



Prisma Cloud Field Guide

How to Deploy Prisma Cloud from Code to Cloud

Updated: September 2024

Table of Contents

Overview	6
Cloud Security	6
Onboarding Best Practices	6
Account Groups	8
Alert Rules	10
Additional Onboarding Information	13
Agentless Scanning	15
Overview	15
Scanning Methodology	15
Required Permissions	15
Onboarding Agentless Scanning	16
Considerations to Choose Agentless or Agent-based Scanning	17
Agentless FAQ	17
VM Image Scanning	22
Cloud Discovery and Exposure Management	23
Attack Path Analysis	26
Config Policy Walkthrough	29
Dashboards and Reporting	37
Compliance in Workloads	44
Vulnerability Policy Strategy - Alerting to Blocking	46
Owners	46
Socialization	46
Change Control	47
Scope	47
Search and Investigate	47
Examples of Common RQL Searches	51
IAM Security	52
Data Security Posture Management (DSPM)	53
Auto-Remediation	60
AWS	62
Azure	63
Google Cloud Platform	65
Runtime Security	67
Defender Deployment Strategy	67
Upgrade/Redeploy the Defenders	68

Runtime Protections	69
Compliance	69
Runtime Models	70
Runtime Policy Configuration	71
WildFire malware detection	72
Custom Runtime Rules	74
Runtime Radar Utilization	74
Cloud Radar View	74
Host Radar View	75
Container Radar View	75
Serverless Radar View	75
Settings	75
Registry Scanning	76
Registry Scan Behavior	76
Scan Scheduling	82
Capacity planning	82
Trusted Images	82
Base Images	83
Web-Application and API Security (WAAS)	84
WAAS Traffic Inspections Modes	84
In-Line (Proxy) Mode	85
Out-of-Band Mode	86
Agentless Mode	87
App Definition (API Protection)	88
Define Rule and Scope	90
Network Controls	90
HTTP Headers	91
Application Profiling	91
Scoping	92
Load Balancers	93
Runtime Collections	93
Collection Functionality Overview in Compute	93
General Best Practices When Implementing Collections	102
Cortex XDR Agent for Cloud	102
What is the Cortex XDR agent for Cloud?	103
Prerequisites	103
Supported Workload Platform	103
Licensing	104
Resource Requirements	104
Deployment Guide	104

Pairing Prisma Cloud and XDR (or XSIAM) tenants	104
Deploy XDR Agent	107
Enable XDR Pro and Vulnerability Scanning	109
Enable Cortex XDR/XSIAM Analytics Engine	112
Review Vulnerability and Compliance Scan result on Prisma Cloud	112
Monitor CDR incident on Prisma Cloud	114
Ingest Prisma Cloud Compute Alerts	117
Application Security	124
Supported Tools & Technologies	124
Deployment Process	124
Enabling Cloud Application Security	127
License information	127
Onboarding Code and Build Providers	128
Scanning	128
Master/Main and Branch Scanning	129
CDK Scanning (CI/CD)	129
Terraform Plan Scanning	129
Third-Party and Private Terraform Modules (CI/CD)	130
Scanning with Checkov CLI	132
Use Cases	133
IDE Integration	133
Drift Detection	133
Enforcement (Scan configurations)	134
Enforcement Categories	134
Vulnerabilities (SCA)	134
License (SCA)	134
Infrastructure as Code (IaC)	135
Secrets	135
Static Application Security Testing (SAST)	135
Prisma Cloud SAST Rule Types	136
SAST Data Flow Terminologies	136
SAST Rules Format	137
Prerequisites for SAST	137
Prisma Cloud with SAST Findings	138
Prisma Cloud with Policy Editor	138
Performing SAST Scanning Using the Search Method	139
Policies Management	142
Out-of-the-Box Policies	142
Import Bridgecrew custom policies into Prisma Cloud	142
Custom Policy for Build-Time Checks	143

Using APIs to manage policies	143
CI Scan for Container Images	145
Best Practices for “Other” Pipeline Integration	146
CI/CD Risks	147
Dashboard	148
Filters	148
CI/CD Risk Inventory Table	148
Kill Chain Graph	149
Suggested Fixes	150
Suppress Events	150
Unsuppress Events	150
Global Suppression	150
Export CI/CD Risk Data	151
Administration	152
Adoption Advisor	152
Managing Prisma Cloud Administrators	153
Single Sign On	154
Access Management for Application Security	158
Set Up Administrator Access for Cloud Application Security	158
Prisma Cloud Roles and Application Security Permissions:	158
Cloud Application Security Roles Configuration	159
Generate API Access Key for Developers (IDE)	162
Generate API Access Key for Application Security Service Account	164
Using CSPM API to manage API Access Keys	168
APIs for Application Security	169
Prisma Cloud Application Security API	169
New Updates and Feature Releases	169
Additional Information	170
Self-Hosted Console	170
Deploy the console	170
Upgrade the console	170
Backup and Restore	170
Supported lifecycle for connected components	171
Third Party Integration	171
Prisma Cloud Enterprise Integration	171
Compute Edition Integration	172
Application Security Notification	173
General Resources and References:	173



Overview

This document is meticulously curated to provide our customers and partners with best-practice implementations pertaining to Prisma Cloud. It is the result of a collaborative endeavor involving our Professional Services and Customer Success teams, who contribute insights from their substantial hands-on experience. This document functions as a valuable resource, facilitating customers in the proficient implementation and deployment of Prisma Cloud capabilities within their Code to Clouds environment.

Cloud Security

Onboarding Best Practices

Onboarding a cloud account at the top hierarchical level is recommended as it saves you time and helps smoothly onboard cloud accounts within the hierarchy (i.e. AWS Organization, Azure Tenant, GCP Organization). Any individual cloud accounts that were onboarded prior to the hierarchy but are within it, will be automatically structured without needing to specify additional information. You have the ability to onboard individual units under the hierarchy, such as eight out of ten Organizational Units (OU) in an AWS Organization. If there are OUs that you would like to exclude, it can be done at the time of onboarding.

At the last step of onboarding, there is a status check that will let you know if Configuration, Cloudtrail, VPC Flow Logs, and the Organization are connecting properly. Validating the statuses to ensure there are no issues (validated by a green check mark) will allow you to ingest data properly from your cloud account or hierarchy. After onboarding is complete, you can navigate to Prisma Cloud's Asset Inventory to review the resources that are being ingested and specific details such as the number of specific resources/assets, configuration details, audit trails, network connectivity, and more.

Cloud Security

Home

Dashboards

Reports

Inventory

Compliance

Alerts

Investigate

Governance

INVENTORY

Applications

Assets

Compute Workloads

API Endpoints

IaC Resources

Data

Date

Most Recent

Add Filter

Cards

Table

Group By: Cloud Type

Data As Of: 26 minutes ago

Cloud

Total

Pass

Fail

Assets with Alerts

Assets with Vulnerabilities

AWS

990K

942,743

47,481

47K

6.4K

23K

17K

23K

11K

137

84

105

116

80

OCI

411K

405,605

5,030

5K

0

2.3K

2.9K

1.7K

3.7K

GCP

121K

100,718

20,047

20K

6

502

1.1K

5K

17K

33

32

33

33

32

AZURE

102K

99,583

2,497

2.5K

24

313

1.1K

1.8K

1K

7

3

7

4

3

IBM

8.5K

7,454

1,014

1K

13

39

966

22

ALIBABA

2.2K

2,086

122

122

0

0

2

117

20

OTHER

168

23

145

145

0

145

0

0

0

0

150

123

148

59

1

When onboarding your cloud account to Prisma Cloud, depending on the organization's security priorities, you can choose the security capabilities grouped as **Foundational** and/or **Advanced** (with a few enabled by default).

- Foundational capabilities are for effectively managing assets in the cloud and on-premises.
- Advanced capabilities are for proactively controlling your cloud operations and identifying and remediating issues before they manifest within your runtime environments.


Here are the detailed step-by-step instructions to onboard your supported CSPs to Prisma Cloud.

- [Onboard Your Amazon Web Services Account](#)
- [Onboard Your Azure Account](#)
- [Onboard Your Google Cloud Platform Account](#)
- [Onboard Your Oracle Cloud Infrastructure Account](#)
- [Onboard Your Alibaba Cloud Account](#)

Account Groups

During onboarding, you will also choose the Account Group that you would like the account/hierarchy to belong to. Account Group creation is important to map to your organization's business and security needs. Prisma Cloud by default comes with a "Default Account Group" that contains all of your cloud accounts. It is tied to the "Default Alert Rule", alert only on the recommended OOTB policies (policies with Prisma_Cloud label) for all of my cloud accounts. Designating certain accounts into Account Groups such as "AWS Development Accounts", "Compliance Team Accounts", "Production Accounts", etc. will allow you to create meaningful roles with least privilege access, create specific Alert Rules based on Account Groups, Investigate, and filter dashboards/views based on those Account Groups.

After specific account groups are created and utilized, the "Default Account Group" may be disabled to reduce noise and incorporate granularity and management of accounts/alerts.



To Add an Account Group

Select Settings > Account Groups > Add Account Group

Add Account Group ×

☐ Make this a parent account group ⓘ

Group By: Cloud Type ▾

Search table data... 🔍 ✕

☐	Cloud Type ↑↓	Cloud Account ↑↓	Parent Account ↑↓
☐	>  (163)	163 Accounts	
☐	>  (64)	64 Accounts	
☐	>  (171)	171 Accounts	
☐	>  (34)	34 Accounts	
☐	>  (2)	2 Accounts	
☐	>  (3)	3 Accounts	

0 Selected | Displaying 1 - 6 of 6

Rows

10 ▾

 Page

1 ▾

 of 1

<

>

Non-Onboarded Account IDs (Optional) ⓘ

Cancel

Save

Alert Rules

Prisma Cloud comes with a “Default Alert Rule” as mentioned previously, where the target is the “Default Account Group” and the alerts are for only the Prisma Cloud Recommended OOTB policies (policies with Prisma_Cloud label) within your tenant. Best practice calls for creating specific custom account groups. Creating specific custom Alert Rules allows you to manage your alerts better. Setup Alert Rules based on organization/security use cases and personas as a starting point. More details about creating alert rules can be found [here](#).


High level Steps to Enabling Alerts after Onboarding:

1. Enable Alerts by adding the cloud account to an account group during the onboarding process.
2. Create an alert rule for run-time checks that correlates with the account group, selecting the policies you'd like to receive alerts on.
3. Verify the alert rule is creating alert notifications in the Alert page.

The last step in the Alert Rule window is the Notification. If you don't select a Notification method but still configure the rule, you will see the alerts in the console. If you'd like to receive it via email or send alerts to a third-party tool, you can select it at this stage.

Examples of custom/specific Alert Rules:

- Tagged resources/alerts - IAM Alerts, Security Groups, etc.
- Teams - DevOps (monitoring build policies and IaC), Compliance (monitoring production or critical accounts that are audited)
- Type of account - Production, Testing, Sandbox, Critical.




The screenshot shows the 'Add Alert Rule' window with the 'Assign Targets' tab selected. The window has a sidebar with 'Add Details', 'Assign Targets', 'Assign Policies', and 'Summary'. The main content area is titled 'Assign Targets' and includes a sub-header 'Select the account groups, compute access groups, regions, and tags that you want to target with this alert rule.' Below this, there are two main sections: 'Account Groups' and 'Compute Access Groups'. The 'Account Groups' section has a green checkmark and a description 'Include these as part of the target for this alert rule.' The 'Compute Access Groups' section has a description 'Include these to target compute policy resources.' Below these, there is a section titled 'Account Groups' with a dropdown menu showing '0 Account Groups Selected'. Further down, there is a section titled 'INCLUDE AND EXCLUDE (OPTIONAL)' with a sub-header 'Select the subset of cloud accounts to exclude and/or regions and resource tags to include.' This section includes three optional fields: 'Exclude Cloud Accounts (Optional)' with a dropdown showing '0 Cloud Accounts Excluded', 'Include Regions (Optional)' with a dropdown showing 'All Regions Selected', and 'Include Resource Tags' with input fields for 'Key (Optional)' and 'Value (Optional)'. At the bottom right, there are 'Previous' and 'Next' buttons.


You can filter your Alerts Overview with Alert Rules, allowing you to focus on a specific set at a time. For example, if you have an “AWS Production CIS Compliance” alert rule, you can focus solely on your AWS Production accounts (assuming you’ve created that Account Group) and be alerted to the CIS compliance-related policies within Prisma Cloud.

There are four steps in an Alert Rule configuration:

1. Details where you create the name and description. Also you have the option to enable Alert Notifications and Auto-remediation. If you enable Alert Notifications, the Configure Notifications step is displayed.



2. Target where you choose the designated Cloud Account Group or individual cloud accounts (utilizing the Advanced Settings to exclude), regions, and resource tags.



Add Alert Rule

Assign Targets

Select the account groups that you want to target with this alert rule.

Account Groups
0 Account Groups Selected

INCLUDE AND EXCLUDE (OPTIONAL)

Select the subset of cloud accounts to exclude and/or regions and resource tags to include.

Exclude Cloud Accounts (Optional)
0 Cloud Accounts Excluded


Include Regions (Optional)
All Regions Selected

Include Resource Tags

Key (Optional) Value (Optional)

Previous Next

3. Select Policies where you can filter and choose specific policies to be alerted on, or choose "Select all policies" if you'd like to be alerted on all of them (increases alerts and noise).



Add Alert Rule

Assign Policies

Select which policies to be alerted on by using the options available below.

Include new policies matching filter criteria ☐ Select all policies ☐

Modifying filter criteria will reset any previous table selections.


Policy Severity Select Policy Severity Cloud Type Select Cloud Type Compliance Standard Select Compliance Standard Policy Label Select Policy Label Show Less

Search...

	Name	Policy Type	Severity	Cloud	Status	Compli
☐	Azure Microsoft Defender for Cloud set to Off for...	Config	Medium	Azure	☒	
☐	AWS Redshift database does not have audit loggi...	Config	Medium	AWS	☒	AUDIT, ACCOL, CONFI, MANAC
☐	AWS ElastiCache cluster not associated with VPC	Config	Medium	AWS	☒	SYSTEM, COMM, PROTE, Safegur
☐	AWS Elastic Load Balancer v2 (ELBv2) load balan...	Config	Medium	AWS	☒	CONFI, MANAC, maintai, configu, cardhol
☐	AWS EC2 instance with IAM write access level	IAM	High	AWS	☒	

Previous Next

4. Alert Notification where you specify how you'd like to be notified on the alerts in this specific Alert Rule. You have the option of third-party integrations, email, or if you do not choose a specific "channel", it will be an Alert Rule contained just in the console which is still helpful to filter when viewing Alerts and data. Note: New accounts must be updated manually in Alert Rules.



Add Alert Rule

- Add Details ✓
- Assign Targets ✓
- Assign Policies ✓
- Configure Notifications**
- Summary

Configure Notifications

Send notifications when alerts are generated through one or more notification channels.

Trigger notification for config alert only after the alert is open for Minutes

Use the following channels to receive notifications on:


Amazon SQS	Not Configured
Amazon S3	Not Configured
AWS Security Hub	Disabled
Azure Service Bus Queue	Disabled
Cortex XSOAR	Not Configured
Emails	Not Configured
Google CSCC	Not Configured
Jira	Not Configured
Microsoft Teams	Not Configured

Previous Next

Additional Onboarding Information


Utilize Filters in *Prisma Cloud Inventory, Governance, Alerts, and Compliance*.

There is typically a lot of data ingested into Prisma Cloud and being able to filter will help you have a more granular view and focus on the relevant information at the time. Understand the use of filters in Policies and Alerts. It helps view more granular/detailed information rather than a flood of Policies and Alerts



The screenshot displays the Prisma Cloud Security console's Alerts page. The top navigation bar includes tabs for Home, Dashboards, Reports, Inventory, Compliance, Alerts, Investigate, and Governance. The Alerts tab is active, showing a list of alerts. The left sidebar contains filters for Time Range (All Time), Alert Status (Alert Opened), and Alert Labels (Prisma Cloud X). The main content area shows a list of alerts, each with a severity level (Critical), a description, and a score. The alerts are sorted by severity, with the highest severity (Critical) at the top. The alerts include details about unauthorized port scanning, privilege escalation, and unauthorized function invocation. The interface also includes a 'Create Alerts Report' button and a 'View Alert Status' button.

Creating meaningful labels for OOTB and custom policies facilitates reviewing data, alerts, and security needs. You will see this option in the first part of a policy, whether it's a default OOTB or a custom one. See the screenshot below.



The screenshot shows the 'Create New Config Policy' dialog box. The dialog has a sidebar with tabs for Add Details, Create Query, Compliance Standards, and Remediation. The 'Add Details' tab is active, showing fields for Policy Name, Description (Optional), Policy Subtype (Run, Build), Severity, and Labels (Optional). The 'Labels (Optional)' field is highlighted with a red border. The 'Next' button is visible at the bottom right of the dialog.



Agentless Scanning

Overview

Agentless scanning is scanning of your cloud resources (Virtual Machines, instances etc.) to inspect for vulnerabilities or risks without actually having an agent installed on them and without actually affecting the execution of any of your instances or VMs.

Agentless scanning is supported on AWS, Azure and GCP hosts for vulnerabilities and compliance. Agentless scanning is also supported for hosts running on Windows OS.

Note: If you have a defender installed on your host, agentless will not scan for that host. Agentless will only scan the host if there is no defender installed on it.

Scanning Methodology

With agentless scanning, the underlying methodology of the scanning is creating snapshots and scanning those points in time images of the resource. The resource will be scanned once in a 24 hour period (you can configure this setting to whatever duration you would like the scan period to be from Manage> System>scan>scheduling) or can manually be kicked off in the UI.

Required Permissions

In order to have agentless scan your resources within your cloud accounts there are a specific set of permissions that are required. When onboarding the account, the pre-flight check will check your permissions. It is always best practice to remediate them first before moving forward.

Here is a list of [permissions](#) that are required for each cloud provider:

AWS permissions (JSON):

<https://redlock-public.s3.amazonaws.com/aws/awsAgentlessPermissions.json>

Azure permissions (JSON):

<https://redlock-public.s3.amazonaws.com/azure/azureAgentlessPermissions.json>

GCP permissions (JSON):

<https://redlock-public.s3.amazonaws.com/gcp/gcpAgentlessPermissions.json>

The agentless scanning is based on creating snapshots of a compute instance and

scanning that image. Therefore IAM policies to create and manage snapshots are essential for agentless scanning. However, for both Azure and GCP, the IAM policies still need to be reduced to least privileged in the Terraform file that onboards cloud accounts. There are IAM policies that are provided that will be reduced in future releases of Prisma Cloud such as the following:

Unset

```
"compute.instances.create",  
"compute.instances.delete",  
"compute.networks.create",  
"compute.networks.delete",  
"compute.subnetworks.create",  
"compute.subnetworks.delete",
```

In the Azure Terraform template, the following IAM policies still need to be reduced to least privileged.

Unset

```
"Microsoft.Network/networkInterfaces/delete",  
"Microsoft.Network/networkSecurityGroups/delete",  
"Microsoft.Network/virtualNetworks/delete",  
"Microsoft.Compute/disks/delete",  
"Microsoft.Compute/virtualMachines/delete"
```

Onboarding Agentless Scanning

Here are the detailed techdoc instructions for onboarding different Public Cloud accounts with agentless scanning.

- [AWS Accounts for Agentless Scanning](#)



- [Azure Accounts for Agentless Scanning](#)
- [GCP Accounts for Agentless Scanning](#)
- [OCI Accounts for Agentless Scanning](#)

Considerations to Choose Agentless or Agent-based Scanning

Agentless	Agents (Defenders)
Simplified visibility into vulnerability and compliance	Continuous, real-time runtime protection
Distributed Access - For distributed environments across different teams, getting access to deploy agents is difficult	Defense in Depth - For sensitive workloads, deploy defenders/agents to gain deep insights into running processes, runtime forensics, behavioral analysis etc.
First step in securing - start with agentless security as soon as you on board accounts for first look into all hosts without need for running deployment scripts per hosts	Advanced security with active prevention - Actively prevent access to sensitive file systems or malware from being deployed with block capabilities
No running management - Agentless uses snapshot based scanning, with no running tools inside your environment to manage	Continuous monitoring - For instances where you require continuous and immediate alerts

Agentless FAQ

Agentless Scanning Mode (TechDoc)	Same Account	Hub Account
How long scanning takes	Depends on the number and size of attached volumes in the EC2 snapshots. Many small files will take longer to scan than fewer large files. The average time is 3 min to scan	same

	26 snapshots	
What happens when scanning is initiated	Snapshot of EC2 is created in the target account(Not in Hub account), then a M5.2xlarge scanner instance is created on the same account to scan the snapshots. If an account spans multiple regions, then there will be at least 1 scanner per region	Snapshot of EC2 is created in the target account and shared with the Hub account, then a M5.2xlarge scanner instance is created on the Hub account to scan the snapshots. If the account spans multiple regions, then there will be at least 1 scanner per region
What happens after completing the scan	Scanner, snapshots, and any other resources get deleted automatically after sending results to the console	same
How many concurrent scans are possible	26 per scanner instance * up to 50 scanners with auto scaling = 1300 concurrent scans	same
How long is the instance in use?	Runs until the scan completes and then the scanner instance gets terminated. The average is 3 min to scan 26 snapshots.	same
What happens if no spot instance is available?	The system tries 3 times for spot instance, if still unavailable, on demand instance gets created	same
Why do I see the system continuously in use	If a scan fails, PC re-tries to scan the same unlimited amount of times	same
What are the most common reasons for scan failure	Unsupported instance, missing permission, SCP conflicting with permission, connectivity	fewer permission errors, but more connectivity errors than same account scanning
Are you attaching any EBS	No	same

volumes?		
Any outbound data transfer to other regions	No customer data is transferred, but the result of the scan is transferred to the console. If an account spans multiple regions, then there will be at least 1 scanner per region so snapshots stay in their region	same
Does agentless support Windows OS	Yes - scans host only and not Windows containers running on them	same
Document link - Permission	https://docs.prismacloud.io/en/enterprise-edition/content-collections/runtime-security/configure/permissions#aws-agentless	same
How to get details of permission?	Setting->connect provider->cloud account->AWS->Fill the tabs, ensure agentless is enabled->give account ID and download CFT	same
After importing CFT, I still get a permission error	Different permission may enforced at the SCP level, SCP permission overwrites any other permissions	same
	Use IAM policy simulator, you need to select individual permission, time consuming	same
AWS ORG account with over 600 Accounts in it. They are requesting that Agentless Scanning be turned on on 6 out of the 600 Accounts	Use a single CFT for the org. Then, turn off the "org scan" flag as part of the onboarding page. This will add the permissions for the org, but then will allow you to selectively turn it on for only	same

	the 6 accounts interested in being scanned. This may be more efficient than having to disable 594 accounts.	
How to enable inactive resources	Under Runtime Security, Manage > Cloud Accounts, edit the account (or bulk edit multiple accounts), advanced settings, enable non-running hosts	same
Is agentless available for both SaaS + Self-Hosted?	yes	same
What happens if a Defender is deployed on a host being scanned by agentless?	Agentless scans will automatically detect the Defender agent and skip any EC2's with Defender running. There will not be any duplicate scans or counts on licensing.	same
Do we scan/have visibility of container images via Agentless?	yes - for Linux containers only	same
Is there a way to scan specific accounts/workloads?	Under Runtime Security, Manage > Cloud Accounts, edit the account (or bulk edit multiple accounts), advanced settings, include or exclude hosts by tags	same
Are new cloud accounts created under the platform auto imported and scanned by Agentless?	Customers can choose to opt-out when onboarding. If agentless is enabled at the Org level, new accounts will be scanned automatically using the same account scanning.	Customers can choose to opt-out when onboarding. If agentless is enabled at the Org level, new accounts will be scanned automatically using the same account scanning and will need to be updated to use hub mode.
Any difference between Defender vulnerability scan and Agentless	The difference in vulnerabilities is that Defender	same

vulnerability scan?	host scan doesn't scan 3rd party packages (python in this case) by design, and Agentless does scan 3rd party packages.	
What scale does it support?	No limit on the number of scanners/scans.	same
What is the performance impact on the customer's machines?	None since we don't install anything on customer's machines.	same
How often are the scans run?	24 hrs, by default. Can be adjusted in 1 hr increments	same
Do we support Windows hosts?	Yes - scans host only and not Windows containers running on them	same
Do we support Agentless for Windows Containers Support?	No	same
How can you disable Agentless scans?	Can only be disabled from the platform cloud accounts page (Settings > Providers) by editing the cloud account and disabling Agentless scans. Automatic scans can be disabled by changing the scan interval from 24hrs to 0 and then you can still run on-demand agentless scans using the UI or API	same
Can you request on demand scans?	Yes, there is a button for on demand scan on top of periodic scans under Compute > Manage > Cloud Accounts	same
Is there an additional risk of using agentless? Is data leaving my cloud environment?	No data is leaving your accounts except metadata that would normally be sent to a security vendor like the	same

	SBOM, vulnerabilities, compliance issues, etc...	
Will agentless scan 'stopped' hosts in addition to running hosts in the cloud account?	Yes. There is a toggle to scan stopped hosts under Agentless configuration.	same
Can you use agentless scanning with your AWS hub accounts to scan encrypted volumes?	Yes. You can use agentless scanning to scan EBS volumes encrypted with KMS. Encrypted volumes using Customer Managed Keys not supported	same
What happens if a scanner fails to communicate back with the Compute Console and is left hanging only to reach timeout (~45 min)?	The scanner may timeout but normally it is not expected. By default the scanner will check for connectivity after deployment. The error will show in the UI and in the logs for further troubleshooting. In OCI this mechanism is still unavailable, and in case of connectivity issues the scanner will remain dangling for the timeout until it is terminated.	same

VM Image Scanning

By allowing the scanning of VM Images (AMIs, VM Images, etc) before VMs are deployed in the public cloud, you will be able to identify vulnerable VM image baselines before they are deployed into your environments as VMs. This is especially important if you have existing workflows that develop custom VM Images. The idea is to be able to scan these VM Images as soon as they are created for any vulnerabilities embedded in the image itself.

The configuration for VM Image scanning works best when you utilize an existing onboarded account (AWS, Azure, or GCP) credentials to establish the scanning. Make sure to add on the additional VM Image scanning permissions required to conduct the scans to your existing IAM role or Service Account, then just re-use those credentials



when configuring the VM Image scanning profile.

Another best practice when configuring VM Image scanning is defining a VPC and Subnet ID for the scanners to spin up and scan the VM Images. When specifying a VPC and Subnet ID, it is important to note that the requirement for this VPC/Subnet combination is that resources being spun up in that particular Subnet must be able to reach out to the internet on HTTPS 443. This will ensure the scanners are able to communicate back to the Prisma Cloud console to report the scan results of the VM Image scans.

One thing to note if trying to restrict certain AWS permissions (Creation/termination of EC2s) is that in the current documentation, it mentions limiting resource creation/termination by using IAM conditional keys based on a tag “twistlock-scan”. This information is inaccurate. The AWS tags associated with any VM image scanning asset will be notated by “prismacloud-scan-*”, not “twistlock-scan-*”.


Cloud Discovery and Exposure Management

Cloud discovery and exposure management in Prisma Cloud helps organizations identify shadow cloud workloads and eliminate the risks they pose. CDEM functionality is a critical component of modern cloud security. It delivers an outside-looking-in view of cloud environments that mimics what threat actors see and enables security professionals to:

- View a complete inventory of managed and unmanaged assets.
- Identify internet exposure risks presented by shadow cloud assets.
- Eliminate exposures to strengthen overall cloud security posture.

Managed Assets

Managed assets are the assets onboarded and are in Prisma Cloud governance.




The screenshot shows the Prisma Cloud Inventory - Assets page. The top navigation bar includes Home, Dashboards, Reports, Inventory (selected), Compliance, Alerts, Investigate, and Governance. The sub-navigation bar includes INVENTORY (selected), Applications, Assets, Unmanaged Assets, Compute Workloads, API Endpoints, IaC Resources, and Data. The main content area shows a table of assets grouped by Cloud Type. The table has columns for Cloud, Total, Pass, Fail, and Assets with Alerts. The data is as of 5 minutes ago.

Cloud	Total	Pass	Fail	Assets with Alerts
AWS	4.5K	4,455	57	57
AZURE	669	617	52	52
OTHER	34	34		

Unmanaged Assets

The **Unmanaged Assets Inventory** page provides information about your unmanaged or shadow IT assets that are publicly exposed on the internet and attributed to your organization. You can use this information to gain valuable insights into your publicly exposed assets and enhance your security posture.



The screenshot shows the Prisma Cloud Inventory - Unmanaged Assets page. The top navigation bar includes Home, Dashboards, Reports, Inventory (selected), Compliance, Alerts, Investigate, and Governance. The sub-navigation bar includes INVENTORY, Applications, Assets, Unmanaged Assets (selected), Compute Workloads, API Endpoints, IaC Resources, and Data. The main content area shows a table of unmanaged assets grouped by Cloud. The table has columns for Cloud and Total. The data is as of 5 minutes ago.


Cloud	Total
AWS	9419
GCP	87

CDEM Workflow

1. When customer enable CDEM Subscription, Prisma Cloud will trigger an API call to Xpanse to provision a tenant
2. Xpanse will fetch customer's entitlement (IP addresses and Domain) from our internal SFDC
3. Xpanse will start scanning customer's exposed assets on internet
4. Xpanse store those scanned assets
5. Prisma Cloud send query every 2 hours for exposed assets based on customer information e.g. IP addresses, domains and certificates
6. Xpanse responds to Prisma Cloud CDEM with data containing unmanaged assets

type: Responsive IP addresses, Domain and Certificates

7. Prisma Cloud CDEM will automatically mapped those assets to specific cloud accounts that it can associate with. Assets that can't be mapped will have error 'account cannot be mapped'



CDEM FAQ

Question - Is there a tag or label added for the converted unmanaged assets to manage?

Answer -None at the moment. You only see a graph of converted assets.

Question - How is license credit charging for CDEM

Answer - After 30 days free trial from enabling CDEM subscription. We charge similar to CIEM. That is 0.25 credits for every managed VM on CPSM

Question - What are the Additional permissions required to onboard unmanaged cloud accounts?

Answer - Please refer to this link to view additional permission required to allow Prisma Cloud to onboard discovered unmanaged cloud assets

<https://docs.prismacloud.io/en/enterprise-edition/content-collections/administration/subscribe-to-cdem#pre-req-for-cdem-aws>

AWS Security Token Service (STS) is an AWS facility for requesting temporary user credentials with limited privileges. It allows you to acquire short-term access to privileged roles in a controlled manner. The role that we create across all member accounts to pull the IP related data. And the above policy will allow org master account to assume the role of member accounts to pull the data. So basically this enables us to use org master account credentials to call APIs on member accounts

Question - What is the frequency of CDEM getting mapping of unmanaged assets from Xpanse?

Answer -every 2 hours

Question - Can we use RQL to query unmanaged assets?

Answer – No

Question - Will there be API queries available for CDEM?


Answer - Yes

<https://pan.dev/prisma-cloud/api/cspm/discovery-and-exposure-management/>

Attack Path Analysis


Your Prisma Cloud tenant will come with default out-of-the-box Attack Path policies. These policies identify the confluence of issues that increase the likelihood of a security breach and are based on a data model that automatically correlates findings across infrastructure misconfigurations, vulnerabilities, excessive IAM permissions, and network exposures that would enable an attacker to target your application. Prisma cloud risk engine helps you identify the attack paths and presents them in a graph view, offering valuable security context to protect against high-risk threats, which requires you to take immediate action.

To view the Attack Path policies, select the **Governance** tab and filter by policy type Attack Path. You can view the status of the policy from the **Status** column. If the status of the policy is disabled, toggle the status to enable each policy. Similar to other policy types in Prisma Cloud, Attack Path Policies need to be added to an [Alert Rule](#) for Prisma Cloud to generate alerts.



Policy Name	Policy Description	Policy Type	Policy Subtype	Cloud	Severity	Status	Actions
Unauthorized account access risk through a publicly exposed AWS Lambda function	This policy identifies AWS Lambda Function that are publicly exposed...	Attack Path	Run	AWS	Critical	Enabled	🔍 🗑️ ⚙️
Potential malicious activity on AWS EC2 instance affected by malware	This policy identifies AWS EC2 instances that are affected by malware...	Attack Path	Run	AWS	Critical	Enabled	🔍 🗑️ ⚙️
Denial of Service (DoS) exploitation attempt on GCP VM with known vulnerabilities	This policy identifies exploitation attempts on GCP VMs that are vuln...	Attack Path	Run	GCP	Critical	Enabled	🔍 🗑️ ⚙️
Remote Code Execution (RCE) exploitation attempt on Azure Virtual Machine	This policy identifies exploitation attempts on Azure Virtual Machine...	Attack Path	Run	Azure	Critical	Enabled	🔍 🗑️ ⚙️
Remote code execution (RCE) exploitation attempt on GCP VM with known vulnerabilities	This policy identifies exploitation attempts on GCP VMs that are vuln...	Attack Path	Run	GCP	Critical	Enabled	🔍 🗑️ ⚙️
Privilege escalation risk due to malware in AWS EC2 instance with S3 access	This policy identifies AWS EC2 instances that are affected by malware...	Attack Path	Run	AWS	Critical	Enabled	🔍 🗑️ ⚙️
Potentially nefarious hacking tools activity detected on a publicly exposed web application	This policy identifies GCP instances that are publicly exposed wit...	Attack Path	Run	GCP	Critical	Enabled	🔍 🗑️ ⚙️
Privilege Escalation risk due to unauthenticated public access and ins...	This policy identifies Azure Function Apps that are publicly exposed...	Attack Path	Run	Azure	High	Enabled	🔍 🗑️ ⚙️
Privilege escalation and lateral movement risk due to a publicly exposed web application	This policy identifies AWS EC2 instances that are publicly exposed w...	Attack Path	Run	AWS	Critical	Enabled	🔍 🗑️ ⚙️
Unauthorized access risk due to a publicly exposed and vulnerable GCP Cloud Function	This policy identifies GCP Cloud Functions which are publicly expose...	Attack Path	Run	GCP	Critical	Enabled	🔍 🗑️ ⚙️
Potentially nefarious ransomware activity detected on Azure Virtual Machine	This policy identifies Azure Virtual Machines that have public serial c...	Attack Path	Run	Azure	High	Enabled	🔍 🗑️ ⚙️

There is a set of Attack Path Rules for each Attack Path policy. Prisma Cloud evaluates those rules and when it finds a match on all of the rules (AND logic), it generates an alert. To view the Attack Path Rules, select the **Governance** tab and filter by policy type Attack Path, click edit icon under the **Action** Column. Also you can view the Attack Path Rules from the **Alerts** tab by clicking on  icon next to the policy name.



The screenshot displays the Prisma Cloud interface. The top navigation bar includes tabs for Alerts, Investigate, and Governance. The Alerts tab is active, showing a list of alerts. A specific alert is highlighted: "Unauthorized access risk due to a publicly exposed and vulnerable AWS EC2 instance". Below the alert details, a table lists the assets associated with this alert.

Asset Name	Alert ID	Alert Time	Account ID	Account	Alert Status	Region	Snoozed For	Snoozed Expires	Reason	Action
jeanclaude-ngin...	A-10088686	6 days ago	886388100502	pcs-aws-lab01	open	AWS Virginia				
jeanclaude-test-v...	A-10086809	6 days ago	886388100502	pcs-aws-lab01	open	AWS Virginia				
jeanclaude-splunk...	A-9995747	6 days ago	886388100502	pcs-aws-lab01	open	AWS Virginia				
i-0a2b2c6b283e...	A-9996226	6 days ago	886388100502	pcs-aws-lab01	open	AWS Virginia				

Below the alerts table, the "Edit Policy" modal is open, showing the "View Rules" tab. It displays the "Attack Path Rules" for the selected policy:

- AWS VM instance in running state that is internet reachable with unrestricted access (0.0.0.0/0)
- Asset has critical or high vulnerabilities

To view the Alerts, Select **Alerts > Overview**, filter the alerts by policy type **Attack Path**, and then click on the policy name. The list of the resources/assets violating the policy will be shown. Click on the asset name to view the Attack Paths/ Evidence details in Graph (default) view. In addition to that, you can view the Findings (all findings) and Vulnerabilities for the selected Asset.

i-0a
EC2 Instance


[View JSON {}](#) ×

Findings Types ?

🔍 Unclassified 🔍 Reconnaissance 🔍 Internet Exposure 🔍 Misconfiguration

Overview **Attack Paths** Alerts (16) Vulnerabilities (94) Findings (13) Package Info Identity Audit

A-9996226 Unauthorized access risk due to a publicly exposed and vulnerable AWS EC2 instance



i-0a
EC2 Instance

[View JSON {}](#) ×

Findings Types ?

🔍 Unclassified 🔍 Reconnaissance 🔍 Internet Exposure 🔍 Misconfiguration

Overview Attack Paths Alerts (16) Vulnerabilities (94) **Findings (13)** Package Info Identity Audit

You are viewing the most recent data about this asset

Type All Type(s) Selected Severity All Severity(s) Selected Source All Source(s) Selected

Search table data... 🔍 ⌂ ☰

Name ↑↓	Description ↓	Source ↑↓	Type ↑↓	Actions
AWS EC2 instance with unres...	This policy identifies EC2 insta...	Prisma Cloud	Internet Exposure	
AWS VM instance in running s...	This policy identifies AWS VM ...	Prisma Cloud	Internet Exposure	
AWS EC2 instance that is inte...	This policy identifies AWS EC...	Prisma Cloud	Internet Exposure	

Config Policy Walkthrough

Your Prisma Cloud tenant will come with some default out-of-the-box policies already enabled. To view your Prisma Cloud policies, click on the Governance menu which will open your policies. There are hundreds of default policies that apply across multiple cloud providers as well as some that are cloud agnostic. Review the enabled policies as well as the disabled ones to determine if they should be enabled or disabled for your organization. Utilize the filter options mentioned earlier in the Setup/Configuration section of the document to filter by different options such as Cloud Type, Compliance Standard, Severity, and more. There is also a list of recommended policies to enable that are aligned to the CIS Benchmarks for each major cloud provider. See the tables below for recommended policies.

Category	AWS Policy	Azure Policy	GCP Policy
Identity Management			
	AWS IAM deprecated managed policies in use by User		GCP IAM Service account has admin privileges
	AWS IAM Groups with Administrator Access Permissions		GCP IAM user have overly permissive Cloud KMS roles
	AWS IAM has expired SSL/TLS certificates		GCP IAM user with service account privileges
	AWS IAM password policy allows password reuse		
	AWS IAM password policy does not have a minimum of 14 characters		
	AWS IAM password policy does not have password expiration period		
	AWS IAM Password policy is unsecure		
	AWS IAM policy allows assume role permission across all services		
	AWS IAM policy allows full administrative privileges		
	AWS IAM Roles with Administrator Access Permissions		
	AWS MFA is not enabled on Root account		
	AWS MFA not enabled for IAM users		

	AWS root account configured with Virtual MFA		
--	--	--	--

Category	AWS Policy	Azure Policy	GCP Policy
Access Management			
	AWS access keys are not rotated for 90 days	Azure Active Directory Guest users found	GCP User managed service account keys are not rotated for 90 days
	AWS Certificate Manager (ACM) has expired certificates	Azure Custom Role Administering Resource Locks not assigned	GCP VM instance configured with default service account
	AWS Customer Master Key (CMK) rotation is not enabled	Azure subscriptions with custom roles are overly permissive	GCP VM instance using a default service account with full access to all Cloud APIs
	AWS KMS customer managed external key expiring in 30 days or less	SQL servers which do not have Azure Active Directory admin configured	
	AWS KMS Key policy overly permissive	Azure Active Directory Security Defaults is disabled	
	AWS KMS Key scheduled for deletion		
	AWS S3 bucket having policy overly permissive to VPC endpoints		
	AWS SQS queue access policy is overly permissive		
	AWS SNS topic policy overly permissive for publishing		
	AWS SNS topic policy overly permissive for subscription		
	AWS EC2 instance not configured with Instance Metadata Service v2 (IMDSv2)		
	AWS IAM policy is overly permissive to all traffic via condition clause		

Category	AWS Policy	Azure Policy	GCP Policy
Data Protection - Encryption at rest			

	AWS EBS snapshot is not encrypted	Azure Key Vault is not recoverable	GCP GCE Disk snapshot not encrypted with CSEK
	AWS EBS volume region with encryption is disabled	Azure SQL Server advanced data security is disabled	GCP KMS encryption key not rotating in every 90 days
	AWS Elastic File System (EFS) with encryption for data at rest is disabled	SQL databases has encryption disabled	
	AWS RDS DB cluster encryption is disabled		
	AWS RDS DB snapshot is not encrypted		
	AWS RDS instance is not encrypted		
	AWS S3 buckets do not have server side encryption		
	AWS SNS topic with server-side encryption disabled		
	AWS SQS server side encryption not enabled		

Category	AWS Policy	Azure Policy	GCP Policy
Data Protection - Encryption in transit			
	AWS Application Load Balancer (ALB) is not using the latest predefined security policy	Azure ACR HTTPS not enabled for webhook	GCP HTTPS Load balancer is configured with SSL policy having TLS version 1.1 or lower
	AWS CloudFront distribution is using insecure SSL protocols for HTTPS communication	Azure App Service Web app doesn't redirect HTTP to HTTPS	
	AWS CloudFront origin protocol policy does not enforce HTTPS-only	Azure App Service Web app doesn't use latest TLS version	
	AWS CloudFront viewer protocol policy is not configured with HTTPS	Azure Application Gateway allows TLSv1.1 or lower	
	AWS CloudTrail logs are not encrypted using Customer Master Keys (CMKs)	Azure Application gateways listener that allow connection requests over HTTP	
	AWS Elastic Load Balancer (Classic) SSL negotiation policy configured with insecure ciphers	Azure CDN Endpoint Custom domains is not configured with HTTPS	

	AWS Elastic Load Balancer (Classic) SSL negotiation policy configured with vulnerable SSL protocol	Azure CDN Endpoint Custom domains using insecure TLS version	
	AWS Elastic Load Balancer v2 (ELBv2) listener that allow connection requests over HTTP	Azure MySQL Database Server SSL connection is disabled	
	AWS Elastic Load Balancer v2 (ELBv2) SSL negotiation policy configured with weak ciphers	Azure PostgreSQL database server with SSL connection disabled	
	AWS Elastic Load Balancer v2 (ELBv2) with listener TLS/SSL is not configured	Storage Accounts without Secure transfer enabled	
	AWS Elastic Load Balancer with listener TLS/SSL is not configured		
	AWS Network Load Balancer (NLB) is not using the latest predefined security policy		
	AWS S3 bucket not configured with secure data transport policy		
	AWS SNS subscription is not configured with HTTPS		
	AWS SNS topic not configured with secure data transport policy		

Category	AWS Policy	Azure Policy	GCP Policy
Public Exposure			
	AWS Amazon Machine Image (AMI) is publicly accessible	Azure Container registries Public access to All networks is enabled	GCP BigQuery dataset is publicly accessible
	AWS Classic Load Balancer is in use for internet-facing applications		GCP SQL database is assigned with public IP
	AWS CloudTrail bucket is publicly accessible		GCP Storage buckets are publicly accessible to all authenticated users
	AWS EBS Snapshot with access for unmonitored cloud accounts		GCP Storage buckets are publicly accessible to all users
	AWS EBS snapshots are accessible to public		Storage Buckets with publicly accessible Stackdriver logs
	AWS EC2 instance allowing public IP in subnets		
	AWS EC2 instances with Public IP and associated with Security Groups have Internet Access		

	AWS RDS database instance is publicly accessible		
	AWS RDS instance not in private subnet		
	AWS RDS snapshots are accessible to public		
	AWS S3 Bucket Policy allows public access to CloudTrail logs		
	AWS S3 bucket publicly readable		
	AWS S3 bucket publicly writable		
	AWS S3 buckets are accessible to any authenticated user		
	AWS S3 buckets are accessible to public		
	AWS S3 Buckets Block public access setting disabled		
	AWS VPC subnets should not allow automatic public IP assignment		
	AWS SNS topic is exposed to unauthorized access		

Category	AWS Policy	Azure Policy	GCP Policy
Network configuration			
	AWS CloudFront web distribution with AWS Web Application Firewall (AWS WAF) service disabled	Azure Cosmos DB IP range filter not configured	GCP Firewall rule allows all traffic on RDP port (3389)
	AWS Default Security Group does not restrict all traffic	Azure Network Security Group allows all traffic on RDP Port 3389	GCP Firewall rule allows all traffic on SSH port (22)
	AWS NAT Gateways are not being utilized for the default route	Azure Network Security Group allows all traffic on SSH port 22	GCP Firewall with Inbound rule overly permissive to All Traffic
	AWS Security Group allows all traffic on RDP port (3389)	Azure Network Security Group having Inbound rule overly permissive to all traffic on any protocol	GCP project is configured with legacy network
	AWS Security Group allows all traffic on SSH port (22)	Azure Network Security Group having Inbound rule overly permissive to all traffic on TCP protocol	GCP VM instances have IP Forwarding enabled
	AWS Security Group Inbound rule overly permissive to all traffic on all protocols (-1)	Azure Network Security Group having Inbound rule overly permissive to all traffic on UDP protocol	GCP VPC Network subnets have Private Google access disabled

	AWS Security Group overly permissive to all traffic	Azure Network Security Group with overly permissive outbound rule	
	AWS VPC allows unauthorized peering	Azure PostgreSQL Database Server 'Allow access to Azure services' enabled	
	Instances exposed to network traffic from the internet	Azure PostgreSQL Database Server Firewall rule allow access to all IPV4 address	
	AWS Application Load Balancer (ALB) not configured with AWS Web Application Firewall v2 (AWS WAFv2)	Azure SQL Servers Firewall rule allow access to all IPV4 address	
		Azure Storage Account default network access is set to 'Allow'	
		Azure storage account has a blob container with public access	
		Azure Virtual Network subnet is not configured with a Network Security Group	
		SQL Server Firewall rules allow access to any Azure internal resources	

Category	AWS Policy	Azure Policy	GCP Policy
Logging			
	AWS Access logging not enabled on S3 buckets	Azure Monitoring log profile is not configured to export activity logs	GCP Project audit logging is not configured properly across all services and all users in a project
	AWS Certificate Manager (ACM) has certificates with Certificate Transparency Logging disabled	Azure Network Watcher Network Security Group (NSG) flow logs retention is less than 90 days	GCP VPC Flow logs for the subnet is set to Off
	AWS CloudFront distribution with access logging disabled	Azure storage account logging for blobs is disabled	
	AWS CloudTrail is not enabled in all regions	Azure storage account logging for queues is disabled	
	AWS CloudTrail logging is disabled	Azure storage account logging for tables is disabled	
	AWS VPC has flow logs disabled		

Category	AWS Policy	Azure Policy	GCP Policy
Others			

	AWS Amazon Machine Image (AMI) infected with mining malware	Azure App Service Web app authentication is off	GCP GCR Container Vulnerability Scanning is disabled
		Azure App Services FTP deployment is All allowed	GCP MySQL instance with local_infile database flag is not disabled
		Azure Application Gateway does not have the Web application firewall (WAF) enabled	GCP SQL database instance is not configured with automated backups
		Azure Security Center Defender set to Off for App Service	GCP VM instance with Shielded VM features disabled
		Azure Security Center Defender set to Off for Azure SQL database servers	VM instances have serial port access enabled
		Azure Security Center Defender set to Off for Key Vault	
		Azure Security Center Defender set to Off for Kubernetes	
		Azure Security Center Defender set to Off for Servers	
		Azure Security Center Defender set to Off for Storage	
		Azure SQL Server ADS Vulnerability Assessment is disabled	
		Threat Detection on SQL databases is set to Off	
		Threat Detection types on SQL databases is misconfigured	


Cloud Security Home Dashboards Reports Inventory Compliance Alerts Investigate Governance Settings 									
GOVERNANCE Overview Incidents Exposure Vulnerabilities Data Custom AI Generated Add View Manage Views Add Policy									
Policy Name Select Policy Severity Select Add Filter									
Group By: None									
Policy Name	Policy Description	Policy Type	Policy Subtype	Cloud	Severity	Category	Actions		
Unauthorized account access risk through a publicly exposed AWS Lambda function with no authentication and excessive IAM permissions	This policy identifies AWS Lambda Functions that are pu... Show More	Attack Path	Run		Critical	Risk			
Copy of Lateral movement and unauthorized IAM actions due to a publicly exposed AWS EC2 instance with IAM management access lev... Show More	This policy identifies AWS EC2 instance publicly expose... Show More	Attack Path	Misconfig	All	Critical	Risk			
Denial of Service (DoS) exploitation attempt on GCP VM with known exploitable vulnerability (WordPress, CVE-2018-6389)	This policy identifies exploitation attempts on GCP VMs... Show More	Attack Path	Run		Critical	Risk			
Remote Code Execution (RCE) exploitation attempt on Azure Virtual Machine with known exploitable vulnerability (Apache Spark, CVE-2022-33891)	This policy identifies exploitation attempts on Azure Vir... Show More	Attack Path	Run		Critical	Risk			
Remote code execution (RCE) exploitation attempt on GCP VM with known exploitable vulnerability (Spring Framework, CVE-2022-22... Show More	This policy identifies exploitation attempts on GCP VMs... Show More	Attack Path	Run		Critical	Risk			
Potentially nefarious hacking tools activity detected on a publicly exposed and vulnerable GCP VM instance with risky service account L... Show More	This policy identifies GCP VM instances that are public... Show More	Attack Path	Run		Critical	Risk			
Critical of AWS S3 bucket publicly readable	This policy identifies the S3 buckets that are public... Show More	Config	Run		Critical	Risk			
raghe-Custom - Config-S3	raghe-Custom - Config-S3	Attack Path	Run		Critical	Risk			
Privilege escalation and lateral movement risk due to a publicly exposed and vulnerable AWS EC2 instance not configured with Metadat... Show More	This policy identifies AWS EC2 instances that are public... Show More	Attack Path	Run		Critical	Risk			
Unauthorized access risk due to a publicly exposed and vulnerable GCP Cloud Function	This policy identifies GCP Cloud Functions which are pu... Show More	Attack Path	Run		Critical	Risk			
Stored Cross Site scripting risk due to publicly exposed and vulnerable Azure Virtual Machine (Jenkins, CVE-2023-27898)	This policy identifies Azure Virtual Machines that are pu... Show More	Attack Path	Run		Critical	Risk			
Denial of Service (DoS) risk due to publicly exposed and vulnerable GCP Cloud Function (Java Spring Framework, CVE-2021-22119)	This policy identifies GCP Cloud Functions that are publ... Show More	Attack Path	Run		Critical	Risk			
Remote code execution (RCE) exploitation attempt on Azure Virtual Machine with known exploitable vulnerability (Log4j, CVE-2021-44228)	This policy identifies exploitation attempts on Azure Vir... Show More	Attack Path	Run		Critical	Risk			
Data destruction risk due to a publicly exposed and vulnerable GCP VM instance with delete permissions	This policy identifies GCP VM instances which are publ... Show More	Attack Path	Run		Critical	Risk			

CIS is a Good baseline compliance standards define cloud-specific technical controls and we develop policies per the standard. Other compliance standards map existing policies to high-level requirements.

Policy Types and Classifications

Category	Class	Type	Subtype
Incident	Behavioral	Anomaly	UEBA
	Behavioral	Anomaly	Network
	Behavioral	Anomaly	DNS
	Privileged Activity Monitoring	Audit Event	Audit
	Network Protection	Network	Network Event
	Runtime	Workload	Run
Risk	Misconfiguration	Config	Run
	Misconfiguration	Config	Build
	Misconfiguration	Data	Data Classification
	Vulnerability	Data	Malware
	Vulnerability	Workload	Run

In Enterprise Settings (Settings → Enterprise Settings), you can select the option for new default policies to be enabled when they are released with Prisma Cloud updates. You can select the severity of the policies to be enabled, for example, if you'd only like new default high severity policies to be enabled, you can select only the “high” option. You can also choose to retroactively enable existing default policies for your chosen severities.



Dashboards and Reporting

Dashboards

Dashboards allow you to gain insights into the security health of your infrastructure, where you should focus in terms of high priority vulnerabilities and risks, and where the biggest security challenges can be found in your Infrastructure as Code (IaC) assets and repositories. Dashboards can be shared across accounts, and custom dashboards can be created to suit custom use cases and needs.

For example, a custom dashboard that compares code vulnerabilities across DevOps teams could help foster improved secure coding practices within those teams. A custom dashboard that shows Compliance violations across business units could help prioritize assets owned by business units that need immediate attention.

The Runtime Security module contains several out-of-the-box dashboards including:


- [Code to Cloud](#)
- [Command Center](#)
- [Vulnerabilities](#)
- [Code Security](#)

- [Discovery and Exposure Management](#)


The Vulnerability dashboard utilizes the Unified Vulnerability Explorer (UVE). The Code Security Dashboard may not be visible by default. However, it can be made visible by selecting the “Manage Dashboards” link in the Dashboards view in the Cloud Security persona. The Discovery and Exposure Management dashboard can only be used if a subscription to [Cloud Discovery and Exposure Management \(CDEM\)](#) has been enabled.

It is possible to create custom dashboards using widgets that are included in Cloud Security. These widgets are used in the out-of-the-box dashboards, making it possible to use them by cloning existing dashboards and customizing them to your needs, or by creating custom dashboards from scratch. Each widget has its own unique set of customizations.


To clone an existing dashboard, click the hamburger menu to the right of any dashboard’s name, and choose “Copy to new Dashboard”.



This will create a cloned dashboard that is named “Copy of [original dashboard name]”. To edit the new dashboard, navigate to it, and click the “Edit Dashboard” button. A list of available widgets will appear in a tray to the right. Widgets can be deleted, resized, moved, and edited as desired.



To create a custom dashboard from scratch, click the “Add Dashboard” link under “Dashboards” in Cloud Security. You will need to name your dashboard to get started. Once your dashboard has been named, a blank canvas will appear. Widgets can be dragged into the canvas from the tray on the right, or added by clicking the “+” in the center of the canvas. Widgets can be resized, moved to different positions, and customized to meet your requirements.



To position a widget, click the [:::] at the top of the widget. To resize a widget, drag the sizing handle in the bottom right corner [//]. To configure a widget, clone it, or delete it, use the icons found in the upper right hand corner of the widget. To learn more about dashboards and how to customize them, visit the [Get Started with Dashboards](#) page in TechDocs.

Reporting

Setting up custom compliance standards can help users to logically group their RQL policies to fit the scope of what they are trying to achieve. An example of a use case could be to have




a compliance standard that is generally used for a specific business unit of a customer environment


There are four different formal reports you can configure from Prisma Cloud; two Alert Reports, Compliance Report, and Cloud Security Report. The [Alert Reports](#) consist of a Cloud Security Assessment Report and a Business Unit Report.


Cloud Security Assessment Report: The Cloud Security Assessment report is a PDF report that summarizes the risks from open alerts in the monitored cloud accounts for a specific cloud type. The report includes an executive summary and a list of policy violations, including a page with details for each policy that includes the description and the compliance standards that are associated with it, the number of resources that passed and failed the check within the specified time period. This report can be useful to share with management, outside third party organizations for assessment purposes, or just a quick review.

To create a Cloud Security Assessment report, navigate to Reports -> Create Report -> Cloud Security Report.



To create a Cloud Security Assessment report,, navigate to Alerts and click the “Create Alert Report” button, and check the “Business Unit Report” radio button.





Create Alerts Report [X]

Name:

Report Type:

☒ **Cloud Security Assessment**
 A PDF report that summarizes the risks from open alerts in the monitored cloud accounts for a specific cloud type.

OR

☐ **Business Unit Report**
 A .csv file that includes the total number of resources that have open alerts against policies for any compliance standard, and you can generate the report on-demand or on a recurring schedule.

Cloud Type:

Account Group (Optional):

Cloud Account (Optional):

Cloud Region (Optional):

Time Range:

Email Address(es) (Optional):

☐ Use custom email notification template


Schedule:


☒ One Time ☐ Recurring

[Cancel] [Save Report]

Business Unit Report: The Business Unit report is a .csv file that includes the total number of resources that have open alerts against policies for any compliance standard, and you can generate the report on-demand or on a recurring schedule. This .csv file allows you to open the report in Microsoft Excel to be able to filter, sort, or utilize any Excel features or you can upload this into a third party tool such as a SIEM tool. You can opt to create an overview report which shows you how you're doing across all your business units, or get a little more granular about each of the cloud accounts you want to monitor. You can also generate the Business Unit report to review policy violations that are associated with specific compliance standards.

To create a Business Unit Report, navigate to Alerts and click the "Create Alert Report" button, and check the "Business Unit Report" radio button.





The 'Create Alerts Report' dialog box contains the following fields and options:

- Name:** A text input field.
- Report Type:** Two radio button options:
 - ☐ **Cloud Security Assessment**: A PDF report that summarizes the risks from open alerts in the monitored cloud accounts for a specific cloud type.
 - ☒ **Business Unit Report**: A .csv file that includes the total number of resources that have open alerts against policies for any compliance standard, and you can generate the report on-demand or on a recurring schedule.
- Account Group (Optional):** A dropdown menu.
- ☐ **Detailed Report** ⓘ
- ☐ **Filter by Compliance Standard**
- Email Address(es) (Optional):** A text input field.
- ☐ **Use custom email notification template**
- Schedule:** Two buttons: **One Time** and **Recurring**.
- Buttons:** **Cancel** and **Save Report** at the bottom right.

Compliance Report: The third type of report that you can generate with Prisma Cloud is a [Compliance Report](#). Prisma Cloud natively includes multiple industry known Compliance Standards such as NIST, CIS, PCI-DSS, HIPAA, and others. You can create compliance reports based on a cloud compliance standard for immediate online viewing or download, or schedule recurring reports so you can monitor compliance to the standard over time.

From a single report, you have a consolidated view of how well all of your cloud accounts are adhering to the selected standard. Each report details how many resources and accounts are being monitored against the standard, and, of those, how many of the resources passed or failed the compliance check. This report can be useful for dedicated Compliance teams, management, or to assist during a security assessment or audit.

Create Compliance Report

Name

Cloud Type

AWS

Date

Select Date

Compliance Standard

Account Group (Optional)

Cloud Account (Optional)

Cloud Region (Optional)

Email Address(es) (Optional)

☐ Use custom email notification template

Schedule

Cancel

Save Report



To create a Compliance Report, navigate to Compliance and click the “Create Compliance Report” button.

Compliance in Workloads

The Center for Internet Security (CIS) publishes recommendations or best practices that should be adhered to when setting up machine images, servers, containers, etc. In addition to these CIS compliance standards the twistlock team has added best practices. As an example, to be HIPAA compliant you may be required to implement certain CIS standards. Computers can fail or block builds to adhere to compliance standards but there is no remediation in place.

Example CIS Kubernetes Compliance Benchmark:

- Ensure admin.conf file permissions are 644 or more restrictive
- Audit: run the stat command to display the permissions
- Remediation: run chmod to set the appropriate permissions

Prisma Compute automates the audit steps checking for misconfigurations in various environments and exposes each benchmark check as a discrete configuration item in a compliance rule allowing for individual lower level checks to be enforced in an environment.

- Critical and high severity alerts only -- medium and low are ignored by default
- Rule creation compliance standard are in the top right corner of creating an alert
- Compliance checks are set by PANW not CIS, customer needs to verify compliance checks based on need
- Anything the customer can script out to run as a compliance check can be integrated as a compliance check (powershell or bash) to be run against each image
- Checks can be organization specific or standard hardening guidelines
- Checks are safely executed against new container instance in a sandbox environment

Trusted Images:

- Allow for specific images, registries, or image base layers to be deemed as trusted
- Allows for choice of which image can be trusted

- Rules can be setup so that the use of untrusted images trigger an alert or the image is blocked by being prevented from launching
- Feature allows organizations to protect against the deployment of arbitrary and potentially malicious images from the internet such as from Docker Hub
- 'Not a get out of jail free card'
 - Even if image is trusted it can still fail at run time based on compliance checks and rules

Cloud Platform Discovery and Compliance

- Check and scan resources based on customer cloud provider credentials
 - Checks if there is a scan in the CSP that coordinates with a Prisma Cloud Scan
 - Scans:
 - Managed Kubernetes
 - Image Registries
 - Serverless Functions
 - Runs set of Twistlock Lab created checks against a cloud environment

Cloud Compliance is defense in depth:

1. Scan vulnerabilities (check critical or high CVEs)
2. Compliance checks for unpublished CVEs
 - a. Available dashboard in UI -- Defend Compliance -- Rule matching is the same (top to bottom)
 - i. Checks high and critical severity -- can create alert, turns off, blocks image spin up, or fails -- checks that nothing is set to block

Manage – Cloud Accounts:

- Used for onboarding cloud providers
- Initially scans for resources to show if there is a protection or not

Compliance Tab:

- Compliance Explorer: shows all non-compliance critical and high alerts
- Clicking on compliance check will show resources that are not compliant
- Container level view shows each container in environment -- name, image, host, cluster, compliance heat map
 - Open up to see best practice within compliance -- make sure to check version number for compliance
 - Check compliance standard vendor with description in Prisma Compliance Container View
 - Compliance vendors will show audits, remediation, etc.
 - Our Audits in Prisma are written in Go so they are not available to see in UI
- Image Level view -- container compliance is different from image level compliance
 - Can be custom, twistlock, CIS (most CIS are written at the container level)

- Every image scan for vulnerability follows compliance at the same time following the compliance rules
- Typically, vulnerabilities are scanned first, the written into rules or policies for compliance
- Function Level View
- Trusted Images: list of images with trust status
- Cloud discovery: Scans all cloud workloads, picks out environments' resources (AWS: EC2, registries, lambda functions, EKS, ECS) count in each region and how many are scanned by twistlock or compute

Vulnerability Policy Strategy - Alerting to Blocking

Having clear visibility into which vulnerabilities to focus on is crucial to identify four items: 1.) the criticality of the alarm, 2.) the owners of an image resource, 3.) the actual image resource that is triggering the vulnerability alarm and 4.) if there is a fix for the vulnerability. In order to facilitate visibility, the [CQ](#) code to pull the images and sort the images by owner tag and image can be found [here](#).

Without measurability, it is impossible to manage in a proper manner. In addition, scanning images in the CI/CD pipeline is crucial in preventing further vulnerabilities from polluting the environment. Once the report is generated by the [CQ](#), it will be clear who the maintainers are, and the images they own. The report will also show the vulnerabilities and their severity levels. When initially generating a report for a customer who has a non-existent vulnerability management policy, their vulnerabilities with critical and high severity levels can be over 1 million alerts. The [CQ](#) code is flexible enough to filter out only specific vulnerability levels and fix dates. When filtered, to simply critical, the list is much more manageable. Also the report will show that across multiple images, the same package is triggering alarms. Therefore updating one package across multiple images could reduce the vulnerabilities significantly.

Owners

It is vital to know who the owners are for resources. Many customers have no visibility into these cloud resources because they do not have an enforcement of tags in their environment. If ownership is not clear, it will be a huge challenge to manage the environment. Also, the tags can be utilized in defining scopes in Prisma Cloud to block vulnerabilities from polluting the environment.

Socialization

It is important to socialize the project of converting from alerting to blocking. Having key stakeholders onboard with the objective of the project is important. To that end, having a regular cadence call multiple times a week is critical to review which images and which packages are triggering alarms. In these meetings, the package updates can be measured to ensure readiness of blocking. If alarms are still not being addressed, during the meeting those issues can be discussed and determined if the blocking for that image is a go or no go for specific business issues.

Change Control

As with any good business process, it is important to have a good documentation system to know what changes affect end consumers of an application. With a change control, it also outlines a backout plan if an outage occurs.

Scope

Defining scopes that you can place the images in to block is important. Having one for alerting only and one for blocking is very important. When converting images to blocking mode causes an outage, move the image back to the scope of alerting and triage the situation later. Once the image has been moved to blocking mode, the CI pipeline scanning is critical in keeping the environment unpolluted.

Search and Investigate

Investigating in Prisma Cloud allows you to look further into security and operational information and investigate alerts and incidents within your cloud environment(s). There are different types of searches; some searches are based on RQLs or Resource Query Language. RQL is similar to the widely-known SQL, and it performs configuration searches into how your cloud resources are deployed. The visibility you gain here saves you time rather than going into your cloud provider portal by helping you understand what is wrong with your cloud resources by investigating directly in the raw data, tracing back in time, and identifying risks across all your cloud service providers.

There are multiple types of queries:


1. Asset Config: Use [Asset Configuration Queries](#) to search for the configuration of the cloud resources.
2. Audit Event: Use [Audit Event Queries](#) to search and audit all the console and API access events in your cloud environment.
3. Permissions: Use [Permissions Queries](#) to gain visibility into the permissions of your cloud resources.
4. Network: Use [Network Flow Queries](#) to search real-time network events in your environment.
5. Network Configuration: Use [Network Configuration Queries](#) to identify assets exposed to the Internet.
6. Asset: Search across all your cloud assets based on findings and vulnerabilities
7. Application: Prioritize security risks, adding application context to findings
8. Application Asset: Identify your software delivery chains and explore engineering attack surface
9. Vulnerability: Identify top vulnerabilities that have been discovered in your cloud environment.

Some questions that might come to mind when securing your cloud environment are:


- Do I have any critical S3 buckets that are publicly accessible?
- Are cloud resources in my environment missing critical patches from vulnerabilities?
- What activities has a root user performed that may not have been necessary?

You can create custom RQL queries as well as search into ones that are already structured, such as within a default policy, or if you navigate to Investigate → Search → Saved Searches, you


will find hundreds of saved searches from not just users within your organization, but also ones that are saved by default within the product that customers and users can find helpful to their organization. See Saved Searches in the following image below.



You can also view the Saved Searches under the **Query Library** Tab. Any helpful Saved Searches can be instantly turned into a policy by selecting the document icon shown in the image below.




You can also use the NLP Query Launcher (Search bar) to get query recommendations from the Query Library by describing what you want using a natural language.



There are now two types of query builders: simple and advanced. Advanced Query Builder uses pure RQL syntax, and Simplified Query Builder is a graphical tool that helps you build or edit a search by selecting and clicking all available parameters based on the search type.

Simplified Query Builder:



Cloud Security

Home Dashboards Reports Inventory Compliance Alerts Investigate Governance

INVESTIGATE Search Query Library Background Jobs

Click here to give your search a name


FIND Asset Configuration Since 24 hours ago WHERE Cloud Type is AWS Cloud Service is Amazon EC2 + Add Reset Query

Simple Advanced Search

Search table data...

Asset Name	Service	Account	Region Name	Last Modified	Deleted	Actions
	Amazon EC2		AWS Virginia	14 hours ago	false	
	Amazon EC2		AWS Virginia	2 days ago	false	
	Amazon EC2		AWS Virginia	8 days ago	false	

Advanced Query Builder



Cloud Security

Home Dashboards Reports Inventory Compliance Alerts Investigate Governance

INVESTIGATE Search Query Library Background Jobs

Click here to give your search a name

FIND Asset Configuration Since 24 hours ago

Simple Advanced Search

config from cloud.resource where cloud.type = 'aws' AND cloud.service = 'Amazon EC2'


Search table data...

Asset Name	Service	Account	Region Name	Last Modified	Deleted	Actions
	Amazon EC2		AWS Virginia	14 hours ago	false	
	Amazon EC2		AWS Virginia	2 days ago	false	
	Amazon EC2		AWS Virginia	8 days ago	false	


Query Type	Simple Mode	Advanced Mode (RQL)
Asset	Yes	N/A
Asset Configuration	Yes	Yes
Application Asset	Yes	N/A
AppSec Asset	Yes	N/A
Vulnerability	Yes	N/A

Permission (IAM)	WIP	Yes
Network Config (CNA)	WIP	Yes
Network	WIP	Yes
Audit Event	WIP	Yes

RQL queries can be created and start with clicking into the search bar, where you're provided options to select the type of the query. Once you select the query type, you can use the simple (filter-based) or advanced (RQL text-based) mode for building the query.



To view the RQL query behind a default policy, navigate to Governance and search for the specific policy you want to look into further. In this example, we'll look at "AWS Security Group allows all traffic on RDP port (3389)". Once you find the policy within the Governance page, click on the edit icon (indicated by the pencil icon on the far right of the policy). You will see the RQL in the second option of the "Edit Policy" popup, and you can click on the Launch Search; you will be forwarded to Investigate > Search, and the RQL will be ready for you.



Edit Config Policy

Add Details ☒ Create Query ☒ Compliance Standards ☒ Remediation ☒

Run Build

New Search Saved Search

AWS Security Group allows all traffic on RDP port (3389)_RL

Launch Search

```
config from cloud.resource where cloud.type = 'aws' AND api.name= 'aws-ec2-describe-security-groups' AND json.rule = isShared is false and (ipPermissions[?any((ipRanges[*] contains 0.0.0.0/0 or ipv6Ranges[*].cidrIpv6 contains ::/0) and ((toPort == 3389 or fromPort == 3389) or (toPort > 3389 and fromPort < 3389)))) exists)
```

Previous Next

Cloud Security Home Dashboards Reports Inventory Compliance Alerts Investigate Governance Settings

INVESTIGATE Search Query Library Background Jobs

AWS Security Group allows all traffic on RDP port (3389)_RL

FIND Asset Configuration Since 1 week ago

Simple Advanced Search

config from cloud.resource where cloud.type = 'aws' AND api.name= 'aws-ec2-describe-security-groups' AND json.rule = isShared is false and (ipPermissions[?any((ipRanges[*] contains 0.0.0.0/0 or ipv6Ranges[*].cidrIpv6 contains ::/0) and ((toPort == 3389 or fromPort == 3389) or (toPort > 3389 and fromPort < 3389)))) exists)

Examples of Common RQL Searches

Policy Description	RQL
AWS: List EC2 instances with a public IP address	config from cloud.resource where api.name = 'aws-ec2-describe-instances' and json.rule = publicIpAddress exists
Find workloads with vulnerability 'CVE-2015-5600'	network from vpc.flow_record where dest.resource IN (resource where finding.type IN ('Host Vulnerability') AND finding.name = 'CVE-2015-5600') and bytes > 0
Azure SQL instances that allow any IP address to connect to it	config from cloud.resource where cloud.service = 'Azure SQL' AND api.name = 'azure-sql-server-list' AND json.rule = firewallRules[*] contains "0.0.0.0"
List Azure Storage accounts (can be used for Azure flow log checks)	config from cloud.resource where cloud.type = 'azure' AND api.name = 'azure-storage-account-list' addcolumn location
List VPCs that do not have Flow Logs enabled	config from cloud.resource where api.name = 'aws-ec2-describe-vpcs' as X; config from cloud.resource where api.name = 'aws-ec2-describe-flow-logs' as Y; filter ' not (\$Y.resourceId equals \$X.vpcId)'; show X;

Useful Documentation for RQL:

- Review the [RQL Example Library](#) for useful queries to run and utilize. It allows you to modify easily since you can edit the existing RQLs
- Review the [RQL Operators](#) document to understand the different capabilities within RQL searches
- Review the APIs that are ingested for reference to build out custom investigate searches and policies
 - [Alibaba APIs Ingested by Prisma Cloud](#)
 - [AWS APIs Ingested by Prisma Cloud](#)
 - [GCP APIs Ingested by Prisma Cloud](#)
 - [Azure APIs Ingested by Prisma Cloud](#)
 - [OCI APIs Ingested by Prisma Cloud](#)
- Visit the [RQL FAQs](#) for additional help and information

IAM Security

The IAM Security Module provides net-effective permissions to cloud infrastructure resources based on policies that are configured out of the box. This allows Prisma Cloud administrators the ability to rightsize these permissions quickly, which reduces the risk of compromise from identity credentials.

The IAM Security module is enabled on the Subscriptions tab in the Prisma Cloud console (click on Learn Mode, and Activate to enable). Once IAM Security has been activated, you will be able to run RQL queries. Verify this on the Investigate tab.

Be sure to save searches that you have created to save time the next time you need to run the query again. If you need to analyze permissions offline you can download the results of a query in CSV format.

Integrate IAM Security with your IdP to calculate permissions for your SSO provider (e.g. Okta, Azure AD).

[Manually remediate IAM security alerts](#) by going to:

1. Alerts > Overview
2. Select the violating policy
3. Policies that can be remediated are indicated by a  icon.
4. Under the Options column, click the Remediate button.
5. Prisma Cloud will make recommendations for CLI commands to run in your CSP.

[Create an Alert Rule for Cloud Infrastructure](#) and follow the instructions for configuring a custom python script on AWS or Azure; this will allow you to manage auto-remediation for IAM alert rules using the messaging queuing service on the respective CSP.

Some best practice guidelines to consider for IAM Security

- Which users have access to resource X?
- What accounts, services and resources does the user name@domain.com have access to?
- What are the cross account permissions between my accounts?
- Can any users outside of group C access resources in region D?
- What roles are not configured according to best practices?
- What resources can be effectively assessed by the public?
- Which compute workloads have permissions that are not actually being used?
- Resolve all over-permissive policies and enforce least-privilege from now on

Data Security Posture Management (DSPM)

Data Security provides protection against sensitive data existing in places where it should not exist, or that is overly accessible to internal and external users. This is accomplished through data classifications (e.g. Secret, Internal) sensitivity, (e.g. PII, PHI etc.) as well as scanning data for sensitive information, objects that are overly permissive, and even objects infected by Malware.

In addition to protections for AWS S3 buckets and Azure Storage blobs, Prisma Cloud DSPM adds capabilities to protect cloud services, including databases (e.g. CloudSQL, RDS), file system services (e.g. EFS and Filestore) and other services (e.g. Snowflake and Microsoft 365).

Prisma Cloud Data Security (DSPM) is an agentless platform that allows for the discovery, classification, protection and governance of sensitive data. By reviewing suspected issues and Risks, Alerts can be generated to provide real time information about policy violations. Additionally, DSPM can be integrated with WildFire to provide defense against malware

This section provides detailed information about using Prisma Cloud DSPM, from onboarding accounts to remediating alerts.

Note: It is important to understand that the DSPM module (Dig) will be deployed to a tenant and will use the Darwin UI/UX language. This may cause some confusion, as there will be elements of the legacy “Data Security” module that will share the same terminology as the DSPM module (e.g. Data Security). For example, subscribing to Data Security by clicking your Persona icon->View Subscriptions will enable the subscription for the legacy Data Security Module and not DSPM.

- Prerequisites


In order to be prepared to onboard your accounts to DSPM and begin using it,

you will need to consider the following:

- **Does it matter which Infrastructure as Code (IaC) language is used?** When performing the onboarding steps, in most cases, you will be able to choose to use CloudFormation or Terraform (whichever is appropriate).
- **Which approach would you like to use to monitor your data assets?** Prisma Cloud DSPM allows you to onboard your organization level account, but it is also possible to onboard individual accounts. Additionally, you can add a new [Orchestrator](#) or use Monitored Accounts (e.g. add to an existing Orchestrator).
- **Be sure to understand the permissions and networking requirements needed:** The list of permissions that Prisma Cloud DSPM needs are extensive. Each Cloud Service Provider (CSP) has its own requirements. Additionally, it is important to review the architecture and remember that, at a minimum, Prisma Cloud DSPM will need to scan data assets and report its findings to the Prisma DSPM console.
 - Review the permissions required for [AWS](#)
 - Review the permissions required for [Azure](#)
 - Review the permissions required for [GCP](#)
 - To get an architectural overview as well as additional requirements, review the [Components](#) of Prisma DSPM.
- **What types of data assets do you hope to monitor using Prisma Cloud DSPM?** The list of supported assets is more extensive than just S3 buckets and storage blobs! But in order to get the most value from DSPM, it is important to understand all of [the assets that it supports](#) so that you have realistic expectations.
- **Account Onboarding**

Onboarding accounts and projects are fairly straightforward. However, there are a few things to take note of for each supported CSP: To begin, choose “Data Security” from the platform drop-down menu at the top left and then choose “Preferences’.

Next, choose the “Integrations” tab. There are various options for onboarding Cloud accounts and services. It is also possible to configure notification services, such as Email, Slack, Webhooks, and SNS here. Additionally, Malware services, such as Crowdstrike and Wildfire can also be configured here.



Note: Once the onboarding process is complete, you will have the ability to enable AI-SPM features and functions when available.

○ AWS

To onboard an AWS account, click the “Configure” button in the AWS tile on the Settings page. Click the “Add” button, and choose whether you want to onboard an Org account, or an individual account. It is important to note that AWS Orgs can only be added once an Orchestrator has been onboarded. On the next page, you will need to decide whether you want to create a new Orchestrator, or add to an existing Orchestrator. You will also need to determine if you want to use Terraform or a CloudFormation Template (CFT) to automate the provisioning. Review the [documentation](#) for more information. Here are some additional notes:

- If Terraform is selected, a link is provided which adds the account information to a Terraform module. Click [this link](#) to download a zip file containing the Terraform that is used for onboarding individual accounts.
- If you wish to onboard an AWS org account, you will need to add the Org account to an existing Orchestrator. Additionally, it is not possible to onboard an AWS Org account using Terraform.
- If onboarding an AWS account using Cloud Formation Templates (CFT), follow the directions in the [documentation](#) and review the stack set before it is applied.
- Once onboarding is complete, the following will have been created in AWS:
 - A Role named *DigSecurityScannerRoleuse*
 - A Permission Policy named *DigSecurityScannerPolicy*
 - An S3 bucket named *dig-security-[numerical string]*

- **Note:** This bucket will not be accessible by the AWS root user.
- **Azure**

Onboarding Microsoft Azure accounts is slightly different than the process of onboarding AWS accounts. While the [documentation](#) is fairly straightforward here are a couple of things to consider:

 - Azure accounts can be onboarded by installing an Azure Enterprise app. This must be done by someone who has Administrative access at the Tenant Level. This should preferably be done by a Global Administrator.
 - Azure accounts can also be onboard at the tenant level. Please keep in mind that the user performing the onboarding should have the required permissions to create roles at the tenant level and approve enterprise applications, preferably holding the Global Admin role.
 - To Onboard your Azure account:
 - Paste in your Tenant name
 - Approve the DSPM enterprise application. Click the “Approve” button will open Azure in another tab. You will need to log in using your Tenant Admin account.
 - Acknowledge the terms and generate the template.
- **GCP**

To onboard a GCP account, click the “Configure” button in the GCP tile on the Settings page. Click the “Add New” button, and choose whether you would like to add a project or an Organization. Note that adding a GCP org is only available once an orchestrator has been onboarded. You can add a project via CloudShell or Terraform. If you are onboarding a project, you will need to have the project name ready. A label for the project is required, so you will need to choose one of the following labels: “Development”, “Staging”, “Testing”, or “Production”. You will also need to determine if you want to use an existing Orchestrator, or if you want to deploy a new Orchestrator. If you are onboarding an Org account, you will need to have the Organization ID. You can only onboard a GCP Org to an existing Orchestrator.

 - [This](#) is a sample of the shell script that runs the CloudShell commands. Note that the project name is “sampletest”
 - [This](#) link will download a zip file containing a sample of the Terraform template used when onboarding a GCP project.
 - Once your project has been onboarded, the following Service Accounts will exist:
 - dig-use1-orchestraotr-<string>@<project name>.dspm.iam.gserviceaccount.com
 - dig-use1-readonly-<string>@<project name>.dspm.iam.gserviceaccount.com
 - dig-use1-scanner-<string>@<project name>.dspm.iam.gserviceaccount.com
 - Once your project has been onboarded, the following Roles will exist:
 - role_dig_use1_collector1_<string_<string>
 - Cusrole_dig_use1_orchestrator1_<string_<string>
 - role_dig_use1_orchestrator2_<string_<string>
 - role_dig_use1_orchestrator3_<string_<string>

- role_dig_use1_orchestrator4_<string_<string>
 - role_dig_use1_readonly1_<string_<string>
 - Custole_dig_use1_scanner1_<string_<string>
 - Custole_dig_use1_scanner2_<string>
 - Custole_dig_use1_scanner3_<string>
- **Onboarding Other Cloud Platforms**
 - **Snowflake**

It is possible to add projects from Snowflake so that they can be scanned by Prisma Cloud DSPM. You can find detailed information [here](#). Please take note of the following before you begin:

 - Snowflake accounts can only be onboarded under existing orchestrators. These can be in Azure, GCP, or AWS.
 - Two accounts will need to be created in Snowflake for DSPM.
 - You will need the following from Snowflake:: Account ID, Organization ID, Username, Login Name and Password.
 - Access to a cloud secrets manager.
 - **Microsoft 365**

It is possible to onboard Microsoft 365, which will allow you to scan OneDrive, Sharepoint, and Enterprise Applications. [This document](#) provides more information. There are a few things to remember:

 - Like with Azure, you will need to approve an Enterprise Application.
 - You can only onboard a Microsoft 365 account under an existing orchestrator.
 - **On-Prem File Shares**

You can use DSPM to scan internal file shares. These are shares located on an internal network that is managed by Active Directory. While [this document](#) contains detailed information, here are a few things to consider:

 - You can only onboard On-Prem Shares under an existing orchestrator.
 - You will need access to Active Directory in order to perform the onboarding process.
 - **Account or Project Exclusions**

If you would like to exclude a project or accounts from an onboarded Organization or CSP, you can do this by going to Settings -> Exclusions. Click on the “Add Now” button to add the Exclusion:

Add new exclusion project

If the project belongs to the organization, you are able to exclude it from DDR alerts and DSPM risks regarding foreign projects. Note that this exclusion will impact all relevant risk rules and DDR policies. You can manage the excluded projects in the settings.

Cloud

GCP

Project ID

mygcpproject

Add Exclusion

Cancel


- **Configuration**

- Licensing

DSPM consumes one credit per data store and 1 credit per 1TB of volume for Snowflake. For more information, see the [credit guide](#).
- RBAC

Prisma Cloud DSPM does not observe the RBAC structure from the rest of the platform (e.g. Settings-> Access Control). You must be a System Administrator in order to view and/or use DSPM.
- Scanner Settings

It is possible to change the services that an orchestrator scans. To do this, Select “Preferences”, and then the cloud account type (e.g. AWS, Azure, or GCP). Select an account by clicking the check box to the left of its name. Then select “Change Scanning Settings” from the “Options” menu. From the sidecar, choose the appropriate service and then select Enable or Disable it:



- Removing Onboarded Accounts
Once you have onboarded an account, it is not possible to delete it. To discontinue use of DSPM in a cloud account or project, remove the roles and service account permissions and groups that were created during onboarding (Review the Terraform or Cloud Formation Templates for this). Additionally, All DSPM created resources are tagged with “dig-security-true”. Removing resources with this tag will effectively discontinue use of DSP.
- Classifications
Out of the box classifications exist (e.g. PII, Sensitive, PII, etc.). However, additional classification labels can be added for files. Additionally, custom configuration rules can be created based on labels. These rules can be configured with additional data classes, file types or permissions. This can all be configured under the “Preferences” tab. New classification rules can also be created based on search filters in “Findings”.
- Monitoring Issues
The “Monitoring Issues” tab under “Settings” displays potential issues scanning resources in cloud accounts. These issues could indicate potential problems with data discovery:

Settings

Integrations Classification Settings Classification Rules Exclusions API **Monitoring Issues** Help

monitoring issues

Status ▾ Open × Filter by

Showing 4 out of 4 results


First Discovered	Issue Name	Issue Source	Cloud	Status
Apr 24, 2024, 4:35 PM	Access denied to read an asset	Resource Policy	aws	Open ▾
May 14, 2024, 11:00 AM	Access denied to list assets	SCP	aws	Open ▾
May 14, 2024, 11:00 AM	Access denied to list assets	Resource Policy	aws	Open ▾
Apr 24, 2024, 5:46 PM	Elastic IP Quota Exceeded	Quota Limit	aws	Open ▾

- API
 - To use API calls, generate a key using “Setting->API” and test using the “Assets” endpoint. At the time of writing, this is the only endpoint supported.
- DDR Policies
 - The DDR policies can be found on the DDR Policies tab. There is not a way to create custom policies.
- Alerts/Notifications
 - Alerts and notifications can be configured by adding integrations. Alerts can be customized so that they are received based on specific findings, severities, etc. while customizing the configuration.

Auto-Remediation

Remediate policies - from Prisma Cloud console can be done manually or by auto-remediation

1. Auto-Remediation requires:
 - a. Onboard the cloud account (or update the onboarded cloud account) with Enabled **Remediation**.



Add Cloud Account

aws Configure Account

Organization Details
Add the necessary details for Prisma Cloud to apply permissions for capabilities and create stack.

Account ID: Account Name:

☐ **Remediation**
Remediate misconfigured resources.

Advance Configuration Settings
Manage advance settings for management and member accounts

Apply permissions for cloud account

[Create IAM Role](#) or [Download IAM Role CFT](#)

[View Steps for more information](#) [View Steps for more information](#)

IAM Role ARN
eg - arn:aws:iam::123456789012:role/abcd


Select All Member Accounts
Prisma Cloud recommends selecting all member accounts (default). Disable to include or exclude specific accounts from monitoring. Enabled

Account Group
[Default Account Group](#)

For product documentation please click [here](#)

[Previous](#) [Next](#)

- b. Policies (default or custom config policies only) are configured with auto-remediation. You can view all of the remediable policies by navigating to Governance and adding the filter “Remediable” and setting it to “True”.




Policy Name	Policy Description	Policy Type	Policy Subtype	Cloud	Severity	Category	Remediable	Class	Actions
AWS EC2 instance with the creation of a new Group with attached policy permission	With access to 'iam:CreateGroup', 'iam:Att...	IAM	Permissions	aws	Critical	Risk	Yes	Exposure	🔍 📄 ⚙️
Char Copy of Critical - AWS S3 Object Versioning is disabled	This policy identifies the S3 buckets whic...	Config	Run	aws	Critical	Risk	Yes	Exposure	🔍 📄 ⚙️
AWS EC2 instance with the creation of a new Role with attached policy permission	With access to 'iam:CreateRole', 'iam:Att...	IAM	Permissions	aws	Critical	Risk	Yes	Exposure	🔍 📄 ⚙️
AWS EC2 instance with the creation of a new User with attached policy permission	With access to 'iam:CreateUser', 'iam:Att...	IAM	Permissions	aws	Critical	Risk	Yes	Exposure	🔍 📄 ⚙️
Critical - AWS S3 Object Versioning is disabled	This policy identifies the S3 buckets whic...	Config	Run	aws	Critical	Risk	Yes	Exposure	🔍 📄 ⚙️
AWS EC2 instance with IAM write access level	This policy identifies IAM write permissi...	IAM	Permissions	aws	High	Risk	Yes	Exposure	🔍 📄 ⚙️
Azure VM instance with Run command permissions	With access to 'Microsoft.Compute/Virtu...	IAM	Permissions	azure	High	Risk	Yes	Exposure	🔍 📄 ⚙️
AWS Lambda Function with data destruction permissions	With access to 's3:DeleteBucket', 's3:De...	IAM	Permissions	aws	High	Risk	Yes	Exposure	🔍 📄 ⚙️
Okta User with IAM permissions management access level	This policy identifies IAM permissions m...	IAM	Permissions	aws	High	Risk	Yes	Exposure	🔍 📄 ⚙️
AWS Lambda Function with IAM permissions management access level	This policy identifies IAM permissions m...	IAM	Permissions	aws	High	Risk	Yes	Exposure	🔍 📄 ⚙️

Displaying 1 - 25 of 226

Rows 25 Page 1 of 10

- c. Alert Rule (a new rule or modify an existing one) with auto-remediation enabled.



2. Understand how auto-remediation works since it pushes CLI commands automatically once an alert is detected and could possibly create unwanted changes.
3. Prisma Cloud Automatically runs the remediation CLI to resolve the policy violations for all open alerts regardless of when they are generated.
4. If you are modifying an existing alert rule (enable auto-remediation) that includes non-remediable policies, those policies will no longer be included in the alert rule.
5. It's helpful to test out auto-remediation in sandbox environments.
6. It's helpful to start with auto-remediation by configuring a new alert rule with auto-remediation enabled, limiting it to an account group, and assigning some policies to it. You can add more account groups and policies to the alert rule after that.

The following lists recommend starting policies for Auto Remediation within the three major cloud providers.

AWS

Typically when first starting out with auto-remediation, you will want to focus on low hanging fruit configurations. Down below are 5 AWS policies which we recommend getting started with:

Policy	Description
--------	-------------

AWS Security Group allows all traffic on RDP port (3389)	Used as the remote access port for Microsoft Windows, it is advised to keep this port open to only trusted IP addresses. This policy checks the configuration of your security groups and ensures that this port is not allowed from any IP address (0.0.0.0/0).
AWS Security Group allows all traffic on SSH port (22)	Most commonly used as the SSH port for Linux, it is highly recommended to lock down access to only trusted IP address ranges. Leaving this port open to 0.0.0.0/0 can expose your instance to brute-force type attacks.
AWS Amazon Machine Image (AMI) is publicly accessible	Unless for very specific use-cases, AMIs should not be made public as they may contain sensitive information. Having a public facing AMI would allow anyone with an AWS account the ability to launch your AMI image.
AWS EBS snapshots are accessible to the public	EBS snapshots are typically used for backups or for security tools. From an EBS snapshot, a volume can be created which can then be attached to an instance, allowing access to the contents of that volume.
AWS CloudTrail logging is disabled	AWS CloudTrail is a service that enables governance, compliance, operational & risk auditing of the AWS account. It is a compliance and security best practice to turn on logging for CloudTrail across different regions to get a complete audit trail of activities across various services.

Azure

Policy	Description
--------	-------------

Azure Network Security Group allows all traffic on SSH port 22	As a best practice, restrict SSH solely to known static IP addresses. Limit the access list to include known hosts, services, or specific employees only. This policy will remove the entry for port 22 which allows access from anywhere from your NSG.
Azure Network Security Group allows all traffic on RDP Port 3389	As a best practice, restrict RDP solely to known static IP addresses. Limit the access list to include known hosts, services, or specific employees only. This policy will remove the entry for port 3389 which allows access from anywhere from your NSG.
SQL databases have encryption disabled	Transparent data encryption protects Azure database against malicious activity. It performs real-time encryption and decryption of the database, related reinforcements, and exchange log records without requiring any changes to the application. This policy will automatically enable encryption on the databases which have this disabled.
Azure Key Vault is not recoverable	The key vault contains object keys, secrets and certificates. Accidental unavailability of a key vault can cause immediate data loss or loss of security functions (authentication, validation, verification, non-repudiation, etc.) supported by the key vault objects.

Google Cloud Platform

Policy	Description
GCP Firewall rule allows all traffic on SSH port (22)	Allowing access from arbitrary IP addresses to this port increases the attack surface of your network. It is recommended that the SSH port (22) should be allowed to specific IP addresses. This policy can remove the entry exposing port 22 to 0.0.0.0/0 from your firewall rules.
GCP Firewall rule allows all traffic on RDP port (3389)	Allowing access from arbitrary IP addresses to this port increases the attack surface of your network. It is recommended that the RDP port (3389) should be allowed to specific IP addresses. This policy can remove the entry exposing port 3389 to 0.0.0.0/0 from your firewall rules.
GCP Firewall rule allows all traffic on SMTP port (25)	This policy identifies GCP Firewall rules which allow all inbound traffic on SMTP port (25). Allowing access from arbitrary IP addresses to this port increases the attack surface of your network. This policy can remove the entry exposing port 25 to 0.0.0.0/0 from your firewall rules.
GCP Firewall rule logging disabled	This policy identifies GCP firewall rules that are not configured with firewall rule logging. Firewall Rules Logging lets you audit, verify, and analyze the effects of your firewall rules. When you enable logging for a firewall rule, Google Cloud creates an entry called a connection record each time the rule allows or denies traffic. This policy can automatically enable this functionality on the firewall rules.
GCP Storage log buckets have object versioning disabled	This policy identifies Storage log buckets which have object versioning disabled. Enabling object versioning on storage log buckets will protect your cloud storage data from being overwritten or accidentally deleted. This policy can enable object versioning features on all storage buckets where sinks are configured.

Some best practices to keep in mind here :

1. Create custom compliance frameworks if needed for specific organizational needs
2. Enable disposition of new default policies that are added with Prisma Cloud product
3. updates by reviewing Enterprise Setting option
4. Setup Trusted IPs and Prisma Cloud Login IPs to reduce false positive alerts
5. Create and manage access keys as needed for certain cloud tools and third party integrations
6. Policies
 - a. What policies are recommended to start with for remediating for customers -
 - i. Steps on Alert Burndown
 1. Understand what alert burndown means - bringing down the number of alerts so it's manageable (important alerts being looked at and resolved regularly)
 2. Create a plan
 - a. Focus on what's important to the organization (high volume alerts or high severity alerts)
 3. Understand alert actions
 - a. Dismiss, snooze, remediate, investigate
 4. Effort vs. result
 - a. Low effort in lowering many alerts (ex. Audit events)
 - b. High volume alerts (think of severity, impact from remediation, complexity of remediation)
 - i. Ex. policies:
 1. Config : AWS Security groups allow internet traffic
 2. Audit Event : IAM configuration updates
 3. Anomaly : Unusual user activity
 4. Network : Internet exposed instances

"none" -m)

- Fargate
 - <https://pan.dev/compute/api/post-defenders-fargate-json/>
- Detailed guides in github for twistcli for deployment of images
 - <https://github.com/PaloAltoNetworks/prisma-cloud-compute-sample-code>

Upgrade/Redeploy the Defenders

After the console upgrade, either SaaS or Self-hosted, the defenders need to be upgraded by the customer.

It is recommended to automate the defender deployment process. With automation, the defender upgrade will be much smoother in a large-scale environment.

If the customer has an automation process for the container deployment, the customer should integrate the defender deployment process into the existing pipeline process.

Defender deployment can be automated with API calls (defender yaml, helm chart, or fargate defender) as referenced.

To upgrade the container defenders manually, you will need to generate and download the yaml file or helm chart from the upgraded console. You will need to update the current yaml or helm chart and apply the new version to upgrade the defenders on your cluster.

If you are using the installation script from the console, we recommend using the uninstall script to remove defenders first. Then you can run the install script copied from the console on the cluster.

Here is the detailed steps for the upgrade:

- Host Defenders:
 - [Defender and Prisma Cloud components Upgrade Process Page](#)
- Kubernetes Defenders:
 - [Upgrade Defender with DaemonSets Page](#)

At times it may become necessary to remove Defender because of a change in the Console or, perhaps, because the Container Defender was installed when the Host Defender was needed, or vice versa. In this case, remove the defender by doing the following:

- On Linux systems that use systemd:
 - Stop the service: `systemctl disable --now twistlock-defender-server.service`
 - Removing the following folders (if present): `/opt/twistlock` and `var/lib/twistlock`
- On Windows servers:
 - Stop the twistlock service
 - Remove the following folder: `c:/Programfiles/Twistlock`
 - Removing the following registry folder and keys:

```
HKEY_LOCAL_MACHINE - System - ControlSet001 - Service  
twistlock-defender-server.service
```
- On Docker
 - Stop the running twistlock defender pod: `docker stop <pod id>`
 - Delete the image: `docker image -rmi <image id>`

Note: It may be necessary to use the -f flag to force the image to be removed.

App-Embedded Defenders and serverless defenders need to be upgraded manually. These defenders are also recommended to integrate to the pipeline of task deployment or serverless script deployment process.


Runtime Protections

Configure, enable, and customize Prisma Cloud policies. Familiarize yourself with and customize compliance requirements.


Compliance

A key goal when it comes to customizing compliance is to work with your customer's Compliance team to mirror their policies inside of Prisma. You want to get the specific frameworks a customer is using to make the alerts meaningful. You can also scope these policies out by using collections. Make sure to work directly with teams and use collections before blocking containers to avoid impact to production.

These are examples of the policy frameworks for deployed containers and images:



These are examples of some the policy frameworks for hosts:



Runtime Models

One key goal is minimizing the amount of work you're required to do to manage runtime defense. Leverage the models that Prisma Cloud can automatically create and manage. Because behavioral learning for model creation is mature technology for Prisma Cloud, in most cases, you won't need to create auxiliary rules to augment model behavior. There will be some exceptions. For example, a long-running container that changes its behavior throughout its lifecycle might need some manually created rules to fully capture all valid behaviors. This is atypical for most environments, however, as containers that need to be upgraded are typically destroyed and reprovisioned with new images.

If you do need to create runtime rules, here are some best practices for doing so:

Minimize the number of rules — Creating static rules requires time and effort to build and maintain; only create rules where necessary and allow the autonomous models to provide most of the protection.

Precisely target rules — Be cautious of creating rules that apply to broad sets of images or containers. Providing wide ranging runtime exceptions can lower your overall security by making rules too permissive. Instead, target only the specific containers and images necessary. Don't use

wildcard (*) in the whitelist or blacklist because it can interrupt the execution of legitimate services.


Name rules consistently — Because rule names are used in audit events, choose consistent, descriptive names for any rules you create. This simplifies incident response and investigation. Also, consider using Prisma Cloud's alert profile feature to alert specific teams to specific types of events that are detected.

Rules in alert action — It is recommended to start configuring the runtime rules in alert action and after you are comfortable with the outputs you can change the action to prevent or block.

Rules testing — In case the customer wants to implement a runtime policy with block or prevent actions. It is recommended to test this policy behavior in a test environment before pushing it to production. For example, testing Kubernetes cluster or test namespace.


Runtime Policy Configuration

- Enable the use of ML models in case of container policy by enabling the automatic runtime learning option.



- Provide a descriptive rule name.
- Avoid using a wide scope based on a cluster name for example and make it more specific by providing a namespace name and image or container value.
- Use **Prisma Cloud Advanced Threat Protection** intelligence feed, to apply malware prevention techniques across processes, networking and filesystem.
- In case the container is running inside a Kubernetes cluster it is better to enable the **Kubernetes attack** option to monitor attempts to directly access Kubernetes infrastructure from within a running container. In case a container is developed for some use case to communicate with Kubernetes API, it needs to be excluded from the selected scope to avoid false positive alarms.

- **Suspicious queries to cloud provider APIs** can be enabled to monitor access to cloud provider metadata API from within a running container. In case a container is developed for some use case to communicate with a cloud provider API, it needs to be excluded from the selected scope to avoid false positive alarms.



- Use **Advanced Malware Analysis** based on Wildfire malware analysis engine, to detect malware. Currently Prisma Cloud Compute uses WildFire for file verdicts only in the following scenarios for Container runtime / CI:
 - ELF files written to a linux container file system in runtime.
 - Shared objects are not examined via WildFire.
 - File must be smaller than 100MB (WildFire limit).
 - You can submit up to 5000 files per day, and get up to 50,000 verdicts on your submissions to the WildFire service.
 - Wildfire is supported on Linux only. Windows containers and hosts aren't currently supported.

WildFire malware detection

- **Use WildFire for runtime protection** — Enable WildFire malware scanning in runtime for containers and hosts.
- **Use WildFire for CI compliance checks** — Enable WildFire malware scanning for containers CI checks.
- Choose the closest WildFire cloud region
- **Upload files with unknown verdicts to WildFire** — Determines whether files with unknown verdict will be sent to WildFire for full analysis. When off, WildFire will only provide verdict for files that have been uploaded to WildFire via a different client.
- **Treat grayware as malware** — Use a more restrictive approach and treat files

with grayware verdict as malware.

WildFire malware detection

Use WildFire integration to enhance malware detection capabilities

Configure wildfire	✓ Active
Enable runtime protection	On ☒
Enable CI compliance checks	On ☒
WildFire cloud region	Global (US) ▼

Advanced configuration	
Upload files with unknown verdicts to WildFire (recommended)	On ☒
Treat grayware as malware	On ☒

- Processes:
 - Review the learned processes in the container model and whitelist or blacklist the process in the rule based on the business need.
 - Configure the Anti-malware and exploit prevention option on alert mode for testing and then change to prevent or block mode.
- Networking:
 - Review the learned networking ports and domains in the container model and whitelist or blacklist it in the rule based on the business need.
 - Configure the Anti-malware and exploit prevention option on alert mode for testing and then change to prevent or block mode
- File System:
 - Review the learned file system paths in the container model and whitelist or blacklist it in the rule based on the business need.
 - Configure the Anti-malware and exploit prevention option on alert mode for testing and then change to prevent or block mode

Custom Runtime Rules

- Precise way to describe and detect specific runtime behaviors.
- Can help fill in a lot of gaps on hosts since our model is more focused on services
 - Example: Preventing writes to a particular file system on a host.
- Make sure to test specific use cases before telling customers what you can and can't do.
- Sample built-in runtime rules are usually good enough in POC.

Defend / Runtime

Container Policy

Host Policy

Serverless Policy

App Embedded Policy

Custom Rules

The following table is a library of available rules that can be individually added to runtime policies.

33 total entries

Search custom runtime rules

Type	Rule Name	Description	Owner	Modified	Used	Actions
network-outgoing	Check if someone is connecting to suspi...	Check if someone is doing connections t...	kyates_paloaltonet...	Feb 3, 2020 8:11:37 AM	Yes	
filesystem	Check if a non-root user is writing into etc.	Checks if a process is writing into the file...	kyates_paloaltonet...	Feb 3, 2020 8:11:37 AM	Yes	
kubernetes-audit	Detect Kubernetes authorization failures	It will detect any Kubernetes authorizati...	kyates_paloaltonet...	Feb 3, 2020 8:11:37 AM	No	
kubernetes-audit	Check for exec or attach to a pod	Audit any attach or exec to a pod	kyates_paloaltonet...	Feb 3, 2020 8:11:37 AM	No	
processes	Check if a non-root user is using nmap?	This rule will check if a non-root user is u...	kyates_paloaltonet...	Feb 3, 2020 8:11:36 AM	Yes	
kubernetes-audit	Twistlock Labs - GKE - Tampering with T...	Audit changes to Twistlock objects, such ...	system	Feb 3, 2020 8:10:28 AM	No	
kubernetes-audit	Twistlock Labs - Tampering with Twistloc...	Audit changes to Twistlock objects, such ...	system	Feb 3, 2020 8:10:28 AM	No	
processes	Twistlock Labs - Running privileged proc...	Detect privileged management tools star...	system	Feb 3, 2020 8:10:28 AM	No	
network-outgoing	Twistlock Labs - Common data exfiltratio...	Detect usage of common data exfiltratio...	system	Feb 3, 2020 8:10:28 AM	No	

Runtime Radar Utilization

Cloud Radar View

Allows user to view geographic locations of onboarded cloud accounts as well as filter these geographic locations by CSP region, CSPs, resources that are protected within each cloud account including (functions, clusters, registries, app embedded and hosts), and individual cloud accounts

Upon selecting a specific location's resources the user can view the compliance posture of the resources within that location with options to protect the resources if available should they not be compliant. Users can select a resource through the list of available resources to view the resource details including (name, version, runtime, ARN, and protected status), as well as the option to view listed compliance violations with detailed descriptions of each violation upon individual selection of the unprotected resource. The view of non-compliant resources includes the resource's onboarded cloud credential, ID, severity, result, title and associated collection



Host Radar View

Allows users to view available hosts that Prisma has access to. There is an option to color code the hosts by vulnerability severity, runtime behavior compliance, and overall compliance. Users can filter by specific collections, clusters, CSP regions, CSP hosts, only connected hosts, and overall severity level

Container Radar View

Allows users to view individual clusters of containers with available egress/ingress connections. Upon selecting a cluster, individual nodes can be viewed in detail. The console provides a view into the risk summary, environment, and networking information

Risk Summary:

Provides view of vulnerabilities, runtime, compliance, and WAAS for individual container selected with a link to that part of the console within Compute

Shows Image, Image ID, Cluster, Namespace, and Service Account

Environment:

Shows containers and hosts of selected cluster's selected node

Networking Information:

Shows the connected inbound and outbound ports and protocols to the node as well as the outbound IP addresses

Serverless Radar View

Allows user to view connected Serverless functions within cloud environments

Upon selection of a node representing a serverless function, the user can view the services that the function has permission to as well as general info including the function's CSP, region, and runtime name.

Upon selection of permitted services associated with the serverless function the user can view the resource association (AWS ARN or equivalent) within the service as well as the associated permissions that the function can perform within the associated services

There is also an option, if applicable, to see the scanning levels that are not natively associated with the serverless function within Prisma Cloud Compute. As an example, if the serverless function is not scanned by vulnerabilities or associated with a compliance standard, there will be an option to protect the function with the natively available infrastructure not yet associated within PCC

Settings

Provides the user with the options to toggle container network monitoring and host network monitoring



Allows the user to configure network objects by adding subnets to scan that include the network object name and the CIDR block of the associated subnet.


Note: If resources do not appear in their expected radars (e.g. container, hosts, cloud etc.) ensure that the appropriate port (8084) is open in all of the appropriate places (firewalls, proxies etc.) between the Defender and Console. To test this connectivity, curl the Console address from the host or container Defenders in question.


Registry Scanning

Prisma Cloud can scan container images in public and private repositories on public and private registries. When you configure Prisma Cloud to scan a registry, you can select the scope of defenders that will be used for performing the scan job.

Registry Scan Behavior

Prisma Console controls the registry scan with assigned defenders. Console scans one registry at a time. If multiple registries are configured to scan, Prisma Cloud console will scan one registry. Once completed the first registry scanning, then move to the following registry. This registry scanning behavior is not configurable.





For example, here are three scenarios to configure multiple registries scan with Prisma Cloud. Depending on the configuration, you can compare the total scanning time. Let's assume the following conditions:

- There are four registries to scan.
- Each registry has eight images.
- It takes 10 seconds to scan the image.


Note: The timing in these scenarios is hypothetical for comparison purposes.

Scenario A:

Build a defender collection for registry scan and assign four defenders. The configures all registry scan profiles with this defender collection. In this scenario, all four defenders scan the images in the registry in parallel. Prisma Cloud would then pick registry A and scan it first. It will take 20 seconds to complete. When the work for registry A is done, Prisma Cloud scans B and moves to C and D. The total time to


finish the four registry scans is 80 seconds.

Scenario A: Assign Defender Collection 1 to the scope of the each registry scan profile.




Scenario B:

Build 4 defender collections with a member of the defender. Configure each registry scan with a specific defender collection. One defender will scan the assigned registry. Prisma Cloud picks the registry A to scan with the defender collection 1. It will take 80 sec to complete the scan. Then Prisma Cloud scans B and moves to C and D. The total time for the scan is 320 seconds.



Scenario C:

Build 4 defender collections with four defenders. Configure each registry scan with a specific defender collection. First, Prisma Cloud picks registry A to scan with defender collection 1. It will take 20 sec to complete the scan with four defenders. Then, Prisma Cloud scans B and moves to C and D. The total time for the scan will be 80 seconds, which is the same as scenario A. However, because Prisma cloud scans registry sequentially, while defender collection 1 is scanning registry A, the other 12 defenders in collections 2, 3, and 4 are running but scanning.



In the above example, scenario A is the most efficient method. It is best to deploy a dedicated defender pool for scanner purposes and create a dedicated scanner collection. It is the best recommendation for small to medium-sized image registries.

Once the scale of the image registry gets large enough and increases scanning time, it is time to begin investigating Defender performance statistics (CPU/MEM & Defender local host activity, etc.) for optimizing their registry scanning throughput. There will be many different aspects to consider depending on the registry environment. Consider increasing the number of defenders in the collection and the number of scanners defined to registry configuration to improve throughput and reduce scan time.

The console will not scan registries simultaneously but sequentially, so creating multiple dedicated defender pools per registry is not generally recommended. However, scenario C can be considered if you have registries with different OS types or reside in various regions.

From the 22.12 release or later, you can add a maximum of 19,999 registry entries for scanning.

Here is more reference for the large scale registry scanning guide:

TechDoc: [Registry with a Large Scale](#)



Scan Scheduling

Depending on the number of images, you may also need to adjust the scanning frequency from something like 24 hours to 72 hours to give the Defenders enough time to complete the scans. If scans roll into the next scanning window, all progress will be stopped on current scans and start from the beginning.

Capacity planning

To ensure optimal performance for Prisma Cloud defenders performing registry scanning, it is crucial to allocate sufficient resources. Each Defender conducting registry scanning should be allocated a minimum of 2GB of RAM and 20GB of storage.

The number of defenders required to scan a registry depends on the registry's size and the frequency of scans. The typical scan speed of a defender is approximately 100MB/s.

For example, scanning a registry with a size of 10TB using a single defender would take approximately $(10\text{TB}/100\text{MB})/(60*60) = 27.7$ hours. To scan it once every 24 hours, a minimum of 2 defenders is required. Please note that these calculations provide general guidance, and actual performance may vary depending on factors such as network conditions.

Trusted Images

As organizations get more familiar with their images and environment, they typically leverage our Trusted Images feature to control developer access to a specific registry or even specific images or layers. [Trusted Images](#) ensures that developers are using verified or approved sources for their images, as well as providing a straightforward way to implement the best practices for container security.

The trusted image function lets you explicitly define which images are permitted to run in your environment. If an untrusted image runs, Prisma Cloud emits an audit, raises an alert, and optionally blocks the container from running.

It is recommended to specify the images it trusts. Declare trust using objects called Trust Groups. Trust Groups collect related registries, repositories, and images in a single entity.


Then, for writing policy rules. It's recommended to use registries or repositories for the trusted groups if they are an organization's standard golden images.

As a best practice, the default rule, Default - alert all should be maintained, and it should be the last rule in your policy as a catchall rule. The default rule matches all clusters and hosts (*). It will alert the images that aren't captured by any other rule in your policy.

Assuming the default rule is in place, the policy is evaluated as follows:

- A rule is matched: The rule is evaluated.
- A rule is matched, but no trust group is matched: The image is considered untrusted. Prisma Cloud takes the same action as if it were explicitly denied.

- No rule match is found: The default rule is evaluated, and an alert is raised for the image that was started. The default rule is always matched because the cluster and hostname are set to a wildcard



The screenshot shows the Prisma Cloud console interface. At the top, there are navigation tabs: Code Repositories, Containers and images, Hosts, Functions, **Trusted images**, Custom, and Cloud platforms. Below the tabs, there are two sub-tabs: **Policy** and Trust groups. The main section is titled "Trusted images rules" and has a toggle switch set to "On". Below this, there is a description: "Create rules to control which sources are trusted and which images are allowed to run. Policy changes take effect immediately. Visibility in Console views requires a rescan." There is a search bar labeled "Filter trust rules" with "1 total entry" next to it. To the right of the search bar are buttons for "Export", "Import", and "+ Add rule". Below this is a table with the following columns: Rule name, Effect, Owner, Scope, Modified, Actions, and Order. The table contains one entry: "Only Allow - Prod Registry" with an effect of "alert", owner "hsong", a blue square icon for scope, and a modified date of "Feb 10, 2022 10:41:48 AM".

Trusted Image Policy

Base Images


Filtering out vulnerabilities whose source is the base image can help your teams focus on the vulnerabilities relevant for them to fix. For Prisma Cloud to be able to exclude base image vulnerabilities, first identify the base images in your environment.

The base image should be specified in the following format: registry/repo:tag. You can use wildcards for the tag's definition. Excluding base image vulnerabilities is currently not supported on Windows images.

It is recommended to select base Images for the most commonly used throughout the different projects/applications and add them to the Base Images.

Once the base images are identified, they can be filtered out from the reports by using the 'Exclude base images vulns' filter.

The maximum number of base images that can be in scope is 50, where each base image is represented by a digest. If there are 50 base images in scope, and the scanner discovers a new base image, the oldest is purged and replaced with the newest.



Web-Application and API Security (WAAS)

WAAS Traffic Inspections Modes

Prisma Cloud supports three different modes to deploy WAAS for containers or hosts. These inspection modes provide great flexibility when you deploy WAAS on your Cloud environment with various security requirements and use cases.

Typically WAAS is deployed with Prisma Cloud Defenders. **In-Line mode**, defenders can operate as a transparent HTTP proxy and inspect and control the traffic to the applications. With **Out-of-Band mode**, the defenders can only inspect the traffic to the applications while not proxying with the traffic flow.

However, in **Agentless mode**, Prisma Cloud monitors the traffic for the remote applications with no Defenders deployed.

Out-of-Band and Agentless modes have no latency cost. These modes can't control the traffic but only send alerts to the Prisma Console.

The table briefly shows the differences among those modes and best use cases.

Prisma Cloud Features	WAAS Inline	WAAS Out of Band	WAAS Agentless

Capability	Inspect all Incoming Traffic Control Traffic	Monitor traffic tapped traffic	Monitor mirrored Traffic
Control Traffic	Yes	No (alert only)	No (alert only)
Defender	Deployment required on Protecting workloads	Deployment required on Protecting workloads	Observer Agent w/ host defender (Prisma Cloud deploys automatically)
Use Case	Enabling best security with complete control of traffic to the web applications and API endpoints.	Without the inline deployment, inspecting traffic to web applications and discovering API endpoints is required.	Without the inline deployment, inspecting traffic to web applications and discovering API endpoints is required.


In-Line (Proxy) Mode

In-line mode inspects all incoming traffic and forwards or blocks them to the protected application per the WAAS rules.

In-line mode provides the best security for web applications and APIs because it has complete control of traffic flows in real-time. However, real-time traffic monitoring may require more resources than Out-Of-Band monitoring. Therefore, testing in-line mode configuration in QA or staging environments before deploying in production is recommended to prevent unforeseen issues.

[Prisma Cloud WAAS Performance Benchmark](#)

In-line mode is required to deploy a Defender to the workload, as shown below.




Here are the detailed instructions for In-Line deployment.

- [Deploy WAAS In-Line for Containers](#)
- [Deploy WAAS In-Line for Hosts](#)

Out-of-Band Mode

Out-of-band mode requires Defender deployment on the workload to monitor the protected applications. Out-of-band mode inspects traffic by tapping the network interface on the relevant application ports. Since Defender passively listens to tapped traffic, out-of-band deployments do not interfere with the customer's application traffic. Another benefit is that this mode can make API discovery without defining an app in the rule. The below diagram shows the traffic flow of out-of-band WAAS mode.




Here are the detailed instructions for Out-of-Band deployment.

- [Deploy WAAS Out-of-Band for Containers](#)
- [Deploy WAAS Out-of-Band for Hosts](#)

Agentless Mode

WAAS Agentless (VPC Traffic Monitoring) can be configured without running Defender on the workload environments. Agentless mode is a good WAAS solution for the Cloud account onboarded to Prisma Cloud with Agentless scanning enabled only. Currently, this mode supports only AWS.

When configuring the WAAS agentless rule, Prisma Cloud will install an EC2 instance called Observer for the VPC traffic mirroring. AWS VPC traffic mirroring feature copies the traffic from the source EC2 instance to the Observer instance within the VPC for WAAS inspection. Auto scaling group is supported for the Observer deployment with the Maxwell version (30.00.140) and higher. The diagram shows the architecture and traffic flow of out-of-band WAAS mode.



Here is the detailed instruction for Agentless deployment.

- [Deploy WAAS Agentless](#)

App Definition (API Protection)

Utilizing Open API / Swagger documents to create the General App Setup and API protection is highly recommended. Not only do Open API and Swagger documents make it easy and consistent to deploy application definitions, they also support the following DevOps philosophy:

- 1.) Open API / Swagger is a good form of documentation on Restful APIs for any given developer group and organization
- 2.) Support automated updating as API and infrastructure changes.
- 3.) Allow for automation of deployment in a consistent manner.
- 4.) Allow versioning for documentation and rollback of configuration if there is a problem with the deployment.

If Open API / Swagger documentation is not available, define the API functions manually. To enable API protection, define the base path and the App port. The App port is the port that the application is listening on rather than the external port that the calling application will use. The more specific to describe your application, there is a better chance of protecting your application. Rather than defining a host, use the scoping mechanism to clearly identify the application to protect.

With the base path, the standard path for RestAPI is “/api/v1”.

App definition

App firewall

DoS protection

Access control

Bot protection

Custom rules

Advanced settings

App ID

app-0986

OpenAPI/Swagger spec

Import

You can define an app by importing an OpenAPI/Swagger spec file or by manually specifying its API endpoints. Importing a spec file will overwrite all previously specified API endpoints, including any that were manually defined.

Endpoint setup

API protection

Description (optional)

test

API endpoint discovery

On

View TLS settings

Protected endpoints

1 total entry

+ Add endpoint

HTTP host	App port	Base path	TLS	HTTP/2	Actions
*	80	*	Off	Off	...

HTTP host

* Add [host];[external port]

App port ?

80

Base path

* Add [base path]

Cancel

Save

Specify your APIs by defining the signature of your endpoint as far as what is an acceptable input parameter. More details about API protection can be found [here](#).

App definition
App firewall
DoS protection
Access control
Bot protection
Custom rules
Advanced settings

App ID

OpenAPI/Swagger spec

i You can define an app by importing an OpenAPI/Swagger spec file or by manually specifying its API endpoints. Importing a spec file will overwrite all previously specified API endpoints, including any that were manually defined.

Endpoint setup
API protection

API protection - Parameter violations

API protection - Unspecified path(s)/method(s)

API resources

1 total entry

Path	Methods	Actions
✓ /	PUT, POST, DELETE, OPTIONS, HEAD, PATCH, GET	🗑

Define Rule and Scope

To properly profile your application, clearly defining a scope is very important. Think broad enough to encompass the application yet specific enough to define the application. When you build a collection for WAAS rules, you have to specify an image minimum to assign to the scope of the WAAS rules. Prisma Cloud will use the collection to identify the application to apply the WAAS rules. If you need behavior to be different based on labels, or a cluster where the container will be running, create a separate collection and apply the given WAAS rule to that collection. More details about WAAS rule resources and application scope can be found [here](#).

Network Controls

Allowing only sources that need access to the application reduces the attack radius of a given application. In order to do this, define a CIDR block that is allowed to access the application, and allow only that CIDR block. Prevent for all other CIDR blocks by setting the “Prevent” action for all others.

[App definition](#) [App firewall](#) [DoS protection](#) **[Access control](#)** [Bot protection](#) [Custom rules](#) [Advanced settings](#)

Network controls [HTTP headers](#) [File uploads](#)

IP access control

Control inbound traffic by IP address. Specify IP addresses in [Network lists](#) On ☒

Blocking mode Allowed ☒ Blocklisted ☐

☒ Allow My CIDR block
Specify network lists

☒ Action for all others Alert Prevent

HTTP Headers

With http header inspection, it is looking for a key value pair to inspect for every single http call to the application. The header name is case insensitive and the value could be either allow or deny. For the values, they can be either explicit or a wildcard character (*) can be used for the following three use cases:

- Begins With
 - E.g., "Mozilla/5.0"
- Contains
 - E.g., "*(X11; Linux x86_64)*"
- Ends With
 - E.g., "*Safari/537.36"

Application Profiling

The best way to protect an application is to be as specific in the application profiling. For example, if you are looking for a person at an airport and you have never met this person before, the more accurate detail that describes the person will clearly identify the person with the least amount of false positives. Likewise with an application, the greater detail the WAAS component has to profile the application, the better it will be in blocking attacks. To enforce specific headers such as a specific token type to retrieve data, you inspect the header for a specific token format by means of a regular expression. For example, if in any given request to an application you wanted the following:

- Specific headers with specific values
 - The headers can be inspected with the exact values
- An application token in the header with specific format (e.g., “s.34x7797bxe3p118923”)
 - The header value can be inspected with a regular expression such as the following would match the above token:
 - `/^s\.[0-9a-zA-Z]+/g`
 - Utilizing an online regex tool such as the following [link](#) will assist in creating an application profile that accurately describes your application.
- Body content
 - The body content can look for specific values using regular expressions to block malicious intent
 - Regular expression can be used to inspect body content to ensure it is a valid web request and reject it if it is not.
 - The body content can look for payload that is only valid
 - The body content can be inspected in combination with an action. For example the payload will only be inspected if the http request action is a POST or PUT and will not be examined if the action is a GET or DELETE.

Scoping

Rules are basic building blocks to enforce a specific action but scopes are used to bind a rule to a specific application. Building proper scopes

i Ban is applied by client IP

Firewall settings

Protection	Mode	Exceptions	Actions
SQL Injection	Disable Alert Prevent Ban		⚙️
Cross-Site Scripting (XSS)	Disable Alert Prevent Ban		⚙️
OS Command Injection	Disable Alert Prevent Ban		⚙️
Code Injection	Disable Alert Prevent Ban		⚙️
Local File Inclusion	Disable Alert Prevent Ban		⚙️
Attack Tools & Vulnerability Scanners	Disable Alert Prevent Ban		⚙️
Shellshock	Disable Alert Prevent Ban		⚙️
Malformed HTTP Request	Disable Alert Prevent Ban		⚙️
Prisma Cloud Advanced Threat Protection	Disable Alert Prevent Ban		⚙️
Detect Information Leakage	Disable Alert Prevent Ban		⚙️
Cross Site Request Forgery Protection	On ☒		⚙️
Clickjacking Prevention	On ☒		⚙️
Remove Server Fingerprints	On ☒		⚙️

Cancel Save

Load Balancers

Configuration for load balancers need to have the HTTP Header X-Forwarded-For in each request for WAAS to be able to determine if the HTTP call needs to be blocked based on source country origin. Most load balancers have that header enabled but without that header country origin blocking will not work.

Runtime Collections

You will need access to an existing SaaS or self-hosted tenant that has the compute capabilities enabled. This means that a user correlated to a role that has sufficient permissions has been assigned to you and you have the ability to authenticate with the console.

Collection Functionality Overview in Compute

In terms of collection scoping within the console, the first step will be to create a role that limits user access to the specific modules that they need. In this example, access has been limited to the Defend, Vulnerabilities and Compliance modules through the role that were created:

Create new role

Access to Console UI
On

Radars
Defend
Monitor
Manage

Please note!
Roles that access policies typically require permissions for [collections](#) and [credentials](#) to work properly.

Vulnerabilities & Compliance Policies

Roles with access to container or host compliance policies require permissions for custom compliance policies
☐ Read
☐ Write

Code Repositories Vulnerabilities Policies	☒	☐
Images/Containers Vulnerabilities & Compliance Policies	☒	☐
Hosts Vulnerabilities & Compliance Policies	☐	☐
Serverless & App-Embedded Vulnerabilities & Compliance Policies	☐	☐
Cloud Platforms Policies	☐	☐
Custom Compliance Policies	☐	☐

Runtime Policies

Roles with access to container, host or app-embedded runtime policies require permissions for custom rules
☐ Read
☐ Write

Cancel
Save

Next a collection was created to associate with the users that would be encompassed by this role, and only included a code repository and an image:

Create new collection

**Please Note**

When creating or updating collections, the set of image resources that belong to a collection isn't updated until the next scan. To force an update, manually initiate a rescan.

Name	Example-Collection-Repos-Containers
Description	Enter a description
Color	
Containers	k8s_dvwa-web_dvwa-web-5db8d745b6-qlfv_dvwa_eb93759c-9c70-4bac-a065-46c931cb9efb_0 × Specify a container
Hosts	* Specify a host
Images	* Specify an image
Labels	* Specify a label
App IDs (App-Embedded)	* Specify an app ID
Functions	* Specify a function
Namespaces	* Specify a namespace
Account IDs	* Specify an account ID
Code Repositories	spring-projects/spring-boot × Specify a repository

Then the collection and the role were created with a new user:

Create new user

Username	Example-User
Authentication method	Local LDAP SAML OAuth 2.0 OpenID Connect
Password	
Role	Example-Role-Repos-Containers ▼
Permissions	■ Example-Collection-Repos-Containers ▼




Please Note

If a role allows access to policies, users will be able to see all rules and all collections that scope rules under the Defend section even if the user's view of the environment is restricted by assigned collections

[Cancel](#)[Save](#)

As you can see by the warning sign, in the Defend module of the console where rules are set, one limitation of the console is that specifically within this module users are able to view all rules regardless of assigned collections. This is only within this module and will be shown in a subsequent step. The other modules are completely filtered by the collections that are assigned to a user.

One thing to note is that this will allow console users to see all rules with associated infrastructure in terms of the associated views within the Defend module. For the Radar and Monitor modules, the views are filtered by the individual collections assigned to a user. Once this role was assigned with a collection, the user got the following view of the console when logged in with the new user:



Code repositories vulnerability policy

Vulnerability rules let you raise alerts when vulnerable packages are found in code repositories. Rules are evaluated at scan-time against all GitHub code repositories in scope, where scope is specified by the table under Code repositories scan scope.

GitHub repositories scan scope


Code repositories scan scope 3 total entries
Specify the git repositories to scan.

Filter Code repository Settings by keywords and attributes

Provider	Type	Credential	Code repositories	Paths to exclude	Actions
GitHub	Public	GitHub Access Token	keylowe/FoodTrucks		
GitHub	Public	GitHub Access Token	spring-projects/spring-boot		
GitHub	Public	GitHub Access Token	PaloAltoNetworks/demo_build_repo_scan		

Webhook settings
Configure GitHub webhooks (with the URL below) to rescan your repositories on push events.

As previously mentioned, in this specific module users, when assigned to have access to this module via role, can see all of the rules regardless of the collection that they are assigned to. In the console's current design the Defend module is meant to be managed and reviewed by vulnerability management teams or managers. Only read access was granted to this module. When we click into editing a scope, we get a non-editable view as can be seen below:



Edit scope

Provider: GitHub

Type: Private Public

Credential: GitHub Access Token

Repositories: spring-projects/spring-boot
Specify the full repository name, i.e. OWNER/NAME

Excluded manifest paths: File paths to exclude (supports pattern matching)


Advanced settings

Close


Missing permission: Credentials store
Depending on your role, you might not see some elements in this page or be able to perform certain actions.


Don't show again for this page

By pivoting into the Images tab within the Defend module, we can see the rules associated with images but, again, since this role has read-only access to this module, the fields are not editable:




This behavior is also mirrored within the Compliance view of the Defend module for code repositories as well as containers:






Next, to show how the Monitor module filters by collection, read access was assigned to the image runtime behavior view of the Monitor module for the user's role that was used to log in..




After logging back in with the new user's account, we can see that the Monitor module and the runtime view have been added to the new user's view of the console:




The screenshot shows the 'Monitor / Events' page in the Prisma Cloud console. The left sidebar has 'Monitor' selected, with 'Events' and 'Runtime' sub-items. The main content area is titled 'Container audits' and shows a 'Container audits: 0' status. Below this, there's a 'Collections:' dropdown menu with 'Example-Collection-Repos-Containers' selected. A line chart titled 'Container audits over time' shows 'There is no data to show'. At the bottom, there's a table with columns: Image, Namespace, Cluster, Audit message, Audit time, and Total. The table also shows 'There is no data to show'.

You can also see that the only collection that is available within this view is the collection assigned to the user. This shows how module access can be limited by the scope of the collections that a user is assigned to. In addition to the events view this collection scope has also limited the runtime view within the Monitor module.



The screenshot shows the 'Monitor / Runtime' page in the Prisma Cloud console. The left sidebar has 'Monitor' selected, with 'Events' and 'Runtime' sub-items. The main content area is titled 'Container models' and shows a 'Container models' status. Below this, there's a 'Collections:' dropdown menu with 'Example-Collection-Repos-Containers' selected. A line chart titled 'Container models over time' shows 'There is no data to show'. At the bottom, there's a table with columns: Image, Cluster, Namespace, OS, Entrypoint, State, Forensi..., Collections, and Actions. The table also shows 'There is no data to show'.

As an additional example the Radar module was added, and the container view within this module to the new user's role. When logging back into the console using the new user's credentials, the radar module was limited in functionality to the collection that was assigned to the user.



To summarize, assigned collections combined with roles can limit the scope of what users are able to view within the console. The limitation of the scope of this filtering is the Defend module which, in the present state of the console, is meant to be only assigned to security managers and vulnerability management teams who set the threshold of the risk appetites of environments through severity thresholds.

Hopefully this information helps with your use cases. In terms of newly defended resources, they will be assigned to the OOTB “All” collection that is accessible by system administrators and can be further assigned to the collection associated with a user, team, or environment by system administrators or roles that are granted write access to the collections view of the System module.

In the console defenders are the primary point of ingesting data from the cloud resources on which they are deployed. They push updates via port 8084 to the console and the results of what are pushed are compared to the console’s threat intelligence stream in addition to the various vulnerability threshold rules, runtime rules, and WAAS rules associated with how the defender is scoped into collections. Defenders are the base level of information ingestion into the console, with typically a one-to-one mapping to various resources deployed in a customer cloud environment, such as containers, hosts, registries, and serverless functions.

The information ingested into the console from defenders is made available on a need to know basis through the role associated with the user that is viewing the console as well as the collection that the infrastructure has the defender is deployed to. The role associated with the user can grant read or write permissions to individual modules of the console, as well as individual views within each module. The collection associated with the infrastructure can further limit the information that each user is privy to within the console. All defenders scan their respective infrastructure that they are deployed to but not all of that information is available to each user.

When a user is associated with a role, collections can be associated with either the role that the user is assigned to or with an individual user. As an example, in a DevOps environment there could be a development, QA/testing, and production environment with different business units interacting with each environment. A role could be created for the development team that is associated only with collections that encompass infrastructure in the development environment. As an additional example, there could be a use case of an audit, where an auditing team would need access to the production environment. In this use case a role could be created with full read access to every collection associated with the production environment. Through creating the scope of what defender streams are able to be viewed in the console through collections console administrators can limit the information available to each business unit's use case.

Defenders are built so that the information ingested to the console is tailored to the type of infrastructure on which the defender is deployed. Associating a collection with a single defender would necessitate that each resource in the collection would have to be the same (i.e. all hosts, containers, registries, or serverless functions). Collections allow for the information ingested by defenders to be available regardless of the resource type. However, in the use case of a defender being associated with a collection, the ability for the collection to encompass multiple types of defended infrastructure would be limited by how the defenders are configured for different types of resources. The granular, one-to-one mapping of defenders with individual cloud resources in customer infrastructure allows for, in the present state of the console, the information ingested by the defenders to be grouped on a need to know basis via collections or roles associated with specific collections.

General Best Practices When Implementing Collections

- Naming conventions (code words, hashes, etc.)
- Designing scope with information streams (upstream and downstream) as well as levels of scope in mind (business unit, team, application, environment)
- Looking at available options with the API
- Creating collections for specific incidents/events

Cortex XDR Agent for Cloud



What is the Cortex XDR agent for Cloud?

Cortex XDR and Prisma Cloud offer a unified Cloud Security Agent for Linux. The XDR agent for the cloud provides [true CDR \(Cloud Detection and Response\)](#) capability, as well as vulnerability and compliance coverage on Linux and container cloud environments.

Cortex XSIAM integrates Cloud Detection and Response to provide a single interface for Security Operation Center (SOC) analysts to monitor and respond to cloud threats.

The Cortex XDR Agent expands to include Prisma Cloud features, eliminating the need for multiple security agents and improving visibility across security programs.

The Cortex XDR tenant will manage and deploy the XDR agent for Cloud. Policy management, data, and alerts are managed first between the Cortex XDR tenant and the Prisma Cloud tenant.

Cortex XDR Agent for Cloud is deployed from the Cortex XDR console, where the SOC team will also be able to view, investigate, and respond to cloud incidents, leveraging extensive visibility through logs and events collected from workloads.

XDR Agent for Cloud will send vulnerability and compliance information to the Prisma Cloud tenant, where the Cloud Security team can view it alongside other information (e.g., misconfigurations, code issues) surfaced by Prisma Cloud.

Prerequisites

To enable the capabilities of the Cloud Security Agent, the Prisma Cloud tenant must be paired with a Cortex XDR tenant. Pairing is one-to-one, with the two tenants being in the same region. Please enable this feature by contacting your Customer Success Engineer (CSE) or a Prisma Cloud Solutions Architect (SA).

- Cortex XDR Cloud per Host, or Cortex XSIAM License
- Prisma Cloud Enterprise Edition
- Cortex and Prisma tenants must be allocated to the same GCP region
- Cortex XDR 3.10 or XSIAM 2.2, Cortex Agent 8.2.1 or above, Kernel 4.18 and above

Supported Workload Platform

Currently, the XDR agent for cloud supports only Linux endpoints, Kubernetes nodes, and Openshift as DaemonSet. X86_64 and ARM64 CPU Architecture are supported models. Here is the tech docs page of the [Cortex Support Matrix](#). Prisma Cloud Defender and XDR agent for cloud should not be deployed simultaneously on a cluster or host.



Licensing

- For Prisma Cloud customers, CDR requires 1 credit per agent, and is also included in the Cloud Security Advanced license plan.
- For Cortex customers, CDR (PAN-XDR-ADV-EP-CLOUD) is licensed per cloud host.


Resource Requirements

Currently, the XDR agent requires 2.3 GHz dual-core CPU, 4GB of RAM ([link](#) to technical specifications). XDR Agent for Cloud requires more resources due to its enhanced security and visibility capabilities

Deployment Guide


Pairing Prisma Cloud and XDR (or XSIAM) tenants

1. Request the SA or CS team to enable the tenant to pair.
2. On the Prisma Cloud Console(not for self-hosted)
 - a. Select Runtime Security
 - b. Go to Manage > System
 - c. Find “Pair Cortex XDR Tenant” and copy the Access Key




The screenshot shows the Prisma Cloud console interface. On the left is a dark sidebar with a navigation menu. The main area is titled 'Manage / System' and has tabs for 'General', 'Intelligence', 'WildFire', 'Scan', 'Forensics', 'Proxy', 'Custom feeds', 'Utilities', and 'Backup & restore'. The 'General' tab is active, showing 'Configure general settings for your environment'. Under 'Log scrubbing', there is a toggle for 'Automatically scrub secrets from runtime events' which is turned 'On'. Below this is a 'Pair Cortex XDR Tenant' section with instructions to pair the Prisma Cloud tenant with a Cortex XDR tenant. It shows an 'Access key' field with a generated key and a 'Pairing status' of 'Not paired'. At the bottom, there is a 'Custom log scrubber' section with tabs for 'Runtime' and 'WAAS', and a filter bar for 'Filter by keywords and attributes' showing '0 total entries'.

3. On the XSIAM or XDR Console
 - a. Go to Settings > Configurations > General > Server Settings
 - b. Find “Prisma Cloud Compute Tenant Paring” and paste the access key from the Prisma Cloud Console.



- c. Click Pair - if there is no region issue, it will make a pairing.



The paired Prisma Cloud tenant information is shown in the XSIAM Console.

4. Review the Prisma Cloud Console to validate the pairing status.

Pair Cortex XDR Tenant

Pair your Prisma Cloud tenant with Cortex XDR tenant to deploy Cloud Security Agents. Copy the Access Token below and insert it in your preferred Cortex XDR tenant, under **Settings > Configurations > Server Settings > Prisma Cloud tenant pairing**

 The generated access key will be valid for the pairing process for one hour until Jan 17, 2024, 3:51:35 PM


Pairing status

 Paired to <https://pcsxiam.xdr.us.paloaltonetworks.com>


[Unpair tenant](#)

Deploy XDR Agent

1. Create a Linux Host or Kubernetes endpoint agent with a higher than 8.3.0. It is recommended to use the latest version. ([Cortex XDR agent installation](#)) Refer to the XDR Agent [End-of-Life](#) support page.
 - a. Linux Host Installation:
 - i. From the XDR/XSIAM console, go to Endpoint > Agent Installations.
 - ii. Create an agent for Linux by selecting Platform “Linux” and selecting version 8.3.0 or later.
 - iii. Select Package Type “Standalone Installer.”




- iv. Once it's created, right-click, select the appropriate architecture “x86_64 installer” or “aarch64 installer” and download the correct installer package to install on the target host. (PC defender should be removed if it's deployed on the host)



CREATION TIME	STATUS	NAME	DESCRIPTION	TYPE	PLATFORM
Sep 22nd 2024 06:38:41	Completed	rs-xdr-cloud-host		Standalone installer	Linux
Sep 20th 2024 23:45:45	Completed	HS-latest		Kubernetes installer	Linux
Sep 18th 2024 03:54:06	Completed	rsingh-eks-csa		Kubernetes installer	Linux
Sep 17th 2024 08:03:36	Completed	ee-k8s-csa		Kubernetes installer	Linux
Sep 17th 2024 05:33:14	Completed	srnimmagaddaxdrtest		Kubernetes installer	Linux
Sep 17th 2024 01:56:36	Completed	pcs-demo-swat-k8s		Kubernetes installer	Linux
Sep 17th 2024 00:38:02	Completed	last_code_with_upgrad		Kubernetes installer	Linux
Sep 16th 2024 11:15:07	Completed	ee-win11	csa	Standalone installer	Windows

b. Kubernetes Installation:

- i. From the XDR/XSIAM console, go to Endpoint > Agent Installations. Create an agent for Linux by selecting Platform “Linux” and selecting version 8.3.0 or later.
- ii. Select Package Type “Kubernetes Installer”. Create with the recommended latest agent version and select the check box “Run on all nodes”.



Create Agent Installation

CREATE NEW INSTALLER

* Name

Description (Optional)

Package Type: **Kubernetes Installer**

Version: **Always deploy the latest agent version**

* Agent Daemonset Namespace: **cortex-xdr**

Node Selector: **key1:value1,key2:value2**

Proxy: **10.10.10.1:8080,example.domain.com:8080**

Cluster Name

☐ Run on master node


☒ Run on all nodes

Deployment Platform: **Standard**

Cancel Create

- iii. Once it's created, right-click and select “Download YAML installer,” then install it on the target cluster (PC Defender should be removed if it's deployed on the cluster).


- c. Once it's created, right-click and select “Download YAML installer,” then install it on the target cluster (PC defender should be removed if it's deployed on the cluster)




CREATION TIME	STATUS	NAME	DESCRIPTION	TYPE	PLATFORM
May 2nd 2024 02:21:04	Completed	TEST_MIGRATION		Standalone installer	Windows
May 2nd 2024 01:03:49	Completed	fy-mac-agent		Standalone installer	macOS
May 3rd 2024 08:53:10	Completed	HS-8-3		Kubernetes installer	Linux
Mar 7th 2024 12:37:26	Completed	derjohnson-windows		Standalone installer	Windows
Apr 21st 2024 17:59:46	Completed	CSA-Linux		Standalone installer	Linux
Mar 31st 2024 16:55:21	Completed	jjeanclaude-gke-eval		Kubernetes installer	Linux
Apr 18th 2024 10:34:47	Completed	teste		Kubernetes installer	Linux
Apr 21st 2024 18:16:30	Completed	rsam-amp-linux-csa-host		Standalone installer	Linux
Apr 15th 2024 11:22:39	Completed	rsingh-csa-pcs-lab-tenant		Kubernetes installer	Linux
Apr 11th 2024 14:06:27	Completed	xdr-test-k8s	owner: Jabar	Kubernetes installer	Linux
Mar 31st 2024 08:53:55	Completed	jjeanclaude-cx-eval	container test	Standalone installer	Linux

Enable XDR Pro and Vulnerability Scanning

1. Create a Policy with enabled “XDR Pro Endpoint” and “Active Vulnerability Analysis”
 - a. From the XDR/XSIAM console, go to Endpoint > Policy Management > Profiles and, create a Linux profile with “Agent Settings,” and click “Next”.



- b. Enable XDR Pro Endpoints for CDR.



Cortex XSIAM
BY PALO ALTO NETWORKS

Profiles / Agent Settings - Linux (HS-test-CSA)

General Information

Disk Space

Alerts Data

XDR Pro Endpoints

XDR Cloud

Response Actions

Revert Endpoint Isolation

Content Configuration

Agent Upgrade

Download Source

Agent Proxy Settings

Agent Operation Mode

Active Vulnerability Analysis

IT Metrics

XDR Pro Endpoints

Enable the Cortex XDR Pro agents capabilities including enhanced data collection, advanced responses, and attached add-ons.

XDR Pro Endpoints Capabilities
Enabled ☐ Use Default (Disabled)

Monitor and Collect Enhanced Endpoint Data
Select ☒ Use Default (Enabled)

Note: Before enabling enhanced endpoint data collection make sure your Strata Logging Service storage capacity and quota allocation can support it.
Please refer to the [Strata Logging Service quota configuration guidelines](#) and [storage calculator](#) for more details.
* This feature is supported from agent version 6.1 and above.

Enable Host Insights capabilities
Select ☒ Use Default (Enabled)

* This feature is supported from agent version 7.1 and above.

Endpoint Information Collection
Select ☒ Use Default (Enabled)


Enable Compliance Collection
Select ☒ Use Default (Disabled)

Additional information

Created By
H.S. Song

Creation time
Jan 17th 2024
14:09:23

- c. Enable Active Vulnerability Analysis (AVA) for vulnerability and compliance scanning with Prisma Cloud policy, then create a profile.



Cortex XSIAM
BY PALO ALTO NETWORKS

Profiles / Agent Settings - Linux (HS-test-CSA)

General Information

Disk Space

Alerts Data

XDR Pro Endpoints

XDR Cloud

Response Actions

Revert Endpoint Isolation

Content Configuration

Agent Upgrade

Download Source

Agent Proxy Settings

Agent Operation Mode

Active Vulnerability Analysis

IT Metrics

Active Vulnerability Analysis

Configure advanced vulnerability and compliance scanning on cloud native environments

Advanced Vulnerability Scanning
Enabled ☐ Use Default (Disabled)

Periodic Scan
Custom ☐ Use Default (24 Hours)

☐ Run weekly at: Sun 00 : 00

☐ Run monthly at: First Sun 00 : 00

☒ Run Every: 6 hours

IT Metrics

Configure IT metrics and measurements collection

Collect IT Data
Select ☒ Use Default (Disabled)

Additional information

Created By
H.S. Song

Creation time
Jan 17th 2024
14:09:23

Modified By
H.S. Song

Modification time
Sep 16th 2024
09:05:51

Profile usage
1

2. From the XDR/XSIAM console, go to Endpoint > Policy Management > Policy Rules
 - a. Create a new policy by clicking Create new policy.
 - b. Select the Agent settings with the profile that created the above step.

Create New Policy

1 **General**
2 *Target*
3 *Summary*

* Platform
Linux

Exploit
Default

Malware
Default

Restrictions
Default


Agent Settings
HS-test-CSA

Exceptions
Default (No Excepti...

- c. Select the endpoint that created step 1.
 - d. Save the policy by clicking "Done."

3. Review the status of connected agents.

Go to Runtime Security > Manage > Defenders > Cloud Security Agents



Manage / Defenders

Defenders: Deployed Defenders: Auto-defend **Cloud Security Agents** Settings

Cloud Security Agents

Your complete inventory of deployed Cloud Security Agents. To edit, delete or take further actions on these agents, please log in to the paired Cortex XDR tenant.

1 (Status: Connected X) Filter by attributes x 9 total entries (filtered)

ID	Creation time	First seen	Status	Name	Platform	Agent
e5de8b5f8bd4838bc9e4c54c4cc5ec	Dec 20, 2023, 8:08:48 AM	Jan 19, 2023, 7:13:12 AM	Connected	ocpcluster-f7jps-worker-b-cxpq7	AGENT_OS_LINUX	8.2.1.
6d223f3567d9425f809e31f9491b953	Dec 20, 2023, 8:00:42 AM	Jan 19, 2023, 7:13:12 AM	Connected	ocpcluster-f7jps-worker-a-bwlf	AGENT_OS_LINUX	8.2.1.
e130ac470f084b508dc612a37ade0dd	Dec 20, 2023, 8:08:01 AM	Jan 19, 2023, 7:13:13 AM	Connected	ocpcluster-f7jps-master-1	AGENT_OS_LINUX	8.2.1.
ade8b176152c44a6911729946e947f9	Dec 20, 2023, 8:01:00 AM	Jan 19, 2023, 7:13:13 AM	Connected	ocpcluster-f7jps-master-0	AGENT_OS_LINUX	8.2.1.
d679b0c40c0a47479d18c08b27dc28be	Dec 20, 2023, 8:18:56 AM	Jan 19, 2023, 7:13:16 AM	Connected	ocpcluster-f7jps-worker-c-h768	AGENT_OS_LINUX	8.2.1.
173516ad97594c2aa7534bd117e4ffc2	Feb 1, 2024, 11:28:58 AM	Feb 1, 2024, 11:28:58 AM	Connected	gle-hs-cluster-1-larger-pool-92f53024-lb02	AGENT_OS_LINUX	8.2.0.
29fb679169754bdfaa85641c667d0772	Feb 1, 2024, 11:28:59 AM	Feb 1, 2024, 11:28:59 AM	Connected	gle-hs-cluster-1-larger-pool-3e41cb58-834e	AGENT_OS_LINUX	8.2.0.
ffc02c0eba4341fb31a6b1f60b9e768	Feb 1, 2024, 11:28:59 AM	Feb 1, 2024, 11:28:59 AM	Connected	gle-hs-cluster-1-larger-pool-92f53024-qank	AGENT_OS_LINUX	8.2.0.
632e1a193f9946b192616040c5939d6	Feb 1, 2024, 11:28:59 AM	Feb 1, 2024, 11:28:59 AM	Connected	gle-hs-cluster-1-larger-pool-3e41cb58-yjdl	AGENT_OS_LINUX	8.2.0.


Displaying 1 - 9 of 9

Rows 20 Page 1 of 1

Enable Cortex XDR/XSIAM [Analytics Engine](#)

The Cortex XDR analytics engine is a security service that utilizes network data to detect and report post-intrusion threats automatically. The analytics engine identifies good (normal) behavior on your network to notice bad (anomalous) behavior.


1. Go to Settings > Configuration > Cortex XDR or XSIAM - Analytics and click Enable..
2. Active Identity Analytics by turning on the toggle.



Review Vulnerability and Compliance Scan result on Prisma Cloud

1. Review the Runtime Security > Monitor > Vulnerabilities > Host

Look for the “Scanned By” = “Cloud security agent” to validate the scanning and reports.



Provider	Region	Resource name	Hostname	Cluster	Scanned by	Vulnerabilities	Risk factors	Collections	Actions
AWS	us-west-1	i-0b0523cd5748600...	ip-192-168-2-249	hs-eks-cl-west1	Cloud security agent	3 14 61 1	12	+45	
AWS	us-west-1	i-0x418c9d9817a3...	ip-192-168-32-182	hs-eks-cl-west1	Cloud security agent	3 14 61 1	12	+45	

Host details

Provider	aws	OS release	AL2
Account ID	146306131214	Kubernetes version	22.5.0
Region	us-west-1	Cluster	hs-eks-cl-west1
Resource name	i-0b0523cd574860016	VM image	ami-06d7578325b5eb8cb
Hostname	ip-192-168-2-249.us-west-1.compute.internal-2f1915a643544186...	Scan time	Sep 22, 2024 6:20:59 AM
OS distribution	Amazon Linux 2		
Scanned by	Cloud security agent		

Vulnerabilities
Compliance
Runtime
Package info
Environment

Filter vulnerabilities by keywords and attributes

Type	Highest severity	Description
go	critical	net/netip version 1.20.12 has 1 vulnerability.
go	high	go.opentelemetry.io/contrib/instrumentation/google.golang.org/grpc/otelgrpc version v0.45.0 has 1 vulnerability.
go	high	go.opentelemetry.io/contrib/instrumentation/google.golang.org/grpc/otelgrpc version v0.42.0 has 1 vulnerability.
OS	high	python2-setuptools (used in python2-setuptools) version 41.2.0-4.amzn2.0.3 has 1 vulnerability.
OS	high	mariadb (used in mariadb-libs) version 5.5.68-1.amzn2.0.1 has 62 vulnerabilities. Affected service: postfix.
go	medium	net/http version 1.20.12 has 1 vulnerability.
go	medium	gopkg.in/square/go-jose.v2 version v2.5.1 has 1 vulnerability.
go	medium	google.golang.org/protobuf/internal/encoding/len... version v1.31.0 has 1 vulnerability.

Close

2. Review the Runtime Security > Monitor > Vulnerabilities > Images > Deployed.

Look for the “Scanned By” = “Cloud security agent” to validate the scanning and reports.

Runtime Security
Home
Dashboards
Reports
Inventory
Compliance
Alerts
Investigate
Governance
Settings

Monitor / Vulnerabilities
Vulnerability Explorer
Images
Hosts
Functions
CVE viewer
VMware Tanzu blobstore

Deployed
Registries
CI

Deployed images

Vulnerability scan reports for deployed images

NEW
Information on Container Images and Hosts is also available on Inventory > Compute Workloads

Filter images by keywords and attributes
38 total entries (filtered)
CSV
Refresh
Columns
Scan

Registry	Repository	Tag	Hosts	Clusters	Scanned by	Vulnerabilities	Risk factors	Collections
docker.io	weaveworksdemos/que...	0.3.1	10 hosts	hs-cluster-1, pcs-smart...	Agentless, Cloud s...	4 70 88 45	12	+22
docker.io	weaveworksdemos/whi...	0.4.8	9 hosts	hs-cluster-1, pcs-smart...	Agentless, Cloud s...	3 57 80 41	12	+22
docker.io	weaveworksdemos/ord...	0.4.7	10 hosts	hs-cluster-1, sock-shop...	Agentless, Cloud s...	2 54 71 40	12	+22
docker.io	weaveworksdemos/carts	0.4.8	10 hosts	hs-cluster-1, sock-shop...	Agentless, Cloud s...	2 54 71 40	12	+22
docker.io	weaveworksdemos/ftro...	0.3.12	10 hosts	hs-cluster-1, luis-cluste...	Agentless, Cloud s...	3 31 63 19	10	+23
gke.gcr.io	metrics-server	v0.5.2-gke.3	33 hosts	hs-cluster-1, dm-gke-te...	Agentless, Cloud s...	17 33 6	12	+23
gke.gcr.io	calico/node		12 hosts	hs-cluster-1, cm-micros...	Agentless, Cloud s...	2 15 17 6	12	+21
gke.gcr.io	calico/typha		7 hosts	hs-cluster-1, cm-micros...	Agentless, Cloud s...	12 15 4	12	+21
gke.gcr.io	prometheus-to-sd		59 hosts	hs-cluster-1, jeanclaud...	Agentless, Cloud s...	1 21 44 5	12	+23

Image details

Image

docker.io/weaveworksdemos/queue-master:0.3.1

sha256:76f0de7a12acaa2bdbda36d1b7ddd0ec42204e9c0f4e8c59c329f07336cfa2dc

Alpine Linux v3.4

3.4.6

sha256:6292d3095f4c7aeed8d863527f8ef6d7a75d3128f20fc61e57f398c100142712

Sep 3, 2024 9:21:24 PM (18 days ago)

Agentless, Cloud security agent (Go to [Environment tab](#) for details)

No issues found - Last analyzed Sep 21, 2024 6:09:34 AM

Vulnerabilities

Compliance

Runtime

Layers

Process info

Package info

Environment

Trust groups

Labels

Malware

Secrets

Filter vulnerabilities by keywords and attributes

34 total entries

Type	Highest severity	Description
Custom	critical	busybox version 1.24.2-r12 has 1 vulnerability
Application	critical	busybox version 1.24.2 has 4 vulnerabilities
jar	critical	tomcat-embed-core version 8.5.4 has 54 vulnerabilities
jar	critical	spring-webmvc version 4.3.2.RELEASE has 3 vulnerabilities
jar	critical	spring-web version 4.3.2.RELEASE has 20 vulnerabilities
jar	critical	spring-core version 4.3.2.RELEASE has 16 vulnerabilities
jar	critical	spring-beans version 4.3.2.RELEASE has 2 vulnerabilities
jar	critical	org.yaml_snakeyaml version 1.17 has 8 vulnerabilities


Close

Monitor CDR incident on Prisma Cloud

Monitor the Security Event for the XDR's report for runtime protection for Hosts and Containers.

Note: The “Security events” are only available for the Prisma Cloud tenant paired with the XDR or XSIAM tenant.

1. Review the Runtime Security > Monitor > Events> Security event.



Security events

Prisma Cloud records every host security event detected by the Cloud Security Agent.

Filter by attributes 124 total entries [Refresh](#) [Columns](#)

Creation Time	Alert name	Action	Host	Host OS	Host IP	Severity
Sep 19, 2024 6:36:40 AM	WildFire Malware	POST_DETECTED	gke-pcs-demo-kub...	Linux	10.200.2.1, 10.4.0...	High
Sep 18, 2024 6:09:15 AM	Local Analysis Malware	REPORTED	ip-192-168-1-202	Linux	192.168.6.209, 19...	High
Aug 11, 2024 8:00:22 AM	WildFire Malware	SCANNED	ip-172-31-88-65.e...	Linux	172.31.88.65	Medium
Aug 1, 2024 10:10:47 PM	WildFire Malware	SCANNED	ip-172-31-88-65.e...	Linux	172.31.88.65	Medium
Jul 23, 2024 6:35:22 PM	WildFire Malware	SCANNED	heist	Linux	10.0.0.5	Medium
Jul 21, 2024 8:19:36 AM	WildFire Malware	SCANNED	ip-172-31-88-65.e...	Linux	172.31.88.65	Medium
Jul 20, 2024 7:00:08 PM	WildFire Malware	SCANNED	heist	Linux	10.0.0.5	Medium
Jul 3, 2024 8:59:28 AM	Reverse Shell - 3230579465	BLOCKED	n1	Linux	192.168.1.0, 10.88...	High
Jul 3, 2024 8:51:55 AM	Reverse Shell - 3230579465	BLOCKED	n1	Linux	192.168.1.0, 10.88...	High
Jul 3, 2024 8:44:34 AM	Reverse Shell - 3230579465	BLOCKED	n1	Linux	192.168.1.0, 10.88...	High
Jul 2, 2024 6:25:42 PM	WildFire Malware	SCANNED	heist	Linux	10.0.0.5	Medium
Jun 28, 2024 2:25:02 PM	Behavioral Threat	REPORTED	ip-10-0-3-121	Linux	10.0.3.121, 172.17...	Critical
Jun 27, 2024 2:25:01 PM	Behavioral Threat	REPORTED	ip-10-0-3-121	Linux	10.0.3.121, 172.17...	Critical
Jun 27, 2024 12:00:21 AM	WildFire Malware	SCANNED	heist	Linux	10.0.0.5	Medium
Jun 26, 2024 2:25:02 PM	Behavioral Threat	REPORTED	ip-10-0-3-121	Linux	10.0.3.121, 172.17...	Critical

WildFire Malware

Severity: High

Description: Suspicious executable detected

Time: Sep 19, 2024 6:36:40 AM

Alert ID: 8407620

Action: POST_DETECTED

Category: Malware

Rule ID:

User: NO

ATT&CK tactics:


ATT&CK techniques:

Host details

Host	gke-pcs-demo-kube-goat-default-pool-e8b144c2-216g
Host IP	10.200.2.1, 10.4.0.36
Host OS	Linux
Cluster	pcs-demo-kube-goat
Cloud account ID	pcs-swaf-resources
Cloud labels	gke-pcs-demo-kube-goat-978e9b0c-node

[Close](#)

2. Review the Runtime Security > Monitor > Events > Containers > Security events.



Security events

Prisma Cloud records every container security event detected by the Cloud Security Agent.

Filter by attributes: 47 total entries

Creation Time	Alert name	Action	Container name	Severity
Sep 20, 2024 5:27:42 AM	Cryptominers Protection - 807956875	REPORTED	xmrig	High
Sep 20, 2024 5:04:43 AM	Cryptominers Protection - 807956875	REPORTED	xmrig	High
Sep 20, 2024 4:59:01 AM	Cryptominers Protection - 17641322	REPORTED	xmrig	High
Sep 20, 2024 4:59:01 AM	Cryptominers Protection - 807956875	REPORTED	xmrig	High
Sep 20, 2024 4:59:01 AM	Cryptominers Protection - 43250279	REPORTED	xmrig	High
Sep 20, 2024 4:59:01 AM	Cryptominers Protection - 17641322	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	Cryptominers Protection - 43250279	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	WildFire Malware	REPORTED	xmrig	Medium
Sep 20, 2024 4:59:00 AM	Cryptominers Protection - 43250279	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	Cryptominers Protection - 807956875	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	WildFire Malware	REPORTED	xmrig	Medium
Sep 20, 2024 4:59:00 AM	Cryptominers Protection - 17641322	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	Cryptominers Protection - 807956875	REPORTED	xmrig	High
Sep 20, 2024 4:59:00 AM	WildFire Malware	REPORTED	xmrig	Medium
Sep 20, 2024 2:52:11 AM	WildFire Malware	REPORTED	xmrig	Medium

Cryptominers Protection - 807956875

Severity: High
 Description: Network activity similar to miners
 Time: Sep 20, 2024 5:27:42 AM
 Alert ID: 8412525
 Action: REPORTED
 Category: Malware
 Rule ID:
 User: NO
 ATT&CK tactics: TA0011 - Command and Control, TA0040 - Impact
 ATT&CK techniques: T1132 - Data Encoding, T1496 - Resource Hijacking

Container details

Container ID: 882e4fc469045481d09608e2cdf7093f82e26d9ed8c303209d...

Image: docker.io/rcmelendez/xmrig:6.16.2

Cluster: pcs-demo-kube-goat

Namespace: xmrig

Container name: xmrig

Cloud account ID: pcs-swaf-resources

Cloud labels: gke-pcs-demo-kube-goat-978e9b0c-node

Host details


Host: gke-pcs-demo-kube-goat-default-pool-e8b144c2-zc2l

Host IP: 10.200.0.1, 10.4.0.34


Host OS: Linux

Close

3. You can find more details about the security event from the XSIAM or XDR console.
 - a. Filter the incidents by Incident Sources from Prisma Cloud Compute and Severity levels.




b. Review Incident Response > Incidents to find the matching event.




Ingest Prisma Cloud Compute Alerts

Create an XSIAM [webhook](#) for Prisma Cloud Compute to send Vulnerabilities and Compliance alerts.


1. Create an XSIAM token and webhook
 - a. Select Setting “Data Sources”, click” Add Data Source.”




- b. Generate an XSIAM Token.



- c. Save the XSIAM Token and Create a Basic Authentication Credential in Prisma Cloud Compute.




- d. Copy the XSIAM API URL for Webhook endpoint configuration in Prisma Cloud Compute.




<https://api-pcsxsiam.xdr.us.paloaltonetworks.com/logs/v1/prisma>


2. Create a Prisma Cloud Compute Alert Profile.
 - a. Use the XSIAM-generated Token for Prisma Cloud Compute Webhook Basic authentication.




- b. Create an Alert Profile for Type Webhook




c. Select triggers to send to XSIAM incoming Webhook.




d. Configure the XSIAM incoming Webhook URL obtained from XSIAM.




- e. Select the XSIAM-generated Token as the credential for sending Alerts to XSIAM.




- f. Review the Summary and click “Send test alert” to verify the settings.



g. “Send test alert” successfully validated.






The screenshot shows the Prisma Cloud console interface. The left sidebar contains navigation options: Access, Custom rules, MONITOR, ATT&CK, Events, Runtime, Vulnerabilities, Compliance, WAAS, MANAGE, Cloud accounts, Logs, Defenders, Alerts (selected), Collections and Tags, Authentication, and System. The main content area is titled 'Alert providers' and includes a 'General settings' section with a dropdown for 'Set the default frequency for all alerts with the exception of the following alert triggers: vulnerability, compliance, and cloud discovery.' set to '10 Minutes'. Below this is a table of alert providers. A filter is applied: 'Name: xsiam-csa-alert'. The table shows one entry: 'xsiam-csa-alert' with provider 'Webhook', owner 'rsam@paloaltonetworks.com', modified 'Sep 22, 2024, 11:47:17 AM', and status 'Connected'. The bottom of the table indicates 'Displaying 1 - 1 of 1'.

Name	Provider	Owner	Modified	Status	Actions
xsiam-csa-alert	Webhook	rsam@paloaltonetworks.com	Sep 22, 2024, 11:47:17 AM	Connected	[Edit] [Delete]

- h. Verify from XSIAM with alerts received from Prisma Cloud Compute via the XSIAM incoming Webhook.



The screenshot shows the Prisma Cloud console interface. The left sidebar contains navigation options: CONFIGURATIONS, Data Collection (selected), Data Sources (selected), Automation & Feed Integrati..., Pathfinder Collection Center, External Alert Mapping, Data Management, Dataset Management, Parsing Rules, Data Model Rules, Compute Unit Usage, Integrations, Threat Intelligence, and External Applications. The main content area is titled 'Data Sources' and includes a 'Status: All' and 'Type: All' filter. A search bar shows 'Prisma Cloud'. The status summary shows '6 Connected', '6 Warnings', and '6 Disabled'. Below this is a section for '3RD PARTY INTEGRATIONS' with two entries: 'Custom - HTTP based Collector' (6 Warnings) and 'Prisma Cloud Compute' (1 Connected). The 'Prisma Cloud Compute' entry shows 'Receives alerts from Prisma Cloud Compute.' and a log entry: 'APP2 Logs Received: Last Hour: 83, Last Day: 83, Last Week: 83 Last Communication: Sep 22, 2024, 11:34:15 AM'.



Application Security

Supported Tools & Technologies

Prisma Cloud Application Security supports a wide range of Cloud DevSecOps and Integrated Repositories, development environments (IDEs), and CI/CD pipelines used to build and deploy code and infrastructure for your organization.

- **Integrated Development Environment (IDE)**

- [VSCode](#)
- [JetBrains](#)

- **Version Control System (VCS)**

- [Azure Repos](#)
- [BitBucket](#)
- [BitBucket Server](#)
- [GitHub](#)
- [GitHub Server](#)
- [GitLab](#)
- [GitLab Self-Managed](#)

- **CI/CD runs**

- [AWS Code Build](#)
- [Checkov](#)
- [CircleCI](#)
- [GitHub Actions](#)
- [Jenkins](#)
- [Terraform Cloud \(Sentinel\)](#)
- [Terraform Cloud \(Run Tasks\)](#)
- [Terraform Enterprise \(Sentinel\)](#)
- [Terraform Enterprise \(Run Tasks\)](#)

- **CI/CD systems**

- [CircleCI](#)
- [Jenkins Serve](#)


Details on the package managers and IaC frameworks supported by Prisma Cloud Application Security can be found in the [Prisma Cloud TechDocs](#).

Deployment Process

The diagram below shows a typical Cloud Application Security deployment process. This consists of four phases:

1. Access (Setup)
2. Feedback
3. Expanded Feedback
4. Guardrails

Rapid time to value




The overall process to deploy Cloud Application Security within the client environment will typically involve the following:

1. Review and understand Prisma Cloud Application Security capabilities
2. Setup RBAC access for Developers, Repositories editors, and Prisma Cloud Code Security configurations (including API Access keys management)
3. Onboard VCS repositories for scanning
4. Embed code scanning in CI/CD pipelines with the specific client requirements
5. Configure default and per-repository (if required) scan enforcement settings (Hard/Soft Fail, Bot comments) and repositories/paths exclusions
6. Establish a process for the VCS Scan results review, PRs and Suppress process, Software Supply Chain review, Exposed Secrets resolution, Drift detection for each repository
 - a. After initial scan
 - b. New scan results
7. Create Cloud Application Security deployment goals for Developers and SecOps and scan resolutions KPIs
8. Setup Prisma Cloud Alert notifications integration and configure notifications for each code repository/category scan results
9. Setup a process to manage out-of-the-box tags and custom tags and tag rules for all resources with the assigned repositories integrated on Prisma Cloud
10. Review out-of-the-box policies and create custom build policies (if required) to match custom security guardrails and rules/standards used by the client.
11. Review Cloud Application Security automation options and create an automated



deployment model for onboarding into Cloud Application Security, custom policies management, alerting and scan results review.

The diagrams below represent the target timelines for a typical deployment of Cloud Application Security.



The diagram below shows a typical approach to Cloud Application Security deployment goals:



Enabling Cloud Application Security

The [activation process](#) is very simple and once activated, the CAS module is free during the initial 30-day trial.

Navigate to **Profile > Subscription > Application Security** to **Subscribe**.

Before you begin adding your development environments and pipelines for scanning, you must first generate access keys to allow permissions for specific users.

If you control outbound Internet connectivity from your cloud workloads and IDE users, make sure to add the Prisma Cloud IP addresses and hostnames for your Prisma Cloud SaaS instance to your Cloud or on-prem FWs allow lists. For more information, see [allow access to the Prisma Cloud Console](#).

If you are using [Prisma Cloud Trusted IP Login IP Addresses Lists](#), make sure the IPs that need access to the portal and APIs are also included there.

License information

Prisma Cloud Application Security is currently available as a part of the [Prisma Cloud Enterprise License](#). The application security license is based on the developer metering method:

A Developer seat refers to an active Git committer, identified as such through their unique Git author email address. Anyone who made a contribution to a code repo protected by Prisma Cloud within the last 90 days can also be considered a Developer.

There are four modules that you can choose to enable for application security.


- Infrastructure as Code (IaC) Security
- Software Composition Analysis (SCA)

- Secrets Security
- CI/CD Security

For more license details, please review [this application security license guide](#).

Onboarding Code and Build Providers

Connecting code and build providers to Prisma Cloud can be done by navigating to the Settings menu in the Prisma Cloud console, selecting Providers, expanding the Connect Provider dropdown menu in the top right corner and selecting Code & Build Providers.



From here, you may select your desired VCS solution and follow the steps to integrate it into Prisma Cloud.

For up-to-date list of supported providers as well as additional information and configuration instructions, refer to the [Connect Code and Build Providers](#) documentation page.

Scanning

By integrating your VCS repos you will put an automatic security guardrail in place, enabling selected repositories to be scanned simply by adding the Prisma Cloud integration.

This provides a single interface to administer repository scans and review policy violations in the Cloud Application Security section of the Prisma Cloud console. For more details, see [Visualizing Checkov Output in Prisma Cloud](#).

However, there are some limitations to this approach. As an alternative, integrating CI/CD pipelines with Code Security scans allows you to customize your code security scanning process during the build time.

Please Note: **New repositories created in your VCS are not automatically added** (except Bitbucket Server) to Code Security configuration after the VCS integration is configured.



It is your responsibility to create a process for onboarding new code repositories after integrating your VCS with Cloud Application Security. This might involve updating all aspects of the existing Cloud Application Security configurations for your new repository, such as creating/updating new/existing RBAC Roles, setting up correct notification channels, and creating new automation workflows.

For testing purposes, you can onboard some sample Git repositories:

- [TerraGoat - Vulnerable by design Terraform Infrastructure](#)
- [Cfngoat - Vulnerable by design Cloudformation Template](#)
- [CdkGoat - Vulnerable by design AWS CDK Infrastructure](#)
- [BicepGoat - Vulnerable by design Bicep and ARM Infrastructure](#)
- [KubernetesGoat - Vulnerable by design Kubernetes Cluster](#)
- [KustomizeGoat - Vulnerable by design Kustomize deployment](#)
- [SupplyGoat - Vulnerable by design SCA](#)

Master/Main and Branch Scanning

When a code repository is onboarded into Prisma Cloud, it will be configured by default to scan the main/master branch. You may select a different branch to scan by navigating to the Settings page and selecting Providers, finding the repository in the table, and selecting the “Set scanned branch” option under the actions menu for that repository (the three dots under the Actions column).

Tagging rules are currently only applicable to the main branch.

Additional info can be found in the [Non-Default Branch Scans TechDocs page](#).

Note that if you are doing a fix on an existing PR scan, it will not open a new PR, but rather commit back to that same branch.

CDK Scanning (CI/CD)

Support for CDKTF for Terraform is provided through Checkov integration.

The main problem with using CDKs is that there is no traceability back to the original CDK code, so PRs or anything bot related are not usable with CDKs.

More info here:

[AWS CDK configuration scanning](#)

[CdkGoat - Vulnerable AWS CDK Infrastructure](#)

Terraform Plan Scanning

Checkov can be used to scan Terraform plans stored in JSON files. Plan scanning provides Checkov with additional dependencies and context that can result in a more complete scan result. For example, it may be necessary to implement plan scanning if the codebase makes heavy use of the [built-in Terraform functions](#).

The code below can be used in a CI/CD pipeline to export the Terraform plan to a JSON file



and subsequently scan it with Checkov. Note that this requires the jq utility to format the output file.

Unset

```
terraform show -json tfplan.binary | jq '.' > tfplan.json  
checkov -f tfplan.json
```

To further enrich the Terraform plan and ensure greater coverage, the deep analysis feature can be used to instruct Checkov to reference the HCL (.tf files) used to generate the plan. This allows Checkov to make graph connections where there is incomplete information in the plan file. For example, local values do not have connections defined in the plan file, but this can be accounted for with deep analysis. Deep analysis can be enabled by specifying the `--deep-analysis` flag in combination with the `--repo-root-for-plan-enrichment` flag, as shown in the example below.

Unset

```
checkov -f plan.json --repo-root-for-plan-enrichment /path/to/tf/files --deep-analysis
```

Note that there are cases where certain values may only be known after the Terraform apply operation has been executed, and as such these values will not be available for Checkov scans. To see if any such values exist in your Terraform module, generate a plan and review the results to see if any values are listed as “(known after apply)”. It is recommended that CWP and/or CSPM policies be implemented to account for these gaps for a defense-in-depth approach.

Third-Party and Private Terraform Modules (CI/CD)

Terraform modules abstract the Terraform configuration away from a regular Checkov scan on the current directory.

To ensure coverage of objects within these modules, you can instruct Checkov to scan the `.terraform` directory (after a `terraform init`) which will contain the third-party modules and any associated `.tf` files.

Unset

```
terraform init  
checkov -d . # Your TF files.
```

```
checkov -d .terraform # Module TF files.
```

Note that when scanning the .terraform directory, Checkov cannot differentiate between third-party and internally written modules.

In the case that third-party modules are stored in a private repository or a private Terraform Cloud registry, you can provide access tokens as environment variables for Checkov to attempt to clone and scan those modules. The relevant environment variables are listed in the table below.

Variable Name	Description
GITHUB_PAT	GitHub personal access token with read access to the private repository.
BITBUCKET_TOKEN	Bitbucket personal access token with read access to the private repository.
TF_HOST_NAME	Terraform registry hostname. Example: gitlab.com
TF_REGISTRY_TOKEN	Private registry access token. Supports Terraform Cloud/Enterprise and third-party registries.
BITBUCKET_USERNAME	Bitbucket username (can only be used in conjunction with BITBUCKET_APP_PASSWORD)
BITBUCKET_APP_PASSWORD	Bitbucket app password (can only be used in conjunction with BITBUCKET_USERNAME)
TFC_TOKEN	Deprecated value, use TF_REGISTRY_TOKEN.

For self-hosted VCS repositories, use the following environment variables:

Variable Name	Description
VCS_BASE_URL	Base URL of the self-hosted VCS. Example: https://example.com
VCS_USERNAME	Username for basic authentication

VCS_TOKEN

Password for basic authentication

For more info, refer to the [Checkov documentation site](#) and [Terraform Plan and External Terraform Module Scanning](#) documentation in the official GitHub repository.

Scanning with Checkov CLI

Integrating Prisma Cloud with the Checkov CLI utility makes it possible for Prisma Cloud Application Security to automatically scan your Infrastructure as Code (IaC) files against policies stored in the Prisma Cloud console.

By specifying the URL and API key for your Prisma Cloud deployment using the `--prisma-api-url` and `--bc-api-key` arguments respectively, Checkov will retrieve the policies from your Prisma Cloud console and evaluate the targeted IaC files against them.

Note that the Checkov CLI utility comes with out-of-the-box policies defined and managed by Bridgecrew. By running the utility without pointing it to your Prisma Cloud deployment, you will only be evaluating your code against these policies (as well as any external policies stored in GitHub or local files specified via additional command line arguments). Conversely, running the utility with your Prisma Cloud deployment specified will use only the policies stored in Prisma Cloud unless specified with additional command line arguments.

The table below shows which policies will be applied to a scan according to the supplied arguments.

Prisma Cloud integrated? *	Include all Checkov policies flag provided? **	External checks provided? ***	Applied policies
Yes	No	No	Prisma Cloud policies
Yes	Yes	No	Prisma Cloud policies Bridgecrew policies
Yes	No	Yes	Prisma Cloud policies Externally defined policies
No	N/A	No	Bridgecrew policies
No	N/A	Yes	Bridgecrew policies Externally defined policies

* `--bc-api-key` and `--prisma-api-url` arguments

** `--include-all-checkov-policies` argument

*** `--external-checks-dir` OR `--external-checks-git` provided

Additional arguments such as the `--check`, `--skip-check` and `--run-all-external-checks` may be used to further scope your scan to the specific subset of policies you would like to



execute in a given environment.

For additional information, refer to the [CLI Command Reference page](#).

Use Cases

The following use cases would necessitate the use of the Checkov CLI utility:

- Review the policies that will be applied to your IaC using the `--list` argument
 - `--output-bc-ids` can also be supplied to show which policies exist in the platform and which are managed by Bridgecrew
- Check a given IaC file or module against a specific subset of policies
- Implement code scanning as part of a CI/CD pipeline
- Scan private Terraform modules (Prisma Cloud support is coming)
- Scan Terraform plans
- Download policies
 - The following one-liner can be used to download Bridgecrew policies to a CSV:

Unset

```
echo 'Provider,Benchmarks,Policy ID,Title,Severity,Category' >
policies.csv; curl
'https://www.bridgecrew.cloud/api/v1/policies/table' -H
'Authorization: $BC_API_KEY'+ (.value|join(", ")) | join("; "),
.id, .title, .severity, .category] | @csv' >> policies.csv
```

IDE Integration

Navigate to **Settings > Repositories > Add Repository** and select your IDE.

The IDE integration supports Microsoft Visual Studio Code and JetBrains IntelliJ:

- [Checkov Extension for Visual Studio Code](#)
- [Checkov Plugin for JetBrains IDE](#)

Drift Detection

Prisma Cloud Application Security supports Drift Detection for your VCS. Drifts are inconsistencies in the configuration that occur when resources are modified directly or manually using the CLI or console, and these modifications from the code are not recorded or tracked.

The inconsistencies in code can either be an addition or deletion of values from the template in the source code.

Cloud Application Security periodically scans your repositories to identify drifts that may occur between the build and deploy phase and enables you with corrective resolutions to



handle traceable configuration changes.

Drift detection is currently available only for resources that are deployed using Terraform and CloudFormation on AWS, GCP and Azure.

Support for Terraform state, Kubernetes, and Unmanaged resource detection is being released after this.

Please check the Release Notes for up-to-date information.

- [Prisma Cloud Release Information](#)

For further details, please review the [setup drift detection](#).

Enforcement (Scan configurations)

Once your code repositories are integrated, you can modify your configuration to specify how Prisma Cloud scans your code.

Periodic scans are performed every 12 hours at 8 a.m./8 p.m. EST and can last up to 60 minutes. You can also perform an on-demand scan by pressing the “Scan Now” button in the project repository view.

On the Prisma Cloud console, there are default parameters, based on best practices, for each code category scanned in your repositories.

Using enforcement, you can configure these default parameters and receive violation notifications only for critical issues, helping you reduce unnecessary noise and optimizing secure productivity.

Here is the detailed instruction to [adjust the enforcement settings](#).

Enforcement Categories

Vulnerabilities (SCA)

Vulnerabilities found in open-source packages and container images.

This will automatically scan repositories for container vulnerabilities, leveraging Prisma Cloud’s twistcli, the CLI tool acquired from Twistlock, helping you identify and remediate vulnerabilities in container images with high accuracy and a low false-positive rate.

For more detailed information, refer to [the vulnerability with SCA](#).

Here is [the latest list of supported package managers](#).

License (SCA)



License compliance issues found in open-source packages and container images. Prisma Cloud scans licenses in parallel with the vulnerability scan for open-source packages.

For more detailed information, review [the license compliance with SCA](#).

Infrastructure as Code (IaC)

Prisma Cloud scans Infrastructure as code files for misconfiguration issues.

Secrets

You can use Cloud Application Security to detect and block secrets in IaC files stored in your IDEs, Git-based VCS, and CI/CD pipelines.

A secret is a programmatic access key that provides systems access to information, services, or assets. Developers use secrets such as API keys, encryption keys, OAuth tokens, certificates, PEM files, passwords, and passphrases to enable their applications to communicate with other cloud services securely.


There is nothing that can be configured for secret scans once enabled. It will scan any non-compressed or non-compiled files.

For more detailed information, see this [secret scanning guide](#).

Static Application Security Testing (SAST)

The definition of SAST is a security testing technique that analyzes the source code or compiled version of an application to identify potential vulnerabilities and security weaknesses. It examines the application's code structure, logic, and dependencies, without actually executing the application.

SAST helps to identify security flaws early in the software development lifecycle (SDLC), enabling developers to address them before the application is deployed.





Prisma Cloud SAST feature is currently in beta and enables developers to check for the common OWASP TOP 10 most critical security risks to web applications as identified by the Open Web Application Security Project.



A01:2021 - Broken Access Control
A02:2021 - Cryptographic Failures
A03:2021 - Injection
A04:2021 - Insecure Design
A05:2021 - Security Misconfiguration
A06:2021 - Vulnerable and Outdated Components
A07:2021 - Identification and Authentication Failures
A08:2021 - Software and Data Integrity Failures
A09:2021 - Security Logging and Monitoring Failures
A10:2021 - Server-Side Request Forgery

Common Weakness Enumeration (CWE) Top 25 is the common root cause for OWASP TOP 10. The CWE Top 25 list demonstrates the currently most common and impactful software weaknesses.

Prisma Cloud SAST Rule Types

- Search (Rule-Based Analysis)
 - Pattern Matching
 - Control Flow Analysis

Multiple rules can be created to identify particular programming patterns linked to malicious actions or instances of systematic misuse that may pose a security threat.

- Data Flow Analysis

Rules focus on inbound data from external sources that traverse the application code and interact with calls accessing outbound requests, logs or data.

The SAST engine tracks the data flow within the application, aiding in the identification of potential security risks associated with the handling of sensitive data.

SAST Data Flow Terminologies

- SOURCE

This is the initial location within an application where sensitive or user-controlled



data is inputted. Examples include form fields, query parameters, and API endpoints.

- SINK

A sink is a critical section of code where data, particularly data originating from a source, is consumed in operations. Such operations include, but are not limited to, database queries, external HTTP requests, and system command execution such as 'exec' command.

- SANITIZER

A sanitizer serves as a defensive coding practice designed to "clean" or validate input data against known exploitation patterns before the data reaches a sink. Through processes such as encoding, escaping, and input validation, sanitizers help to ensure that incoming data does not contain malicious content that could exploit vulnerabilities in the application.

Source and Sink are used in a data flow analysis. The source is where data comes from and the sink is where it ends. Source and Sink are frequently used for taint analysis.

Data is "tainted" if it comes from an insecure source such as a file, network or user.

SAST Rules Format

- Rules are written in YAML using text editor or Prisma Cloud Policy Editor
- The three main sections of the rule
 - metadata
 - description of the rule
 - scope
 - specify the programming language
 - definition
 - SAST logic

Prerequisites for SAST

Linux Virtual Machine

- Ubuntu 22.04 or higher
- Python 3
- Checkov
 - apt install python3-pip
 - pip install -U checkov
- BC_API_KEY [Prisma Cloud Access Key and Secret Key]
- Prisma Cloud API URL

Prisma Cloud

- Contact the customer's SA or CSM to join the SAST extended beta program



- Require to onboard repositories for SAST scanning

Per Developer credits

- No credits will be consumed during the beta evaluation
- SAST will be 4 credits with General Availability (GA)


Prisma Cloud with SAST Findings

The screenshot shows the Prisma Cloud Application Security interface. The 'Weaknesses' tab is selected in the top navigation bar. The dashboard displays a summary of 60 issues detected in the code scope. Below the summary, a table lists specific findings with details on repository, policy, CWE, line numbers, labels, git user, and first detected date.

Repository	Policy	CWEs	Line(s)	Labels	Git User	First Detected
(sohnaeo/badcode: java/main.java)	CSRF protection is di...	CWE-352: Cross-Sit...	33	CWE top 25, OWASP top 10		April 22, 2024
	Encryption keys with...	CWE-326: Inadequa...	63	OWASP top 10		April 22, 2024
	Unsafe DES algorith...	CWE-327: Use of a ...	55	OWASP top 10		April 22, 2024
	File path not validate...	CWE-22: Improper L...	147	CWE top 25, OWASP top 10		April 22, 2024
	Cookie stored for an ...	CWE-539: Use of Pe...	80-81	OWASP top 10		April 22, 2024

SAST findings with OWASP top 10 and CWE top 25 labels are included in the Application Security Weakness tab dashboard.

Prisma Cloud with Policy Editor



Prisma Cloud Policy Editor supports the following policy type using Build Policy:

- **IaC policies** can be created via the code and visual editor. Policies can be tested against onboarded repositories.
- **Secrets policies** can be created via the code and visual editor. Policies can be tested against code snippets in the visual editor.
- **SAST policies** can be created via the code editor. Policies can be tested against multiple code snippets that can be marked as violating and non-violating snippets.



Prisma Cloud Code Editor supports auto-completion fields.

Performing SAST Scanning Using the Search Method

Method 1 - Prisma Cloud Config Policy

Create New Config Policy

Add Details

Create query

Compliance Standards

Remediation

Create query

Build

SAST Policy

Test

```

1 metadata:
2   guidelines: "find-exec-code"
3   category: "sast"
4   approach: "define failing"
5   scope:
6     languages:
7       - "python"
8   definition:
9     patterns:
10      or:
11        - exec(<ANY>)
12        - exec("<ANY>")
13

```

Test 1/1

Include code snippets that should or should not violate your custom policy to validate the policy is written correctly.

☒ Violating snippet

```

1 // insert your custom code
2 print("hello sast")
3 exec("rm -rf .")
4 print("goodbye!")
5 exec(a)
6
7

```

An example above was created for Config Policy with Build Subtype using Prisma Cloud Policy Editor. Click on the Test button to show the violated command syntax highlighted in red under the code snippets section.

Application Security Home Dashboards Inventory Investigate Governance Settings 62%									
GOVERNANCE Overview Vulnerabilities Custom Menu Add Policy									
Policy Name find-exec-code Policy Severity Select Add Filter Group By: Cloud, Policy Type, Policy Subtype Search table data...									
Cloud	Policy Type	Policy Subtype	Policy Name	Status	Severity	Category	Actions		
▼ All (1)			1 Policies	1 active policies	0 1 0 0 0	1			
	▼ Config (1)		1 Policies	1 active policies	0 1 0 0 0	1			
		▼ Build (1)	1 Policies	1 active policies	0 1 0 0 0	1			
find-exec-code					High	Risk			

The custom SAST Config Policy will be used by Prisma Cloud Console to scan against the on-boarded repositories for any violation.



Method 2 - Checkov CLI Tool

```
export BC_API_KEY= $PRISMA_ACCESS_KEY::$PRISMA_SECRET_KEY

checkov -s -d appcodes/ --framework sast_python --repo-id cli/find-exec-code
--prisma-api-url https://api.prismacloud.io/ --external-checks-dir sast_policy
```

The example above shows how to execute Prisma Cloud SAST scanning using Checkov tool.

The Python application source codes (.py and .yaml) are stored under the directory appcodes. The SAST rule is stored under the directory sast_policy (yaml).

Obtain the BridgeCrew (BC) API key required by the Checkov tool by generating Access Key or Service Account Key from Prisma Cloud Console.

Assumption: Tenant on app.prismacloud.io with SAST feature enabled.

```
checkov
By Prisma Cloud | version: 3.2.106
Update available -> 3.2.109
Run pip3 install -U checkov to update

sast_python scan results:
Failed checks: 4, Skipped checks: 0
Source code files scanned: 1, Policies found: 58

Check: 1060149136915596288_SAST_1716823837358: "find-exec-code"
Severity: HIGH
FAILED for file - app.py:2
Guide: find-exec-code
2 | exec("rm -rf .")

Check: 1060149136915596288_SAST_1716823837358: "find-exec-code"
Severity: HIGH
FAILED for file - app.py:4
Guide: find-exec-code
4 | exec(a)

Check: CKV3_EXEC_RULE: "Unchecked exec command is bad"
Severity: CRITICAL
FAILED for file - app.py:2
['CWE_XXX: Untrusted Execution']
2 | exec("rm -rf .")

Check: CKV3_EXEC_RULE: "Unchecked exec command is bad"
Severity: CRITICAL
FAILED for file - app.py:4
['CWE_XXX: Untrusted Execution']
4 | exec(a)
```

The output from the Checkov scan result.



Policies Management

Out-of-the-Box Policies

Prisma Cloud Application Security default policies are being updated once a month, policies deleted, new ones created and existing policies are updated. Please check the release information for up to date information:

- [Prisma Cloud Release Information](#)

Prisma Cloud includes out-of-the-box policies that enable you to detect misconfigurations and provide automated fixes for security issues seen across your integrated code repositories and pipelines. You can review this list of Configuration policies with a filter for subtype **Build** on the Prisma Cloud administrative console **Policies**.

Please note: It is currently (August 2022) not possible to select the Policies view filter to show build only policies. FYI - all build policies are of type - Config.

To view the Custom created policies choose the policies created by your Prisma Cloud users listed in the "Last modified by" column. All Prisma Cloud default policies are modified by "Prisma Cloud System Admin". It might be easier to use the API to do this, more info on this later.

Cloud Application Security offers thousands of out-of-the-Box policies for scanning various types of resources. You can find more about supported policies here:

- [Alibaba Policies](#)
- [AWS Policies](#)
- [Azure Policies](#)
- [Build Integrity Policies](#)
- [CI/CD Risk Policies](#)
- [Docker Policies](#)
- [GCP Policies](#)
- [Kubernetes Policies](#)
- [OCI Policies](#)
- [OpenStack Policies](#)
- [Secrets Policies](#)

Bridgecrew default policies are listed on github here:

<https://github.com/bridgecrewio/checkov/tree/master/docs/5.Policy%20Index>

Checkov Policy Index:

<https://www.checkov.io/5.Policy%20Index/all.html>

Import Bridgecrew custom policies into Prisma Cloud



If you are an existing Bridgecrew user and have created custom policies then it will be possible to import Bridgecrew custom policies into the Cloud Application Security module. Please use the following script to do this:

https://github.com/PaloAltoNetworks/prisma_channel_resources/blob/main/prisma_bash_to_olbox-main/export_bridgecrew_custom_yaml_policies_and_load_into_ccs.sh

Custom Policy for Build-Time Checks

The following document provides good details on the custom policy definition - [Custom Policy definition](#)

This document gives good examples of custom policies - [Examples - YAML-Based Custom Policies](#)

For example: Can a user create a custom policy that could prevent the usage of certain resources as a cost saving feature?

```
metadata:
  name: "don't create s3 buckets"
  guidelines: "save some monies"
  category: "general"

  severity: "critical"
  scope:

  provider: "aws"

  definition:
    cond_type: "attribute"
    resource_types:
      - "aws_s3_bucket"

    attribute: "bucket" # pick any required attribute
    operator: "exists"
```

[How to add custom policies for build time.](#)

It is recommended to utilize Labels for your custom build policies - for example Code-Security-Build-AWS/GCP/Azure-AppX-01. Creating a Policy Label standard for your organization and the appropriate labels assignment to your custom policies will enable you to filter the policies list to your specific build policy set on the Prisma Cloud portal and will significantly simplify your custom policy operational management.

Using APIs to manage policies

You can use the [CSPM Policy API](#) to create and manage Cloud Application Security build policies.

The following API calls are also available to manage build policies in Prisma Cloud Application Security:

- `POST /code/api/v1/policies/definition/{queryId}`
- `POST /code/api/v1/policies`
- `GET /code/api/v1/policies/table/data`
- `POST /code/api/v1/policies/{policyId}`
- `DELETE /code/api/v1/policies/{policyId}`
- `POST /code/api/v1/policies/preview`
- `POST /code/api/v1/policies/clone/{policyId}`
- `POST /code/api/v1/remediations/buildtime`
- `GET /code/api/v1/remediations/buildtime/{fixId}`
- `GET /code/api/v1/remediations/buildtime/baseFile/{filename}`

PolicyId is auto-generated.

Here is a sample custom policy definition in JSON

```
{
  "cloudType": "aws",
  "name": "Sample API build policy",
  "policyType": "config",
  "rule": {
    "name": "Sample API build policy",
    "parameters": {
      "withIac": "true",
      "savedSearch": "false"
    },
    "type": "Config",
    "children": [
      {
        "criteria":
        "{\"category\":\"General\",\"resourceTypes\":[\"aws_s3_bucket\"],\"conditionQuery\":{\"attribute\":\"bucket\",\"operator\":\"equals\",\"value\":\"abc\",\"cond_type\":\"attribute\"}}",
        "type": "build",
        "recommendation": "Get good"
      }
    ]
  }
}
```

```
},  
  
  "severity": "low"  
}
```

It *might* be technically possible to send more complex policy definitions in the Prisma API payload, allowing you to create more complex policies than can be expressed in the visual editor. This will result in the policy being visible on the Prisma policies page, but it will not be fully viewable or editable, because the definition cannot be rendered in the visual editor. This is officially not supported,

More info on Application Security API usage here: <https://prisma.pan.dev/api/cloud/code/>

Instructions on how to setup the Postman Collections and Environments relating to Prisma Cloud (including Compute Console) API requests - <https://github.com/PaloAltoNetworks/pcs-postman>

You can use the following API call to retrieve all build policies in JSON format

```
https://{api-endpoint}/code/api/v1/policies
```

Not all policies are YAML policies. YAML policies are the only OOTB policies that have definitions that are visible and editable (via cloning) in the platform. Python policies in the platform are sort of "black box" policies and do not have visible definitions (but the definitions are visible in Checkov).

Please read this article to find out how to convert JSON to YAML format using jq/yq. This is useful if you want to use the YAML Policy Code Editor to be able to create new custom policies with the Policy Code Editor using the existing policies as a template.

<https://stackoverflow.com/questions/53315791/how-to-convert-a-json-response-into-yaml-in-bash>

Custom Compliance Standard

Prisma Cloud includes an extensive list of out-of-the-box compliance standards..

You can also create your own Compliance standards and assign the chosen default policies and/or your custom policies for compliance checks for this custom compliance standard.

Create one from scratch or use "clone existing standard" as a framework for your new custom compliance standard. Please note that you will need to edit your policies in order to assign them to a specific compliance standard.

CI Scan for Container Images

You can integrate Prisma cloud scanning in the pipelines for these technologies:

- Azure DevOps
- Gitlab CI/CD
- Jenkins
- Any other pipeline that we don't integrate with CI plugins can be integrated with Twistcli

[CI plugin Guide](#)

Best Practices for “Other” Pipeline Integration

For CI/CD pipeline tools that we do not natively support integration with, TwistCLI can be used instead and should be embedded in the pipeline at a stage that is before the container/image is deployed. This allows for the container/image to be scanned for vulnerabilities and blocked as needed per the CI rules in Prisma Cloud Compute.

The following Code Block below was pulled from a Gitlab pipeline task and demonstrates how to pull down and configure TwistCLI on the pipeline job runner agent, run the image scans using credentials and URL pertinent to the customer's tenant, and publish the results both in the pipeline output and to the console.

```
prisma-cloud-compute-scan:
  stage: build
  variables:
    prisma_cloud_compute_url: ""
    prisma_cloud_compute_username: ""
    prisma_cloud_compute_password: ""
    prisma_cloud_scan_image: node:lts-alpine
  before_script:
    - apk update && apk add --no-cache docker-cli
    - docker version
    - apk --no-cache add curl
    - apk add --no-cache --upgrade bash
    - |
      if ! /tmp/twistcli --version 2> /dev/null; then
        echo "Download twistcli binary file ..." curl
        -k -u
        ${prisma_cloud_compute_username}:${prisma_cloud_compute_password} \
          --output /tmp/twistcli
        ${prisma_cloud_compute_url}/api/v1/util/twistcli
        chmod +x /tmp/twistcli
      fi
    - |
      /tmp/twistcli --version
    - |
      echo "Create image scan helper script image_scan.sh ..."
      cat > ./image_scan.sh << EOF
      #!/bin/bash
      set +e
      /tmp/twistcli images scan --details --address \${prisma_cloud_compute_url} \
      \
      --user=\${prisma_cloud_compute_username}
      --password=\${prisma_cloud_compute_password} \
      --output-file twistcli.json \${prisma_cloud_scan_image}
      rc=\${?}
      if [ -f twistcli.json ]; then
```



```
mkdir -p report/image_scan
touch report/image_scan/results.xml
docker run --rm \
  -v \${PWD}/twistcli.json:/tmp/twistcli.json \
  -v \${PWD}/report/image_scan/results.xml:/tmp/results.xml \
  redlock/pcs-sl-scanner pcs_compute_junit_report
fi
exit \${rc}
EOF
chmod +x ./image_scan.sh
script:
  # if script is defined in extended job, make sure below command is added
  - bash
  ./image_scan.sh
artifacts:
  when : always
  paths:
    - report/image_scan/results.xml
reports:
  junit:
    - report/image_scan/results.xml
tag
- shell
```

[Additional CI platform integrations sample code](#)

CI/CD Risks

CI/CD Risks are a set of predefined rules that aim to identify vulnerabilities in the CI/CD pipeline by analyzing the security settings and configurations of various systems, as well as their interconnectivity. The risks are classified based on different security categories, including attack vectors, misconfigurations, and bad practices throughout the CI/CD pipeline.

Through 'CI/CD Risks,' you can view an inventory and description of the CI/CD risks detected in a scan, including their severity and the potential impact of the risk. You can remediate risk events by applying suggested solutions or suppress them if you do not want to actively address them. In addition, you can view a Kill Chain graph, which visualizes the stages of a cyber attack on the repository CI/CD pipeline.

To access CI/CD Risks, in Application Security select **Home > CI/CD Risks**.

The CI/CD Risks page includes the following components:

- Filters: Narrow a search for a CI/CD risk by using filters.
- Dashboard: Displays charts representing a visual display of CI/CD risks by system and category



- CI/CD risk inventory table: A list of all risks detected in your CI/CD pipelines.

Dashboard

The 'Dashboard' provides a visual display of CI/CD risks categorized by system and category. When you select a value, such as a GitHub Actions, in the 'Risks by System' chart, both the 'Risks By System' and the 'Risks By Category' charts will update to show results relevant to your selection. In addition, the inventory table will be filtered to display the risks associated with your selection.

Filters

The following filters allow you to narrow a search for a CI/CD risk:

- Status: Filter by status. Values: 'Open', 'Suppressed' and 'Fixed'.
- Severity: Filter by severity. Values: 'Critical', 'High', 'Medium', 'Low', 'Informative'.
- Risk Priority: Filter by the potential impact of the risk and ease of remediation. Risk priority is determined through automatic risk analysis. Values: 1-4.
- System: Filter by the type of technologies found in the system, such as VCSs, CI, container registry and so on.
- Risk: Filter by a risk detected in the pipeline
- Category: Filter by system-based categories detected in your organization
- Repository: Filter by a repository in your environment
- Group By: For details, see CI/CD Risk Inventory Details below
The System filter only displays technologies in which CI/CD risks have been detected.
- Risk: Filter by CI/CD risk detected in the organization pipelines.
- Category: Filter pipeline risks by system-based categories detected in the organization.
By default, Prisma Cloud assigns a category to a risk. The category cannot be modified.
- Repository: Filter by repository which includes detected risks.

CI/CD Risk Inventory Table

The CI/CD risk inventory table includes the following details.

- Group By: Group risks in the table by values corresponding to the inventory table headers. This makes it easier to find and understand specific information, allows for a more systematic data analysis, and helps identify patterns, trends, and relationships.
- Risk Name: The name of the CI/CD risk
- Open Events: The number of open events of a risk detected in the CI/CD pipeline. See below for more on open events.
An 'event' is a particular instance of a CI/CD risk. For example, if the CI/CD risk: 'Possible command

injection attack using crafted Issue user event on [REPO_NAME] by [USER]' (details protected by confidentiality) is detected in your GitLab account, it is considered a single event or instance of the 'Possible command injection detected in user event' risk.

- **Severity:** The severity of the CI/CD risk indicated by a color. Values: 'Critical,' 'High,' 'Medium,' 'Low' and 'Informative'.
CI/CD risk severity levels are set by Prisma Cloud.
- **System:** The system (such as GitHub, Jenkins and so on) containing the CI/CD risk
- **Risk Priority:** For details, see the Risk Priority filter description above
- **System:** The system (such as GitHub, Jenkins and so on) containing the CI/CD risk
- **CI/CD Category:** The risk category that the CI/CD risk is assigned to allows organizations to focus efforts to secure their CI/CD ecosystem. Includes a link to the [Top 10 Security Risks](#).
For more information on CI/CD risk categories, refer to [OWASP's Top 10 CI/CD Security Risks](#).
- **Last Event:** The latest instance of the CI/CD risk detected in your environment
Selecting a CI/CD risk in the inventory opens the resource explorer, displaying additional information about the selected risk.

Tab Information

- The Overview tab opens as the default view, displaying metadata about the CI/CD risk, including a detailed description, severity, graph view of the location of the risk in the delivery chain, the status of the CI/CD risk events (the number of open, closed or suppressed instances of the risk), the system and category in which they were detected, and when last calculated. In addition, remediation is provided through Steps to Solve. See [Suggested Fixes](#) for more on remediation.
- The Open Events tab provides a list and details of the open events, including the name and description of the event, and when detected. You can suppress an event (see [Suppress Events](#) below for more on suppressing events). In addition, you can access the kill chain graph, when applicable to the event by selecting the icon under Actions. See the [Kill Chain Graph](#) below for more.
- The Suppressed Events tab provides a list and details of suppressed events, including an option to unsuppress an event.
- The Fixed Events tab provides a list and details of fixed events.

Kill Chain Graph

The Kill Chain graph visualizes the stages of a cyber attack on the CI/CD pipeline, from the initial reconnaissance phase to the final objective of the attack, resulting from multiple misconfigurations across various systems. The graph is used to illustrate the various steps that an attacker takes to penetrate the pipeline, and it can help identify potential vulnerabilities in the pipeline.



The Kill Chain graph currently supports the Direct Poisoned Pipeline Execution and Direct Poisoned Pipeline Execution by external collaborators CI/CD security risk categories. For more on CI/CD security risk categories, see [OWASP Top 10 CI/CD Security Risks](#).

To view the kill chain graph for a risk, select a risk > Open Events.

The Kill Chain graph includes nodes and edges, describing the connections between them.

For more actions you can take on the graph, refer to the [Repository Application Graph](#).

Suggested Fixes

Prisma Cloud provides suggested solutions to fix instances of CI/CD risks detected in your system: Select a risk from the inventory > in the Details tab, and scroll down to Steps to Solve.

Suppress Events

An 'event' represents a particular instance of a CI/CD risk. By suppressing an event, you intentionally choose not to actively address the event. This can be useful if the error is known and does not require immediate attention or if the error is expected and does not impact the functionality or stability of the system in which it was detected.

To suppress events, select a risk from the inventory table > choose the required events under the Open Events tab of the resource explorer > Suppress. The selected events will be removed and displayed under the Suppressed Events tab.

Unsuppress Events

Unsuppress an event or multiple events in order to take action on it when the event requires attention.

To unsuppress an event, select the risk from the inventory table > choose the required events under the Suppressed Events tab > Unsuppress. The status of the selected events will be restored as 'Open', and will be displayed under the Open Events tab.

Global Suppression

Disable policies globally to exclude calculating issues (risks) detected during a scan in order to reduce overall scan time, prevent unnecessary policies from being scanned, and help reduce false positives:

1. On the Prisma Cloud console, select Governance.
2. Select the relevant policy > in the Status column, toggle OFF.
3. If the "Status" column is missing, you can add it by accessing the table menu and selecting Status.
4. Click Confirm in the Confirm Disabling Policy Status popup that is displayed.
Disabling policies automatically resolves open events. You cannot re-enable the policy for the next 4 hours.

For more information on disabling policies, refer to [Manage Prisma Cloud Policies](#).

Export CI/CD Risk Data

You can export all CI/CD risk data or the data relating to an open, suppressed, or fixed event as a CSV file:


- To export all CI/CD risk data: Select the Download icon found on the top right of the CI/CD risk inventory.
- To export open, suppressed, or fixed event data: Select the Download icon in a corresponding sidecar tab when selecting a risk in the inventory table.

Administration

First steps are to inventory cloud accounts to start ingesting platform logs to begin the process of analysis by Prisma Cloud. Administrators need access to the platform based on their roles and responsibilities. Enabling the Adoption Advisor is highly recommended, as it allows customers to view progress in configuring initial critical configuration tasks, to get the most out of the platform and its various modules.

Adoption Advisor

Prisma Cloud's Adoption Advisor (AA) is a tool that helps you see how far you've explored the tool's capabilities. It allows you to view the various tasks to perform in order to adopt Prisma Cloud – providing you visibility into security areas that you have not discovered yet. Adoption Advisor is available for Cloud Security, Runtime Security, and Application Security. It groups the tasks into three categories – Foundations, Intermediate, and Advanced. There's a percentage that's associated with how much you've adopted Prisma Cloud so far, and completing tasks under the three categories will further increase the percentage shown. You will see AA in the top right of your screen with the percentage showing. Your team should utilize AA to get the most of the Prisma Cloud tool and discover/learn capabilities that you may have not configured yet but are beneficial to have within your organization. Complete the different actions where you'll see a task, description, summary, and then clicking into it will have the tool walk you through how to complete that specific task.



ADOPTION ADVISOR

Great Progress, You've Got This! 🎉

10/17 10/24 11/01 11/08 11/15 10/17 10/24 11/01 11/08 11/15

● Critical ● High

Recommended Next Steps

To gain visibility into your operationalization journey so you know where you are, what to do, and why.

[View By: Code & Build, Deploy, Runtime](#)

Task	Progress	Category
Enable Your Team	6 of 6 Tasks Completed	100%
Code & Build	12 of 15 Tasks Completed	80%
Deploy	5 of 8 Tasks Completed	63%
Runtime	18 of 25 Tasks Completed	72%
Remediation	0 of 1 Task Completed	Intermediate
Workload Scanning	1 of 3 Tasks Completed	Foundational
Access Control	2 of 3 Tasks Completed	Advanced
Vulnerability Detection	2 of 2 Tasks Completed	Intermediate
Data Exposure Visibility	1 of 3 Tasks Completed	Foundational
Discovery and Exposure Management	Not Enabled	Intermediate
Workload Defense	2 of 3 Tasks Completed	Intermediate
Drift Detection	✓	Intermediate



Configuring and spending time in the initial onboarding process will help in the long run when it comes to other Prisma Cloud configuration tasks, such as Alert Rules, RBAC, alert monitoring, and more.


Managing Prisma Cloud Administrators

Review Prisma Cloud role permissions, and create organization-specific Roles tied to the appropriate [Permission Groups](#) (System Admin, Account Group Read-Only, Cloud Provisioning Admin, etc.). It is best practice to not make every user a System Administrator and to tie the least amount of access needed to each user/team. A common role used across organizations is an Account Group Administrator or Account and Cloud Provisioning Administrator. These two roles allow a user to utilize the Prisma Cloud tool but only access the Account Group(s) they're responsible for.

Here are some key concepts to consider:

- What is the function of the user?
- Does the user need to access the Prisma Cloud Portal, or will automation/integration provide what they need?
- If the user does not require access to the Prisma Cloud Portal, will an emailed report be sufficient enough?
- Does the user need to do more than just consume asset/security data?
- Is there asset/security data the user should NOT have access to?
- Is there a capability the user needs to access that cannot be done with a read-only role?

See the below screen capture showing Access Control settings. Here, you can access "Permission Groups", "Roles", "Users", "Access Keys", and "SSO" configurations.



The screenshot shows the Prisma Cloud Security console. The top navigation bar includes links for Home, Dashboards, Reports, Inventory, Compliance, Alerts, Investigate, and Governance. The 'Settings' button is highlighted with a red box. The left sidebar shows the 'CONFIGURE' section with 'Access Control' highlighted. The main content area displays the 'ACCESS CONTROL' settings, with a sub-tab for 'Permission Groups'. A table lists various roles with their names, associated roles, types, last modified by, last modified date, and actions.

Name	Roles	Type	Last Modified By	Last Modified	Actions
Appsec Admin	appsec	Default	template@redlock.io	5 years ago	
Cloud Provisioning Admin	Giang CPA, Test role with Clou...	Default	template@redlock.io	5 years ago	
Developer	sca test, test-developer-role, de...	Default	template@redlock.io	5 years ago	
JIT Requester	jlt-requester	Default	template@redlock.io	5 years ago	
JIT Admin	jlt-admin, JIT Admin new	Default	template@redlock.io	5 years ago	
Build and Deploy Security	MG Build and Deploy Security, ...	Default	template@redlock.io	5 years ago	
NetSecOps	Netsecops Read Only, sawyer...	Default	template@redlock.io	5 years ago	
Account and Cloud Provisioning Admin	ACPA, 1672985036.1021314, ...	Default	template@redlock.io	5 years ago	
Account Group Read Only	TempRo, aarnav AGRO, Repo-r...	Default	template@redlock.io	5 years ago	

Single Sign On

In setting up authentication management, [SSO integration](#) is recommended as best practice. You can enable SSO using the Identity Provider (IdP) that supports Security Assertion Markup Language (SAML) or OpenID Connect (OIDC).

Examples include Okta, Microsoft Active Directory Federation Services (AD FS), Azure Active Directory, Google, or OneLogin. Note, you can only configure a single IdP across the cloud accounts that Prisma monitors. Organizations can add administrative users on Prisma Cloud to create their local account when SSO is enabled or utilize Just-In-Time (JIT) provisioning on the SSO configuration if you'd like to have the accounts created locally.

Note: It is important to provide at least two admin users who can bypass the third-party SSO - a setting under the SSO settings/configuration. This is needed in the event there are SSO issues, such as a SSL certificate expiration or an IdP problem.

Some Examples of SSO Roles

- Create roles based on user personas (DevOps, SecOps, SOC/IR, Compliance, App Developers)
- Attach roles to different account groups based on cloud account ownership

If you are facing issues with user authentication and SSO configuration, you can find a list of the last five SAML login failures by navigating to the bottom of the SSO configuration page.

SAML Troubleshooting common errors

1. Mandatory Fields: Check to ensure that mandatory fields are filled out correctly. IdP

Issuer and Certificate are the two required fields. If you're using HIT, the additional fields must also be filled correctly.

2. SSO Not Enabled: Enable SSO under the main Prisma Cloud settings.
3. Authentication Failed Errors: If a user experiences an authentication failure when they try to log in, you can investigate the issue using a SAML browser plugin to capture the assertion that's being sent to the user's browser. SAML Assertion is the XML document that the IdP sends to the service provider that contains the user authorization. It is important to remember that the URL or certificate information in the asset may not match the Prisma Cloud configuration.
4. JIT Authentication Failed Errors: The URL, certificate, or JIT user parameters may not be correct and can be analyzed from the Assertion XML document. There may be missing attribute values in that the Prisma Cloud SSO config may have an incorrect attribute key name.

⚙️ ACCESS CONTROL | Permission Groups | Roles | Users | Access Keys | **SSO**

Reset Save

SSO Settings

SAML SSO Enabled
 With SAML SSO enabled, all users can access Prisma Cloud console only via the IdP configured URL. To avoid account lockout, Direct User Authentication may be enabled for individual users to login using Prisma Cloud console credentials.

Direct User Authentication
 In addition to logging in via the SSO provider, the following list of selected users can also access the Prisma Cloud console directly by using their email and password registered with Prisma Cloud.

Users (Optional)

SAML **OIDC**

☒ Enabled

Audience URI (SP Entity ID)
 https://app0.prismacloud

Identity Provider Issuer
 http://www.okta.

Certificate
 -----BEGIN CERTIFICATE-----

Generate a SAML SP Metadata file

By default, Cloud Security Module can't generate a metadata file which means that we need to generate it using a third party tool like <https://www.samltool.com> or create it manually.

1. Go to <https://www.samltool.com>

Build SP Metadata

Build the XML metadata of a SAML Service Provider providing some information: EntityID, Endpoints (Attribute Consume Service Endpoint, Single Logout Service Endpoint), its public X.509 cert, Nameid Format, Organization Info and Contact info.

This metadata XML can be signed providing a public X.509 cert and the private key.

 CLEAR FORM FIELDS

EntityId 	SP X.509 cert (same cert for sign/encrypt) (Optional) 
Attribute Consume Service Endpoint (HTTP-POST) 	
Single Logout Service Endpoint (HTTP-REDIRECT) (Optional) 	
Nameid Format (Optional) urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified	AuthnRequestsSigned (Optional) False
	WantAssertionsSigned (Optional) False

- Identify the following values on your Prisma Cloud Tenant:

EntityId → Audience URI (SP Entity ID) →

`https://app<TENANTNUMBER>.prismacloud.io/customer/<CUST_ID_HERE`

> Attribute Consume Service Endpoint (HTTP-POST) →

`https://api<TENANTNUMBER>.prismacloud.io/saml`

WantAssertionsSigned → True

- Fill out the above data on <https://www.samltool.com> and click on “Build SP METADATA”

Build SP Metadata

Build the XML metadata of a SAML Service Provider providing some information: EntityID, Endpoints (Attribute Consume Service Endpoint, Single Logout Service Endpoint), Its public X.509 cert, NameID Format, Organization Info and Contact Info.

This metadata XML can be signed providing a public X.509 cert and the private key.

 CLEAR FORM FIELDS

EntityId https://app.prismacloud.io/customer/ Attribute Consume Service Endpoint (HTTP-POST) https://api.prismacloud.io/saml Single Logout Service Endpoint (HTTP-REDIRECT) (Optional) NameId Format (Optional) urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified	SP X.509 cert (same cert for sign/encrypt) (Optional) AuthnRequestsSigned (Optional) False WantAssertionsSigned (Optional) True
--	--

SP Metadata Output example:

```
<?xml version="1.0"?>
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata"
    validUntil="2022-10-20T16:32:31Z"
    cacheDuration="PT604800S"
    entityID="https://app<TENANTNUMBER>.prismacloud.io/customer/<CUST_ID_HERE>">
    <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:Name
        IDFormat>

        <md:AssertionConsumerService
        Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"

        Location="https://api<TENANTNUMBER>.prismacloud.io/saml"

        index="1" />

    </md:SPSSODescriptor>
</md:EntityDescriptor>
```

Access Management for Application Security

Administrator access is the same process as for CSPM. You create [roles](#), [users](#) and [access keys](#) via the Prisma Cloud Settings. When creating roles, it is important to note which [Permission Groups](#) are relevant to the CCS module. For instance, only the SYS ADMIN and the ACCOUNT AND CLOUD PROVISIONING ADMIN permissions allow to add, update or delete repositories, while the DEVELOPER role provides the least privileged permissions.

Navigate to **Settings > Access Control** and select **Add > Role**.

Navigate to **Settings > Access Control** and select **Add > User**.

Navigate to **Settings > Access Control** and select **Add > Access Key**.

Administrators can also access the CAS module using any Single Sign-On integrated with Prisma Cloud. No extra configuration is required.

Set Up Administrator Access for Cloud Application Security

You need to enable administrative access for all the DevSecOps and Security teams who need to add code repositories or pipelines, create policies and review scan results on Prisma Cloud. To know more see - add [Prisma Cloud Administrators and role permissions](#). You can also see: - [add administrative users](#).

Prisma Cloud Roles and Application Security Permissions:

The following reference page provides Application Security list of the access privileges associated with each role for the different parts of the Prisma Cloud administrative console.

- [Application Security Administrator Permissions](#)

The following are the available options for creating Users with access to the Cloud Application Security module:


1. Use the existing System Admin Role for full access to the Cloud Application Security features.
2. Set up Developer role for Developers with view access to the relevant repositories and access to issue Fix and Submit changes to the relevant repositories.

For CI/CD scanning add the relevant Developer Service account with the relevant Developer and repository access.

3. Create Account and Cloud Provisioning Admin role to allow repositories edit access to the relevant repositories (Alternative is to using the existing SysAdmin roles for this)
4. Create new Cloud Application Security specific roles or update existing Account Group Read Only role to allow view access to relevant repositories (if required)

Cloud Application Security Roles Configuration

1. Create **Developer** role for Developers and Service Accounts with view/submit fix access to the relevant repositories.



Create New Role ✕

Name
Developer

Description (Optional)

Permission Group [View Permissions](#)
Developer

Repositories (Optional)

Search


63 Results [Select All](#) [Deselect All](#)

☐	aks	cli (806775482162247680_cli_repo)
☐	aws-containers-task-definitions	Github (hssong97)
☐	Azure	Github (afdsioh234)
☐	AzurePAVM	Github (hssong97)
☐	cfngoat	Github (PCS-LAB-ORG)
☐	cfngoat	Github (github-hadi)
☐	ckurepo01	Github (PCS-LAB-ORG)

2. Create **Account and Cloud Provisioning Admin** role with permissions to edit a specific repository.

Example to create a new IaCodeEditRepositories role with the "Account and Cloud Provisioning Admin" permission group and assign it to the relevant users.

This role permissions can be enabled for all onboarded repositories and/or filtered to specific repositories:



Create New Role ✕

Name
IaCodeEditRepositories

Description (Optional)

Permission Group [View Permissions](#)
Account and Cloud Provisioning Admin

Account Group (Optional)

Resource List (Optional)

Repositories (Optional)

Search

☐ Select All (3 results)

☐ AST-IV

☐ prismapoc

☐ SbD_IaC_POC_Modulescan

3. Create New Cloud Application Security specific or update existing **Account Group Read Only** permissions role to allow view access to relevant repositories (if required)

Create New Role

×

Name

Prisma_Cloud_Read_Only

Description (Optional)

Permission Group

View Permissions

Account Group Read Only

▼

Account Group (Optional)

▼

Resource List (Optional)

▼

Repositories (Optional)

^

Search

Q

63 Results

Select All

Deselect All

☐	aks	cli (806775482162247680_cli_repo)
☐	aws-containers-task-definitions	Github (hssong97)
☐	Azure	Github (afdsioh234)
☐	AzurePAVM	Github (hssong97)
☐	cfngoat	Github (PCS-LAB-ORG)
☐	cfngoat	Github (github-hadi)
☐	ckurepo01	Github (PCS-LAB-ORG)

Generate API Access Key for Developers (IDE)

Prisma Cloud uses Access Keys to integrate with the environments where you host your templates, source code, or pipelines.

For CI/CD integration and API use, it is recommended to use Service Account keys (No User Access Account keys). Please see the next section. Service accounts do not provide access to the Prisma Cloud portal.

There should be no Cloud Application Security user API Access Keys present unless you are using IDE integration.

The user must have the “Allow user to create API Access Keys” ticked.

User Information ×

First Name

Developer-with-IDE

Last Name

Email

Assign Roles

DeveloperTest-OK ▼

Default Role

DeveloperTest-OK ▼

☒ Allow user to create API Access Keys

Cancel

Save and add another


Save and close

Access keys are specific to a user and they enforce the role and permissions assigned to the specified user. Use the key expiry date for each key (you can have up to 2).


A user can have more than one Prisma Cloud Role assigned to them. The default Role (assigned when the user is created or updated later) will be used when accessing Prisma Cloud with the user API Access Key.

Note: When you are prompted to add an API Token on any plugin, make sure to provide the relevant User/Service account Prisma Cloud access key ID and secret as the input.

1. Select Settings > Access Control > Access Keys Tab > Add Access Key.




2. Enter an access key **Name** and **Save**.

The screenshot shows the 'Add Access Key' dialog box. It has a title bar with a close button (X). The dialog contains a 'Name' input field, an 'Enable Expiration' toggle switch (currently turned off), and two buttons at the bottom: 'Cancel' and 'Save'.

3. Enable Expiration and set the **Key Expiry**.

As a best practice, set an expiration time for the validity of your access key



Add Access Key [X]


Name

Enable Expiration ☒

Access Key Expiration
 [v]

[Cancel] [Save]

4. Copy and save your new **Access Key ID** and **Secret Key** in a secure location.



Access Key ID
 [Copy]

Secret Access Key
 [Copy] [Eye icon]

[Download .csv file]

[Done]

You can select **Download .csv file** to download this information.

Save your secret key once it is generated, as **you cannot view it again on Prima Cloud**.


Generate API Access Key for Application Security Service Account

Prisma Cloud uses Access Keys to **integrate with CI/CD pipelines or API usage**.

It is recommended to set up a service account with API Access key access only for CI/CD onboarding and API use.

It is also best practice to use dedicated service account/API keys for each CI/CD tool or software deployment pipeline and use key expiry date for each key (you can have up to 2) **And don't forget to test and document API key rotation process.**

1. Select Settings > Access Control > Users > Add > Service account.



We will be using the **Developer** role created previously with access to all repositories (please note: it is recommended to create a dedicated role for access to each security domain with access to dedicated IaC repositories), here are the details.

Edit Role

×

Name

Developer-All-Repos

Description (Optional)

Permission Group

View Permissions

Developer

▼

Repositories (Optional)

aks cli (806775482162247680_cli_repo),aws-containers-task-definitions Github (hssong97)...

▼

Submit

2. Name your service account and select the role that you created in the previous step. This will provide this access key with API access to the role that was created. This means that the API key will have access to the same data that the role provides access to if a user was provisioned with this role.

Add Service Account

Account Details

Access Key Details

Service Account Details

A service account is a special Prisma Cloud identity used to access Prisma Cloud programmatically via API. To create a service account provide a descriptive name and fill in the additional details carefully because you cannot modify these inputs once the account is created.

Service Account Name

Role

Developer-All-Repos

Next

3. Define a name for the access key to be provisioned as well as a date for the access key to expire. It is best practice to name the key with something that references that this is a service account key and if possible what it has access to. Since all keys will fall under the Access Keys section of access control, this will enable the user with a quick glance to identify which keys are used for service accounts (API access only) and which keys belong to a user.

Add Service Account

Account Details

Access Key Details

Access Key Details

Enter a meaningful name for the access key. As a best practice, set an expiration date and rotate the key frequently.

Access Key Name

CodeSecurityServiceAccount-All-Repos

Access Key Expiration

Apr 17, 2024, 12:00:00 AM

Previous

Save & Create (1 of 2)

4. The API key will be created and will state how many keys can be created. Note only 2 access keys can be generated per service account.

Access Key Results ✕

1 of 2 total available access keys has been successfully generated.

This is the only time that the secret access key can be viewed or downloaded. Be sure to

i *download and save these details in a secure location. You cannot recover your secret access key later but you can create a new access key, as needed. As a best practice, do not share the access key id and secret access key with anyone.*

Access Key ID

b1710162-ace9-4202-8ed3-afdb79767025 

Secret Access Key

.....  

[Download .csv file](#)

[Done](#)

Download API key credentials and store them in a safe place like Vault where you can retrieve them securely as part of your CI/CD pipeline run.

Save your secret key once it is generated, as **you cannot view it again on Prisma Cloud**. This is the key that will be used for CI/CD scan onboarding.

Using CSPM API to manage API Access Keys

It is possible to use API endpoints to create and update API access keys at scale:

<https://prisma.pan.dev/api/cloud/cspm/access-keys>

Here is an example on how to use the endpoints to write a script to rotate users and service account keys:

https://github.com/PaloAltoNetworks/prismacloud-api-python/blob/main/scripts/pcs_rotate_service_account_access_key.py

The Prisma Cloud Python API library also includes methods for the other endpoints:



https://github.com/PaloAltoNetworks/prismacloud-api-python/blob/main/prismacloud/api/cspm/_endpoints.py

APIs for Application Security

Prisma Cloud Application Security API

The Prisma™ Cloud Application Security API enables you to check your Infrastructure-as-Code resources against Prisma Cloud out-of-the-box and custom security policies programmatically. The Cloud Application Security API enables you to:

- Initiate Cloud Application Security scans of repositories you've added to Prisma Cloud
- View the repositories you've connected to Cloud Application Security
- Manage Cloud Application Security suppression rules
- Fix or suppress Cloud Application Security policy violations

Use the [CSPM Policy API](#) and [Prisma Cloud Terraform provider](#) to create and manage Cloud Application Security build policies.

New Updates and Feature Releases

Reviewing the Prisma Cloud release notes to learn about all the new features and known issues periodically is recommended. Every release will add new policies, permissions, and APIs to catch up with ever-evolving cloud services.

- [Prisma Cloud Release Information](#)

Additional Information

Self-Hosted Console

Deploy the console

The Console must be deployed first. The initiation of the Console will ensure that only the Defenders that a Console deploys will only be controlled by the Console. Automation can be used to deploy Compute. The Console must be running and licensed before any further configuration can be implemented (e.g. deploy Defenders). The most common methods of deployment are:

- [Onebox](#) - simple single docker host deployment. The `twistlock.sh` (`$ twistlock.sh -sy onebox`) bash script that is included within the release tar will deploy a Console and Defender on the node. Note: Onebox will only deploy the Console on a RHEL node running podman.
- [Kubernetes](#) - Most flavors of K8s ([ACK](#), [ACS](#), [AKS](#), [EKS](#) and [IKS](#)) are supported. Every release is tested on several versions of K8s and they are listed [here](#).
- [OpenShift](#) - basically it is Kubernetes but there are some slight nuances (e.g. external router, OpenShift internal registry, `<=v3.11` - docker, `>=v4.0` - cri-o)

Upgrade the console

You should have kept good notes when initially installing Prisma Cloud. The configuration options set in `twistlock.cfg` and the parameters passed to `twistcli` in the initial install are used to generate working configurations for the upgrade.

Prerequisites: Document and save all options set in `twistcli.cfg` and parameters passed to `twistcli` during the install.

The console upgrade is the one way to upgrade to the higher version, meaning you cannot downgrade to the current version. Due to the restore only support to the same version, it's strongly recommended to test the upgrade with the dev or secondary console first and validate the all utilized functionality.

To safely upgrade the console, it is required to keep the "Default - ignore Twistlock components", in the vulnerability to defend the policy. If this rule is disabled or deleted, there is a chance that an upgrade will fail.

Backup and Restore

Prisma Cloud automatically backs up all data and configuration files periodically. You can view all backups, make new backups, and restore specific backups from the Console UI. You can also restore specific backups using the `twistcli` command line utility. You can also manually backup any point of the time from the console.



Prisma Cloud is implemented with containers that cleanly separate the application from its state and configuration data. To back up a Prisma Cloud installation, only the files in the data directory need to be archived. Because Prisma Cloud containers read their state from the files in the data directory, Prisma Cloud containers do not need to be backed up, and they can be installed and restarted from scratch.

You can only restore Console from a backup file whose version exactly matches the current running version of Console. If the console is unresponsive, you can use `twistcli` to restore the console.

Supported lifecycle for connected components

Any supported version of Defender, `twistcli`, and the Jenkins plugin can connect to Console. Prisma Cloud supports the latest release and the previous two releases (n, n-1, and n-2).

There are some exceptions to this policy as we roll out this new capability. Please find the details of the lifecycle [here](#).

Third Party Integration

Prisma Cloud Enterprise Integration

[Configure integrations](#) with third party security tools and SOC workflow tools. Configure alert workflows for notifications and remediation. Organizations already have multiple processes in place when it comes to managing security in the cloud, whether that's a SIEM tool, logging tool, or ticketing system such as ServiceNow.

1. Setup integrations - helps to monitor alerts and send alert notifications to security processes that already exist in an organization or a new process can be created to manage large amounts of alerts/data. Integrations help with the process of keeping Prisma Cloud alerts manageable.

- a. 3rd party integrations (ex. w/Splunk, ServiceNow, and other native integrations)
 - b. Webhooks (non-native integrations such as a SIEM tool)
2. Create Alert Rules
 - a. This is used to generate specific Alerts - specify the target accounts/account groups, the specific policies (or all), as well the notification channel if one is needed
3. Set up alert profiles and integrations to test alert notification functionality with different types of policies that generate alerts
 - a. This helps to create a parallel to a DevOps pipeline with alert workflows so that instead of having to test how the alert generates information on the channel that will be used, users can test sending alert information in a standardized way through alert profiles and integrations that mirror the configuration of the production level information streams

Compute Edition Integration

Prisma Cloud Compute supports the following third-party tool [integrations](#) out of the box. You can find detailed instructions to integrate those supported tools below.:

- [AWS Security Hub](#)
- [Cortex XDR](#)
- [Cortex XSOAR](#)
- [Email](#)
- [Google Cloud Pub/Sub](#)
- [Google Cloud SCC](#)
- [IBM Cloud Security Advisor](#)
- [JIRA](#)
- [PageDuty](#)
- [ServiceNow Security Incident Response](#)
- [ServiceNow Vulnerability Response](#)
- [Slack](#)
- [Splunk](#)
- [Webhooks](#)

Best practice is to first determine if a customer is using any of the above tools as a SOC tool that can ingest Prisma Cloud alert data. If there is no available out-of-the-box integration for your customer's tool, try to figure out if their tool supports Webhooks ingestion. If Webhooks ingestion is not a probability, there is a possibility that Professional Services can help set up a custom integration with their tool using our APIs.



Application Security Notification

You can enable notifications to the external integrations you have configured in Prisma Cloud Enterprise.. JIRA, ServiceNow, Microsoft Teams, Slack, Splunk and Webhook are supported. Notifications are disabled by default.. See [Configure External Integrations on Prisma Cloud](#) to set up an integration. Then configure [Application Security Notifications](#).

General Resources and References:

- [Prisma Cloud: Customer Corner Monthly Videos](#)
- [Prisma Cloud Live Community](#)
- [Prisma Cloud Official Documentations](#)