Owner

The Account Owner (Primary Administrator) can assign a different user to be the Account Owner. Then, the former Account Owner is automatically assigned the Super User role.

1. With the Account Owner credentials, log in to the CloudGuard portal and navigate to the **Users** page in the **Settings** > **Users & Roles** menu.
2. Select the user that is necessary to set as the Account Owner and click **Set as account owner** on the menu bar.

Unlock a user

Users who enter an incorrect password more than a set number of times when logging in are locked out of their account. Their account can be unlocked by a Super User, on the **Users** page.

- To unlock the user, select the user and click **Reset password** on the menu bar.

Notifications

CloudGuard can send notifications in an email or through an integration with a third-party platform. Notifications show CloudGuard findings and security scores that CloudGuard assigns to your environments.

Notes:

- The Code Security feature has its own third-party integrations and sends its own notifications. To configure Code Security to send notifications, see *"Code Security Integrations" on page 739*.
- The Toxic Combinations feature sends its own notifications. The Toxic Combinations feature uses the same third-party integrations as other CloudGuard features. Only some third-party integrations are supported for Toxic Combinations. For more information, see *"Action Hub" on page 93*.

Notification Types

You can send these types of notifications:

- **Summary Report** shows you the security score for each of your environments and compares it to the results in the previous report. In addition, it shows an aggregated result for all your accounts.
- **Executive Summary Report** shows the status of your environments and assets based on the results of the last test that CloudGuard performed. This report focuses on a specific ruleset for multiple environments on one cloud platform. The report includes:
  - The environments with the highest number of severity findings
  - The distribution of assets that passed or failed the test
  - The test score
  - The number of failed tests, sorted by the severity of the rule
- **Detailed Report** shows details for each failed test. It also shows the current status of findings from the previous report. This provides a complete picture of the compliance posture of your cloud environments and an indication of progress in resolving open issues.
- **Immediate Notification** sends information about a specific finding immediately after CloudGuard generates the finding.

How to Configure a Notification

Step 1: Create an integration to receive the notification from CloudGuard

1. In the CloudGuard UI, from the left menu, click **Integration Hub**.
2. Click the external service you want to integrate (for example: Microsoft Teams).

   A sliding window opens.
3. In the sliding window, click **Add**.
4. Enter the required information from the external service.
5. Click **Save**.

For more information, see *"Integration Hub" on page 796*.

Step 2: Create a notification

Note - The Toxic Combinations feature sends its own notifications. Notifications for Toxic Combinations do not need to be added to a policy. To configure Toxic Combinations to send notifications, see *"Action Hub" on page 93*.

1. Navigate to **Settings** > **Configuration** > **Notifications**.

   A list of notifications appears.
2. Click **Add**.

   The **Create New Notification** window opens.
3. Enter a unique **Name** and a **Description** for the notification.
4. To schedule CloudGuard to send scheduled reports, in the **Schedule Report** section select **Email scheduled reports** and fill the relevant fields. To schedule CloudGuard to send reports on a custom schedule, see *"Appendix: How to schedule reports on a custom schedule" on page 855*
5. To configure CloudGuard to send information findings as soon as CloudGuard detects changes in your environment:

   a. In the **Immediate Notification** section, select a notification type. Use the Filter bar to send notifications only about certain kinds of findings. For example, you can create an immediate notification only for Critical findings.

      Note - Some notification types apply for all findings (for example: **Email notification per newly created finding).** Other notification types apply only to specific types of findings (for example**: CSPM- Summary report to Teams channel** applies only to CSPM findings).

b. Select a configuration of an integration.

Note - If there is no configuration of a specific integration (for example: there is no configuration of a Microsoft Teams integration), select **Add new configuration**. Then, create the configuration in the sliding window. For more information, see *"Integration Hub" on page 796*.

Note - It is not possible to select more than one configuration of the same integration type. For example, it is not possible to select more than one Microsoft Teams configuration.

c. **Optional -** Select more notification types and integrations to add to the notification.

6. Click **Save**.

The new notification appears in the list of notifications.

Step 3: Add the notification to a policy

1. From the left menu, navigate to **one** of these screens:
   - **CSPM** > **Continuous Posture**
   - **CIEM** > **Policies**
   - **Workload Protection** > **Admission Control** > **Policies**
   - **CDR > Manage Policies**
2. Do **one** of these:
   - To create a new policy, in the top right click **Add Policy** > select the policy type.
   - To edit a policy, select the checkbox to the left of the policy > click **Edit**.

   The **Add Policy** or **Edit Policy** wizard window opens.
3. **Optional -** To create a new notification, in the **Notifications Selection** step of the wizard, click **Add Notification.** In the **Create New Notification** window that opens, follow the procedure in *"Step 2: Create a notification" on the previous page*.
4. In the **Notification Select** step of the wizard, select one or more notifications.

   Important - Select only notifications that are relevant to the policy. For example, **CIEM Notifications - Email** is relevant to CIEM, but it is not relevant to CSPM.
5. Finish the wizard.

Appendix: How to schedule reports on a custom schedule

In the **Create New Notification** window > **Schedule Report** section, you can use the dropdown menus to configure CloudGuard to send reports daily, weekly, or monthly at a specific time of day. You can use a cron expression to configure CloudGuard to send reports on a custom schedule.

A 7-digit cron expression contains seven fields. To leave a field blank, enter an asterisk (*).

Cron Expression Fields (from left to right)

Field	Allowed Values
Second	0-59
Minute	0-59
Hour	0-23
Day of Month	1-31
Month	1-12
Day of Week	0-6
Year	1970-2099

Cron Expression Special Characters

Special Character	Meaning
*	Any value
,	Separates a list of values
-	Range of values

Cron Expression Examples

Example	Meaning
0 0 * * * *	Sends a report at the beginning of every hour
0 0 9 * * *	Sends a report every day at 09:00:00 UTC.
0 30 16 * * 1-5 *	Sends a report from Monday through Friday at 14:30:00 UTC.

Example	Meaning
0 0 11 * * 6,0 *	Sends a report on Saturdays and Sundays at 11:00:00 UTC.

To use a Cron Expression to schedule a report:

1. In the **Create New Notification** window > **Schedule Report** section, select **Custom**.
2. In the **Enter cron expression field**, enter a cron expression.
3. Finish configuring the notification.

Sending All Alerts

You can manually send all reports and notifications for a policy immediately. This is useful to do a security investigation or to test integrations. The **Send all alerts** action is supported for these policies:

- **CSPM** > **Continuous Posture**
- **Workload Protection** > **Admission Control** > **Policies**
- **Workload Protection** > **Vulnerabilities** > **Policies**

To send all alerts immediately

1. In the CloudGuard UI, navigate to one of the supported policies.
2. Select the policy that you want to synchronize and click **Send all alerts**.
3. Select the notification type and name from those attached to the policy and click **Send**.

Broken Notifications

If CloudGuard detects a misconfiguration or failure in an integration, it blocks the integration for six hours. After six hours, CloudGuard tries to send new notifications to the integration. Then, if CloudGuard detects a misconfiguration or failure, it blocks the integration again.

To resolved a misconfigured notification

1. In the CloudGuard UI, navigate to **Settings** > **Configuration** > **Notifications**.

   A yellow exclamation point icon appears in the **Status** column to show that a notification is misconfigured.

2. Click the name of the notification to open it.

   The problem is highlighted in red.

3. Click **Open Configuration**.

   A sliding window opens for the integration (for example: Microsoft Teams).

4. Select the relevant configuration of the integration (for example: a Microsoft Teams configuration that you named *teams_integration_1*).

5. Click **Test**.

   If the test fails, an error message describes the problem.

6. Fix the problem.

   CloudGuard test the integration.

7. After a successful test, click **Save**.

8. In the notification window, click **Validate**.

9. Click **Save**.

Security and Authentication

You can configure CloudGuard security and authentication settings on the **Security & Authentication** page of the **Settings** menu.

Security

Dome9 Account Lockout for Failed Password

Super Users can set an account lockout threshold for failed attempts to log in. If a user enters an incorrect password more than the configured amount of times, the account is locked. To unlock their account, the users must reset their password.

Super User can unlock a locked user account.

To set the number of failed attempts:

1. In CloudGuard, navigate to **Settings > Security & Authentication**.
2. In the Security section, set the Dome9 account lockout threshold.

Session Timeout

Super User can configure a timeout for idle CloudGuard sessions.

To set the inactivity period:

1. In CloudGuard, navigate to **Settings > Security & Authentication**.
2. In the Security section, set the period in the **Session idle timeout in minutes (15 min - 12 hours)** field.

Multi-Factor Authentication for Dome9 Accounts

You can configure your CloudGuard account to use Multi-Factor Authentication (MFA).

You can use the applications below to create one-time authentication codes to sign in to CloudGuard:

- Google Authenticator
- Twilio Authy
- Other equivalent applications

To enable MFA:

1. In CloudGuard, navigate to **Settings** > **Security & Authentication**.
2. In the Multi-Factor Authentication section, move the slider to **ON** to configure an application. The configuration window opens.
3. Install an applicable application and use it to scan the QR code. The application displays a 6-digit code.
4. In the MFA configuration window, enter the code below step 3 - **Verify your authenticator**.
5. Click **Verify and Save**. On successful verification, the status changes to **ON**.
6. Account administrators can enforce MFA for all users, when they select this option. In this case, all users receive an email notification that they must enable MFA in 30 days. Users that do not configure MFA cannot log in to CloudGuard by the end of the period.

   Note - When the account owner or super users enforce MFA for all users, super users must activate the MFA for themselves. For the account owner, the MFA activation is optional.

To disable MFA:

1. In CloudGuard, navigate to **Settings** > **Security & Authentication**.
2. In the Multi-Factor Authentication section, move the slider to **OFF**.

To disable MFA for a specific user, see *"Disabling MFA for other users with Dome9 accounts" on page 851*.

Email Notifications

CloudGuard can send notifications on its system events and audit logs to your email. The **Settings > Email Notifications** page shows available auditing groups.

Select the type(s) of events from each group to receive notifications about:

- Network Security
- Assets
- Administration
- Identity (if enabled)
- Posture Management

By default, you receive email notifications for all environments onboarded to your account. For Network Security and Identity, click **All environments selected** on the section header to select only specific environments for which you want to receive the emails.

In addition to the email notifications, you can send this information to an SNS (Simple Notification Service) topic in your AWS account.

More Links:

- *"Sending System Notifications to AWS SNS" on page 806*
- *"System Audit Logs" on page 134*

Workloads Settings

Image Assurance

You can set the period after which CloudGuard can delete inactive images, separately for Kubernetes and ShiftLeft images. The period ranges from 1 to 365 days. CloudGuard considers Kubernetes or ShiftLeft images inactive if none of their related workloads are running. To learn more about your images, see *"Images" on page 425*.

To set the period:

1. In CloudGuard, navigate to **Settings** > **Configuration** > **Workloads**.
2. In the Image Assurance section, enter the number of days for Kubernetes and ShiftLeft images.
3. Click **Save changes**.

Agentless Workload Protection

You can configure some AWP parameters individually for each account onboarded to AWP. For more information, see *"Viewing AWP Details" on page 491*.

To configure these parameters for all onboarded environments, use the Workloads Settings.

Caution - These settings are global and, when applied, override existing AWP settings in all onboarded accounts. To set custom AWP parameters for an environment, change the parameters individually in each specific environment after you finish to configure an OU.

To set the scan preferences for OU:

1. Navigate to **Settings** > **Configuration** > **Workloads** and see the **Agentless Workload Protection** section.
2. To set the period between consecutive scans of the same VM, enter a number between 24 and 1000 (by default, **24**) in **Scan interval (hours).**
3. To set the maximal number of simultaneous scans in the same region, enter a number between 1 and 20 (by default, **20**) in **Max Concurrent scans (scans per region)**.
4. To configure a custom tag, click **Add** and enter the tag's key and value. For more information on custom tags, see *"Custom Tags" on page 491*.
5. For AWS accounts with In-Account mode only, to enable the scanning of AWS

Marketplace licensed images, select the **Scan licensed images** option.

Note - Scanning of AWS Marketplace licensed images may result in additional charges based on AWS pricing policies. Initially, AWP attempts to scan with the regular machine types. If this fails, it reverts to the original machine types, subject to some cost limitations. Review your cloud storage plan to understand potential costs.

6. Click **Save changes**.

Filter and Search

You can filter views in CloudGuard with the filter bar at the top of the page or view. You can search for a specific entity by name in the search bar.

Save filters for future use and share them with other users.

Filter

You can find the filter bar on the top of the CloudGuard pages.

To open the filter selection, click **Add Filter** and then select a filter object. A secondary list is shown with filter options for the selected object.

Select the specific filter options for the selected object.

In addition, you can change (or remove) a selected filter option. Open the list of options for a filter object and change or add to the selected options.

You can save a filter to use it again or make it available to other users.

1. Select the filters in the filter block.
2. Click **Saved Filters**.

   The window shows all existing public and private filters.
3. In the **Save Filter** section, enter a name for the new filter.
4. Select **Public** to share this filter with other CloudGuard users.
5. Click **Save**.

Search

The search area is at the top of the page, on the right of the filter area.

Enter text in the search field to search for a specific object from the list. The search is incremental, and search results are updated as you enter more text. The search text is not case-sensitive.

System Search (ALT-/ shortcut)

You can open a universal search window anywhere in CloudGuard, and search in it for an entity, alert, policy, or a different object, by name or other text. For example, you can search for an asset by its name or for alerts by severity.

Click **ALT-/** anywhere in CloudGuard to open the search window.

Enter the search text and select the entities on which to apply it from a list.

Solutions

Terraform

Terraform is an open-source Infrastructure as code (IaC) tool. With the CloudGuard Dome9 Provider in Terraform, you can onboard environments from in Terraform, create Continuous Posture policies, and manage Security Groups.

See https://www.terraform.io/docs/providers/dome9

CloudGuard Terraform Provider

CloudBots

CloudGuard CloudBots provide automatic remediation for issues discovered in your environments by CloudGuard compliance assessments. CloudBots are an open-source project that you deploy in your environment and trigger from the compliance policy when issues are found. There are CloudBots for AWS, Azure, and GCP environments.

See *"Automatic Remediation with CloudBots" on page 317*

SDKs

CloudGuard has a number of open-source SDK projects that you can use.

Python API SDK

The Python API SDK is an open-source Python wrapper for the CloudGuard API. You can use this SDK to provide programmatic automation for a number of CloudGuard operations in Python, such as onboarding accounts, running assessments, and setting protection modes.

Go SDK

The CloudGuard Go SDK is an open-source Go wrapper for the CloudGuard API.

Tools

These open-source tools do utility functions related to your CloudGuard and environments.

S3 Logger

This script configures your CloudGuard account to send log information (for example, compliance findings) to an AWS S3 bucket.

Single Sign-On

This chapter describes configuring SSO for Dome9 accounts. To configure SSO for the Infinity Portal accounts, see *Infinity Portal Administration Guide*.

Single Sign-On (SSO) provides a means for enterprises to centrally manage and control users authentication and authorization.

Using SSO, organizations reduce the administrative overhead of managing multiple authentication tokens for each user. A user logs in with a single ID and password to gain access to a connected system or systems without using different usernames or passwords.

CloudGuard supports Single Sign-On based on SAML 2.0.

When SSO is enabled for a CloudGuard account, each account user can be configured to use SSO authentication (default), or a built-in user authentication.

Important - The SAML response generated by the Identity Provider (IdP) must be utilized within 24 hours before it expires. Each SAML response is valid for a single use only.

Users with SSO

Users configured to use SSO:

- Have the password managed by the SSO identity provider, so a password reset in CloudGuard direct the users to reset the password on the IdP (SSO Provider)
- Have MFA enabled and managed with the SSO solution provider, so MFA is disabled for these users in CloudGuard portal

Best Practice - Make sure to delete unnecessary SSO users when they are deactivated or no longer need access to CloudGuard (see *"Deleting users" on page 850* for more information).

A CloudGuard Account Owner cannot be configured for SSO. This restriction is a fail-safe in order to allow at least one user to be able to log in to the CloudGuard system if something goes wrong with the SSO identity provider.

Single Sign-On using Just-In-Time (JIT) Provisioning

With JIT provisioning, there is no need to create users for logging in. The Identity Provider provisions (creates or updates) them when the user attempts to access the service.

The provider allocates permissions based on the groups to which the user belongs. Roles, with specific permissions, must be defined and associated with the groups.

The provider generates temporary tokens for the user to access the service, so no actual user is created.

To enable JIT:

1. Configure SSO as usual according to the instructions for your IdP.
2. In SAML configuration, enter a meaningful attribute name, for example, *JIT-for-CloudGuard*.
3. In CloudGuard, navigate to **Settings** > **Security & Authentication** and configure SSO.
4. Select **Allow** for the **Just-in-time provisioning for the account** option.
5. In **Attribute name in SAML for just-in-time role**, add the name that you entered, *JIT-for-CloudGuard* (by default, *memberOf*).

Configure SSO

Before you use SSO, make sure your configuration meets these prerequisites:

- The organization has a SAML 2.0 SSO infrastructure in place
- Users are provisioned in the identity provider's SSO application
- A CloudGuard user with the same user identity email is provisioned in CloudGuard (when not using JIT)
- The CloudGuard user is assigned permissions in CloudGuard

SSO End User Login

An end user configured for SSO can log in to CloudGuard in two ways:

- Access the CloudGuard portal with the URL *https://secure.dome9.com/sso/yourcompanyname*, which redirects the user to log in with the SSO solution provider login page and, once successfully authenticated there, redirects the user back to the CloudGuard portal (Service-Provider-initiated)
- Log in through the SSO provider login page (IdP-initiated) and from there select the CloudGuard application

To log in with the SP-initiated flow:

1. Navigate to *https://secure.dome9.com/sso/yourcompanyname*, where your company name is the **Account ID** identifier configured in the SSO settings page.
2. You are redirected to the SSO provider's login page.
3. Log in to the SSO provider's site.

4. You are redirected back to the CloudGuard portal, with an authenticated session with the CloudGuard user corresponding to the user on the SSO site (with the same user email).

To log in with the IdP-initiated flow:

1. Navigate to the login page for the SSO provider and log in there with the SSO user name.
2. Select CloudGuard as the destination site.
3. You are redirected to the CloudGuard portal, with an authenticated session with the CloudGuard user corresponding to the user on the SSO site (with the same user email).

Actions

Disabling SSO for an Account

You can disable SSO for a CloudGuard account. If this is done, SSO is disabled for all users in the account, and a password reset invitation is issued to all SSO users.

1. Navigate to the **Security & Authentication** page in the **Settings** > **Configuration** menu.
2. Click **Disabled**.

Configuring new users with SSO

1. Log in to the CloudGuard portal with a super user account.
2. Navigate to the **Users** page in the **Settings** > **Users & Roles** menu.
3. Click **Add User**.
4. Enter details for the user. Note that SSO is enabled by default on accounts that have SSO enabled.
5. Click **CREATE**.

Changing the SSO configuration for users

1. Log in to the CloudGuard portal with a super user account.
2. Navigate to the **Users** page in the **Settings** > **Users & Roles** menu.

To disconnect a user from SSO:

3. Select the user that you want to disconnect from SSO and click **Disconnect from SSO** on the menu bar.

   When SSO is disabled, an email is sent to the user to reset the password.

To connect a user to SSO:

4. Select the user that you want to connect to SSO and click **Connect to SSO** on the menu bar.

   When SSO is enabled, an email is sent to the user to indicate they must use SSO and to specify their SSO provider to log in to CloudGuard.

SSO Configuration Troubleshooting

Most of the SSO issues are caused by wrong configuration. To troubleshoot the issues, navigate to **Events** > **Operational** > **System Audit Logs** page in CloudGuard.

If the System Audit log contains an *SSO Login failed* record, it means that there are specific configuration errors.

If there is no SSO Login records, it means that the SAML request is not configured to target a valid environment of the IdP.

SSO Login Failed record

The description alerts on the cause of the failure.

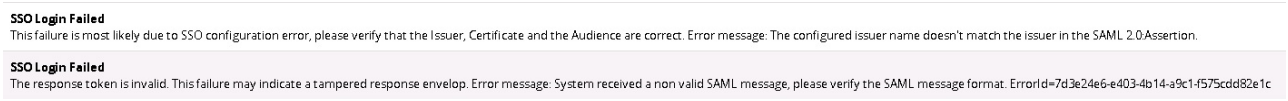

If you see one of the two CloudGuard audit log messages above, verify these:

- Ensure the Token-signing certificate is **SHA256**.
- Ensure the Token-signing certificate public key is entered in the CloudGuard SSO configuration in **Base-64 encoded X.509** format.
- Ensure your **Issuer** field is correct.
  If the **Issuer** URI field is incorrect, your browser does not forward to your IdP login page, but instead forwards you to the standard CloudGuard login page. To solve this, adjust the Issuer field on the CloudGuard SSO configuration.
- With the second audit log entry on invalid response token, verify that the certificate is valid, not encrypted and not expired.

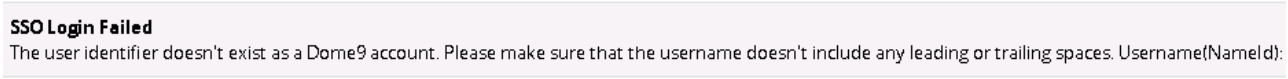

This record mentions that the user that tried to log in using SSO does not exist on CloudGuard system. Verify that the user name is correct.

SSO Failed and no Audit on CloudGuard portal

This situation indicates that Relying Party Trust is not created correctly or the associated claim rules are not configured correctly in the Claims Issuance Policy. This could also indicate that the metadata file contains a wrong Account Identifier.

Configure SSO JIT Provisioning with ADFS

With JIT provisioning, users log in to ADFS with their ADFS credentials and select to log in to CloudGuard configured as an SSO application. SAML authenticates the credentials and transfers them to CloudGuard to create users based on their email address and AD group membership that are mapped to the existing CloudGuard roles.

Configuring Active Directory

To map dynamically AD groups to CloudGuard roles, create AD groups with CloudGuard or other relevant prefix (for example, *CG-Admins*, *CG-Users*). Add users to these AD groups. Check Point recommends to add a user to only one CloudGuard AD group.

Create AD Group and users

1. Create an Active Directory security group for JIT provisioning to CloudGuard (for example, *CG-jit-auditors*) in your domain environment. Add to the group the domain users who have to log into CloudGuard.
2. Add users to the group.

Configuring ADFS

On the ADFS side, ensure that you have a working ADFS configuration. For this, log into the main page

https://<server.domain.com>/adfs/ls/idpinitiatedsignon.htm

where *<server.domain.com>* is your ADFS WAP server. If you need a valid SSL certificate, you can get one for free from *Let's Encrypt*.

Then you need to add CloudGuard as a Relying Party Trust to be a consumer of the ID provider.

Adding a Relying Party Trust

1. In ADFS, add a new **Relying Party Trust** with **Claims aware** option and follow the steps in the Relying Party Trust Wizard.
2. Copy the **Service Provider Metadata XML** text and save it. Replace <*your-company-name*> to match the **Account ID** string that you configure in step 3(a) below in *Configuring CloudGuard*. If necessary, update the property values of **validUntil** and **cacheDuration**.

XML Template for Service Provider Metadata:

```
<?xml version="1.0"?>
  <md:EntityDescriptor
xmlns:
md="urn:oasis:names:tc:SAML:2.0:metadata" validUntil="2025-09-
03T06:43:37Z"
 cacheDuration="PT604800S" entityID="https://secure.dome9.com">
    <md:SPSSODescriptor
AuthnRequestsSigned
="false"

WantAssertionsSigned
="true"

protocolSupportEnumeration
="urn:oasis:names:tc:SAML:2.0:protocol">
      <md:NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-
format:unspecified</md:NameIDFormat>
      <md:AssertionConsumerService
Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-
POST" Location="https://secure.dome9.com/sso/saml/<your-
company-name>" index="1" />
    </md:SPSSODescriptor>
  </md:EntityDescriptor>
```

3. On the Select Data Source page, import data from the **Service Provider Metadata XML** file created above.
4. On the Specify Display Name page, in the **Display name** field, enter **CloudGuard**.
5. Click **Next** on each screen until the end of the Wizard.

Adding Claims Issuance Policy rules

1. From the **Actions** panel, select **Edit Claim Issuance Policy**.
2. Click **Add Rule** and edit the **Get Email** rule as the picture shows.
3. Click **OK**.
4. Click **Add Rule** to create another rule.
5. Set the rule name to **Convert Email to NameID** and edit the rule.
6. Select **Transform an Incoming Claim** and click **Next**. Set these fields:

a. **Incoming claim type** - **E-Mail Address**

b. **Outgoing claim type** - **Name ID**.

c. **Outgoing name ID format** - **Email**.

7. Click **OK**. You have two rules:
8. Click **Apply**.

Configuring CloudGuard

When you continue in CloudGuard, log in with your super user credentials.

Configuring CloudGuard

1. In CloudGuard, navigate to the **Security & Authentication** page in the **Settings** menu.
2. In the SSO section, click **Edit**. The SSO Configuration window opens.
3. Complete the SSO Configuration form with these details:
   a. **Account ID** - Use the Account ID from the XML file above.
   b. **Issuer** - Set to *http://<your.server.com>/adfs/services/trust*
   c. **Idp endpoint url** - Set to *http://<your.server.com>/adfs/ls*
   d. **X.509 Certificate** - Paste from your PFX file, including the `--BEGIN CERTIFICATE--` and `---END CERTIFICATE---` text. This is the certificate from your ADFS web server.
4. Check **Allow** for **Just-in-time provisioning for the account**, and leave **Attribute name in SAML for just-in-time role** as **memberOf**.
5. Click **Save**.

Testing ADFS Single Sign-On

Testing sign in to CloudGuard using ADFS SSO

1. Log in to the ADFS sign-in page with AD credentials and click **Sign in to one of the following sites** - CloudGuard or other Display Name.
2. Click **Sign in**. This action logs you into CloudGuard as a JIT user, matching the AD group to a role with the same name in CloudGuard.
3. Check the System Audit Trail to see the JIT event:

Configure CloudGuard SSO with Microsoft Entra ID

To enable Single Sign-On (SSO) with Microsoft Entra ID (formerly Azure AD), start the configuration on the Microsoft Azure portal. You can then log in to CloudGuard as *super user* to configure SSO and redirect users' login requests to the Identity Provider login page. After it, CloudGuard administrative users can log in with SSO.

Adding the CloudGuard Application

In these steps, you add the CloudGuard application from the gallery to your list of managed SaaS applications.

To add CloudGuard from the gallery:

1. In the Azure portal, on the left navigation panel, click **Microsoft Entra ID**.
2. Navigate to **Enterprise applications** and select **All applications**.
3. Click **New application** to add a new application.
4. In the search field, type **CloudGuard** and select the **Check Point CloudGuard Posture Management (Microsoft Entra ID)** application from the results panel.

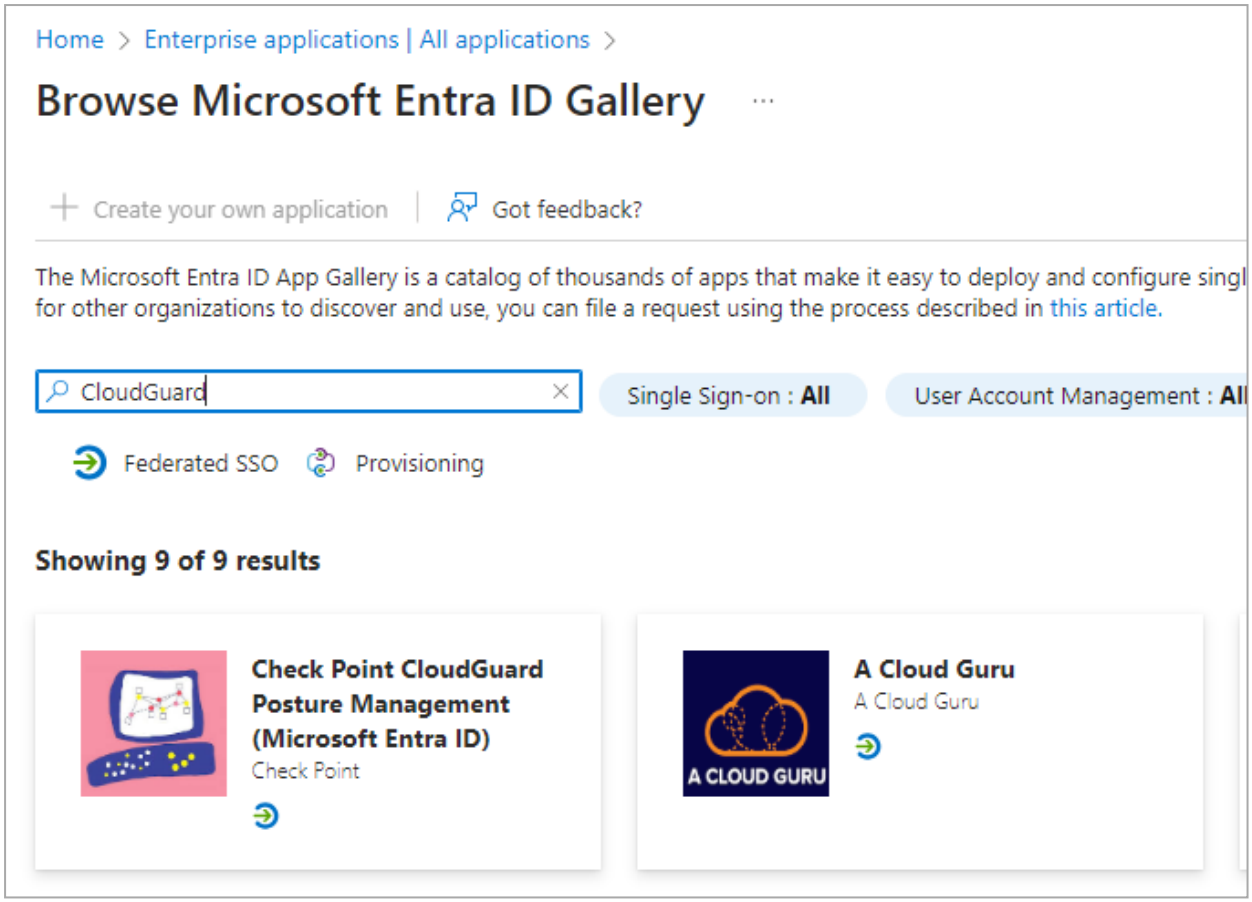

5. Click **Create** to add the application.

Configuring Microsoft Entra ID Single Sign-On

In these steps, you configure SAML as single sign-on method in Microsoft Entra ID.

To configure Microsoft Entra ID SSO with CloudGuard:

1. In the Azure portal, on the **CloudGuard** application integration page, select **Single sign-on** from the **Manage** menu.
2. From the gallery of available methods, select **SAML**.

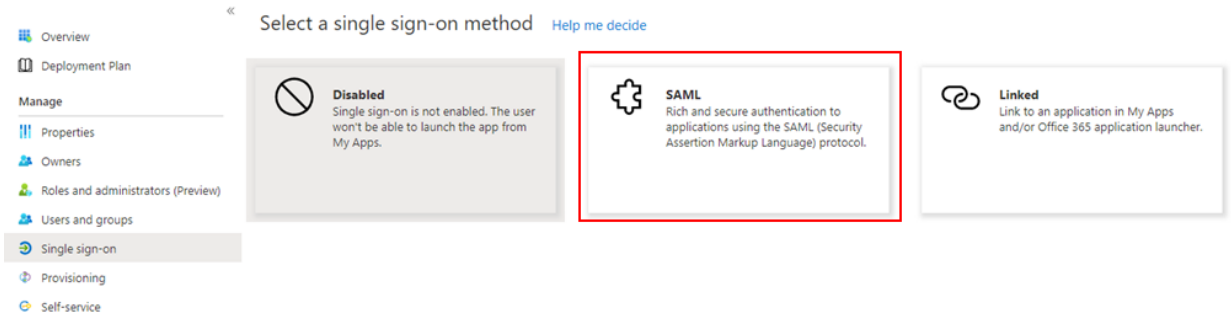

3. On the **Set up single Sign-On with SAML** page, click the pencil icon of **Basic SAML Configuration** to edit the settings.
4. To configure the application in the IDP-initiated mode:
    a. In the **Identifier** field, enter your data center address, for example, `https://secure.dome9.com/`.

    Note - In this and similar fields below, replace the generic address (https://secure.dome9.com/) with your regional (*YYY*) address:
        - United States - `www.dome9.com`
        - Europe - `www.eu1.dome9.com`
        - Singapore - `www.ap1.dome9.com`
        - Australia - `www.ap2.dome9.com`
        - Canada - `www.cace1.dome9.com`
        - India - `www.ap3.dome9.com`

    b. In the **Reply URL** field, enter a URL with the pattern: `https://secure.YYY.dome9.com/sso/saml/<yourcompanyname>`.

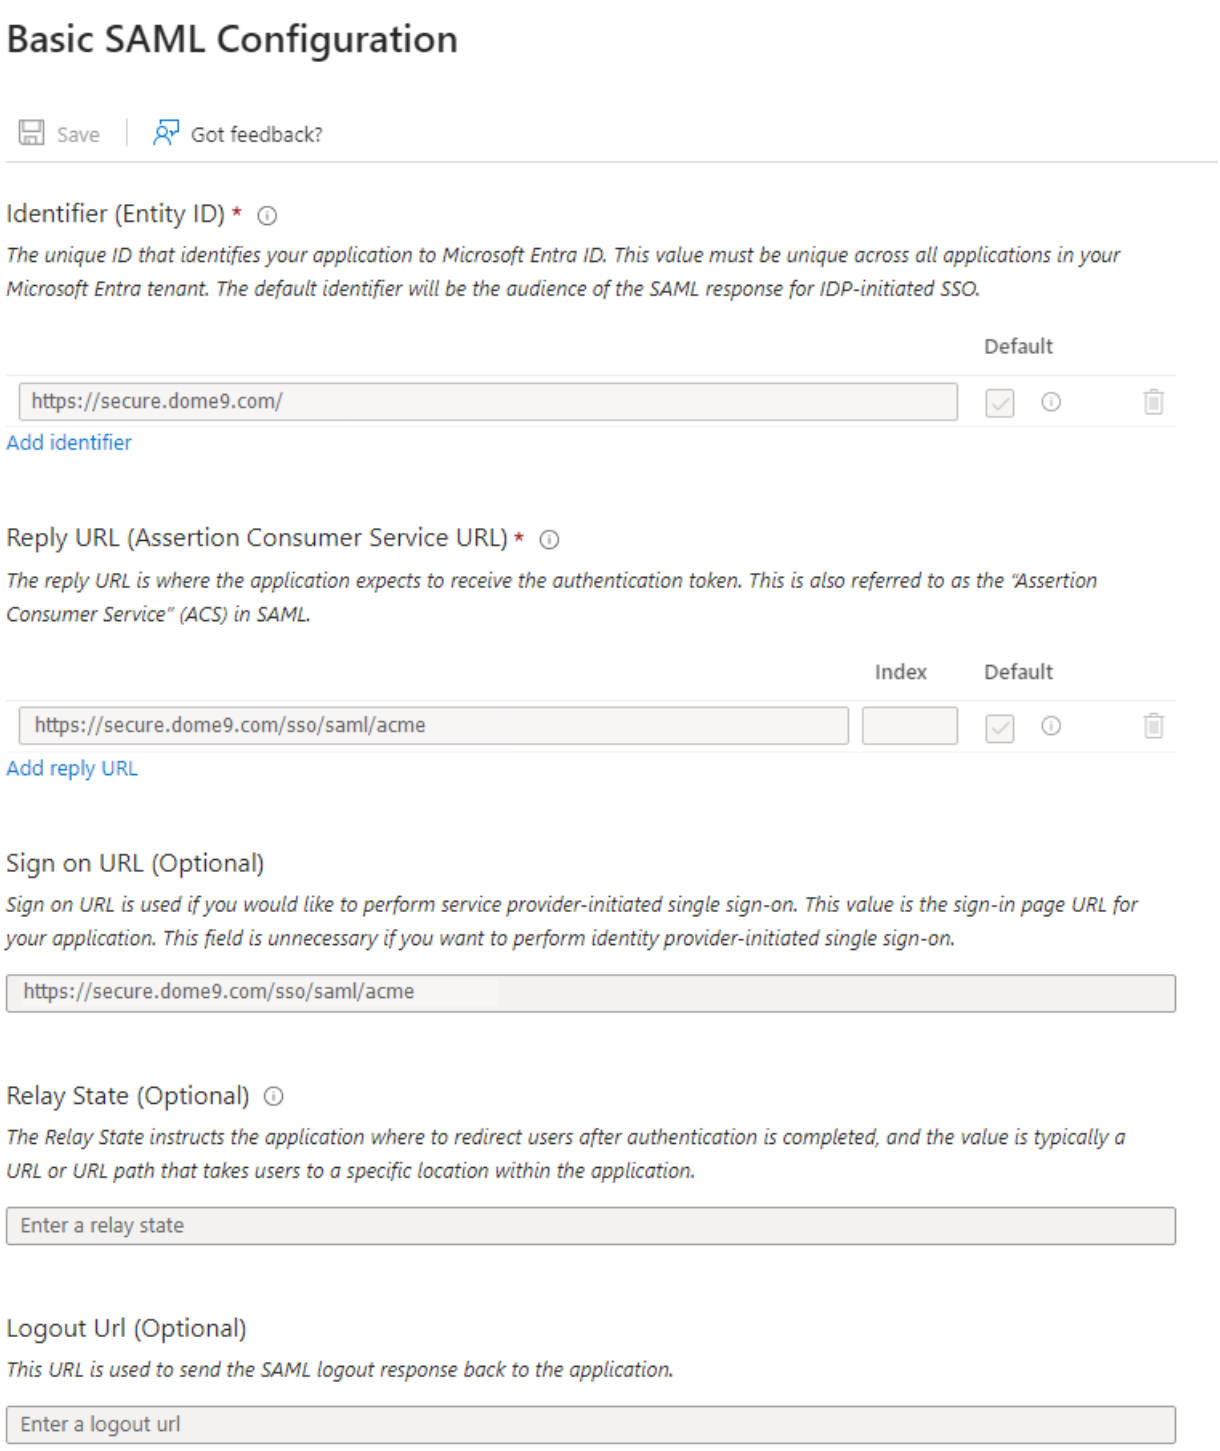

Note - Enter your company name in the CloudGuard portal as explained in step **4a** of *"Configuring Single Sign-On in CloudGuard" on the next page*.

5. To configure the application in the SP-initiated mode, select **Show advanced URL settings** and, in the **Sign on URL** field, enter a URL with the pattern `https://secure.YYY.dome9.com/sso/saml/<yourcompanyname>`.

6. Add custom attribute mappings to your SAML token attributes configuration. See the image below for the list of required attributes. The *memberof* attribute is non-default, so you have to search it to add. All other attributes are default and immediately available.

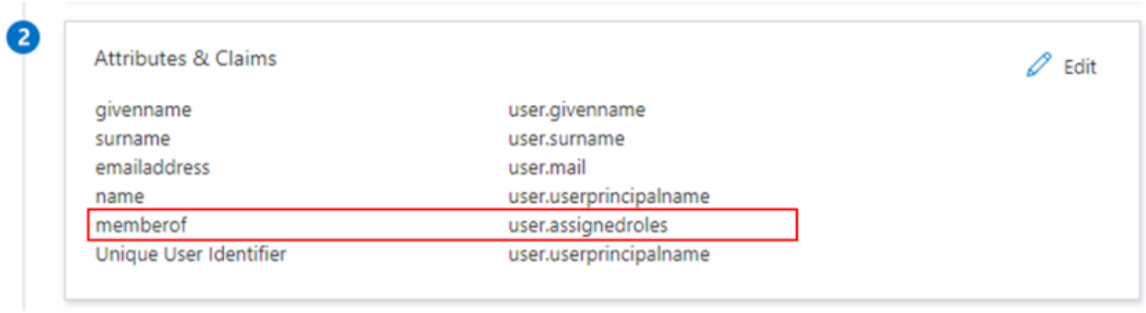

7. In the **SAML Signing Certificate** section, click **Download** for **Certificate (Base64)** and save the certificate file on your hard drive.

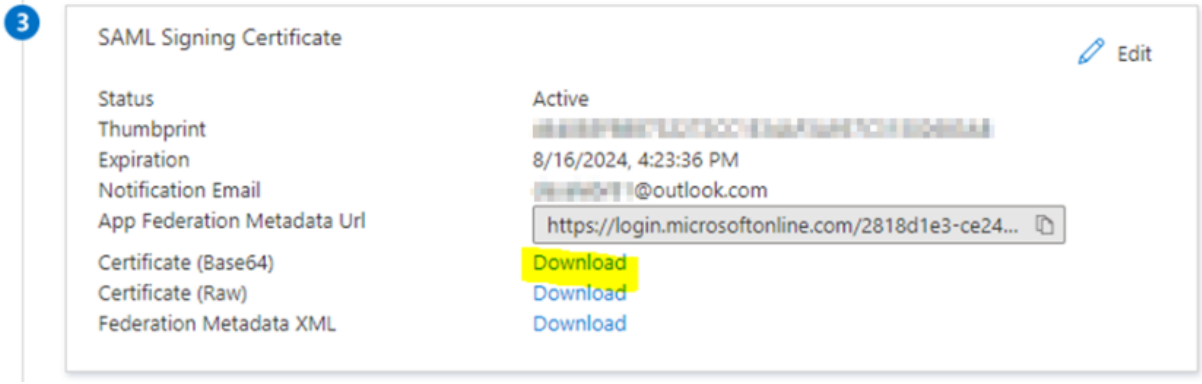

8. If you want to automate the application configuration, click **Install the extension** to install the **My Apps Secure Sign-in** browser extension. If you want to configure it manually, go to step **2** in *"Configuring Single Sign-On in CloudGuard" below* below.
9. After you add the extension, click **Set up Check Point CloudGuard Posture Management (Microsoft Entra ID)** to open the CloudGuard portal.

Configuring Single Sign-On in CloudGuard

1. Log in to the CloudGuard portal with the super user credentials. The browser extension installed in step **8** of the previous section automatically configures the application and automates steps **2-5** below.
2. If you want to set up the application manually, in a new web browser tab, log in to the CloudGuard portal with super user credentials and go to **Settings** > **Configuration** > **Security & Authentication**.
3. In the **SSO** section, click **ENABLE**.
4. In the **SSO Configuration** window, enter these details:

a. **Account ID** - Enter the company name. It appears in the **Reply URL** and **Sign on URL** fields of the **Basic SAML Configuration** section at the Azure portal.

b. **Issuer** - Paste the value of **Microsoft Entra ID** from the Configuration URLs on the Azure portal.

c. **IDP endpoint URL** - Paste the value of **Login URL** from the Configuration URLs on the Azure portal.

d. **X.509 certificate** - In the Notepad, open the certificate your downloaded in step **7**, copy its content into your clipboard, and then paste in this field.

e. Optionally, below **Just-in-time provisioning for the account**, click **Allow** to enable just-in-time provisioning. The **Attribute name ...** field appears.

f. In **Attribute name in SAML for just-in-time role**, make sure it contains the same **memberOf** name that you defined in step **6** of the Microsoft Entra ID configuration. See *"Just-in-Time (JIT) Provisioning with Microsoft Entra ID SSO" below* below.

5. Click **Save**.

Selecting Users for SSO

1. In Microsoft Entra ID, select **Allow select users to authenticate directly with CloudGuard** to configure several users that can access CloudGuard directly, with their email address and password registered with CloudGuard. This is an important provision to ensure that there are users that can always log in locally and not be locked out of the portal in case of SSO misconfiguration.

   Note - To access data through API, you have to authenticate to CloudGuard directly.

2. Select the users that can access CloudGuard with local (CloudGuard) authentication credentials or with SSO.

3. Verify user access with SSO. Administrative users with enabled SSO must log in to Microsoft Entra ID and click the CloudGuard application icon to access CloudGuard.

Just-in-Time (JIT) Provisioning with Microsoft Entra ID SSO

If you configure the option in steps **4e** and **4f** of *"Configuring Single Sign-On in CloudGuard" on the previous page*, you can enable JIT provisioning for users in the CloudGuard SSO account. These users cannot have their registration to CloudGuard with a user name and password. JIT users must log in to Microsoft Entra ID and access the CloudGuard portal with SSO.

Configure CloudGuard SSO with Okta

Use the Okta Administrator Dashboard to add an application and view the values that are specific for your organization.

Then you can log in to CloudGuard as a super user, configure SSO and redirect the login requests to the Okta login page so that your administrative users can log in using SSO.

Okta Configuration for SSO

1. Log in to Okta as an Administrator, go to **Applications** and click **Create App Integration**.

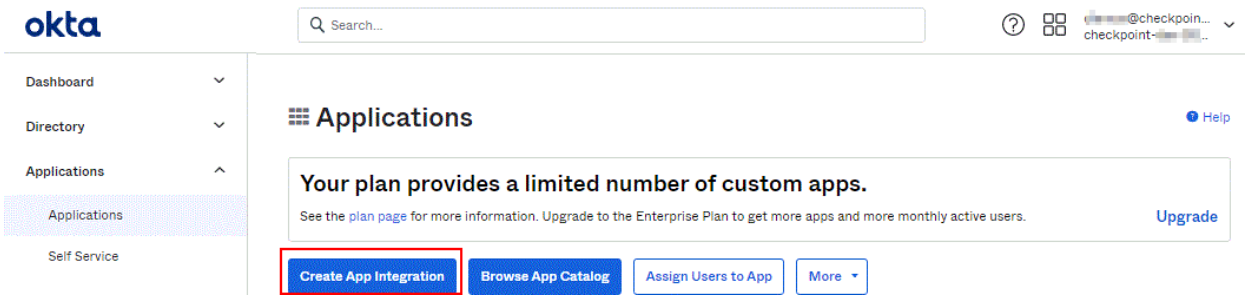

   The App integration wizard opens.

2. In the **Create a new app integration** wizard, for the Sign-in method, select **SAML 2.0** and click **Next**.

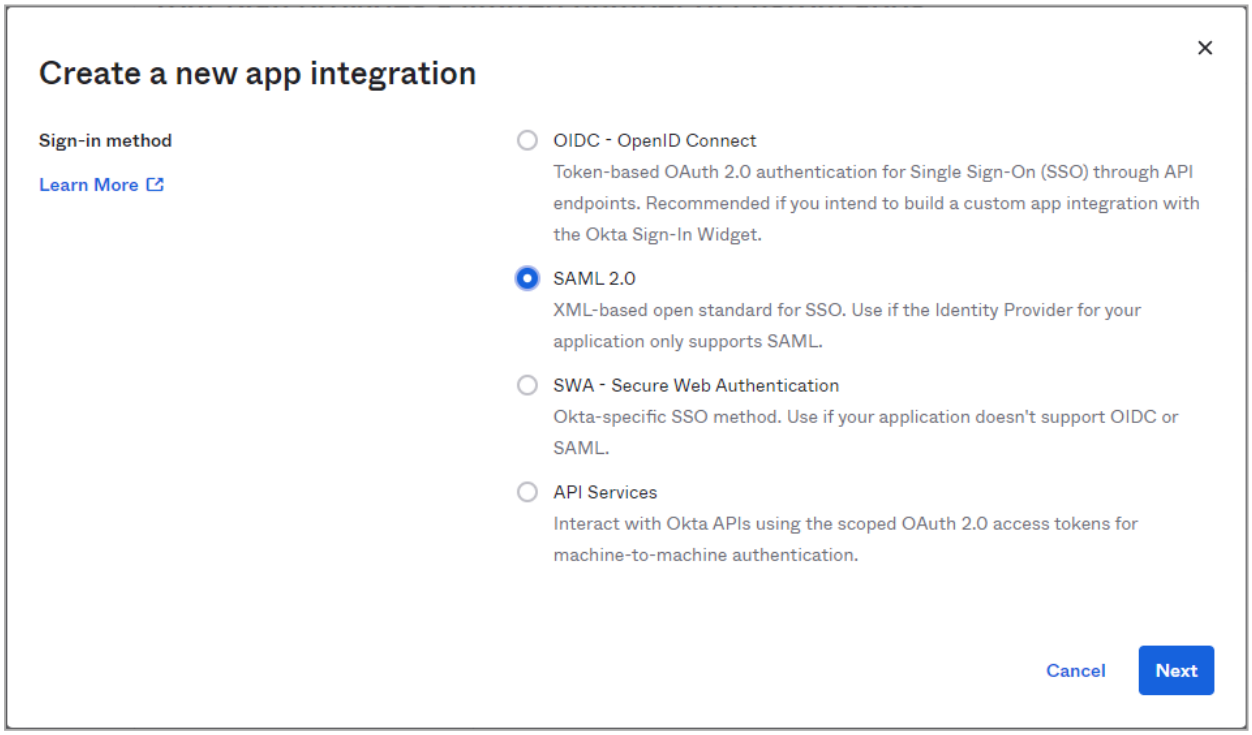

3. In the **General Settings** section of the **Create SAML Integration** window, set these values:

- **App name**: Select the app name, for example, CloudGuard
- **App logo**: Optionally, upload the CloudGuard logo
- Make sure to clear the options under **App visibility**

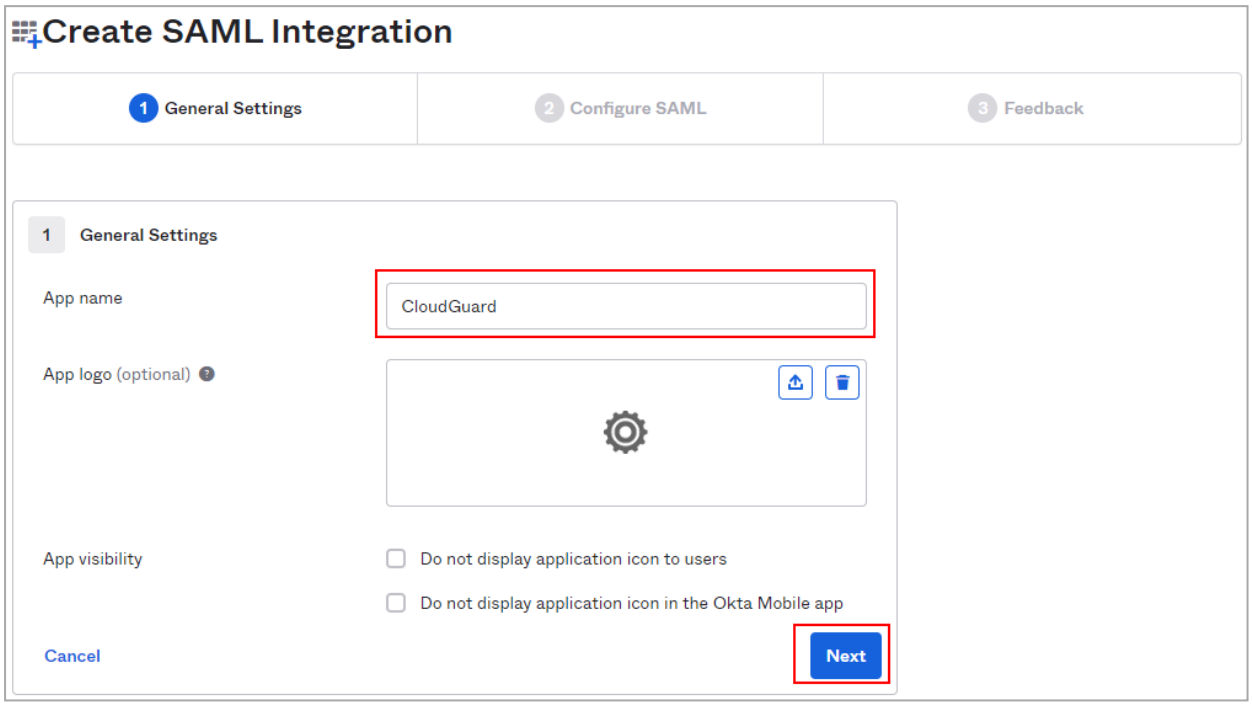

4. Click **Next**. The **Configure SAML** page opens.
5. In the **General** section of **SAML Settings**, set these values:
   - **Single sign on URL**: Enter **https://secure.dome9.com/sso/saml/***yourcompanyname*, where *yourcompanyname* is the Account ID string used in the CloudGuard SSO configuration. Make sure to select the option to **Use this for Recipient URL and Destination URL**.
   - **Audience URI (SP Entity)**: Enter **https://secure.dome9.com**
   - **Name ID format**: Select **EmailAddress**
   - **Application username**: Select **Okta username**
   - Leave default values for Advanced settings

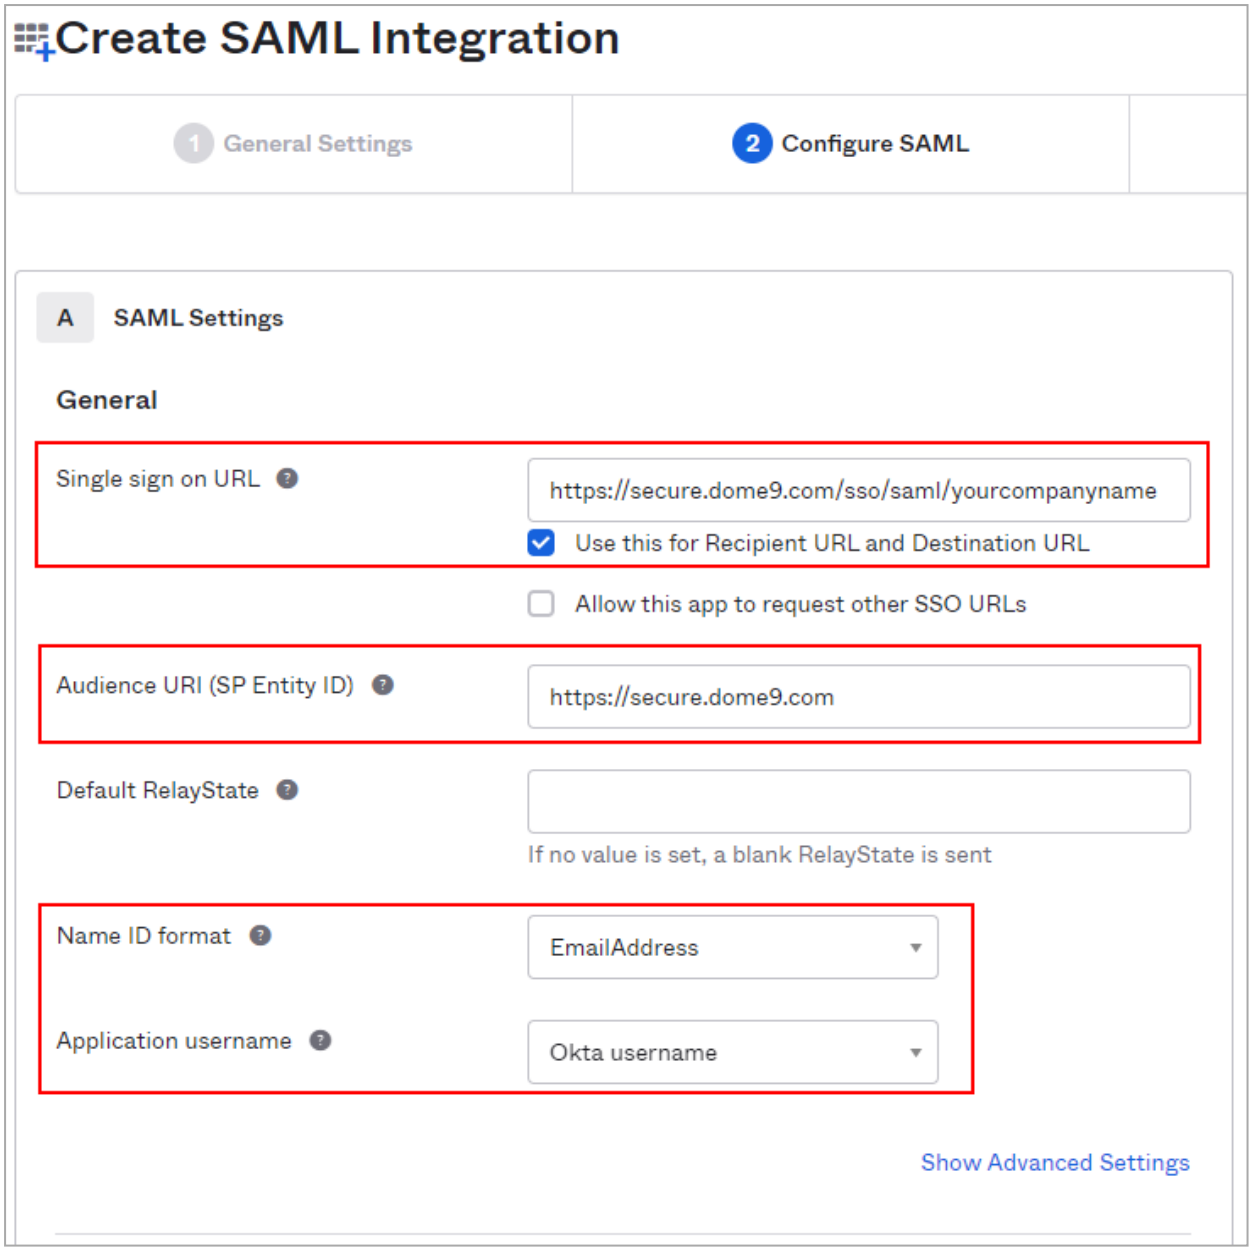

6. Click **Next**.
7. Click **Finish**.

   In your newly created application, complete the configuration of the SAML 2.0 settings.
8. On the **Sign On** tab of your Application, under **Settings**, find the SAML 2.0 section and click **View Setup Instructions**.

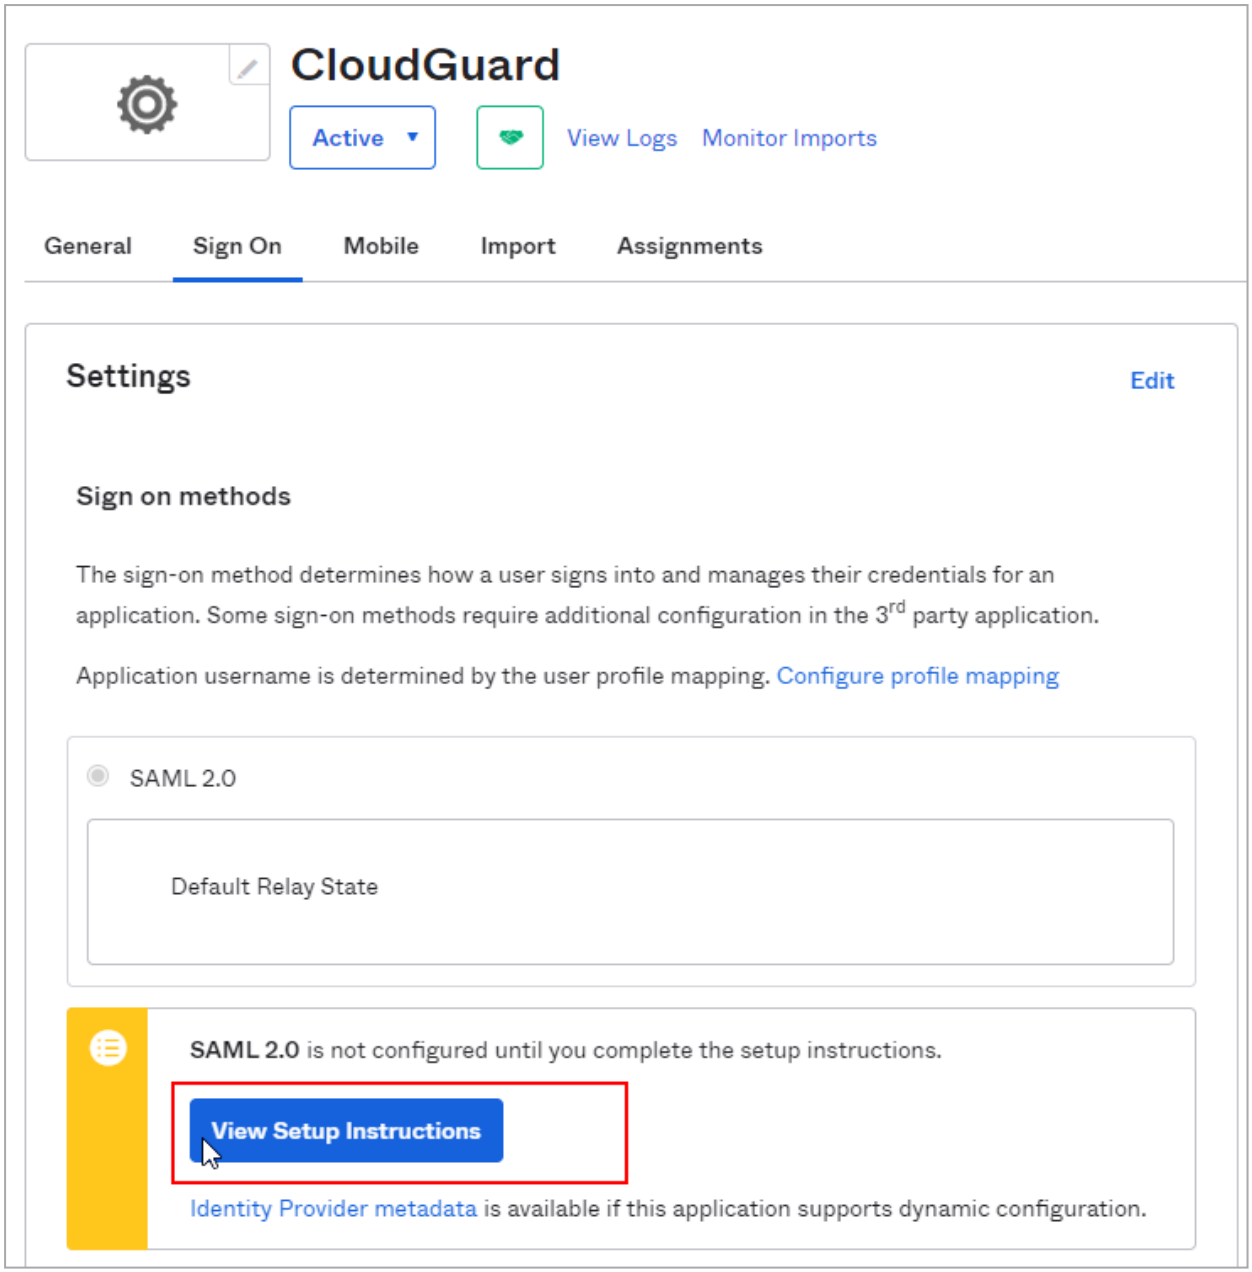

The Setup Instructions open in a separate window.

9. Copy and save for future use the values under 1 and 2: **Identity Provider Single Sign-On URL** and **Identity Provider Issuer**.
10. To get the **X.509 Certificate** under 3, click **Download certificate** and save the file on your computer.

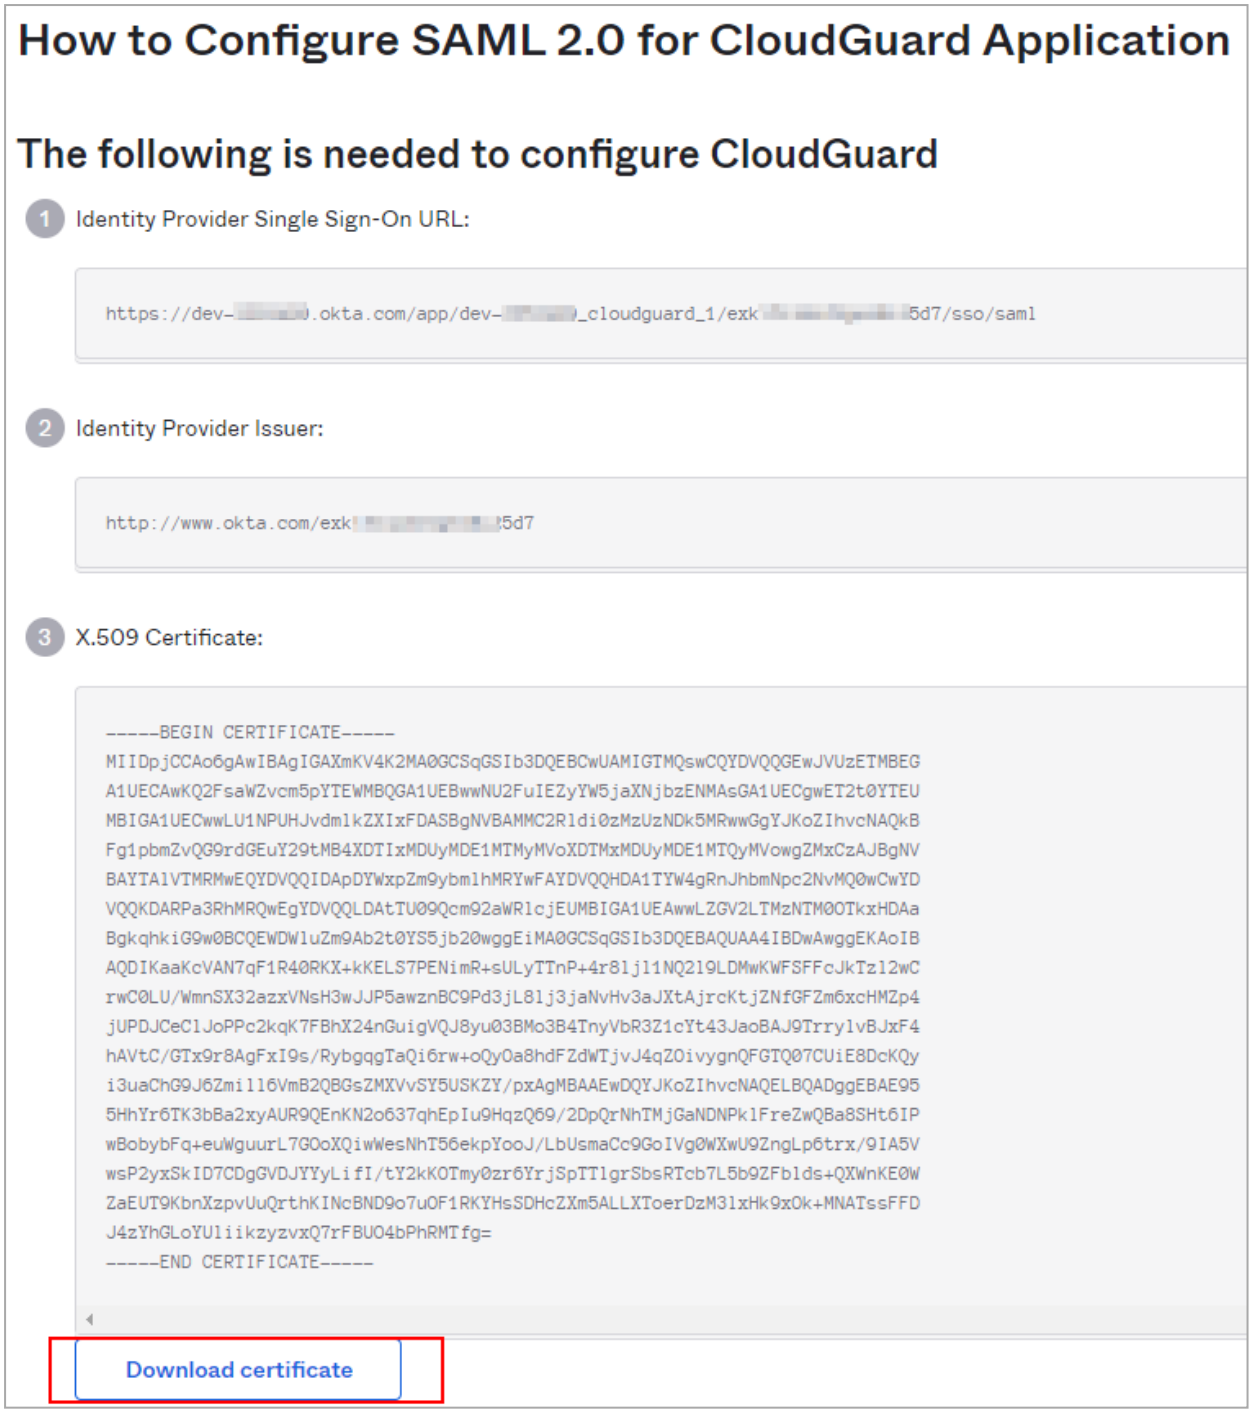

11. Go back to your application and open the **Assignments** tab to assign users or groups of users to the application. This enables the users to log in to CloudGuard with SSO.
12. Click **Assign** and select **Assign to People** to add individual users or **Assign to Groups** to add groups of users.

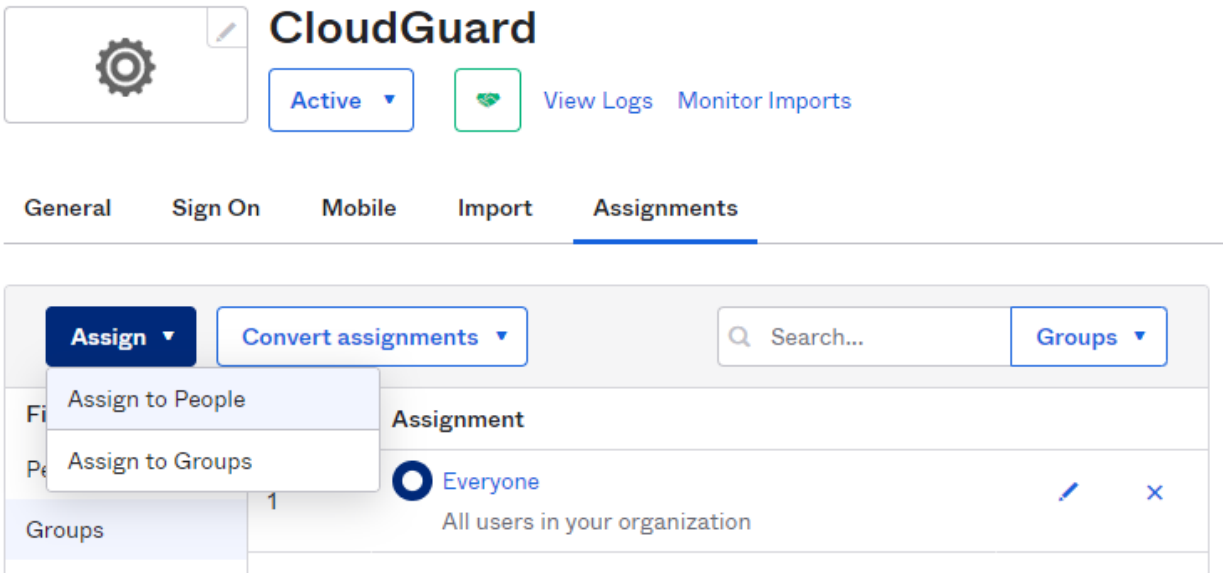

CloudGuard Configuration for SSO

1. Log in to CloudGuard with a super user account.
2. Navigate to **Settings** > **Security & Authentication**.
3. In the **SSO** section, select **Enabled**.
4. In the **SSO Configuration** window that opens, enter the below values and click **Save**.
   - **Account ID** - enter a name without spaces that can serve as your **company name identifier**
   - **Issuer** - enter the URL that you saved from the Setup Instructions, step 9 above, item 2 **Identity Provider Issuer**
   - **Idp endpoint url** - enter the URL that you saved from the Setup Instructions, step 9 above, item 1 **Identity Provider Single Sign-On URL**
   - **X.509 certificate** - paste the contents of the file that you saved from the Setup

Instructions, step 10 above, item 3 **X.509 Certificate**

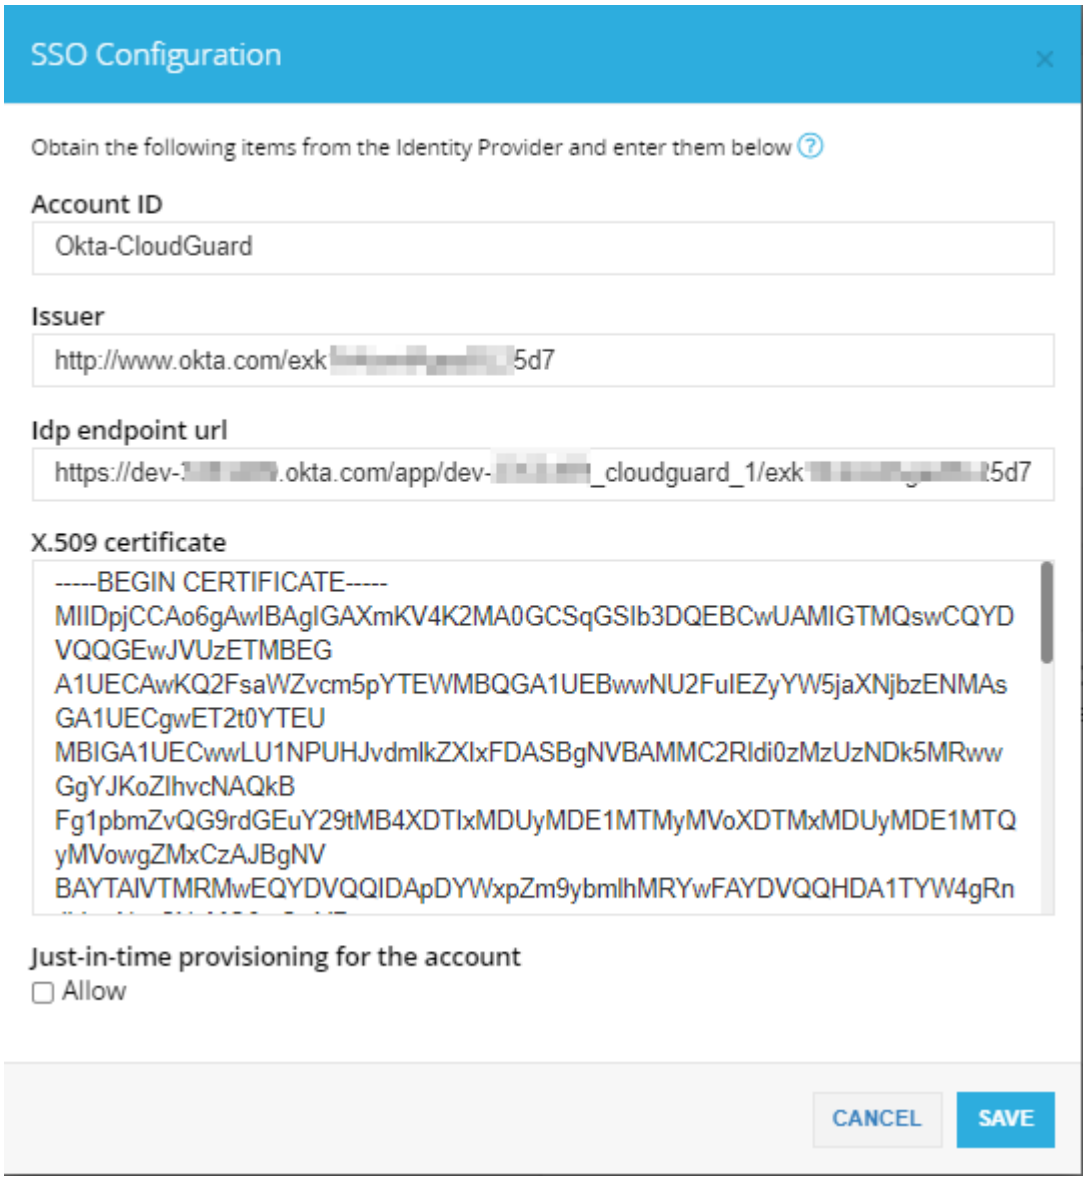

Note - SP-initiated flows and IdP-initiated flows are supported.

Configure SSO JIT Provisioning on Okta

Okta Configuration

1. In Okta, go to the Admin panel.

2. In the Directory menu, select **Groups**.

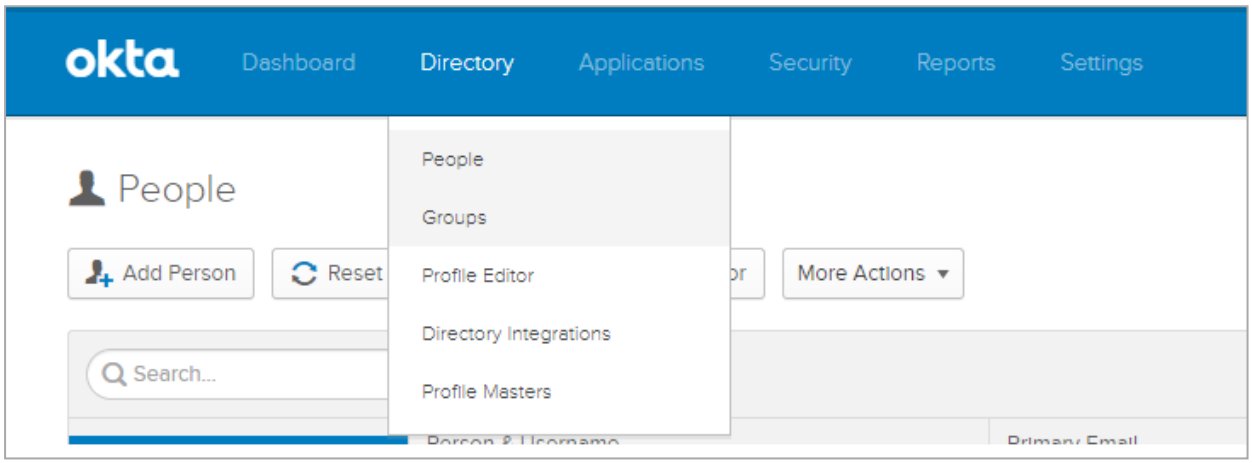

3. Click **Add Group**.
4. Enter a name and description for the group (remember the name as you will need it later), and the click **Add Group**.

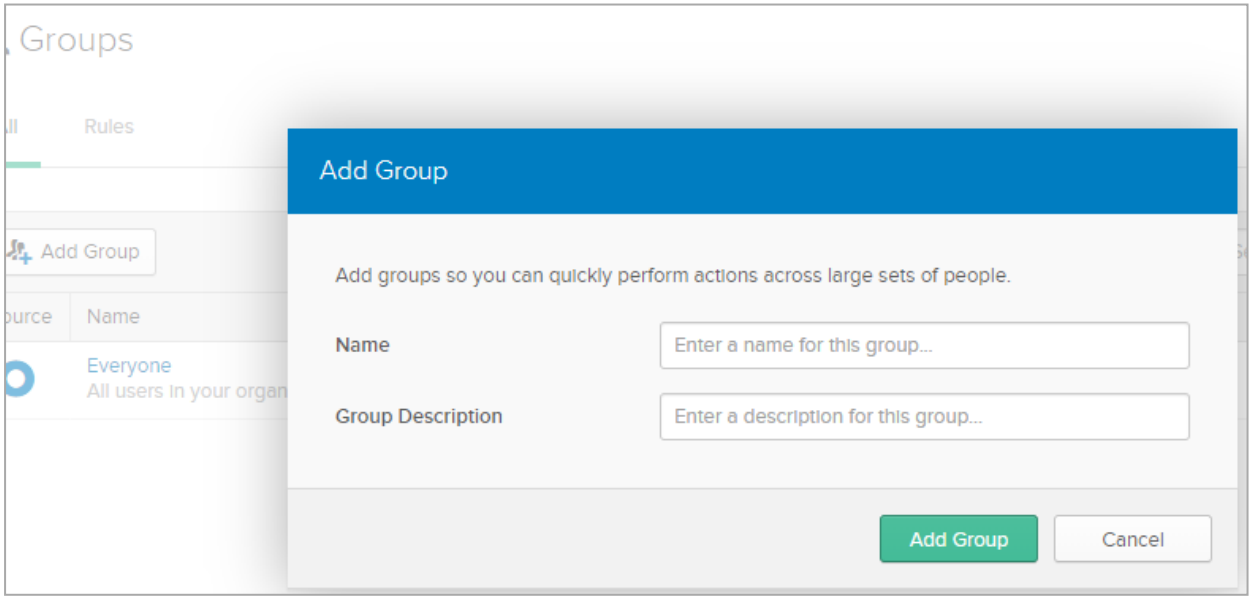

5. In the Application menu, select **Application**.

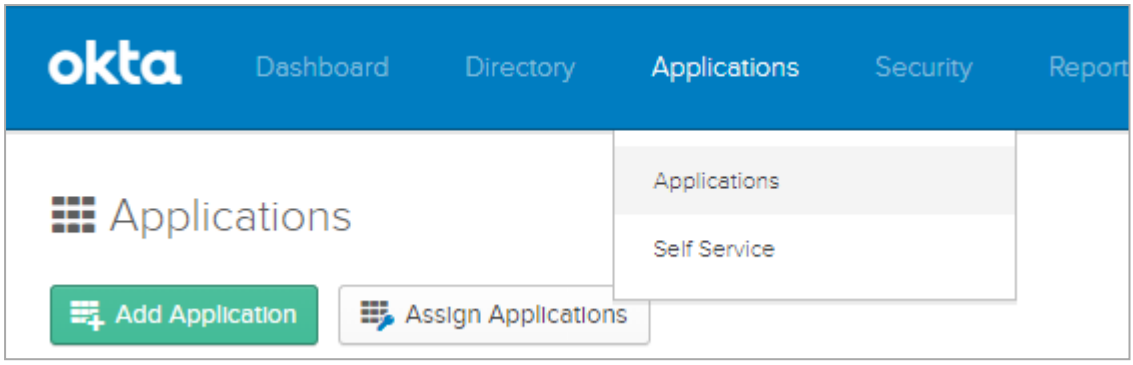

6. Click **Create New App**.

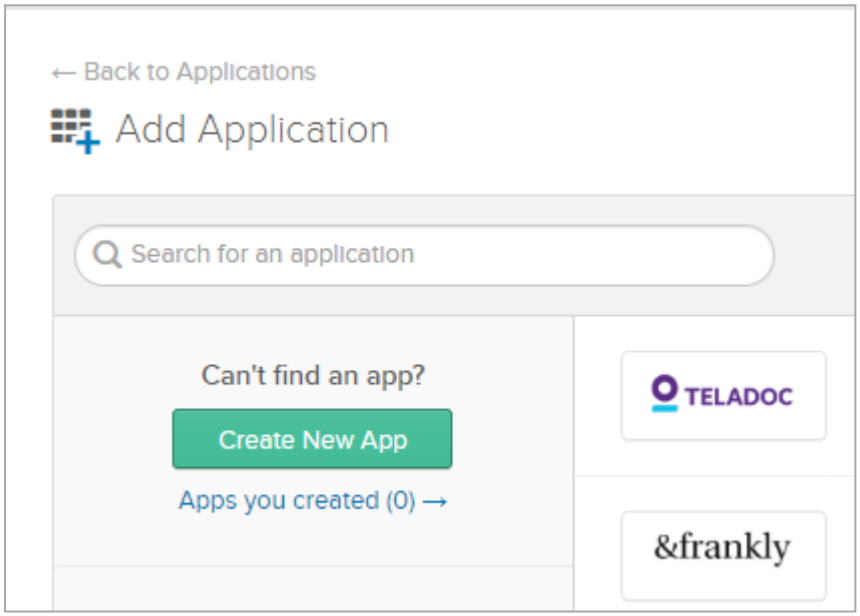

7. Select the following, and then click **Create**:

   Platform: Web

   Sign on method: SAML 2.0

8. Set the App name, then click **Next**.
9. Set the following parameters:

- The "Name-up-select" can be changed to any name.

- The Name in the "GROUP ATTRIBUTE STATEMENTS" (memberOf) can be set to any name you choose.

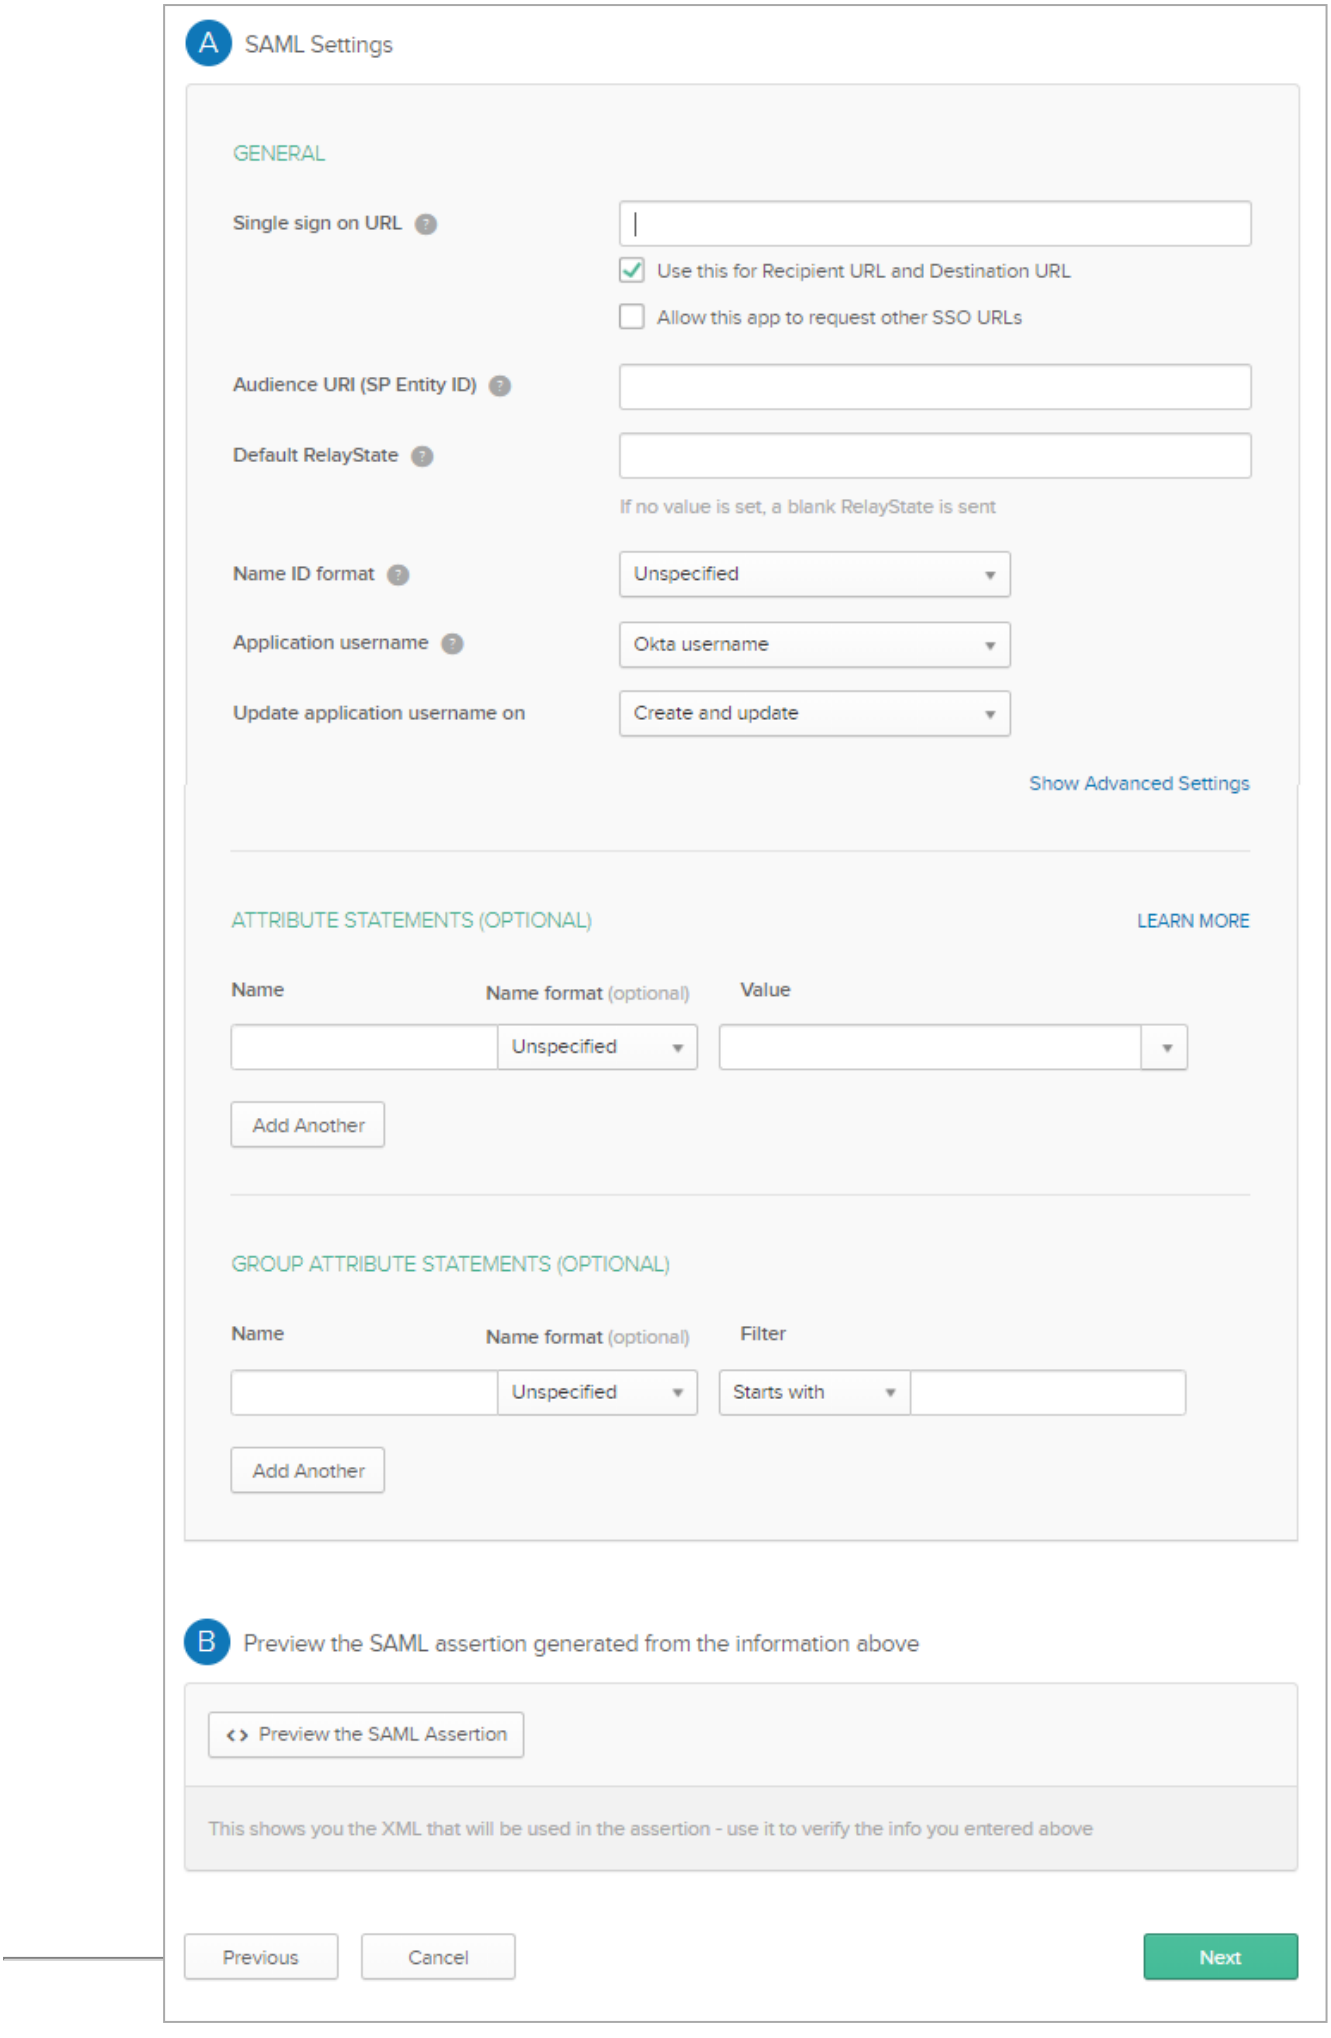
A
SAML Settings
GENERAL
Single sign on URL
Use this for Recipient URL and Destination URL
Allow this app to request other SSO URLs
Audience URI (SP Entity ID)
Default RelayState
If no value is set, a blank RelayState is sent
Name ID format
Unspecified
Application username
Okta username
Update application username on
Create and update
Show Advanced Settings
ATTRIBUTE STATEMENTS (OPTIONAL)
LEARN MORE
Name
Name format (optional)
Value
Unspecified
Add Another
GROUP ATTRIBUTE STATEMENTS (OPTIONAL)
Name
Name format (optional)
Filter
Unspecified
Starts with
Add Another
B
Preview the SAML assertion generated from the information above
Preview the SAML Assertion
This shows you the XML that will be used in the assertion - use it to verify the info you entered above
Previous
Cancel
Next

10. Click **Next** and then **Finish**.
11. Click **View Setup Instructions**.

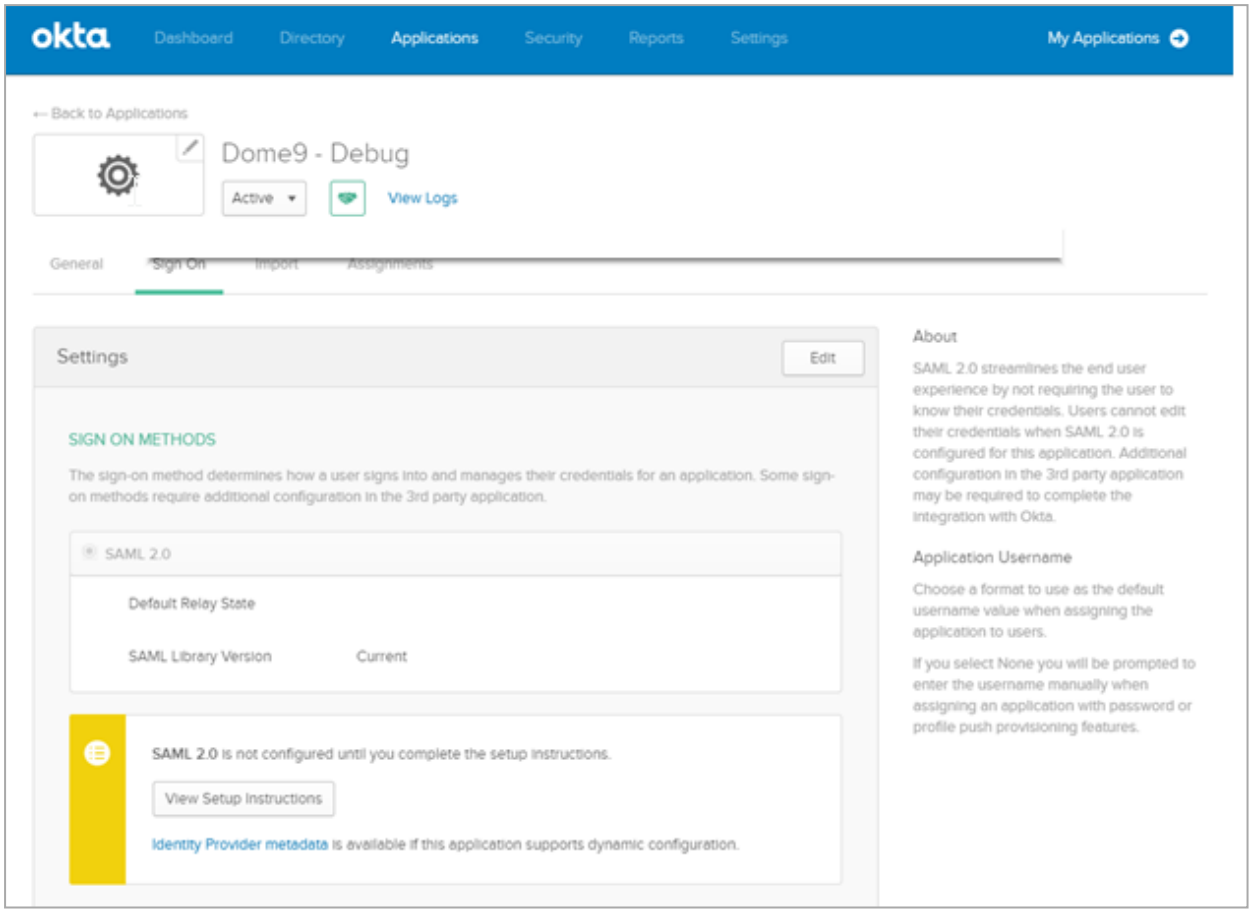

CloudGuard Configuration

1. In CloudGuard, navigate to the **Authentication** page in the **Settings**.
2. In the **SSO** section, click **Enabled**.
3. Click **Edit**, to open the **SSO Configuration** box.
4. Enter the following details:
   - **Account ID** - the value that you entered instead of "Name-up-select".
   - **Issuer** - the Identity Provider Issuer from Okta.
   - **Idp endpoint url** - the Identity Provider Single Sign-On URL from Okta.
   - **X.509 Certificate** - the X.509 Certificate from Okta.
   - Check **Just-in-time provisioning for the account** checkbox.

- **Attribute name in SAML for just-in-time role** - add the name that you entered instead of the "member Of"

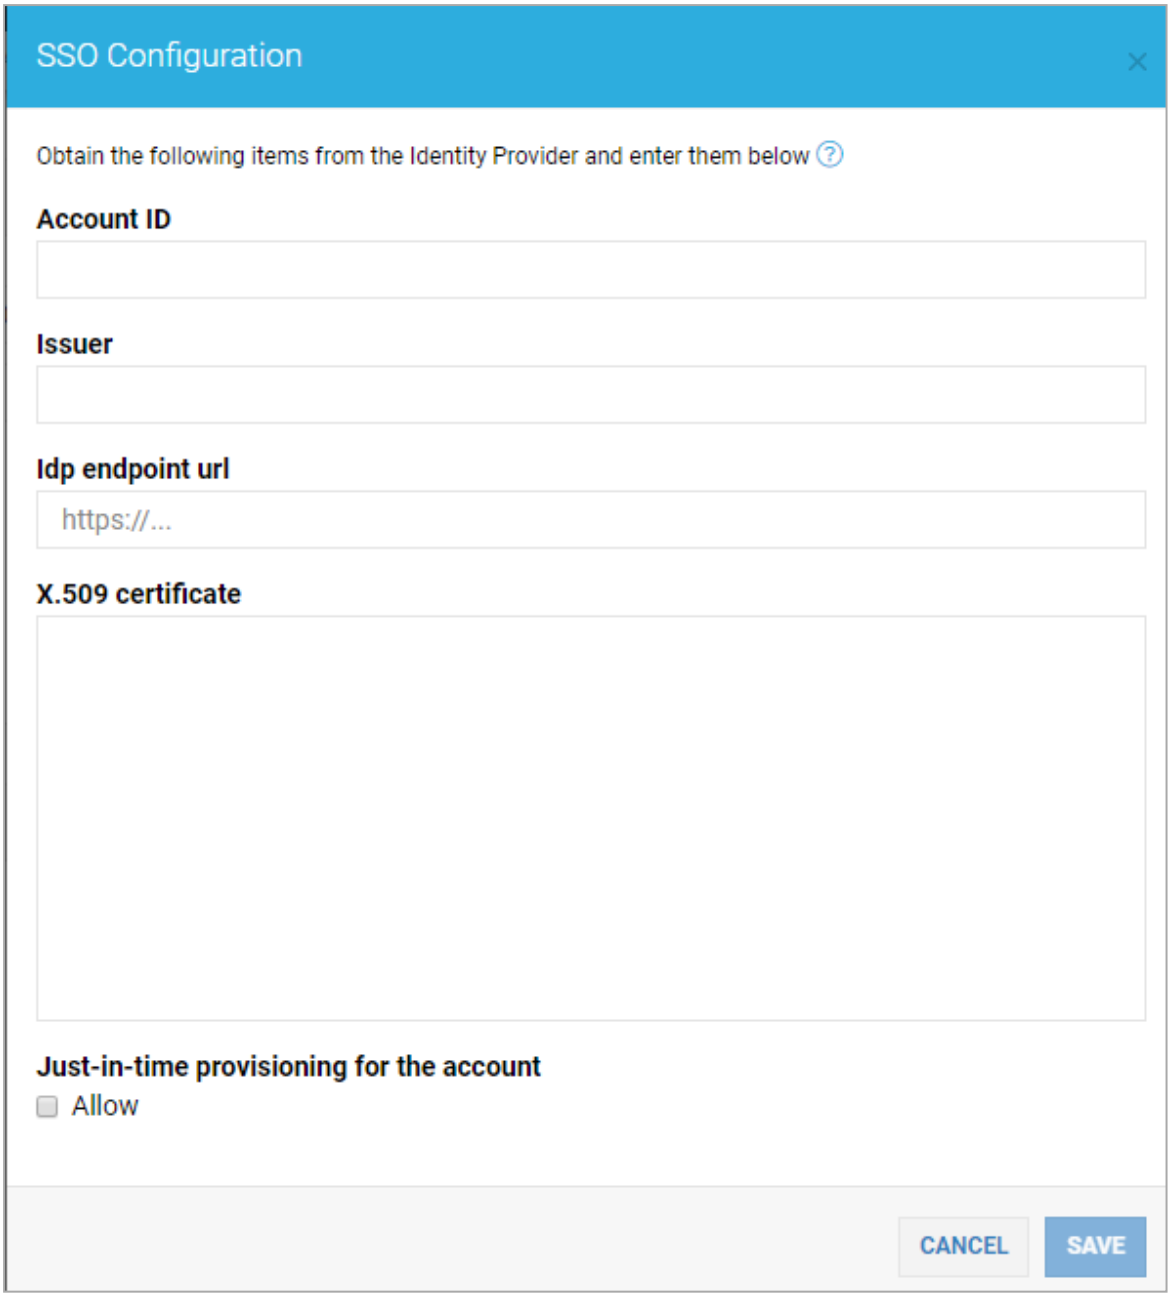

5. Click **Save**.
6. Assign the group that you created in step 4 to the application.
7. Navigate to the **Roles** page in the **Settings** menu.
8. Create a role with the same name as the name of the group that you created in Okta.

Configure CloudGuard SSO with JumpCloud

Based on JumpCloud documentation

Single Sign-On (SSO) with JumpCloud

Prerequisites

To successfully complete the integration between JumpCloud and CloudGuard, you must use an owner account in CloudGuard.

Notes:

1. CloudGuard does not support automatic new user provisioning via SSO. Prior to attempting SSO, all users must have a CloudGuard account that uses the same email as their JumpCloud account.
2. To prevent account lockout, CloudGuard does not allow the account owner to use single sign-on.
3. This instruction assumes that the JumpCloud administrator that performs the integrations understands the process of generating private keys in addition to public certificates. See below the generation of signed certificates on Linux as an example. For generation of keys on other operating systems, refer to these operating systems documentation.
   - Create a private key `opensslgenrsa -out private.pem 2048`
   - Create a public certificate for that private key:
     `openssl req -new -x509 -key private.pem -out cert.pem -days 1095`

To restrict access to a smaller group of users:

1. Notice the **IdP URL** name for this app in the Application details, for example, https://sso.jumpcloud.com/saml2/ **ConnectorName**.
2. Create a new Tag and name it **SSO-ConnectorName**. Important: This tag name is case sensitive.
3. Add users to this Tag who should be given access to CloudGuard via Single Sign-On. Any other users who are not in this tag will be denied access.

   Important - If the Tag does not exist, all users in your organization will be authorized to access CloudGuard.

Step 1: Configure CloudGuard for JumpCloud SSO

1. Log in to CloudGuard with a super user account.
2. Navigate to **Settings** > **Security & Authentication**.

3. In the **SSO** section, select **Enabled**.
4. In the **Account ID** field, enter a unique value (no spaces) that is later used to identify your company's SSO configuration with CloudGuard (your company name is a good value to use here) and copy this value.

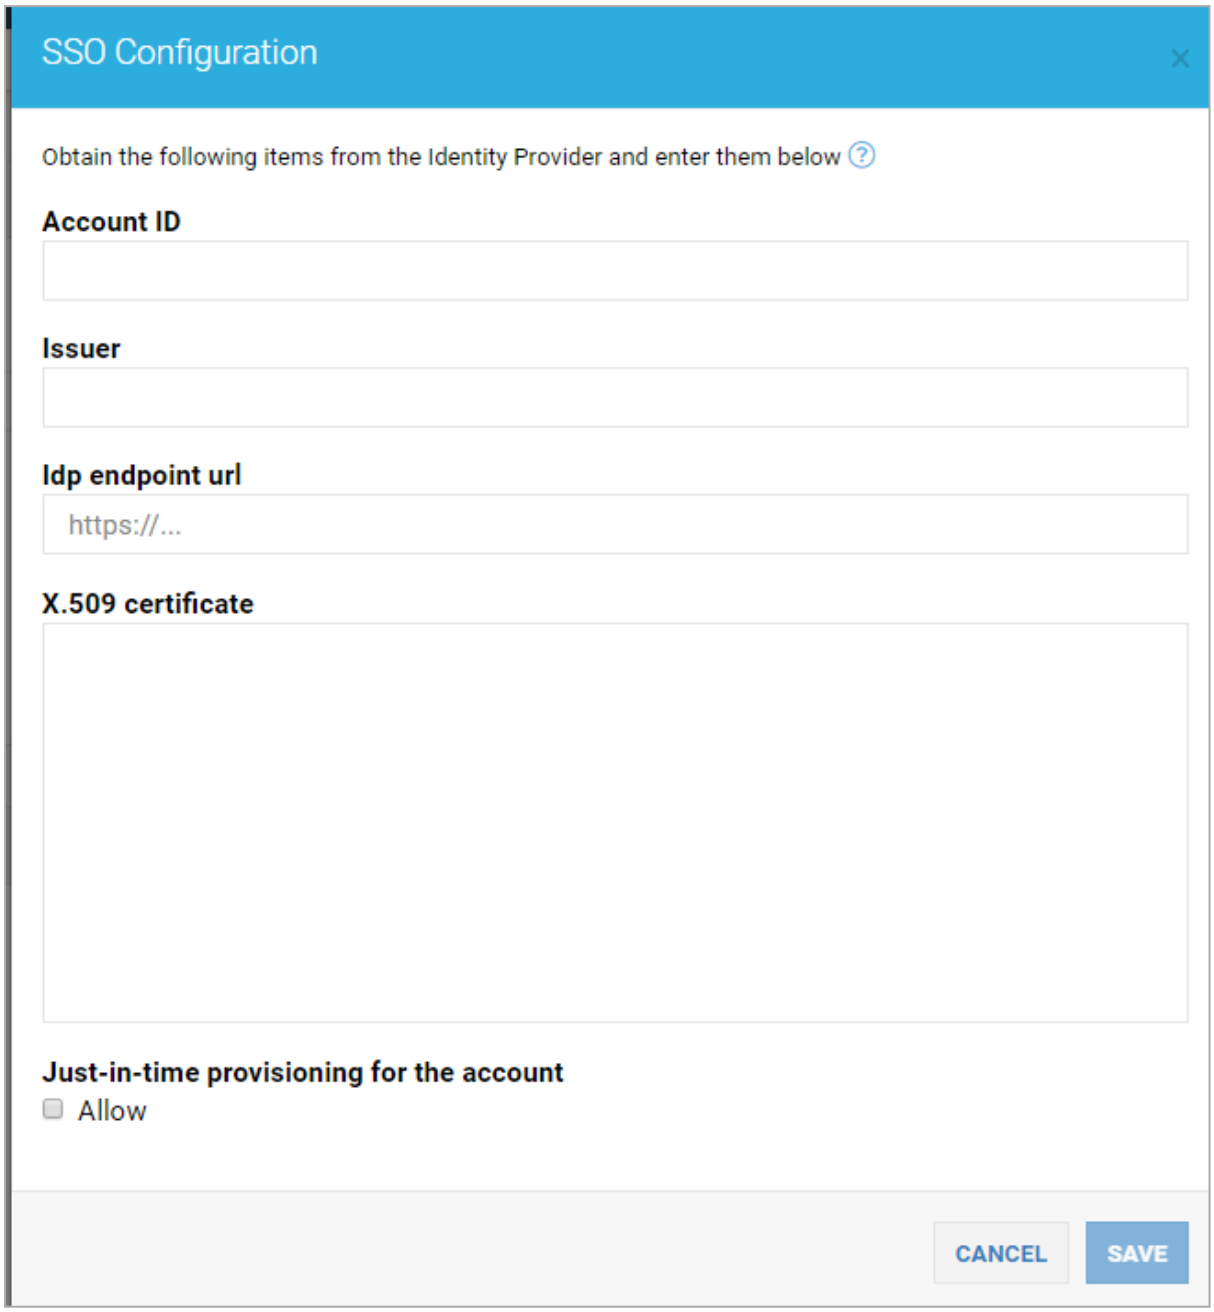

5. In the **Issuer** field, enter https://YOURDOMAIN.com (replace YOURDOMAIN with your company's unique domain).
6. In the **Idp Endpoint Url** field, enter https://sso.jumpcloud.com/saml2/dome9.
7. In the **X.509 Certificate** field, paste your entire **public certificate** (see Note 3 above).
8. Click **Save**.

9. Create a test user to test your configuration as appears in *"Adding a New User in the Dome9 Portal" on page 849*.
10. Fill in the necessary fields to create the user and ensure that **SSO User** is toggled to **On**.
11. To enable a pre-existing user to sign in via SSO, see *"Connecting a user to SSO for Dome9 accounts" on page 850*.

Step 2: Configure JumpCloud SSO for CloudGuard

1. Log into the JumpCloud Admin console at https://console.jumpcloud.com.
2. Click **Applications** in navigation pane on the left.
3. Click the green **+** icon in the upper left corner and find CloudGuard in the list.
4. Click **Configure**.
5. In the **IdP Entity ID** field, enter https://YOURDOMAIN.com/(this should be the same value that you entered in the Issuer field in CloudGuard, with "/" at the end).
6. Click **Upload Private Key** and upload your private key (see Note 3 above).
7. Click **Upload IdP Certificate** and upload your public certificate (see Note 3 above).
8. In the **ACS URL** field, enter https://secure.dome9.com/sso/saml/ACCOUNT_ID (replace ACCOUNT_ID with the value that you entered in the Account ID field in CloudGuard).
9. Click **Activate**.

Test the SSO Configuration

Case 1 - IdP-Initiated Flow

1. Log into the JumpCloud User Console with the email you used to create a test user in CloudGuard (or another email used by a CloudGuard account that does not have owner privileges, see Note 2 above).
2. Click the **CloudGuard** icon.
3. You should automatically be logged in to CloudGuard.

Case 2 - SP-Initiated Flow

1. In your Web browser, navigate to https://secure.dome9.com/sso/ACCOUNT_ID.
2. If necessary, log into the JumpCloud User Console as the appropriate user (see Note 2 above).
3. Now you automatically log in to CloudGuard.

Configure SSO using SAML from Google Workspace

Configure your CloudGuard account to use Single Sign-On (SSO) from Google Workspace using SAML

1. In the Google Workspace Admin console, navigate to SAML apps.
2. Click **+** to add a new service.
3. Click **SETUP MY OWN CUSTOM APP**.
4. Download the Certificate. We will use it in a later step.
5. Keep the Google Admin Console open on this page.
6. In a new tab, open the CloudGuard portal and navigate to **Settings** > **Configuration** > **Security & Authentication**.
7. In the **SSO** section, click **Enable**.
8. On the CloudGuard **SSO Configuration** page set the following:
   - **Account ID** – This can be any text you want.
   - **Issuer** - Copy the "Entity ID" field from Step 2 of the G Suite page and paste it here.
   - **Idp Endpoint URL** - Copy the value from the **SSO URL** field from Step 2 of the G Suite page and paste it here.
   - **X.509 certificate** - Using a text editor, open the certificate file you downloaded earlier and copy the full contents. Paste it in this field.
   - **Just-in-time provisioning for the account** – This option allows for CloudGuard users to be created and deleted when a Google Workspace user is created or deleted.
9. Click **Save**. After the page refreshes, it shows the enabled status.

   Leave this page open.
10. Switch back to Google Workspace Console and click **Next**.
11. Fill in the details as you like. These are details that users will see.
12. Click **Next**. Fill in the following fields.

- **ACS URL** - Copy this URL from the "Login Page" field of the CloudGuard SSO configuration. Add: */saml* after the */sso* (the full URL should look like this:

- **Entity ID** - This is always https://secure.dome9.com/
- **Name ID Format** - Change to **EMAIL**.

13. Continue to click **Next** until you are back at the SAML apps page.
14. Click on the newly created CloudGuard SAML app.
15. Click **Edit Service**.
16. Choose to turn ON/OFF for your organization (or specific groups).
17. Switch back to the CloudGuard portal and navigate to **Settings** > **User & Roles** > **Users**.
18. In the **Actions** menu bar, select **Connect to SSO** to enable the user to log in with SSO.

Notes:

1. When SSO is enabled, creating new users will enable SSO by default for the user.
2. Connecting a user to SSO will disable the normal login method for that user.
3. When disconnecting SSO from a user, the user will need to re-enable MFA in the CloudGuard portal (if MFA was originally used).

Log in with SSO

If a user has permissions in Google Workspace and in CloudGuard to use SSO, when logged in to their Workspace account, the user can click the menu in Google and select CloudGuard from the list of apps.

Configure CloudGuard SSO with Centrify

Please click on this link to learn how to configure SSO with Centrify

Configure CloudGuard SSO with Centrify

Configure CloudGuard SSO with OneLogin

OneLogin Configuration

1. On OneLogin, navigate to **Apps**, and select **Add App**.
2. Search for Dome9 (CloudGuard).
3. Click the app icon.
4. Enter a Display Name and click **Save**.
5. In the configuration page, enter a **Dome9 (CloudGuard) Client ID** and click **Save**.

CloudGuard Configuration

1. Log in to the CloudGuard portal, with a super user account.
2. Navigate to the **Authentication** page in the **Settings**.
3. In the **SSO** section, click **Enabled**.
4. Click **Edit** to open the **SSO Configuration** window.
5. Enter details for the SSO provider as follows:

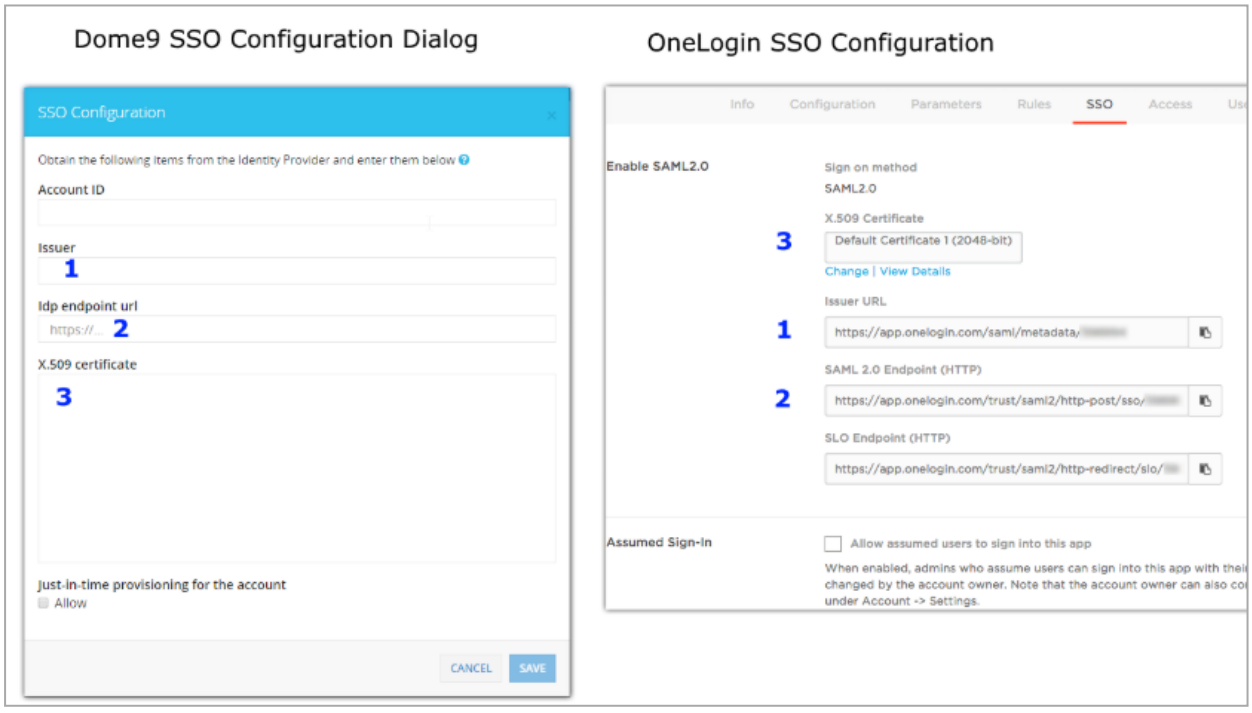

6. In the **Account ID**, enter the **CloudGuard Client ID**, created above in OneLogin (step 5).
7. Click **Save**.

To enable Just-In-Time provisioning, see the *"Configure SSO JIT Provisioning on OneLogin" on page 902*.

Configure SSO JIT Provisioning on OneLogin

OneLogin configuration

1. In OneLogin, in the **Applications** tab, add your app.
2. In the **Users** tab, add relevant users (previously configured).
3. In the CloudGuard portal, navigate to the **Roles** page in the **Settings** menu.
4. Create a new role.
5. In OneLogin, navigate to **Users -> Roles**.
6. Create a new role, with the same name as the role created in CloudGuard above.
7. Edit the new role.
8. Create a new app **SAML Test Connector (IdP)**.
9. Select **Apps** in the menu, and click **Add App**
10. Search for the newly created app (*SAML Test Connector (IdP)*)
11. Set the name and click **Save**.
12. In the **Configuration** tab, set the following:
    - **CloudGuard-onelogin-SSO** - set to the SSO account ID configured in CloudGuard
    - **RelayState** - https://secure.dome9.com
    - **Audience** - https://secure.dome9.com
    - **Recipient** - https://secure.dome9.com/sso/saml/CloudGuard-onelogin-SSO
    - **ACS URL** - https://secure.dome9.com/sso/saml/CloudGuard-onelogin-SSO
13. In the **Parameters** tab, click **Add parameter**.
14. In the **name** field, enter *memberOf* (or another name).
15. Click **Save**.
16. Navigate to the **Security & Authentication** page in the **Settings** menu.
17. In the **SSO** section, click **Enabled**.
18. Enter these details for the SSO configuration:
    - "Account ID" - enter the Value that you entered instead of "Name-up-select"
    - "Issuer" - enter the "Issuer URL" from OneLogin.

- "Idp endpoint url" - enter the Identity Provider Single Sign-On URL from OneLogin.
- "X.509 Certificate" - enter the X.509 Certificate from OneLogin.

19. Select the **Just-in-time provisioning for the account** option.
20. In **Attribute name in SAML for just-in-time role**, add the name that you entered instead of the *member Of*, above (step 14).
21. Click **Save**.

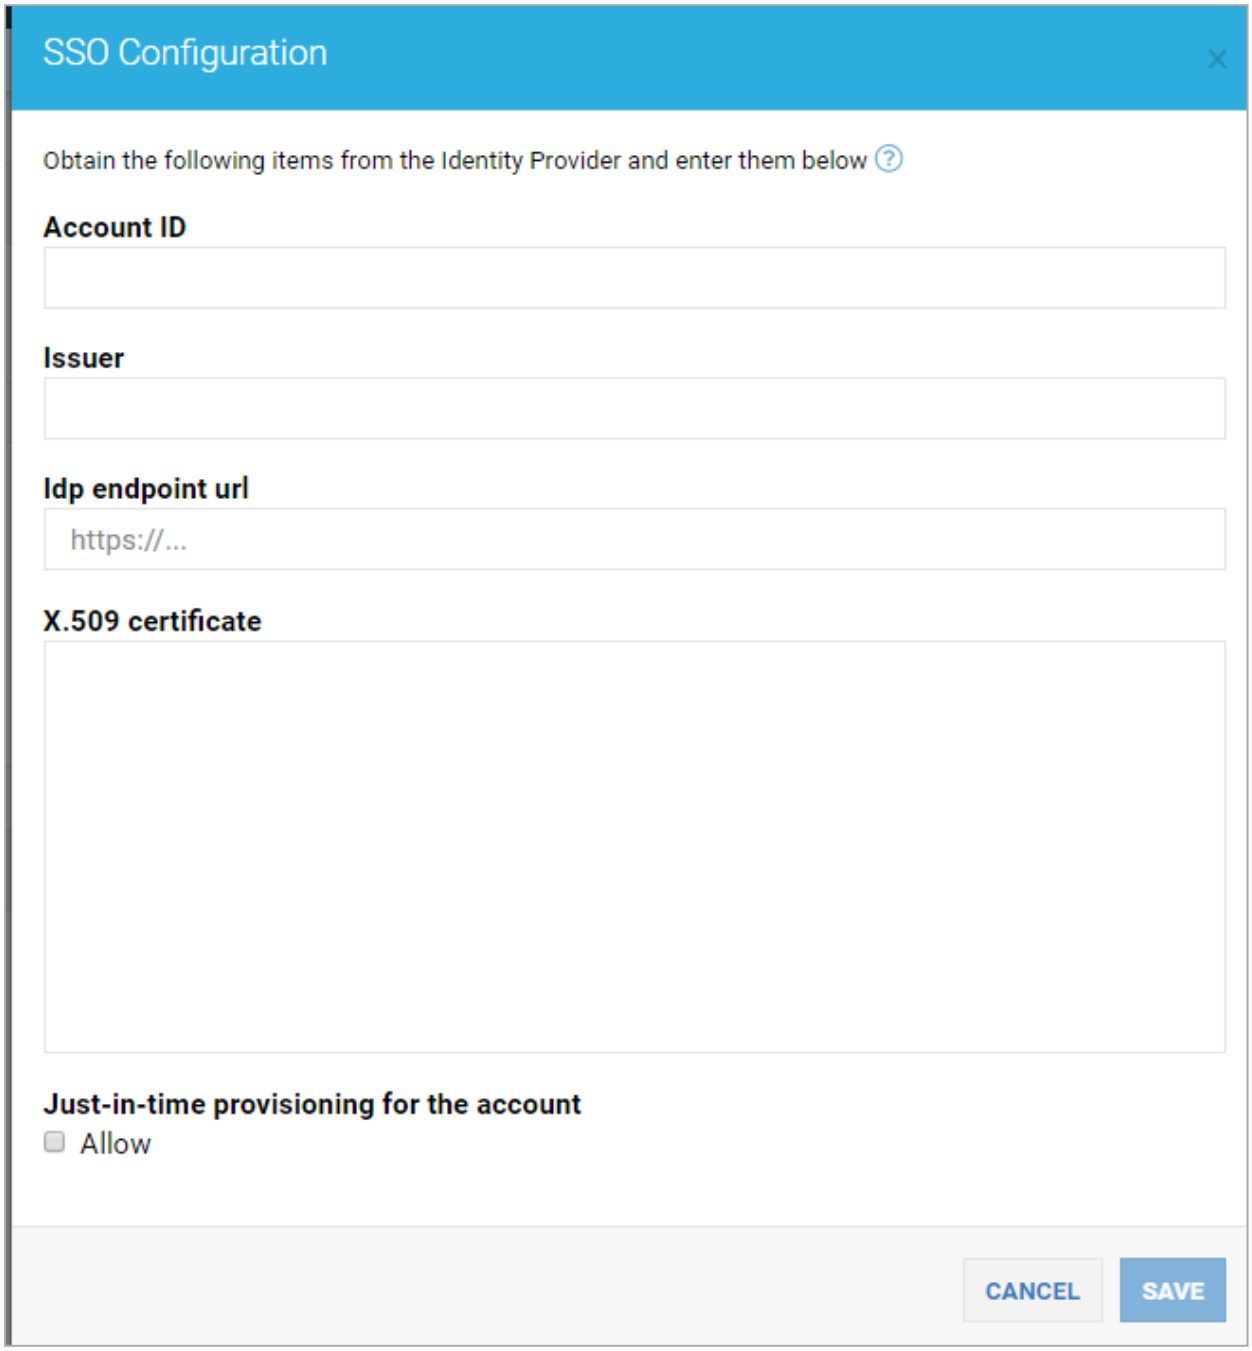

22. Navigate to the **Roles** page in the **Settings** menu.
23. Create a role with the same name as the name of the Role that you created in OneLogin.

24. If the mail address user for OneLogin is already known in CloudGuard, add another user in OneLogin, with the role from the previous step.

    Note - JIT Provisioning is created for a user who does NOT exist in CloudGuard, but belongs to a CloudGuard SSO account.

Configure CloudGuard SSO with a Generic / Custom Configuration

You can configure CloudGuard SSO with each Identity Provider that supports SAML 2.0 using a generic (custom) configuration. In this configuration, you create a custom SAML connection. Each SSO Identity Provider requires specific information for creation and configuration of the new connection. Usually, the required information differs by the Identity Provider.

To create a generic SAML connection, you need the Identity Provider Metadata URL that is available from your Enterprise customer's SAML instance.

Configuring the CloudGuard account for SSO

1. Log in to the CloudGuard portal with a super user account.
2. Navigate to the **Security & Authentication** page in the **Settings** menu.
3. Select **Enabled**.

   The **SSO Configuration** window opens.
4. Enter the details for the SSO provider as follows:
   - Account ID (For example, for OneLogin, the Account ID is the OneLogin Client ID for CloudGuard)
   - Issuer
   - IDP endpoint URL

SSO Configuration Dialog

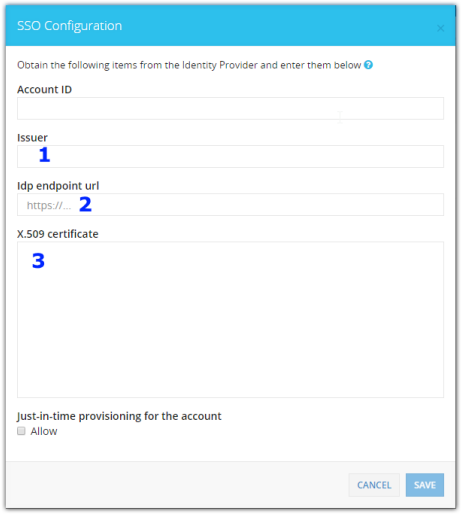

OneLogin SSO Configuration

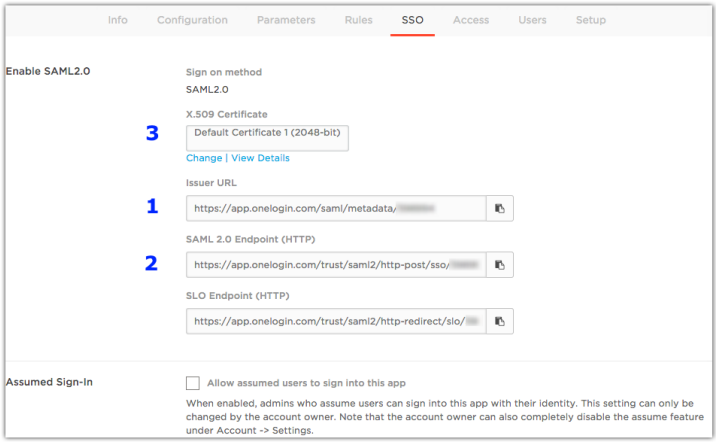

5. Copy the content of the IDP's X.509 certificate to the **X.509 certificate** text area in the CloudGuard **SSO Configuration** window. To view the certificate details in your IDP SSO setting application, click **View Details** below the certificate context and see the BASE64 representation of the certificate.
6. Click **SAVE**.

Configuring the Identity Provider Custom Connector

When you use an IDP custom application connector, specify these parameters:

1. Set the SSO URL / ACS URL to: **https://secure.dome9.com/sso/saml/yourcompanyname**, where **yourcompanyname** is the Account ID string used in the CloudGuard SSO configuration.
2. Set the Audience /Entity ID field to: **https://secure.dome9.com**.
3. Make sure the assertion element is signed but not encrypted. Encryption is handled by the transport layer.

   Below is an example of the Okta custom connector.

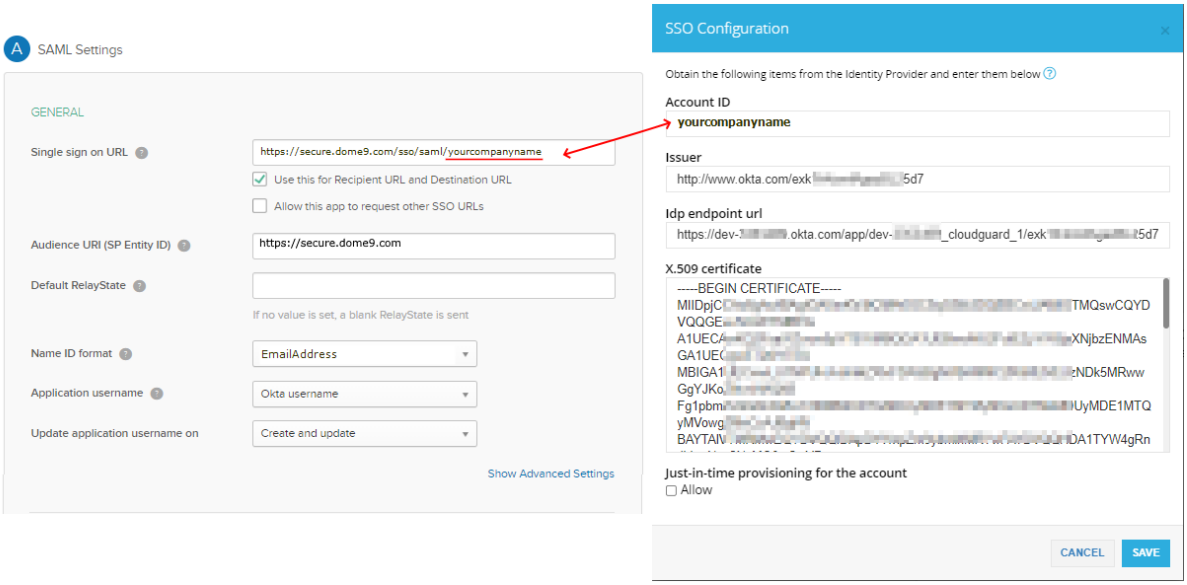

REST API

You can use CloudGuard programmatically, with the REST API, to onboard accounts, manage security groups, retrieve findings, and run posture assessments. Application developers can get access to these functions with RESTful HTTP requests. The resources and methods listed in the API include the CloudGuard functionality that the application developers need to onboard and manage the environments in CloudGuard.

The resource groups are CloudGuard resources and Inventory resources. CloudGuard resources include functional features (Intelligence, Compliance/Posture Management, Alerts), and entities (access leases, rulesets, and rules, CloudGuard users and roles). Inventory resources include entities such as Security Groups, instances, regions, and VPCs.

The API is based on HTTP requests and responses and uses JSON blocks.

The base URL for the CloudGuard API depends on your Data Center location (**Settings** > **Account Info**):

Location	For Dome9 accounts	For Infinity Portal accounts
United States	https://api.dome9.com/v2/	https://api.us1.cgn.portal.checkpoint.com
Europe	https://api.eu1.dome9.com/v2/	https://api.eu1.cgn.portal.checkpoint.com
Australia	https://api.ap2.dome9.com/v2/	https://api.ap2.cgn.portal.checkpoint.com
Canada	https://api.cace1.dome9.com/v2/	https://api.cace1.cgn.portal.checkpoint.com
India	https://api.ap3.dome9.com/v2/	https://api.ap3.cgn.portal.checkpoint.com
Singapore	https://api.ap1.dome9.com/v2/	N/A

To create an API key:

- in the Dome9 portal, see *"V2 API" on page 840*
- in the Infinity Portal, see *Infinity Portal Administration Guide* > Account Settings > API Keys.

For the full API reference guide, see https://docs.cgn.portal.checkpoint.com/reference. All examples in the reference guide use the base URL for the United States as a default.

Managed Service Providers

This section describes how MSPs and resellers can use the CloudGuard MSP Portal to create and manage customer accounts in CloudGuard Dome9 Portal.

To configure MSSP for the Infinity Portal accounts, see *Infinity Portal Administration Guide*.

MSP Portal

The CloudGuard MSP portal allows MSPs and CloudGuard resellers to create and manage their customers' CloudGuard Dome9 accounts.

MSPs and MSSPs

Managed Service Providers (MSP) provide services to business customers to manage their IT needs. This can include procurement, setup, and ongoing operational monitoring. For customers using cloud-based IT, or a cloud-based web presence, these services are in the cloud (on providers such as AWS and Azure). Resellers are MSPs that sell CloudGuard and cloud platform services to customers.

Managed Security Service Providers (MSSP), in addition, provide network security services to their customers. This could include configuring a secure network, monitoring their security posture, and responding to security events. These additional services can be applied to the cloud if the customer's presence is located there.

How CloudGuard can help MSPs for cloud-based computing

CloudGuard provides cloud security and compliance services for customers with a presence on AWS, Azure or GCP. This includes analysis of a customer's current security posture, or compliance, ongoing monitoring, and corrective actions to remedy problems.

MSPs that work through CloudGuard can offer these services on to their customers. Further, they can use different use-case models with their customers. These are described in the next section.

Resellers can use the CloudGuard MSP portal to create CloudGuard accounts for their customers.

Work Modes

MSPs can work with CloudGuard in different ways:

- In one model, the MSP provides a full service to the customer:
  - Create accounts for them with the cloud provider and with CloudGuard
  - Have full access to the customer's CloudGuard account to act on their behalf
  - Generate reports for them
- In another model, the MSP or Reseller provides the accounts and bills the customer for them, while the customer manages these accounts on their own.
- In a third model, a large enterprise, with a number of business units, can work with CloudGuard as an MSP. Each individual business unit is a separate account. Each account is managed by one overall MSP account (in either of the above models).

The MSP can build this flexibility in managing customer accounts into their business and pricing models. They can offer full services for customers who do not want to be bothered with the day-to-day management of their network, or reduced services for customers who do want to manage their own account, but yet want to procure all their services from one provider. They can also choose which of the CloudGuard services they wish to offer their customers.

CloudGuard Account Types

There are two types of CloudGuard accounts that you can manage:

- **Enterprise** accounts are for regular CloudGuard enterprise customers.
- **Reseller/MSP/Distributor** accounts are for customers that have and manage enterprise customers of their own. You must sign-on to the portal with this type of account.

Roles

An MSP account itself is a Super User, with full permissions over itself and its enterprise accounts. You can sign-on to a managed account in CloudGuard with any of the roles that are defined for the account, including as a Super User. You can define roles with specific permissions for each account (see *"Adding a Role" on page 850*).

Cross-Account Trust Capability

When you, as an MSP, create enterprise accounts for your customers, you can choose to allow the MSP account to sign-on to them and assume a role on them. This is called cross-account trust (also referred to as federated access), and it allows you, as the MSP, to actively manage customer accounts in CloudGuard.

Use the MSP Portal

This section explains how to use the CloudGuard MSP portal with Dome9 accounts.

Signing in to the MSP Portal

You must use an MSP account to sign-on to the MSP portal. Contact *Check Point Support Center* to change your account to an MSP type.

To sign in to the MSP portal:

1. Sign in to the CloudGuard portal (secure.dome9.com) with your MSP account name and password.
2. To redirect to the MSP portal, change the URL in the browser address bar to the MSP address according to your account region:
   - United States - msp.dome9.com
   - Europe - msp.eu1.dome9.com
   - Australia - msp.ap2.dome9.com
   - Canada - msp.cace1.dome9.com
   - India - msp.ap3.dome9.com
   - Singapore - msp.ap1.dome9.com

Actions

You can perform the following actions from the CloudGuard MSP Portal, or from the CloudGuard Console. Some actions can also be performed using the CloudGuard API.

Viewing your Accounts

The home page of the MSP portal shows all your accounts. Enterprise accounts are grouped under their MSP account (the top row, an MSP account, is the MSP account you with which you signed on). For each account you can see the CloudGuard modules selected for it, as well as the current number of users, and billable items.

On the left is a list of all distinct account names including the one you signed in with. Select one of these names to filter the list to show only these accounts.

Adding a Customer Account

1. To add a CloudGuard customer account, you must select one of the MSP accounts in the list (the top row is the account with which you signed on to portal). The new account will be managed by the selected MSP account, and will appear beneath it in the list.
2. Click **Add Account**.
3. In the pop-up window, select the type of account and then fill in the remaining details for the customer, including the email (which is used as the CloudGuard sign-on name).
4. Select the CloudGuard modules that the customer account can use (from Network, IAM Safety, and Posture Management).
5. Select whether the account has Enterprise capabilities and whether the account has FIM capabilities.
6. Select **Trust** if you (the MSP) want to be able to access (sign-on) the customer's account and act on their behalf in CloudGuard (see cross-account trust below).
7. Select the number of CloudGuard users for the account (or select UNLIMITED).
8. Click **Save** to add the account.

   It appears later in the list of accounts. A message is sent to the email address you entered.
9. Open the email message and follow the link to activate the new account.

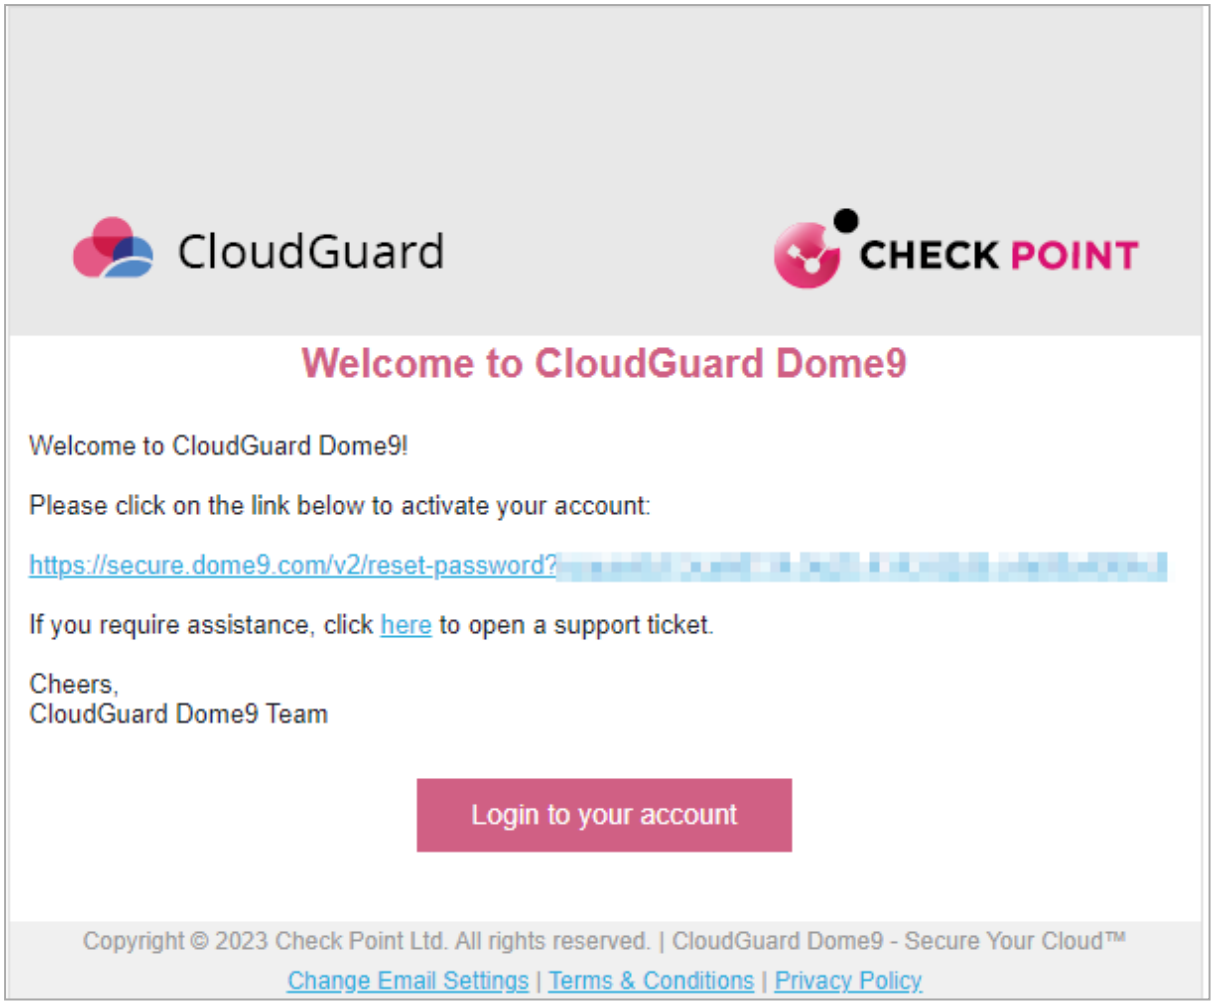

10. Enter a password for the account.

Changing Customer Account Details

This changes details for a CloudGuard customer account.

1. On the portal home page, click the menu ⋮ at the right on the line for the account you wish to change, and select **Edit account**.
2. In the pop-up window, change any of the details for the account, as necessary. You can change the plan (type of account), name, and modules for the account.
3. Click **Save** to save the changes for the account. The list of accounts on the portal home page will show the updated details for the account.

Deleting a Customer Account

This action deletes a CloudGuard customer account. It does not delete any cloud environments associated with it.

1. On the portal home page, click the menu ⋮ at the right on the line for the account you wish to change, and select **Delete account**.
2. Confirm the deletion; the account is deleted and removed from the list on the home page.

Exporting MSP information

You can export information for managed accounts to a CSV file.

1. Select a view of the managed accounts.
2. Click **Export to CSV** in the upper right and then select whether to export actual data or average data.

Switching to the Customer Account on CloudGuard

Connect to the CloudGuard portal with your MSP account and then switch to one of your managed accounts:

1. Sign in to the CloudGuard portal (secure.dome9.com) with your MSP account username and password.
2. Open the user option menu from the top bar, and select **Switch role** > **More**.
3. In the Switch Role window, select one of the listed accounts (these are your managed accounts) and then, from the adjacent list, select a role. Click **Switch Role**.

   This connects you to CloudGuard in the selected account and role. Your account appears shaded in the upper right of the screen to indicate that you have switched accounts.
4. To switch back to your original CloudGuard account, open the user option menu again and select **Back to** [name_of_your_account].

Using the CloudGuard API

You can establish a cross-account trust relationship between an MSP account and a customer account using the CloudGuard API instead of the MSP Portal. The accounts (one of them an MSP account) must be created first.

Configuring a Cross-account Trust Relationship Between Accounts

This procedure establishes a cross-account trust relationship between an MSP account and one or more customer accounts.

On the MSP account:

1. In the CloudGuard portal for the MSP account, select **Account Info** in the **Settings** menu.
2. Select the **Cross Account Access** tab.
3. Click **GENERATE**. This generates an account ID. Save the value for use in the next step.

For each customer account:

1. In the CloudGuard portal for the customer account, select **Settings** in the menu.
2. Create an API key as described in *"Credentials" on page 840*. This generates a unique API Key and Secret. Copy the value of the secret; it cannot be displayed again when you close the window.
3. Use this AccountTrust method in the API to establish the cross-account trust as in the example below:

Cross-account trust

```
curl -X POST --user '<api-key-id>:<api-key-secret>' -H "Content-Type: application/json" -d '{ "sourceAccountId": "<cross-account-identifier>", "description": "Grant access for MSP account", }' "https://api.dome9.com/v2/AccountTrust"
```

Where:

- **api-key-id** and **api-key-secret** are the API Key and secret, generated in the previous step
- **cross-account-identifier** is the account ID generated for the MSP account (above).

Configuring Role Restrictions for Cross-Account Trust

You can configure access to a customer account for specific roles only. Use this, for example, if the MSP will access the customer account with restricted permissions.

Add the following snippet to the method:

"restrictions": { "roles": ["Role1","Role2"]}}

This allows the MSP account to connect only as Role1 or Role2 (the specific role is selected when the MSP signs in to the account).

The URL would then appear like this:

Cross-account with restrictions

```
curl -X POST --user "<api-key-id>:<api-key-secret>" -H "Content-
Type: application/json" -d '{ "sourceAccountId" : "<cross-
account-identifier>", "description" : "Grant access for MSP
account", "restrictions": { "roles": ["Role1","Role2"]}}'
https://api.dome9.com/v2/AccountTrust
```

Frequently Asked Questions

Intelligence

I have a firewall on the storage that holds my logs. Which IP or IP range Intelligence uses that I can allow in the firewall configuration?

You can allow these IP addresses based on your Data Center location (appears in **Settings** > **Account Info** > **Data Center**):

- For US Data Center: **34.197.175.194**
- For EU Data Center: **54.228.172.67**
- For Canada Data Center: **15.222.80.139**
- For Mumbai (Asia-Pacific) Data Center: **13.233.191.132**
- For Singapore (Asia-Pacific) Data Center: **54.169.243.71**
- For Sydney (Asia-Pacific) Data Center: **52.64.4.254**

Can CloudGuard provide a cross-account traffic pattern analysis and view for two VPCs in different accounts?

Yes. These two accounts must be onboarded to CloudGuard, and the VPCs must have a peering connection. For more information on VPC peering, see AWS documentation.

CloudGuard Connectivity

Which CloudGuard endpoints do I have to allow on my network?

The endpoints depend on your Data Center location (region). See below the URLs that you must allow in your environments for network traffic to transit your security perimeter.

United States (US)

M odule	Dome9 Portal URL	Infinity Portal URL
Portal	https://secure.dome9.com	https://portal.checkpoint.com
API server	https://api.dome9.com https://api-cpx.dome9.com	https://api.us1.cgn.portal.checkpoint.com

M odule	Dome9 Portal URL	Infinity Portal URL
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com	
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://us-gw.sg.iaas.checkpoint.com	
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query	
Container Registry	https://quay.io/checkpoint	
Code Security Spectral	https://spectral-us.dome9.com/ https://dl.spectralops.io/	https://spectral-us.checkpoint.com/
Code Security Spectral OSS	https://eiffel.spectralops.io/	
Intelligence	https://magellan.dome9.com	https://magellan.us1.cgn.portal.checkpoint.com
Kubernetes Intelligence	https://validator-prod-k8s.s3.amazonaws.com	
MSP	https://msp.dome9.com	-

Europe (EU)

Module	Dome9 Portal URL	Infinity Portal URL
Portal	https://secure.eu1.dome9.com	https://portal.checkpoint.com
API server	https://api-cpx.eu1.dome9.com https://api.eu1.dome9.com	https://api.eu1.cgn.portal.checkpoint.com
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/	

Module	Dome9 Portal URL	Infinity Portal URL
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://eu-gw.sg.iaas.checkpoint.com	
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query	
Container Registry	https://quay.io/checkpoint	
Code Security Spectral	https://spectral.eu1.dome9.com/ https://dl.spectralops.io/	https://spectral-eu.checkpoint.com/
Code Security Spectral OSS	https://eiffel.spectralops.io/	
Intelligence	https://webserver.logic.eu1.dome9.com	https://webserver.logic.eu1.cgn.portal.checkpoint.com
Kubernetes Intelligence	https://validator-prod-533924475734-k8s.s3.eu-west-1.amazonaws.com	
MSP	https://msp.eu1.dome9.com	-

Australia (AU)

Module	Dome9 Portal URL	Infinity Portal URL
Portal	https://secure.ap2.dome9.com	https://ap.portal.checkpoint.com
API server	https://api.ap2.dome9.com https://api-cpx.ap2.dome9.com	https://api.ap2.cgn.portal.checkpoint.com

Module	Dome9 Portal URL	Infinity Portal URL
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/	
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://au-gw.sg.iaas.checkpoint.com	
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query	
Container Registry	https://quay.io/checkpoint	
Code Security Spectral	https://dl.spectralops.io/	
Code Security Spectral OSS	https://eiffel.spectralops.io/	
Intelligence	https://webserver.logic.ap2.dome9.com	https://webserver.logic.ap2.cgn.portal.checkpoint.com
Kubernetes Intelligence	https://validator-prod-583664506098-k8s.s3.ap-southeast-2.amazonaws.com	
MSP	https://msp.ap2.dome9.com	-

Canada (CA)

Module	Dome9 Portal URL	Infinity Portal URL
Portal	https://secure.cace1.dome9.com	https://ca.portal.checkpoint.com
API server	https://api.cace1.dome9.com https://api-cpx.cace1.dome9.com	https://api.cace1.cgn.portal.checkpoint.com

Module	Dome9 Portal URL	Infinity Portal URL
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/	
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://ca-gw.sg.iaas.checkpoint.com	
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query	
Container Registry	https://quay.io/checkpoint	
Code Security Spectral	https://dl.spectralops.io/	
Code Security Spectral OSS	https://eiffel.spectralops.io/	
Intelligence	https://webserver.logic.cace1.dome9.com	https://webserver.logic.cace1.cgn.portal.checkpoint.com
Kubernetes Intelligence	https://validator-prod-052001227150-k8s.ca-central-1.amazonaws.com	
MSP	https://msp.cace1.dome9.com	-

India (IN)

Module	Dome9 Portal URL	Infinity Portal URL
Portal	https://secure.ap3.dome9.com	https://in.portal.checkpoint.com

Module	Dome9 Portal URL	Infinity Portal URL
API server	https://api.ap3.dome9.com https://api-cpx.ap3.dome9.com	https://api.ap3.cgn.portal.checkpoint.com
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/	
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://in-gw.sg.iaas.checkpoint.com	
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query	
Container Registry	https://quay.io/checkpoint	
Code Security Spectral	https://dl.spectralops.io/	
Code Security Spectral OSS	https://eiffel.spectralops.io/	
Intelligence	https://webserver.logic.ap3.dome9.com	https://webserver.logic.ap3.cgn.portal.checkpoint.com
Kubernetes Intelligence	https://validator-prod-573281234161-k8s.s3.ap-south-1.amazonaws.com	
MSP	https://msp.ap3.dome9.com	-

Singapore (SG)

Module	Dome9 Portal URL
CloudGuard Portal	https://secure.ap1.dome9.com

Module	Dome9 Portal URL
API server	https://api.ap1.dome9.com https://api-cpx.ap1.dome9.com
Image Assurance	https://rpm-serv.sg.iaas.checkpoint.com https://shiftleft.portal.checkpoint.com/
Image Scan	https://shiftleft-prod-bucket.sg.iaas.checkpoint.com For agent ver. 2.28.0 and lower: https://sg-gw.sg.iaas.checkpoint.com
Runtime Protection	https://storage.googleapis.com/cos-tools https://rep.checkpoint.com/file-rep/service/v2.0/query
Container Registry	https://quay.io/checkpoint
Code Security Spectral	https://dl.spectralops.io/
Code Security Spectral OSS	https://eiffel.spectralops.io/
Intelligence	https://webserver.logic.ap1.dome9.com
Kubernetes Intelligence	https://validator-prod-155213570047-k8s.s3.ap-southeast-1.amazonaws.com
MSP	https://msp.ap1.dome9.com

Known Limitations

Alibaba Cloud Accounts

See *"Onboarding Alibaba Cloud Accounts" on page 75*.

- You cannot apply *"Organizational Units" on page 297* actions to Alibaba Cloud accounts.
- *"Automatic Remediation with CloudBots" on page 317* for Alibaba Cloud does not exist, so active remediation can be limited.
- Some Alibaba Cloud assets are not supported.
- Some dashboard widgets (see *"Dashboards" on page 86*) can fail to show Alibaba Cloud account data.
- Alibaba Cloud accounts have three tabs:
  - Protected Assets
  - Compliance Policies
  - Assessment History

CloudBots

See *"Applying a CloudBot immediately (Fix it)" on page 127*.

The **Fix it** option is not applicable to GCP environments.

CDR

AWS

See *"Onboarding AWS Environments to Intelligence" on page 572*

- AWS allows you to set only one event notification on a specific prefix with a specific event type. You cannot onboard an S3 bucket if the bucket has the event notifications set, and one of the notifications meets two conditions below:
  - The notification has an empty prefix filter set (that is, all of the bucket) or an explicit AWS log-prefix set.
  - The event type is *PutObject* or all object-created events.
- For Intelligence, you can onboard an S3 bucket through one SNS topic only.

- Intelligence cannot analyze the involved IAM policies: S3 bucket policy, existing SNS topic policy, and policies of the CloudGuard IAM trust role. Therefore, permission-related issues can occur when you create the onboarding stack.

Kubernetes

See *"Intelligence for Kubernetes Containers" on page 645*

- Kubernetes Intelligence takes about 5 minutes to start the traffic visualization in the CloudGuard portal. For new assets, the process can take 10 minutes, and while it is in progress, you can see assets by their IP addresses and not the names.
  - CloudGuard receives Flow Logs every half a minute, while inventory update occurs each 5 minutes. This is why the traffic of the new assets can be seen as traffic that originated from an IP address (not enriched) and not from a Pod (enriched).
  - It takes time for CloudGuard to handle (enrich/store) the data.
- The CloudGuard portal does not show traffic for pods in the host networks. Therefore, for example, Flow Logs agent pods cannot be seen on the Graph.
- Kubernetes Intelligence agents support Linux kernel v4.1 and higher.
- Kubernetes Intelligence does not identify the connection's direction and treats each connection as bidirectional.
- Kubernetes Intelligence categorizes IP addresses as *Private IP* based on RFC-1918.

  The applicable ranges are **10.0.0.0/8**, **172.16.0.0/12**, and **192.168.0.0/16**.

  This means that a workload that uses one of the IP addresses is considered a Private IP, even if it is external to the cluster.

Kubernetes Clusters

AWS EKS Support

See *"Image Assurance on AWS Fargate" on page 435*

CloudGuard supports Posture Management, Image Assurance, and Admission Control on AWS Fargate clusters.

To use Runtime Protection and Threat Intelligence with AWS EKS, install them on your non-Fargate clusters.

Note - For Runtime Protection and Threat Intelligence, CloudGuard does not support mixed (EC2 and Fargate compute) clusters.

GKE Autopilot

Runtime Protection is not supported on Autopilot clusters.

Runtime Protection

Kubernetes Runtime Protection agents support Linux kernel v4.1 and higher.

Vulnerabilities

Images

See *"Images" on page 425*.

- Images used by short-lived pods may not be visible to Image Assurance.
- The **Request Scan** usage is limited to 200 requests in an hour.
- Requests for a scan of *inactive* images are not available.
- On-demand scanning is not supported for ShiftLeft images and environments.

Container Registry Scanning

See *"Onboarding Container Registries" on page 204* and *"Container Registry Scanning" on page 464*.

- By default, CloudGuard adds to Protected Assets and scans only 10 recent images of each repository. You can change the default value with the API call (maximal number is 1000 for a JFrog Artifactory and Sonatype Nexus). For more information, see the API Reference Guide.
- Scanning Windows container images is not supported.
- For JFrog Artifactory, it can take about 20 minutes that the images start to show for the first time.
- For JFrog Artifactory and Sonatype Nexus, the maximal number of tags per repository is 1000. Container images from the repositories with more than 1000 tags are neither shown as protected assets, nor scanned. The number is limited due to extensive API calls and performance considerations.

ImageScan Findings

See *"Image Scan Findings" on page 449*.

- Sometimes, the ImageScan category is not available in the filter when you create a notification. This happens with newly onboarded environments where CloudGuard has not finished yet to scan images for the first time. Wait approximately 5-10 minutes to let it finish and try again.
- The remediation length is limited to 25,600 symbols. The remediation that exceeds this length is truncated to 25,600 symbols.

AWP

For a detailed list of limitations, see *"Known Limitations" on page 492*.

Infinity Portal

CloudGuard was integrated with the Infinity Portal, and some features existing in the standalone version are not supported, as of this writing.

- Managed Service Providers (MSP) are not available.
- Usernames do not allow aliases such as `johndow+demo@mycompany.com`.
- SSO is limited to IDPs that integrate with the Infinity Portal.
- The Singapore Data Center is not available in the Infinity Portal.