



# USER MANUAL

Detailed guide to walk you through the  
cloud security journey

Protect every cloud native  
application, Protect everywhere

# -=AccuKnox Manual=-

## Table of Contents

<b>Getting Started Guide</b>	<b>8</b>
<b>1. Assets Count</b>	<b>8</b>
1.1 Cloud	8
1.1.1 AWS	8
1.1.2 Azure	11
1.1.3 GCP	14
1.2 Container Images Counts	16
1.2.1 DockerHub	16
1.2.2 AWS ECR	16
1.2.3 GCR	16
1.2.4 ACR	17
<b>2. CSPM Prerequisites</b>	<b>18</b>
2.1 AWS	18
2.2 AZURE	20
2.3 GCP	29
<b>3. Cloud Onboarding</b>	<b>36</b>
3.1 AWS Onboarding	36
3.2 Azure Onboarding	39
3.3 GCP Onboarding	42
3.4 Cloud Account Deboarding	44
<b>4. CWPP Prerequisites</b>	<b>45</b>
4.1 Minimum Resource required	45
4.2 AccuKnox Agents	45
<b>5. Cluster Onboarding</b>	<b>51</b>
<b>6. Cluster Offboarding</b>	<b>58</b>
<b>7. VM Onboarding with Systemd/Docker Mode</b>	<b>60</b>
7.1 Systemd	60
7.1.1 Container Protection Requirements (Optional)	61
7.1.2 Resource Requirements	61
7.1.3 Network Requirements	62
7.1.4 Onboarding	63
7.1.5 Onboarding Control Plane	64
7.1.6 Onboarding Worker Nodes	65



7.1.7 Deboarding	66
<b>7.2 Docker</b>	<b>67</b>
7.2.1 Resource Requirements	67
7.2.2 Network Requirements	67
<b>7.3 Onboarding</b>	<b>70</b>
7.3.1 Install Knoxctl/accuknox-cli	70
7.3.2 Onboarding Control Plane	70
7.3.3 Onboarding Worker Nodes	71
<b>7.4 Deboarding</b>	<b>72</b>
<b>8. Registry Onboarding</b>	<b>73</b>
<b>8.1 Azure Container Registry</b>	<b>73</b>
8.1.1 Steps to generate credentials for onboarding ACR	73
8.1.2 Steps to onboard the registry on AccuKnox SaaS	75
<b>8.2 Harbor Registry</b>	<b>77</b>
8.2.1 Prerequisites for Harbor Registry Onboarding in Accuknox:	77
8.2.2 Steps to Onboard Harbor Registry on Accuknox:	79
<b>8.3 Deboarding a Registry</b>	<b>83</b>
<b>9. AccuKnox CNAPP Dashboard Widgets</b>	<b>84</b>
<b>9.1 CWPP Widgets</b>	<b>84</b>
9.1.1 Top 5 cluster findings Widget	84
9.1.2 Findings by Asset Categories Widget	84
9.1.3 K8S Security Metrics Widgets	85
9.1.4 Workload Alerts Widgets	85
9.1.5 Workloads without any Policy Applied Widget	86
9.1.6 K8s Resource SummaryWidget	86
9.1.7 Cluster Connection Status Widget	87
9.1.8 Workloads without Network Policies Widget	87
9.1.9 Top 5 K8s CIS Findings Widget	88
9.1.10 Block based Policies with Associated Alerts Widget	88
<b>9.2 CSPM</b>	<b>89</b>
9.2.1 Top 3 cloud accounts with failed controls Widget	89
9.2.2 Top 10 risk associated to cloud accounts Widget	89
9.2.3 Findings Trends Widget	90
9.2.4 Cloud Accounts Widget	90
9.2.5 Findings Widget	91
<b>9.3 KIEM</b>	<b>91</b>
9.3.1 Kiem Risk Assessment Widget	91

9.3.2 KIEM Findings by Asset type Widget	92
9.3.3 Top 5 most critical findings Widget	92
<b>9.4 Cloud Misconfiguration Widget</b>	<b>93</b>
Cloud Account Risk Assessment Widget	93
<b>9.5 Container images Widgets</b>	<b>93</b>
9.5.1 Image Severity Distribution Widget	93
9.5.2 Image Risk Assessment Widget	94
<b>9.6 Tickets Widgets</b>	<b>94</b>
Tickets by status Widget	94
<b>10. CSPM (Cloud Security Posture Management)</b>	<b>95</b>
<b>10.1 Asset Inventory</b>	<b>95</b>
10.1.1 How to find a particular asset	95
10.1.2 How to group assets	97
10.1.3 How to search asset by label	100
<b>10.2 Misconfigurations</b>	<b>101</b>
10.2.1 Where to find misconfigurations	101
10.2.2 How to group by Asset and find misconfiguration	102
10.2.3 How to group by Findings	105
10.2.4 How to group by criticality and Status	107
10.2.5 How to create a Ticket	111
<b>10.3 Issues/Findings</b>	<b>116</b>
10.3.1 Group findings by source and severity	116
10.3.2 How to group by Findings and severity	117
10.3.3 How to group by Asset and severity	118
<b>10.4 Baselines</b>	<b>120</b>
10.4.1 How to create a Baseline out of a data source	120
10.4.2 How to compare baselines	122
<b>10.5 Compliance</b>	<b>124</b>
10.5.1 How to get Compliance for Cloud Assets	124
<b>10.6 Remediation - Fix Problems/Create Tickets</b>	<b>126</b>
<b>10.7 CSPM Reports</b>	<b>128</b>
<b>10.8 Rules Engine</b>	<b>130</b>
<b>11. ASPM (Application Security Posture Management)</b>	<b>133</b>
<b>11.1 SAST</b>	<b>133</b>
11.1.1 Integrating SonarQube SAST with AccuKnox in a GitLab CI/CD Pipeline	133
11.1.2 Pre-requisites	133
11.1.3 Steps for Integration	133

11.1.4 Initial CI/CD Pipeline Without AccuKnox Scan	137
11.1.5 CI/CD Pipeline After AccuKnox Integration	137
11.1.6 View Results in AccuKnox SaaS	137
<b>11.2 DAST</b>	<b>139</b>
11.2.1 Gitlab DAST Scan	139
11.2.2 Pre-requisites	139
11.2.3 Steps for Integration	139
11.2.4 Initial CI/CD Pipeline Without AccuKnox Scan	142
11.2.5 CI/CD Pipeline After AccuKnox Scan Integration	142
11.2.6 View Results in AccuKnox SaaS	142
<b>11.3 IaC GitLab Scan</b>	<b>145</b>
11.3.1 Integrating IaC with AccuKnox in a GitLab CI/CD Pipeline	145
11.3.2 Pre-requisites	145
11.3.3 Steps for Integration	145
11.3.4 Initial CI/CD Pipeline Without AccuKnox IaC Scan	149
11.3.5 CI/CD Pipeline After AccuKnox IaC Scan Integration	149
11.3.6 View Results in AccuKnox SaaS	149
<b>11.4 GitLab IaC Scan via Accuknox</b>	<b>152</b>
11.4.1 Prerequisites	152
11.4.2 Configuring Code Source in Accuknox	153
11.4.3 Setting Up IaC Configuration	154
11.4.4 Viewing and Managing IaC Findings on Accuknox	155
<b>11.5 Container Scan Use Case</b>	<b>156</b>
11.5.1 Scenario Before Integrating AccuKnox	156
11.5.2 Scenario After Integrating AccuKnox	157
11.5.3 Remediation and Rescan	159
11.5.4 Steps needed to be taken for integration	160
<b>12. KSPM (Kubernetes Security Posture Management)</b>	<b>165</b>
12.1 CIS Benchmarking Compliance Scan Onboarding	165
12.2 Cluster Misconfiguration Scanning	171
12.2.1 Remediation	173
12.2.2 Vulnerability Management Lifecycle	174
12.3 Kubernetes Identity and Entitlement Management (KIEM)	177
12.3.1. Onboarding Process	177
12.3.2 Install KIEM Agents	177
12.3.3 Post-Onboarding Steps	177
12.3.4 Permissions Overview	178

12.3.5 Key Queries	179
12.3.6 Full-text Search	180
12.3.7 Entity Exploration	181
12.3.8 Interactive Visualization	184
<b>13. CWPP (Cloud Workload Protection Platform)</b>	<b>185</b>
13.1 Cloud Workloads	185
13.1.1 How to find graph view of clusters	185
13.1.2 How to find list view of clusters	185
13.1.3 How to find details on cluster	186
13.1.4 How to get Compliance for Cloud Workload	188
13.2 App Behavior	189
13.2.1 How to interpret network graph	189
13.2.2 How to see App Behavior Telemetry	193
13.3 Runtime Protection w/ Policy Management	195
13.3.1 How to understand discover policies	195
13.3.2 How to understand Hardening policies	198
13.3.3 How to Audit application and get alerts for that	201
13.3.4 When do we say policies are stable?	203
13.3.5 What if something changes in Application?	204
13.3.6 How to create a custom Policy	208
13.3.7 How to enforce Policies and see anomalies	215
13.3.8 How to perform bulk operation on applying policies	222
13.3.9 How to Find Nodes of a VM cluster	224
<b>14. Host Security</b>	<b>226</b>
14.1 Host Scan	226
14.2 Prerequisites for Nessus Integration	226
14.3 Asset Inventory	227
14.4 Vulnerability Management	227
<b>15. Admission Controller Support Using Knoxguard</b>	<b>230</b>
15.1 Introduction	230
15.2 Prerequisite for Knoxguard Admission Controller	231
15.3 Deployment of Knoxguard	232
15.4 Policy Enforcement	233
15.5 Policy Violation and Alerts	235
15.6 Pod Security Admission Controller	236
15.7 Enabling Pod Security Admission (PSA)	237
15.8 PSA Protection Example	240

<b>16. CWPP Report Generation</b>	<b>242</b>
16.1 Regex	242
16.1.1 Rules for Regular Expression	242
16.2 Reports Configuration	244
16.2.1 On Demand Report Configuration	244
16.2.2 Scheduled Report Configuration	245
<b>17. Integrations</b>	<b>248</b>
17.1 Integrate SIEM tools	248
17.1.1 Splunk	248
a. Prerequisites:	249
b. Steps to Integrate:	249
17.1.2 AWS Cloudwatch	251
a. Prerequisites	251
b. Steps to Integrate:	251
c. Configuration of Alert Triggers:	252
d. Logs Forwarding:	253
17.1.3 Azure Sentinel Integration	253
a. Prerequisites:	253
b. Steps to Integrate:	253
17.1.4 Creating webhook using the Azure Logic App	254
a. About the logic app:	254
b. To see Logs in the Sentinel:	254
17.1.5 Rsyslog	254
a. Prerequisites:	254
b. Steps to Integrate:	255
17.2 Integrate Notifications Tools	256
17.2.1 Slack	256
a. Prerequisites:	256
b. Steps to Integrate:	256
17.3 Integrate Ticketing Tools	257
17.3.1 Jira Integration	257
a. Prerequisites	258
b. JIRA integration for CWPP:	258
17.3.2 JIRA integration for CSPM:	259
17.3.3 ServiceNow Integration	262
a. Prerequisites	262
b. Steps for integration	262

17.3.4 Freshservice Integration	262
a. Prerequisites	262
b. Steps to Integrate:	263
17.4 Creating Ticket Configuration	267
17.5 Email Integration	269
<b>18. User Management</b>	<b>271</b>
18.1 Inviting a New User	271
18.2 User Receives Invitation	273
18.3 User Login Options	273
Option A: Traditional Login	273
Option B: Single Sign-On (SSO) with Google	274
18.4 Assign RBAC	275
18.5 Create Roles and Assign Users	276
<b>19. Ticketing Procedures</b>	<b>278</b>
19.1 How to raise an AccuKnox support ticket?	278
19.2 How to track the issue resolution status?	281
<b>20. FAQs</b>	<b>283</b>
20.1 AccuKnox FAQs	283
20.2 Bonus Questions	289
<b>References:</b>	<b>291</b>

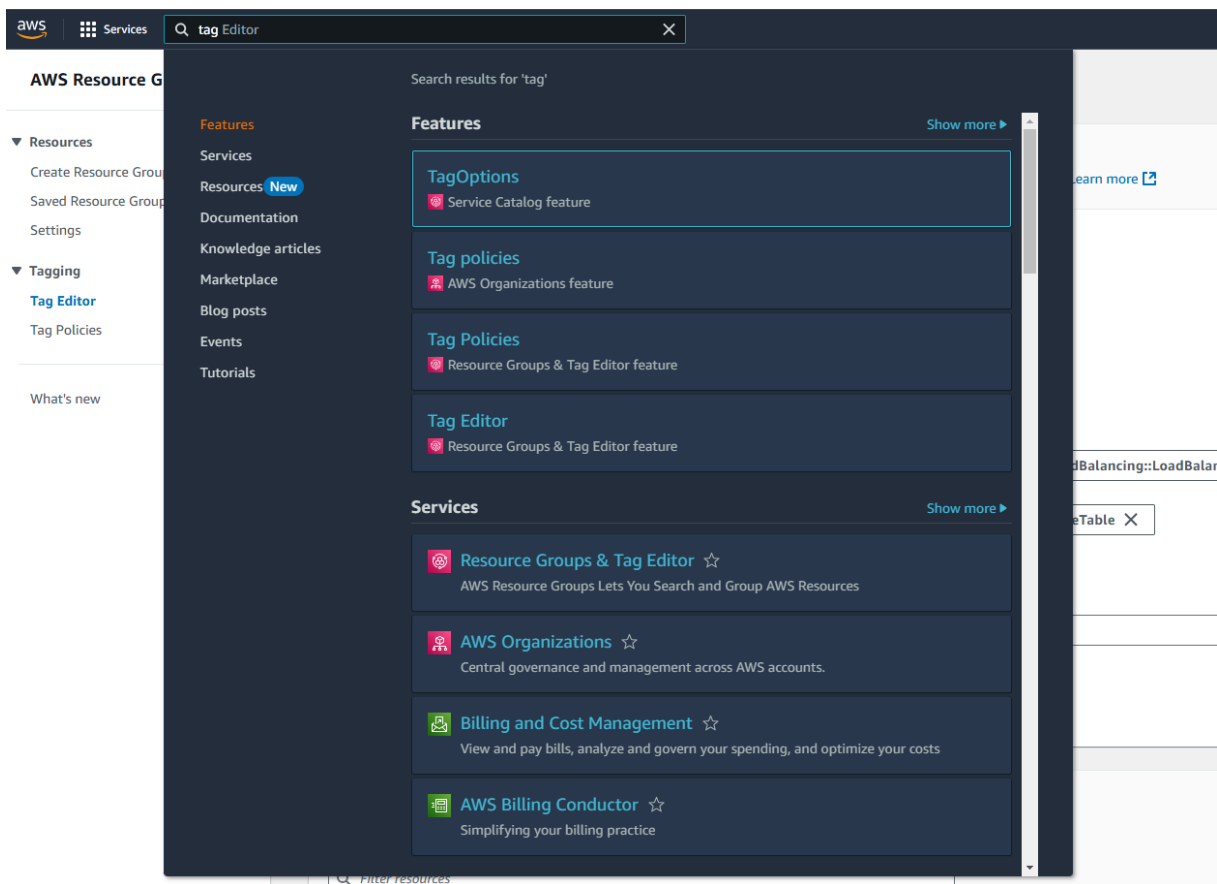
# Getting Started Guide

## 1. Assets Count

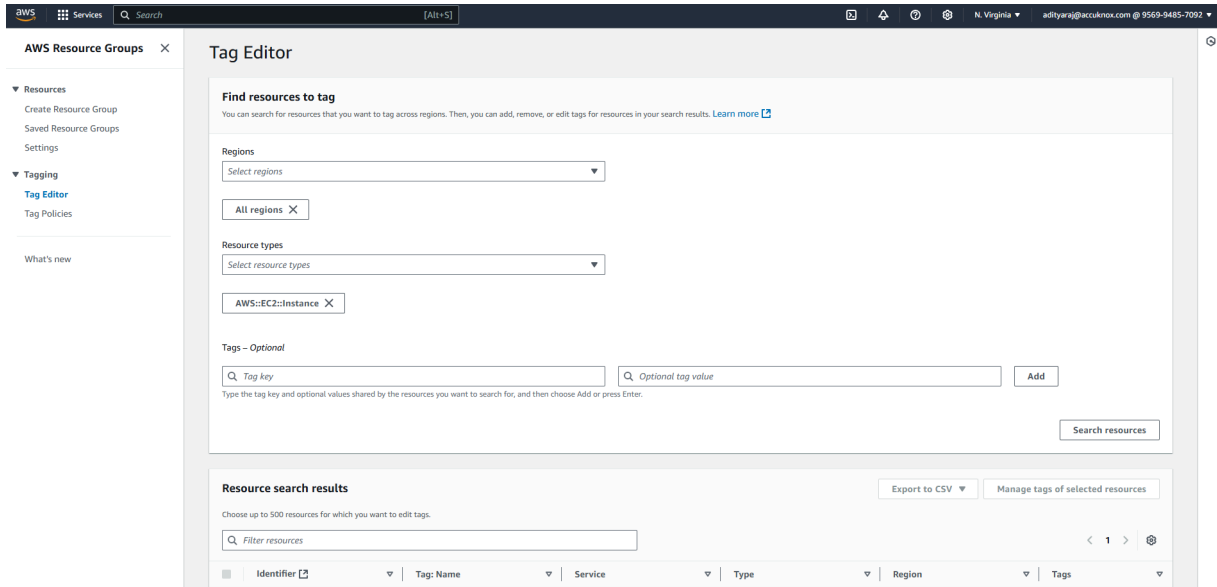
### 1.1 Cloud

#### 1.1.1 AWS

##### Step 1: Search for Tag Editor



##### Step 2: In the tag editor screen, filters can be applied by passing them in the URL



Paste the following snippet at the end of the URL in the browser:

```
#query=regions:!('AWS::AllSupported'),resourceTypes:!('AWS::Lambda::Function','AWS::CloudFront::Distribution','AWS::CloudFront::StreamingDistribution','AWS::IAM::InstanceProfile','AWS::IAM::ManagedPolicy','AWS::EC2::VPC','AWS::EC2::Subnet','AWS::EC2::NetworkAcl','AWS::EC2::NetworkInterface','AWS::ElasticLoadBalancingV2::LoadBalancer','AWS::ElasticLoadBalancing::LoadBalancer','AWS::EC2::EIP','AWS::EC2::SecurityGroup','AWS::EC2::RouteTable','AWS::Route53::Domain','AWS::Route53::HealthCheck','AWS::RDS::DBSubnetGroup','AWS::EC2::Instance','AWS::ECS::Cluster','AWS::EKS::Cluster','AWS::ElastiCache::CacheCluster','AWS::ElastiCache::Snapshot','AWS::S3::Bucket','AWS::EC2::Volume','AWS::EC2::Snapshot','AWS::RDS::DBInstance','AWS::RDS::DBCluster'),tagFilters:!(,),type:TAG_EDITOR_1_0
```

Eg:

```
https://us-east-1.console.aws.amazon.com/resource-groups/tag-editor/find-resources?region=us-east-1#query=regions:!('AWS::AllSupported'),resourceTypes:!('AWS::Lambda::Function','AWS::CloudFront::Distribution','AWS::CloudFront::StreamingDistribution','AWS::IAM::InstanceProfile','AWS::IAM::ManagedPolicy','AWS::EC2::VPC','AWS::EC2::Subnet','AWS::EC2::NetworkAcl','AWS::EC2::NetworkInterface','AWS::ElasticLoadBalancingV2::LoadBalancer','AWS::ElasticLoadBalancing::LoadBalancer','AWS::EC2::EIP','AWS::EC2::SecurityGroup','AWS::EC2::RouteTable')
```



```
able', 'AWS::Route53::Domain', 'AWS::Route53::HealthCheck', 'AWS::RDS::DBSubnetGroup', 'AWS::EC2::Instance', 'AWS::ECS::Cluster', 'AWS::EKS::Cluster', 'AWS::ElasticCache::CacheCluster', 'AWS::ElasticCache::Snapshot', 'AWS::S3::Bucket', 'AWS::EC2::Volume', 'AWS::EC2::Snapshot', 'AWS::RDS::DBInstance', 'AWS::RDS::DBCluster'), tagFilters: !(), type: TAG_EDITOR_1_0
```

**Resource types**

Select resource types ▼

AWS::Lambda::Function X

AWS::CloudFront::Distribution X

AWS::CloudFront::StreamingDistribution X

AWS::IAM::InstanceProfile X

AWS::IAM::ManagedPolicy X

AWS::EC2::VPC X

AWS::EC2::Subnet X

AWS::EC2::NetworkAcl X

AWS::EC2::NetworkInterface X

AWS::ElasticLoadBalancingV2::LoadBalancer X

AWS::ElasticLoadBalancing::LoadBalancer X

AWS::EC2::EIP X

AWS::EC2::SecurityGroup X

AWS::EC2::RouteTable X

AWS::Route53::Domain X

AWS::Route53::HealthCheck X

AWS::RDS::DBSubnetGroup X

AWS::EC2::Instance X

AWS::ECS::Cluster X

AWS::EKS::Cluster X

AWS::ElasticCache::CacheCluster X

AWS::ElasticCache::Snapshot X

AWS::S3::Bucket X

AWS::EC2::Volume X

AWS::EC2::Snapshot X

AWS::RDS::DBInstance X

AWS::RDS::DBCluster X

**Tags - Optional**

Type the tag key and optional values shared by the resources you want to search for, and then choose Add or press Enter.

---

**Resource search results (1938)**

Choose up to 500 resources for which you want to edit tags.

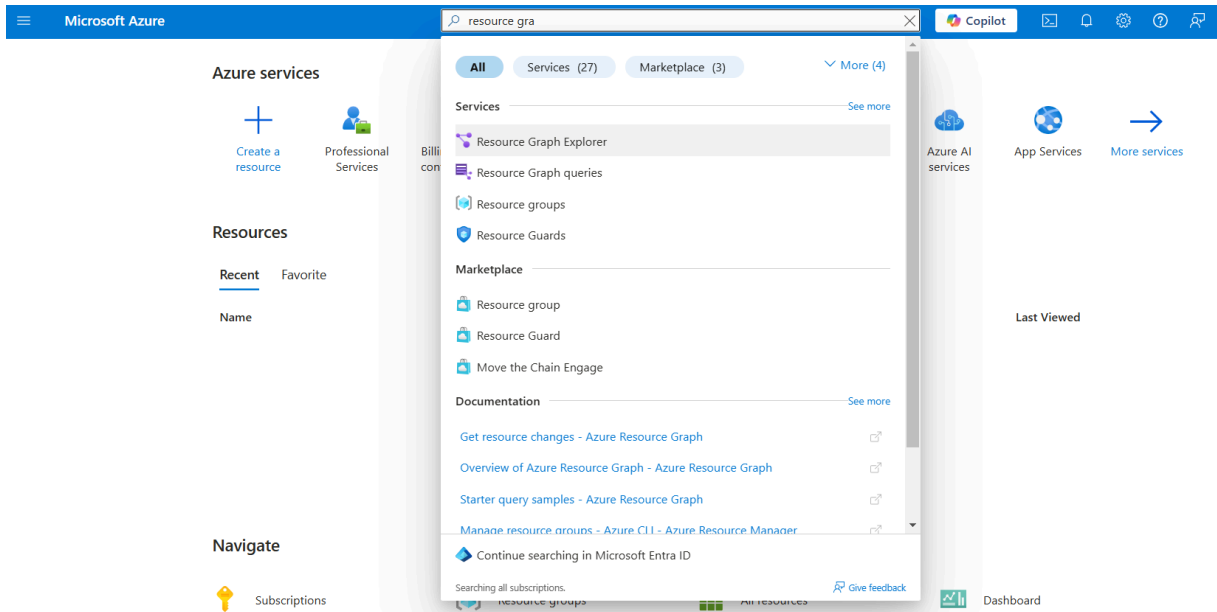
< 1 2 3 4 5 6 7 ... 20 > ⚙

<input type="checkbox"/>	Identifier <span style="font-size: 8px;">🔗</span>	Tag: Name	Service	Type	Region	Tags
<input type="checkbox"/>	<a href="#">vpc-00c8ce89bf3cce433</a>	(not tagged!)	EC2	VPC	ap-southeast-1	-
<input type="checkbox"/>	<a href="#">vpc-077e684bb6773b467</a>	(not tagged!)	EC2	VPC	ap-northeast-3	-

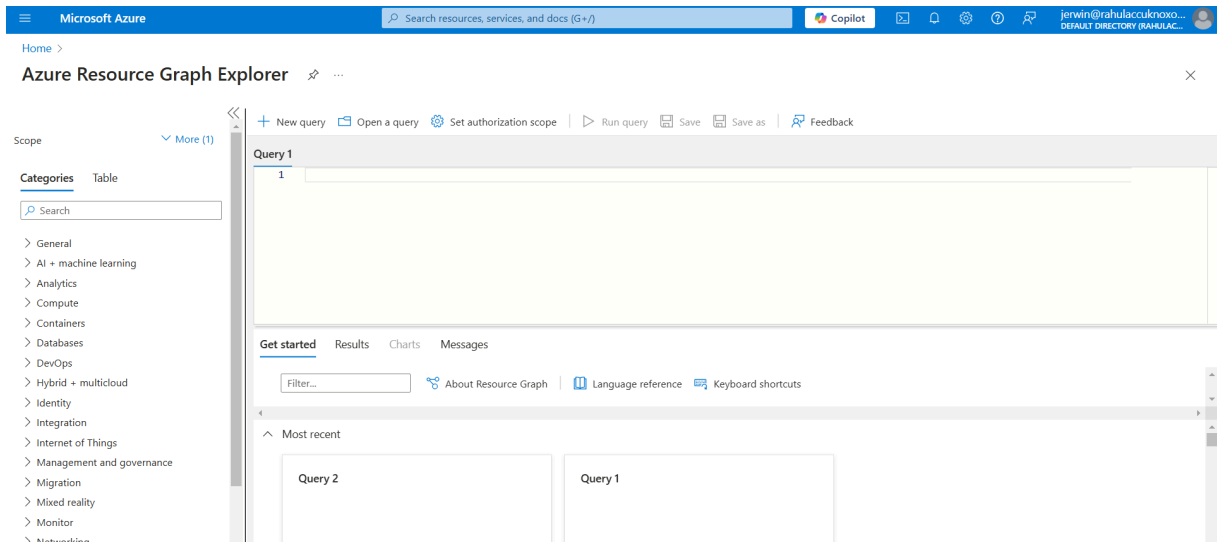
**Step 3:** Click on Export to CSV.

## 1.1.2 Azure

### Step 1: Search for Resource graph explorer



### Step 2: In the resource graph explorer screen, create a new query

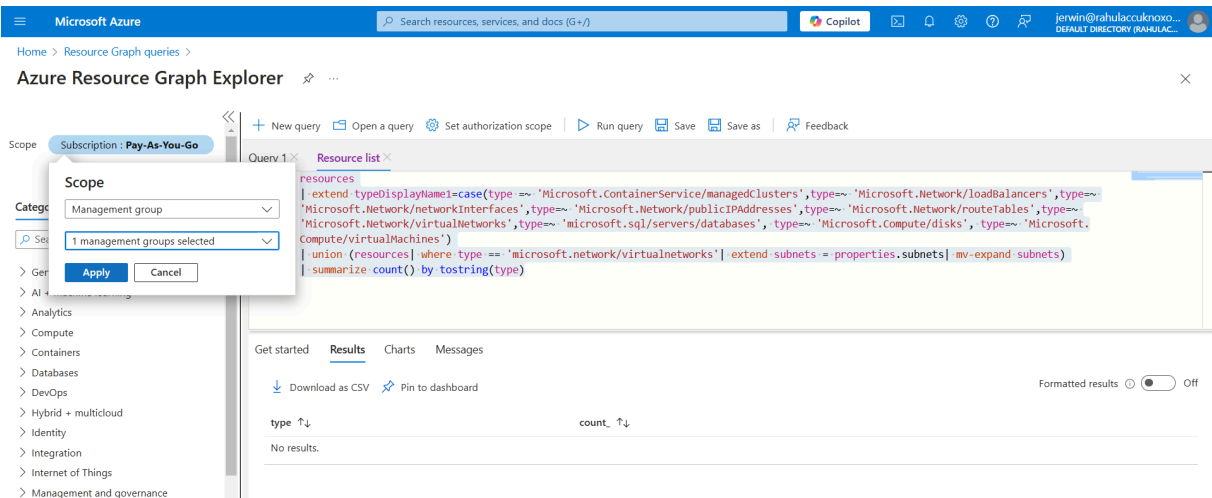


Add the following query to the resource graph explorer:

```
resources
| extend typeDisplayName1=case(type =~
```

```
'Microsoft.ContainerService/managedClusters', type=~
'Microsoft.Network/loadBalancers', type=~
'Microsoft.Network/networkInterfaces', type=~
'Microsoft.Network/publicIPAddresses', type=~
'Microsoft.Network/routeTables', type=~
'Microsoft.Network/virtualNetworks', type=~
'microsoft.sql/servers/databases', type=~ 'Microsoft.Compute/disks',
type=~ 'Microsoft.Compute/virtualMachines')
| union (resources | where type == 'microsoft.network/virtualnetworks' |
extend subnets = properties.subnets | mv-expand subnets)
| summarize count() by tostring(type)
```

**Step 3:** Click on **More** at the left side and set the scope for the query as required



The screenshot shows the Azure Resource Graph Explorer interface. A 'Scope' dialog box is open, showing 'Management group' selected. The background shows a query editor with the same query as in Step 1, and a results table that is currently empty.

**Step 4:** Click on run query to view the number of assets by their type

Microsoft Azure | Search resources, services, and docs (G+)

Home > All resources > Azure Resource Graph Explorer

Scope: More (1)

Categories: Table

Search: [ ]

- > General
- > AI + machine learning
- > Analytics
- > Compute
- > Containers
- > Databases
- > DevOps
- > Hybrid + multicloud
- > Identity
- > Integration
- > Internet of Things
- > Management and governance
- > Migration
- > Mixed reality
- > Monitor

Query 1 x Resource list x Query 2 x

```

1 resources
2 | extend typeDisplayName1=case(type =~ 'Microsoft.ContainerService/managedClusters',type~='Microsoft.Network/loadBalancers',type~='Microsoft.Network/networkInterfaces',type~='Microsoft.Network/publicIPAddresses',type~='Microsoft.Network/routeTables',type~='Microsoft.Network/virtualNetworks',type == 'microsoft.network/virtualnetworks',type~='microsoft.keyvault/vaults',type~='microsoft.sql/servers/databases',type~='microsoft.cache/redis',type~='microsoft.network/applicationsecuritygroups',type~='microsoft.network/networksecuritygroups',type~='microsoft.insights/activitylogalerts')
3 | union (authorizationresources | extend typeDisplayName2=case(type='microsoft.authorization/roledefinitions',type~='microsoft.authorization/roleassignment',type~='microsoft.authorization/roleassignment'))
4 | summarize count() by tostring(type)

```

Get started | Results | Charts | Messages

Download formatted results as CSV | Pin to dashboard | Formatted results: On

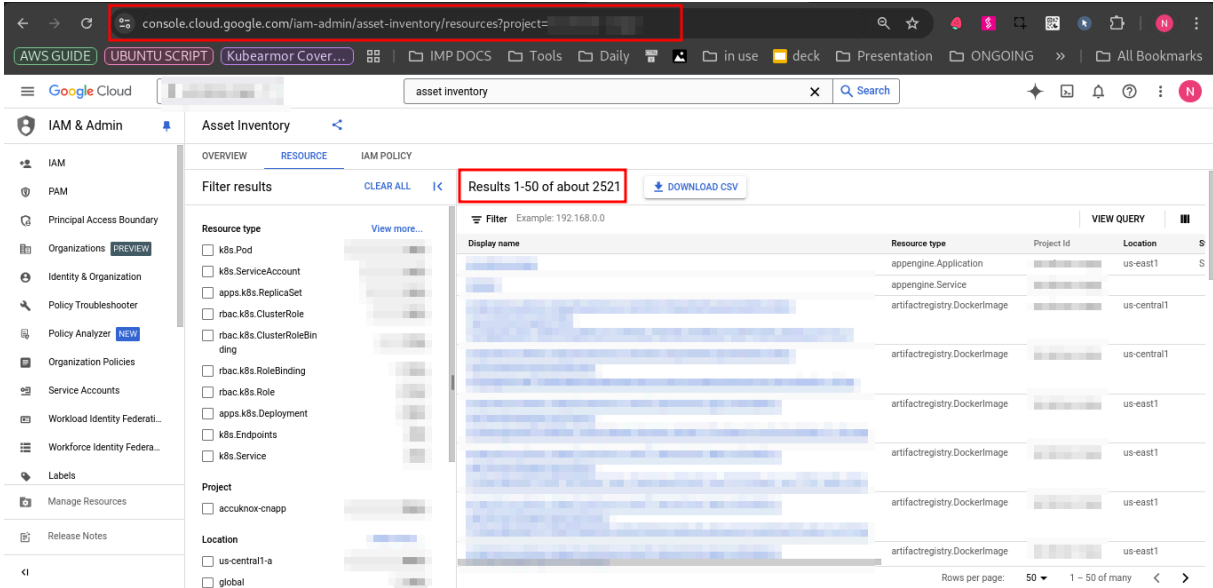
Type	count	
microsoft.authorization/classicadministrators	1	<a href="#">See details</a>
microsoft.authorization/roleassignments	19	<a href="#">See details</a>

< Previous | Page 1 of 1 | Next >

The results can be downloaded as CSV

## 1.1.3 GCP

**Step 1:** Navigate to the GCP Asset Inventory (IAM & Admin > Asset Inventory)



The screenshot shows the GCP Asset Inventory console. The browser address bar contains the URL: `console.cloud.google.com/iam-admin/asset-inventory/resources?project=`. The console interface shows the 'Asset Inventory' page with a 'Filter results' section. The filter results are highlighted in red, showing 'Results 1-50 of about 2521'. The table below shows a list of resources with columns for 'Display name', 'Resource type', 'Project Id', and 'Location'.

Display name	Resource type	Project Id	Location
appengine.Application	appengine.Application		us-east1
appengine.Service	appengine.Service		us-east1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-central1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-central1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-east1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-east1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-east1
artifactregistry.DockerImage	artifactregistry.DockerImage		us-east1

**Step 2:** Paste this snippet at the end of the url on the Browser to apply the required filters:

```
&pageState=("assetTypeFacets":("f":"compute.googleapis.com%2FSubnetwork,compute.googleapis.com%2FRoute,compute.googleapis.com%2FFirewall,logging.googleapis.com%2FLogBucket,serviceusage.googleapis.com%2FService,compute.googleapis.com%2FDisk,compute.googleapis.com%2FInstance,container.googleapis.com%2FCluster,storage.googleapis.com%2FBucket,compute.googleapis.com%2FNetwork"))
```

**Note:** Sometimes the page might need reload to reflect the filter changes.

The results can be downloaded as CSV.

## 1.2 Container Images Counts

### 1.2.1 DockerHub

To get the count of dockerhub images please use the following command after connecting your dockerhub repository to the commandline using dockerdesktop application.

```
docker images <repository-name>
```

**Note:** Replace the <repository-name> with your repository name.

### 1.2.2 AWS ECR

To get the count of the ECR Repository images the users need to connect the AWS account using AWS CLI and use the following command for getting the image count in each repository

```
aws ecr describe-images --repository-name <repository-name> --query "length(imageDetails[])"
```

**Note:** Replace the <repository-name> with your repository name.

### 1.2.3 GCR

To get the count of images stored in the GCR registry using the gcloud command line tool use the following command

```
gcloud container images list-tags gcr.io/<PROJECT_ID>/<REPOSITORY_NAME>  
--format='get(digest)' | wc -l
```

**Note:** Replace the <PROJECT\_ID> with your Google Cloud project ID and <REPOSITORY\_NAME> with the name of the GCR repository you want to count images.

## 1.2.4 ACR

To get the count of images stored in an Azure Container Registry (ACR) using Azure CLI use the following command

```
az acr repository show-tags --name <ACR_NAME> --repository <REPOSITORY_NAME> --output json --query "length(@)"
```

Note: Replace <ACR\_NAME> with the name of your Azure Container Registry and <REPOSITORY\_NAME> with the name of the ACR repository you want to count images.




## 2. CSPM Prerequisites

### 2.1 AWS

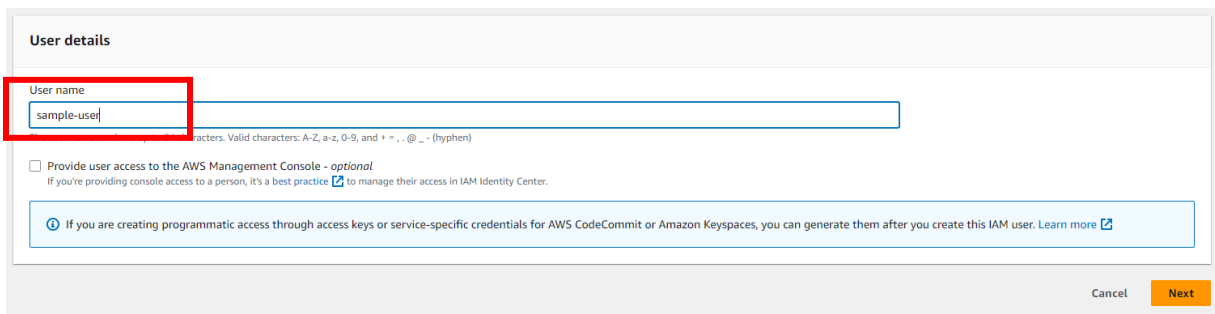
#### AWS IAM User Creation

Please follow the following steps to provide a user with appropriate read access:

**Step 1:** Navigate to IAM -> Users and click on Add Users



**Step 2:** Give a username to identify the user



**Step 3:** In the "Set Permissions" screen:

- a. Select "Attach policies directly"
- b. Search "ReadOnly", Filter by Type: "AWS managed - job function" and select the policy

Step 2  
Set permissions

Step 3  
Review and create

Permissions options

- Add user to group  
Add user to an existing group, or create a new group. We recommend using groups to manage user permissions by job function.
- Copy permissions  
Copy all group memberships, attached managed policies, and inline policies from an existing user.
- Attach policies directly  
Attach a managed policy directly to a user. As a best practice, we recommend attaching policies to a group instead. Then, add the user to the appropriate group.

Permissions policies (1/1116)  
Choose one or more policies to attach to your new user.

Filter by Type: AWS managed - job function (1 match)

Policy name	Type	Attached entities
<input checked="" type="checkbox"/> ReadOnlyAccess	AWS managed - job function	0

c. Search "SecurityAudit", Filter by Type: "AWS managed - job function" and select the policy

Permissions policies (2/1116)  
Choose one or more policies to attach to your new user.

Filter by Type: AWS managed - job function (1 match)

Search: security

Policy name	Type	Attached entities
<input checked="" type="checkbox"/> SecurityAudit	AWS managed - job function	0

▶ Set permissions boundary - optional

Cancel Previous Next

**Step 4:** Finish creating the user. Click on the newly created user and create the Access key and Secret Key from the Security Credentials tab to be used in the AccuKnox panel

Permissions | Groups | Tags | **Security credentials** | Access Advisor

Console sign-in Enable console access

Console sign-in link: <https://864316920010.signin.aws.amazon.com/console>

Console password: Not enabled

Multi-factor authentication (MFA) (0)  
Use MFA to increase the security of your AWS environment. Signing in with MFA requires an authentication code from an MFA device. Each user can have a maximum of 8 MFA devices assigned. [Learn more](#)

Remove Resync Assign MFA device

Device type	Identifier	Certifications	Created on
No MFA devices. Assign an MFA device to improve the security of your AWS environment.			

Assign MFA device

Access keys (0)  
Use access keys to send programmatic calls to AWS from the AWS CLI, AWS Tools for PowerShell, AWS SDKs, or direct AWS API calls. You can have a maximum of two access keys (active or inactive) at a time. [Learn more](#)

Create access key

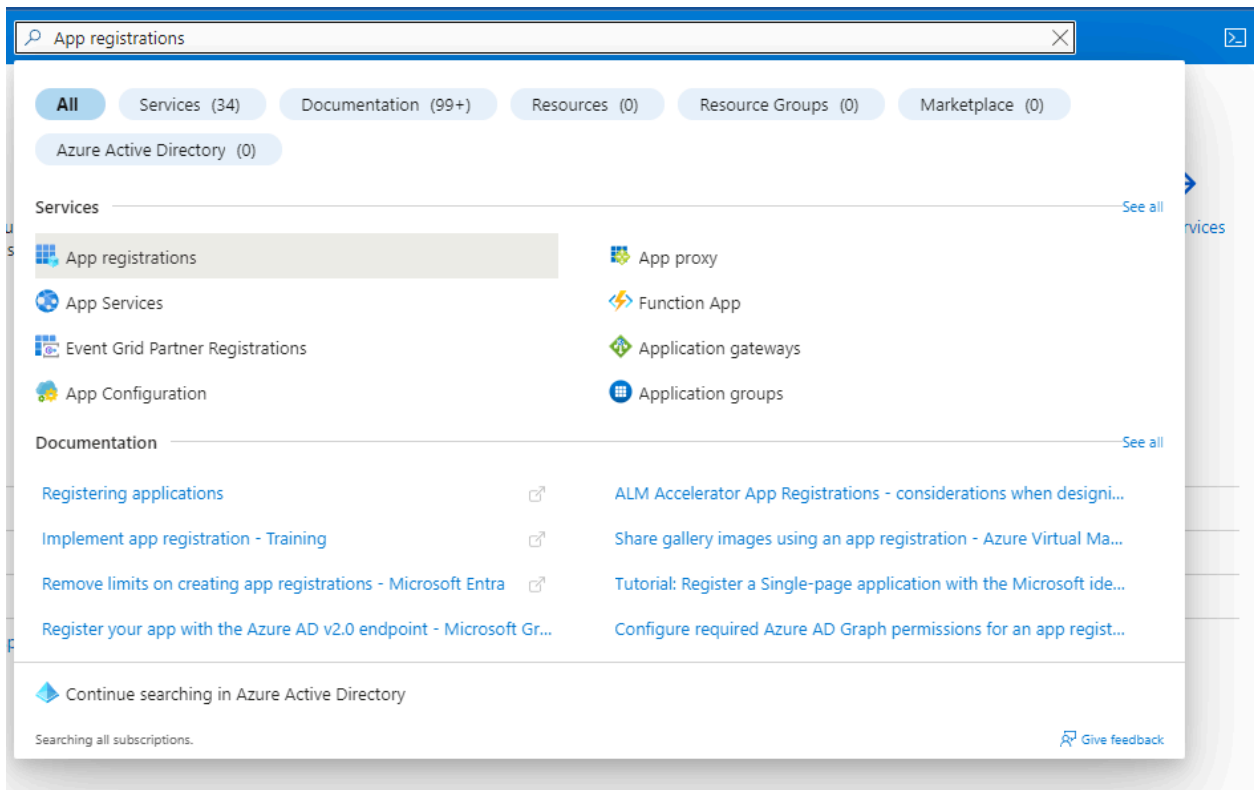
No access keys  
As a best practice, avoid using long-term credentials like access keys. Instead, use tools which provide short term credentials. [Learn more](#)

Create access key

## 2.2 AZURE

For Azure Onboarding it is required to register an App and giving Security read access to that App from the Azure portal.

- Go to your Azure Portal and search for App registrations and open it




- Here click on New registration

Home >

## App registrations


[+ New registration](#)
[🌐 Endpoints](#)
[🔧 Troubleshooting](#)
[🔄 Refresh](#)
[⬇️ Download](#)
[📺 Preview features](#)
[🗨️ Got feedback?](#)

 Starting June 30th, 2020 we will no longer add any new features to Azure Active Directory Authentication Library (ADAL) and Azure AD Graph. We will continue to support existing applications and Microsoft Graph. [Learn more](#)

[All applications](#)
[Owned applications](#)
[Deleted applications](#)

[Add filters](#)

7 applications found

Display name 

- Give your application a name, remember this name as it will be used again later, For the rest keep the default settings

Home > App registrations >

## Register an application

### \* Name

The user-facing display name for this application (this can be changed later).



### Supported account types

Who can use this application or access this API?

- Accounts in this organizational directory only (Default Directory only - Single tenant)
- Accounts in any organizational directory (Any Azure AD directory - Multitenant)
- Accounts in any organizational directory (Any Azure AD directory - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)
- Personal Microsoft accounts only

[Help me choose...](#)

### Redirect URI (optional)


We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

Register an app you're working on here. Integrate gallery apps and other apps from outside your organization by adding from [Enterprise applications](#).

By proceeding, you agree to the [Microsoft Platform Policies](#)

[Register](#)

- Now your application is created, save Application ID and Directory ID as they will be needed to for onboarding on Accuknox SaaS and then click on 'Add a certificate or secret'

3  ...

[Delete](#) [Endpoints](#) [Preview features](#)

---

^ Essentials

Display name : [Accuknox-may-2023](#) Client credentials : [Add a certificate or secret](#)  
 Application (client) ID : 0aaaf206-7336- Redirect URIs : [Add a Redirect URI](#)  
 Object ID : e3dcd617-e4b3- Application ID URI : [Add an Application ID URI](#)  
 Directory (tenant) ID : 57650de0-d901- Managed application in I... : [Accuknox-may-2023](#)  
 Supported account types : [My organization only](#)



**i** Welcome to the new and improved App registrations. Looking to learn how it's changed from App registrations (Legacy)? [Learn more](#)


**i** Starting June 30th, 2020 we will no longer add any new features to Azure Active Directory Authentication Library (ADAL) and Azure AD Graph. We will continue to provide technical support and security updates but we will no longer provide feature updates. Applications will be upgraded to Microsoft Authentication Library (MSAL) and Microsoft Graph. [Learn more](#)

[Get Started](#) [Documentation](#)

- Click on a new client secret and enter the name and expiration date to get secret id and secret value, save this secret value as this will also be needed for onboarding.

Home > App registrations > Accuknox-may-2023

 **Accuknox-may-2023** | Certificates & secrets  ...

Search  Got feedback?

Overview **i** Got a second to give us some feedback? →

Quickstart

Integration assistant

Manage

Branding & properties

Authentication

**+** Certificates & secrets

Token configuration

API permissions

Expose an API

App roles

Owners

Roles and administrators

Manifest

Support + Troubleshooting

Troubleshooting

New support request

Credentials enable confidential applications to identify themselves to the authentication service when receiving tokens at a web addressable location (using an HTTPS scheme). For a higher level of assurance, we recommend using a certificate (instead of a client secret) as a credential.

**i** Application registration certificates, secrets and federated credentials can be found in the tabs below.

Certificates (0) **Client secrets (1)** Federated credentials (0)

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

+ New client secret

Description	Expires	Value	secret ID
may-2023	5/8/2025	zXc9Q-	72e13faf-

**Copied**

- Next, go to the API permissions tab and click on Add a permission

Home > App registrations > Permission-screen

Permission-screen | API permissions

Search Refresh Got feedback?

Overview  
Quickstart  
Integration assistant

Manage  
Branding & properties  
Authentication  
Certificates & secrets  
Token configuration  
**API permissions**  
Expose an API  
App roles  
Owners  
Roles and administrators  
Manifest

The "Admin consent required" column shows the default value for an organization. However, user consent can be customized per permission, user, or app. This column may not reflect organizations where this app will be used. [Learn more](#)

Configured permissions

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. [Learn more about permissions and consent](#)

+ Add a permission ✓ Grant admin consent for Default Directory

API / Permissions name	Type	Description	Admin consent requ...	Status
Microsoft Graph (1)				
User.Read	Delegated	Sign in and read user profile	No	...

To view and manage consented permissions for individual apps, as well as your tenant's consent settings, try [Enterprise applications](#).

- On the screen that appears, click on Microsoft Graph

Home > App registrations > Permission-screen

Permission-screen | API permissions

Search Refresh Got feedback?

Overview  
Quickstart  
Integration assistant

Manage  
Branding & properties  
Authentication  
Certificates & secrets  
Token configuration  
**API permissions**  
Expose an API  
App roles  
Owners  
Roles and administrators

The "Admin consent required" column shows the default value for organizations where this app will be used. [Learn more](#)

Configured permissions

Applications are authorized to call APIs when they are granted permissions by users/admins as part of the consent process. The list of configured permissions should include all the permissions the application needs. [Learn more about permissions and consent](#)

+ Add a permission ✓ Grant admin consent for Default Directory

API / Permissions name	Type	Description	Admin consent requ...	Status
Microsoft Graph (1)				
User.Read	Delegated	Sign in and read user profile	No	...

To view and manage consented permissions for individual apps, as well as your tenant's consent settings, try [Enterprise applications](#).

**Request API permissions**

Select an API

Microsoft APIs APIs my organization uses My APIs

Commonly used Microsoft APIs

**Microsoft Graph**

Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Microsoft Entra ID, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.

**Azure DevOps**

Integrate with Azure DevOps and Azure DevOps server

**Azure Service Management**

Programmatic access to much of the functionality available through the Azure portal

**Azure Storage**

Secure, massively scalable object and data lake storage for unstructured and semi-structured data

**Dynamics CRM**

Access the capabilities of CRM business

**Intune**

Programmatic access to Intune data


**Office 365 Management APIs**

Retrieve information about user, admin,

- Next, select Application Permissions and then search for Directory.Read.All and click on Add permissions

## Request API permissions ×

[← All APIs](#)

 Microsoft Graph  
<https://graph.microsoft.com/> [Docs](#) [↗](#)

What type of permissions does your application require?

Delegated permissions

Your application needs to access the API as the signed-in user.

Application permissions

Your application runs as a background service or daemon without a signed-in user.

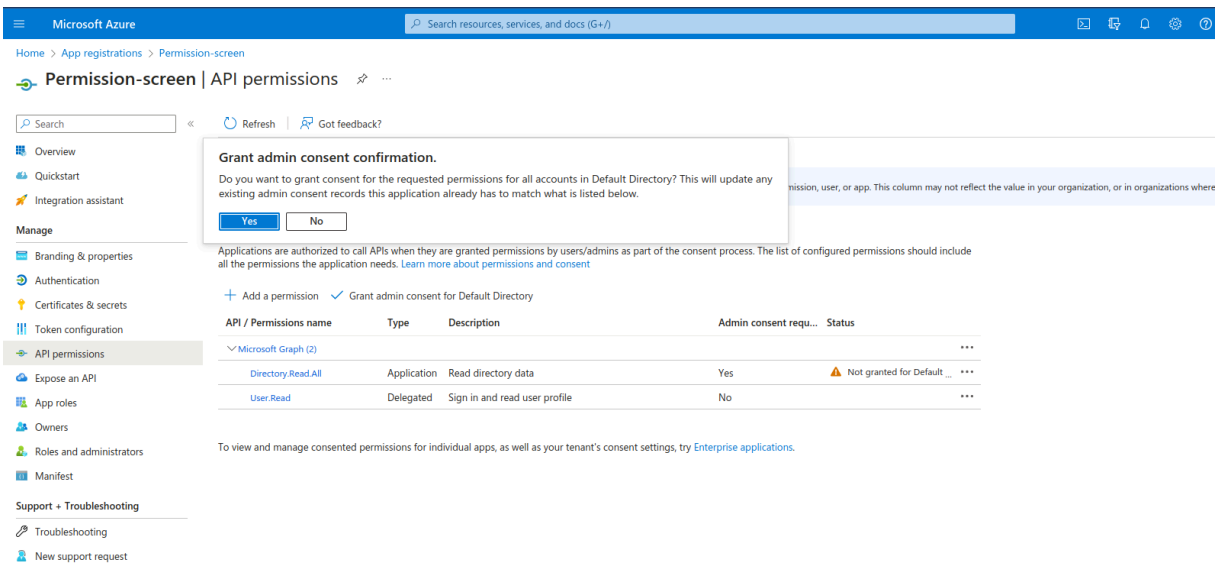
Select permissions [expand all](#)

×

Permission	Admin consent required
<div style="margin-left: 10px;"> <span>▼</span> Directory (1)                 </div> <div style="margin-left: 10px; background-color: #f0f0f0; padding: 5px;"> <input checked="" type="checkbox"/> Directory.Read.All <span>ⓘ</span>                      Read directory data                 </div>	Yes

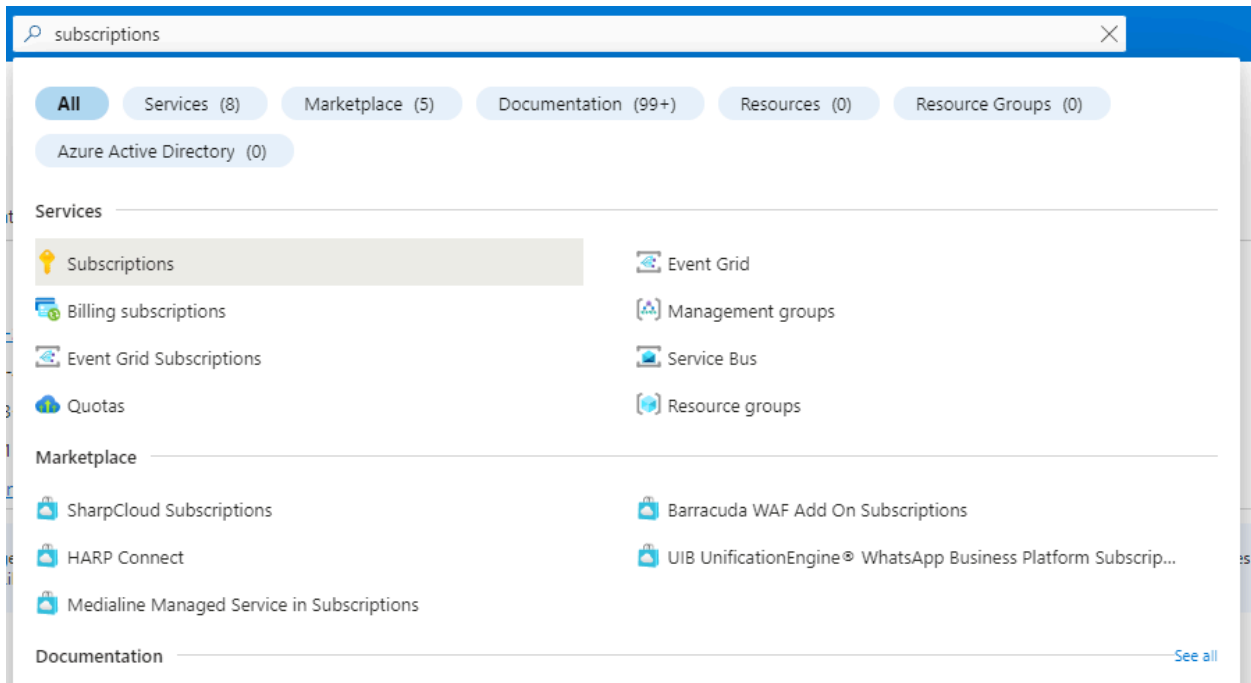
Add permissions
Discard

- Select 'Grant Admin Consent' for Default Directory and click on 'Yes'

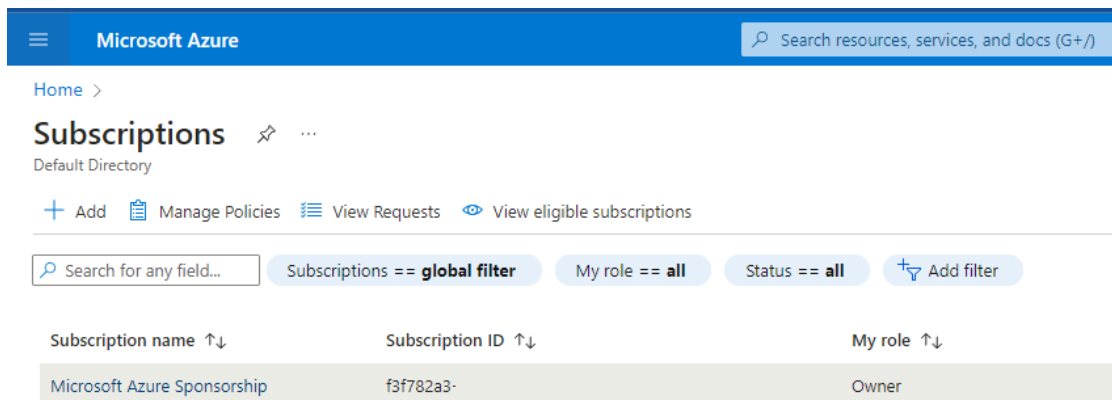


The screenshot shows the Microsoft Azure portal interface. At the top, there's a navigation bar with 'Microsoft Azure' and a search bar. Below it, the breadcrumb trail reads 'Home > App registrations > Permission-screen'. The main heading is 'Permission-screen | API permissions'. A left-hand navigation pane lists various management options like 'Overview', 'Quickstart', 'Integration assistant', 'Manage', 'Branding & properties', 'Authentication', 'Certificates & secrets', 'Token configuration', 'API permissions' (which is selected), 'Expose an API', 'App roles', 'Owners', 'Roles and administrators', 'Manifest', 'Support + Troubleshooting', 'Troubleshooting', and 'New support request'. The main content area shows a table of API permissions for 'Microsoft Graph (2)'. The table has columns for 'API / Permissions name', 'Type', 'Description', 'Admin consent required', and 'Status'. One permission, 'Directory.Read.All', is highlighted in blue and has a status of 'Not granted for Default Directory'. A modal dialog box titled 'Grant admin consent confirmation.' is overlaid on the screen, asking 'Do you want to grant consent for the requested permissions for all accounts in Default Directory? This will update any existing admin consent records this application already has to match what is listed below.' The dialog has 'Yes' and 'No' buttons.

- Now we need to give Security read permissions to this registered Application , to do that go to subscriptions

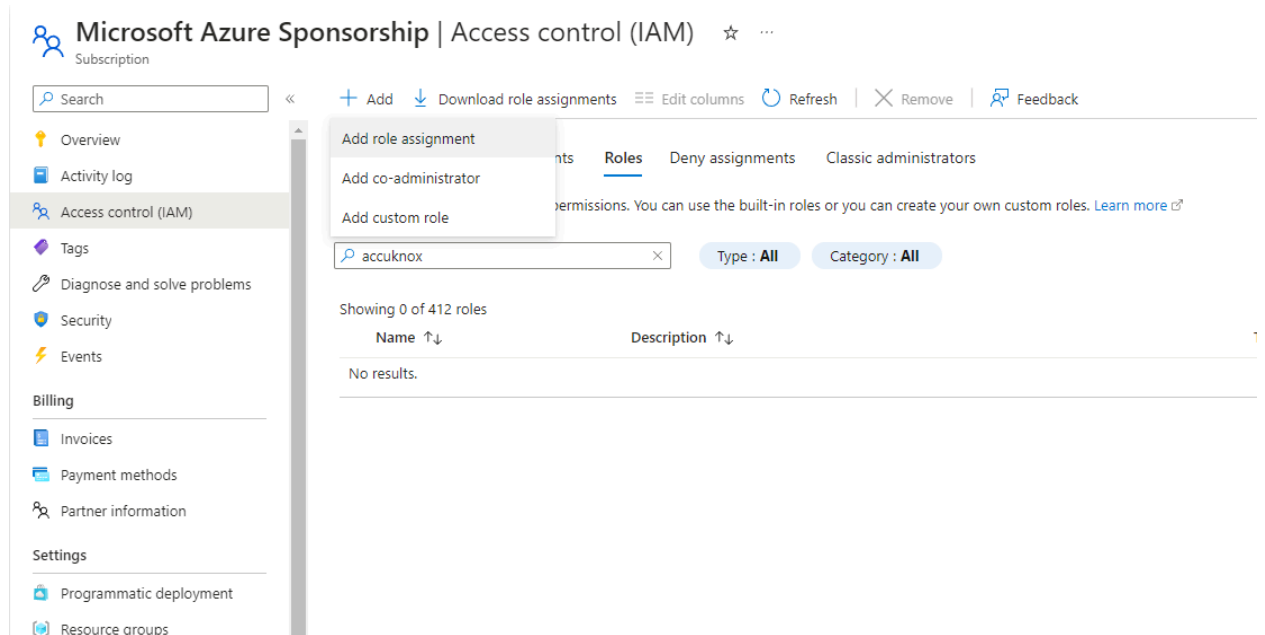


- First save the subscription ID and click on the subscription name , here it is “Microsoft Azure Sponsorship“



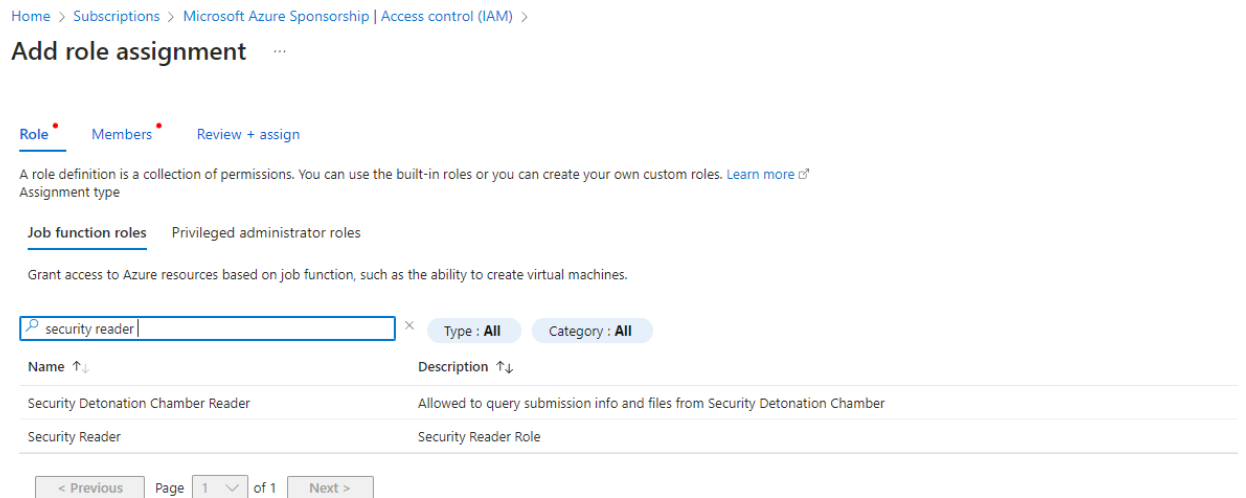


- Navigate to Access control(IAM) and go to Roles , here select Add and Add role assignment



The screenshot shows the 'Add role assignment' dialog in the Microsoft Azure portal. The search bar contains 'accuknox'. The 'Type' and 'Category' filters are both set to 'All'. The results table is empty, showing 'Showing 0 of 412 roles' and 'No results.' The left sidebar shows the navigation menu with 'Access control (IAM)' selected.

- Search for “Security Reader” Job function Role, select it and press next



The screenshot shows the search results for 'security reader' in the 'Add role assignment' dialog. The search bar contains 'security reader'. The 'Type' and 'Category' filters are both set to 'All'. The results table shows two roles:

Name ↑↓	Description ↑↓
Security Detonation Chamber Reader	Allowed to query submission info and files from Security Detonation Chamber
Security Reader	Security Reader Role

At the bottom, there is a pagination control showing 'Page 1 of 1'.

- In the member section click on Select members it will open a dropdown menu on the right hand side

## Add role assignment

Role **Members** Review + assign

**Selected role** Security Reader

**Assign access to**  User, group, or service principal  
 Managed identity

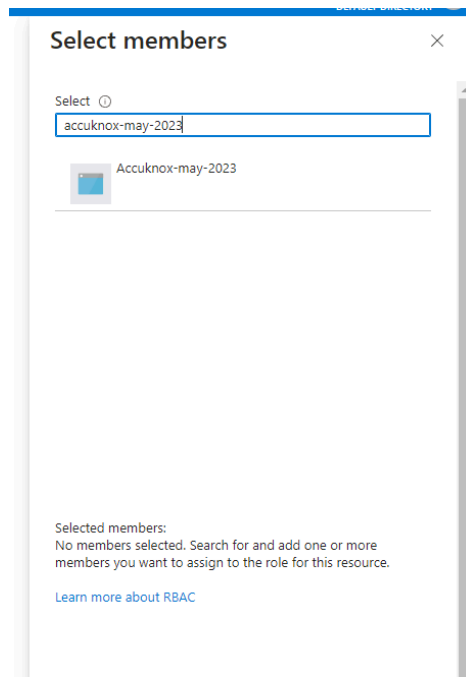
**Members** [+ Select members](#)

Name	Object ID
------	-----------

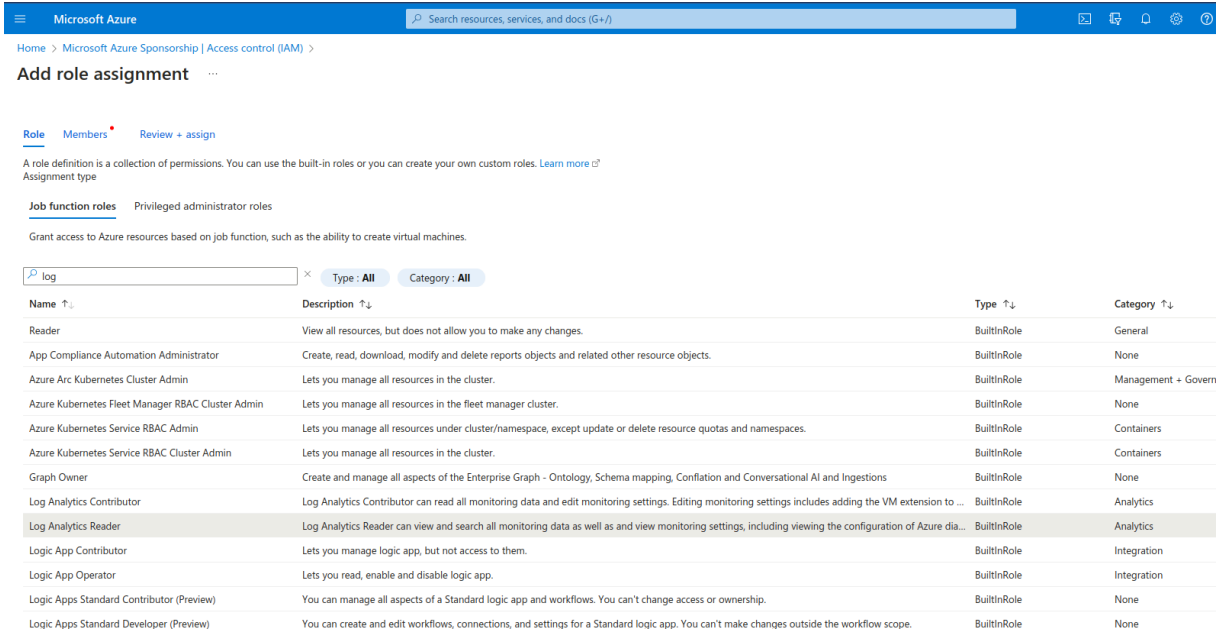
No members selected

**Description** Optional

- Here search for the Application that you registered in the beginning , select the application and click on review and assign.



- Similarly, we have to add another role. This time, search for Log Analytics Reader. Select it and click next



Microsoft Azure | Search resources, services, and docs (G+/)

Home > Microsoft Azure Sponsorship | Access control (IAM) >

### Add role assignment

Role **Members** Review + assign

A role definition is a collection of permissions. You can use the built-in roles or you can create your own custom roles. [Learn more](#)

Assignment type

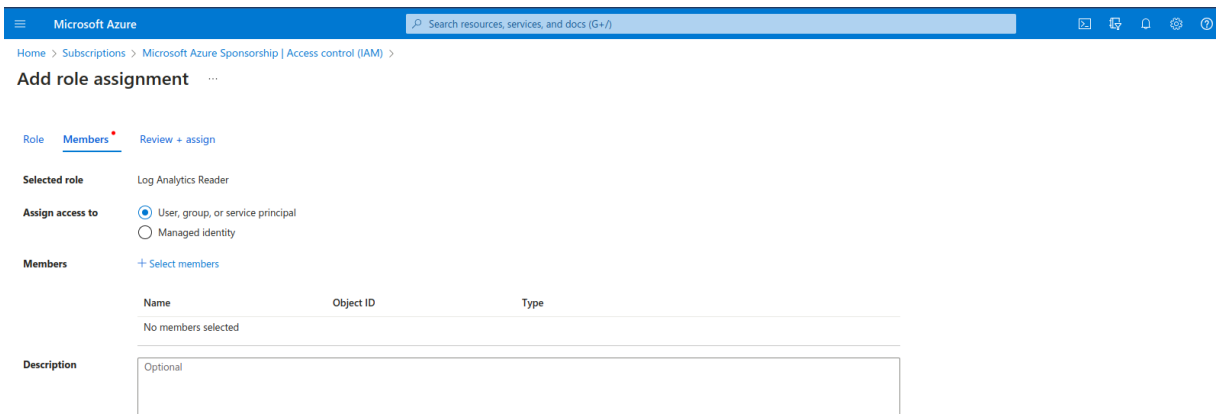
**Job function roles** Privileged administrator roles

Grant access to Azure resources based on job function, such as the ability to create virtual machines.

log × Type: All Category: All

Name ↑↓	Description ↑↓	Type ↑↓	Category ↑↓
Reader	View all resources, but does not allow you to make any changes.	BuiltInRole	General
App Compliance Automation Administrator	Create, read, download, modify and delete reports objects and related other resource objects.	BuiltInRole	None
Azure Arc Kubernetes Cluster Admin	Lets you manage all resources in the cluster.	BuiltInRole	Management + Govern
Azure Kubernetes Fleet Manager RBAC Cluster Admin	Lets you manage all resources in the fleet manager cluster.	BuiltInRole	None
Azure Kubernetes Service RBAC Admin	Lets you manage all resources under cluster/namespace, except update or delete resource quotas and namespaces.	BuiltInRole	Containers
Azure Kubernetes Service RBAC Cluster Admin	Lets you manage all resources in the cluster.	BuiltInRole	Containers
Graph Owner	Create and manage all aspects of the Enterprise Graph - Ontology, Schema mapping, Conflation and Conversational AI and Ingestions	BuiltInRole	None
Log Analytics Contributor	Log Analytics Contributor can read all monitoring data and edit monitoring settings. Editing monitoring settings includes adding the VM extension to ...	BuiltInRole	Analytics
<b>Log Analytics Reader</b>	<b>Log Analytics Reader can view and search all monitoring data as well as and view monitoring settings, including viewing the configuration of Azure dia...</b>	<b>BuiltInRole</b>	<b>Analytics</b>
Logic App Contributor	Lets you manage logic app, but not access to them.	BuiltInRole	Integration
Logic App Operator	Lets you read, enable and disable logic app.	BuiltInRole	Integration
Logic Apps Standard Contributor (Preview)	You can manage all aspects of a Standard logic app and workflows. You can't change access or ownership.	BuiltInRole	None
Logic Apps Standard Developer (Preview)	You can create and edit workflows, connections, and settings for a Standard logic app. You can't make changes outside the workflow scope.	BuiltInRole	None

- Now, click on select members, select the application that was created similar to the previous role. Finally, click on Review and Assign.



Microsoft Azure | Search resources, services, and docs (G+/)

Home > Subscriptions > Microsoft Azure Sponsorship | Access control (IAM) >

### Add role assignment

Role **Members** Review + assign

**Selected role** Log Analytics Reader

**Assign access to**  User, group, or service principal  Managed identity

**Members** + Select members

Name	Object ID	Type
No members selected		

**Description** Optional

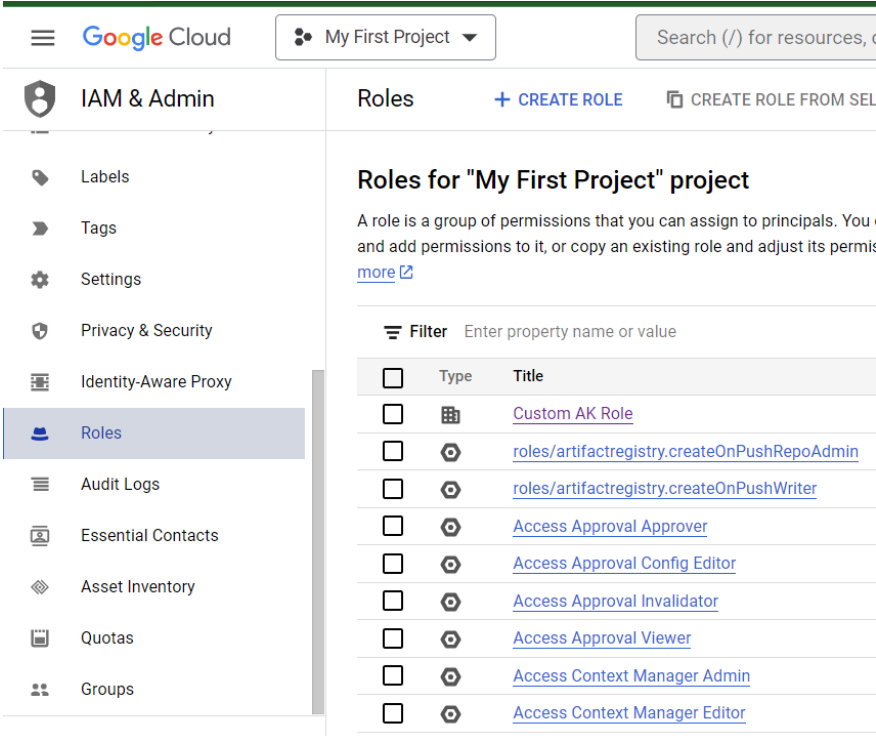
## 2.3 GCP

**Note:** Make sure the Below API Library is enabled in your GCP Account for onboarding into AccuKnox SaaS:

1. Compute Engine API
2. Identity and Access Management (IAM) API
3. Cloud Resource Manager API
4. Cloud Functions API
5. KMS API
6. Kubernetes API
7. Cloud SQL Admin API

For GCP there is a requirement for IAM Service Account Access.

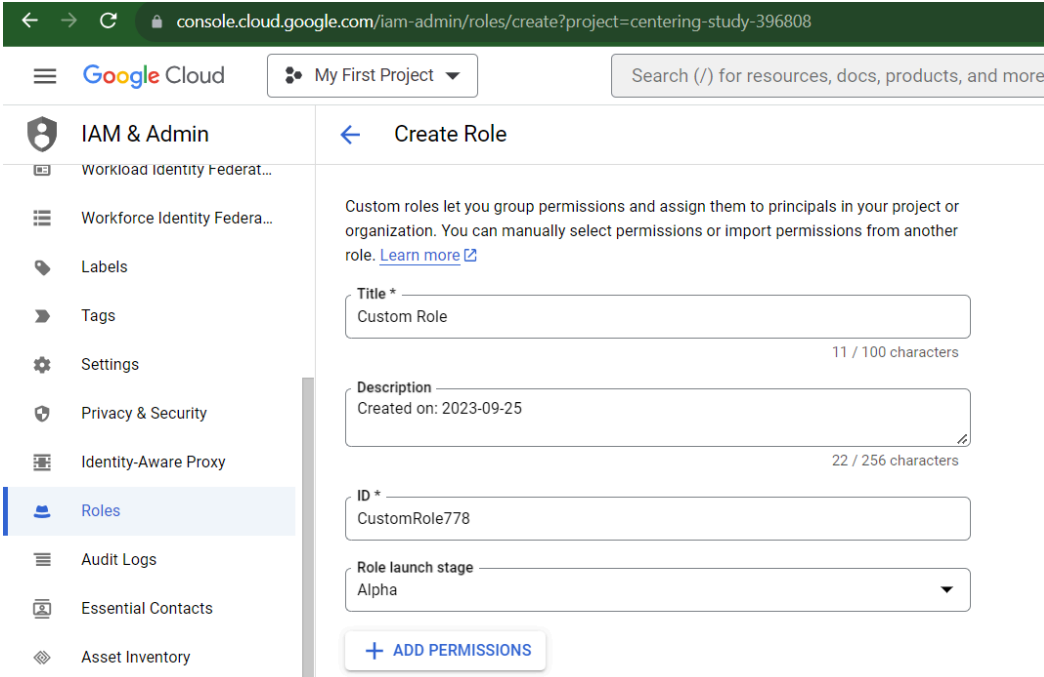
**Step 1:** Log into your Google Cloud console and navigate to IAM & Admin, choose “Roles” and Click “Create Role”



The screenshot shows the Google Cloud IAM & Admin console interface. The left sidebar contains navigation options: IAM & Admin, Labels, Tags, Settings, Privacy & Security, Identity-Aware Proxy, Roles (highlighted), Audit Logs, Essential Contacts, Asset Inventory, Quotas, and Groups. The main content area is titled 'Roles for "My First Project" project' and includes a '+ CREATE ROLE' button and a 'CREATE ROLE FROM SEL' button. Below this, there is a filter input field and a table listing various roles.

<input type="checkbox"/>	Type	Title
<input type="checkbox"/>	Custom	<a href="#">Custom AK Role</a>
<input type="checkbox"/>	System	<a href="#">roles/artifactregistry.createOnPushRepoAdmin</a>
<input type="checkbox"/>	System	<a href="#">roles/artifactregistry.createOnPushWriter</a>
<input type="checkbox"/>	System	<a href="#">Access Approval Approver</a>
<input type="checkbox"/>	System	<a href="#">Access Approval Config Editor</a>
<input type="checkbox"/>	System	<a href="#">Access Approval Invalidator</a>
<input type="checkbox"/>	System	<a href="#">Access Approval Viewer</a>
<input type="checkbox"/>	System	<a href="#">Access Context Manager Admin</a>
<input type="checkbox"/>	System	<a href="#">Access Context Manager Editor</a>

## Step 2: Name the “Role” and Click “Add Permission”



console.cloud.google.com/iam-admin/roles/create?project=centering-study-396808

Google Cloud My First Project Search (/) for resources, docs, products, and more

IAM & Admin

- Workload Identity Federat...
- Workforce Identity Federa...
- Labels
- Tags
- Settings
- Privacy & Security
- Identity-Aware Proxy
- Roles**
- Audit Logs
- Essential Contacts
- Asset Inventory

Create Role

Custom roles let you group permissions and assign them to principals in your project or organization. You can manually select permissions or import permissions from another role. [Learn more](#)

Title \*  
Custom Role 11 / 100 characters

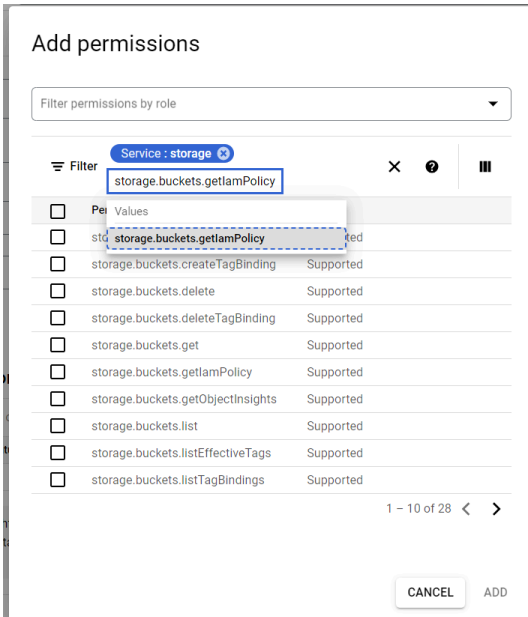
Description  
Created on: 2023-09-25 22 / 256 characters

ID \*  
CustomRole778

Role launch stage  
Alpha

+ ADD PERMISSIONS

## Step 3: Use the Service: storage filter, then value as “storage.buckets.getIamPolicy”



Add permissions

Filter permissions by role

Filter Service : storage

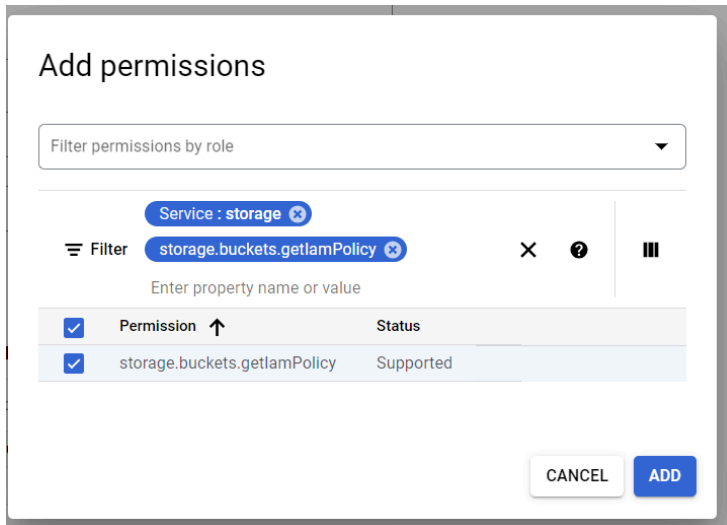
storage.buckets.getIamPolicy

Filter	Values	Supported
<input type="checkbox"/>	storage.buckets.getIamPolicy	Supported
<input type="checkbox"/>	storage.buckets.createTagBinding	Supported
<input type="checkbox"/>	storage.buckets.delete	Supported
<input type="checkbox"/>	storage.buckets.deleteTagBinding	Supported
<input type="checkbox"/>	storage.buckets.get	Supported
<input type="checkbox"/>	storage.buckets.getIamPolicy	Supported
<input type="checkbox"/>	storage.buckets.getObjectInsights	Supported
<input type="checkbox"/>	storage.buckets.list	Supported
<input type="checkbox"/>	storage.buckets.listEffectiveTags	Supported
<input type="checkbox"/>	storage.buckets.listTagBindings	Supported

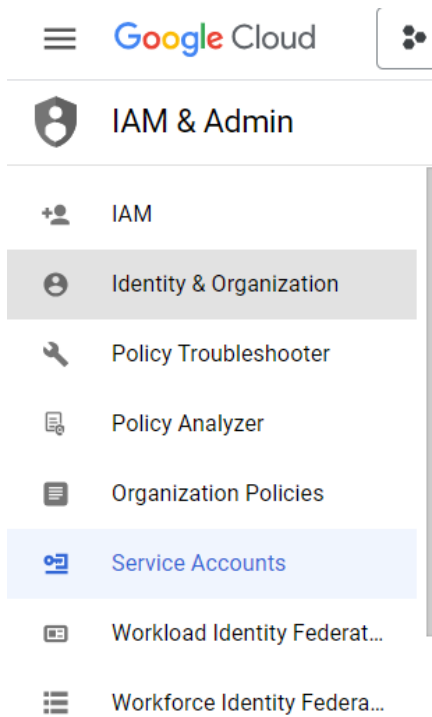
1 - 10 of 28 < >

CANCEL ADD

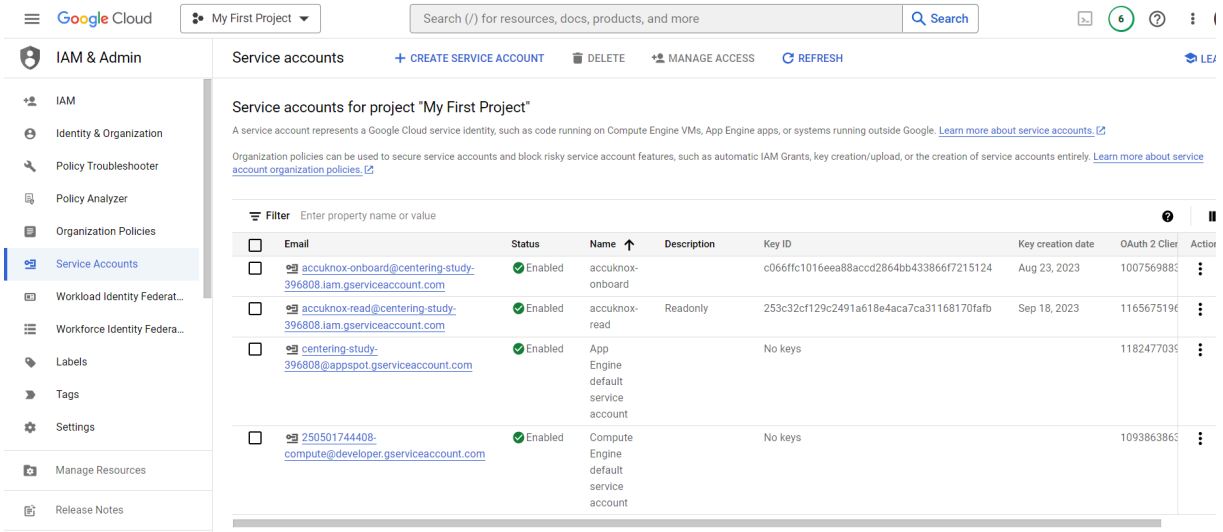
**Step 4:** Choose the permission and Click “Add” then Click Create in the same page.



**Step 5:** In the Navigation Panel, navigate to IAM Admin > Service Accounts.



## Step 6: Click on "Create Service Account"



The screenshot shows the Google Cloud IAM & Admin console. The left sidebar is expanded to 'Service Accounts'. The main content area displays 'Service accounts for project "My First Project"'. Below this, there is a table of existing service accounts:

Email	Status	Name	Description	Key ID	Key creation date	OAuth 2 Client	Action
<a href="mailto:accuknox-onboard@centering-study-396808.iam.gserviceaccount.com">accuknox-onboard@centering-study-396808.iam.gserviceaccount.com</a>	Enabled	accuknox-onboard		c066ffc1016eea88acc2864bb433866f7215124	Aug 23, 2023	1007569883	⋮
<a href="mailto:accuknox-read@centering-study-396808.iam.gserviceaccount.com">accuknox-read@centering-study-396808.iam.gserviceaccount.com</a>	Enabled	accuknox-read	Readonly	253c32cf129c2491a618e4aca7ca31168170fab	Sep 18, 2023	116567519c	⋮
<a href="mailto:centering-study-396808@appspot.gserviceaccount.com">centering-study-396808@appspot.gserviceaccount.com</a>	Enabled	App Engine default service account		No keys		1182477035	⋮
<a href="mailto:250501744408-compute@developer.gserviceaccount.com">250501744408-compute@developer.gserviceaccount.com</a>	Enabled	Compute Engine default service account		No keys		1093863865	⋮

**Step 7:** Enter any name that you want on Service Account Name.

**Step 8:** Click on Continue.

### 1 Service account details

Service account name

Display name for this service account

Service account ID \*  ✕ ↻

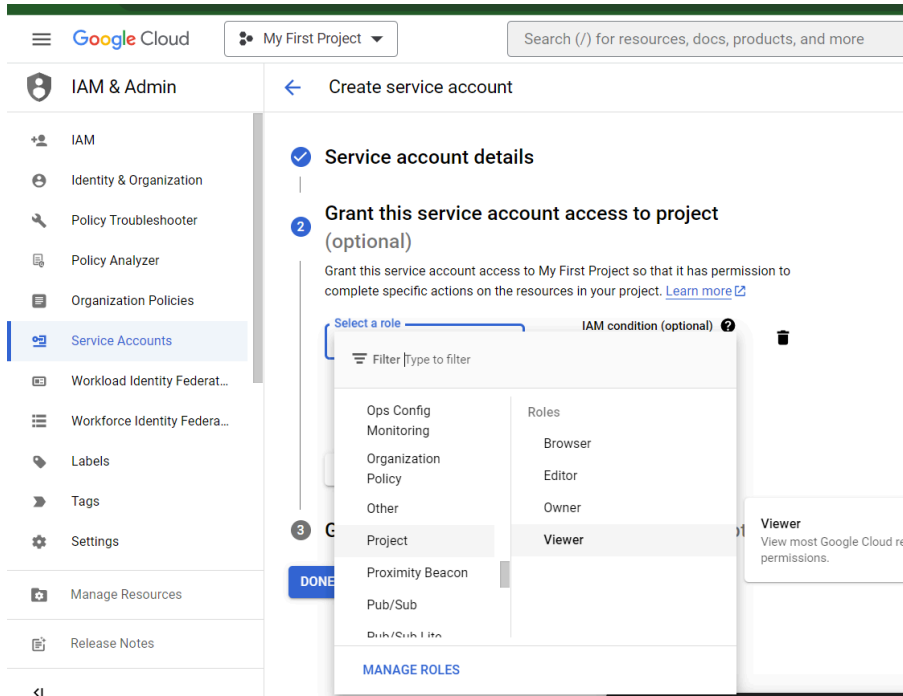
Email address: [ak-test@centering-study-396808.iam.gserviceaccount.com](mailto:ak-test@centering-study-396808.iam.gserviceaccount.com) 📄

Service account description

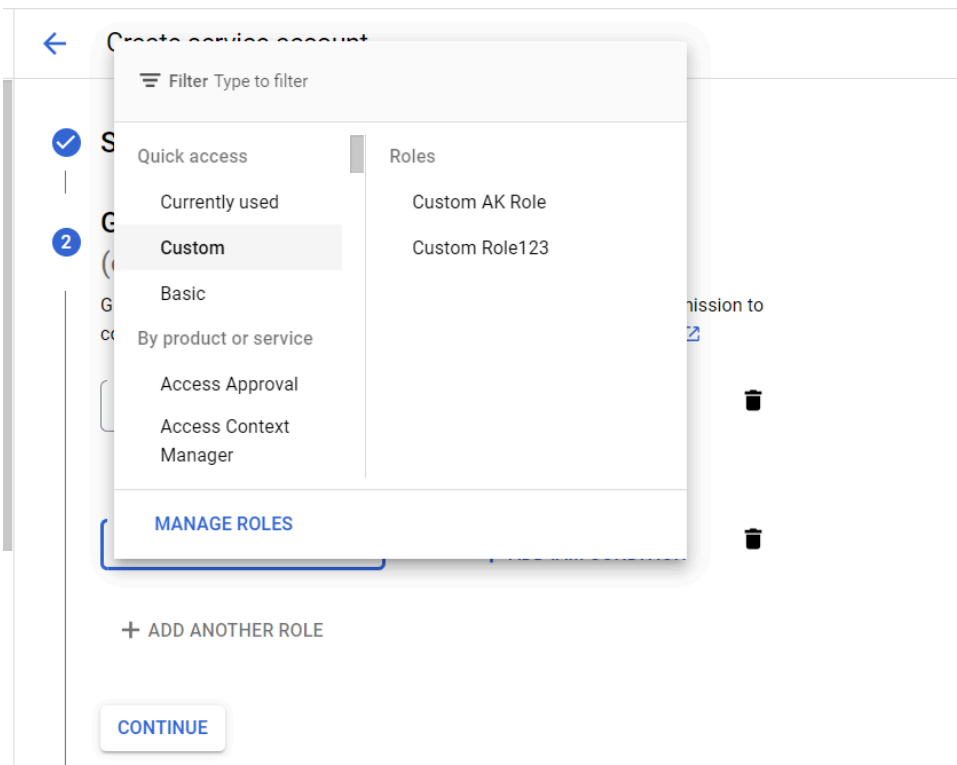
Describe what this service account will do

[CREATE AND CONTINUE](#)

**Step 9:** Select the role: Project > Viewer and click Add another Role.



**Step 10:** Click “Add Another Role” Choose “Custom” Select the created Custom Role.







## Step 11: Click on “Continue” and ”Done”

### ✓ Service account details

#### 2 Grant this service account access to project (optional)

Grant this service account access to My First Project so that it has permission to complete specific actions on the resources in your project. [Learn more](#)


Role Viewer	IAM condition (optional) ? <a href="#">+ ADD IAM CONDITION</a>	
View most Google Cloud resources. See the list of included permissions.		
Role Custom Role123	IAM condition (optional) ? <a href="#">+ ADD IAM CONDITION</a>	
Created on: 2023-09-25		










[+ ADD ANOTHER ROLE](#)

[CONTINUE](#)

#### 3 Grant users access to this service account (optional)

## Step 12: Go to the created Service Account, click on that Service Account navigate to the “Keys” section.


 IAM & Admin

-  IAM
-  Identity & Organization
-  Policy Troubleshooter
-  Policy Analyzer
-  Organization Policies
-  Service Accounts
-  Workload Identity Federat...
-  Workforce Identity Federa...
-  Labels

← AK-test

DETAILS
PERMISSIONS
 KEYS
 METRICS
LOGS

### Keys

 Service account keys could pose a security risk if compromised. We recommend you read [about the best way to authenticate service accounts on Google Cloud](#).

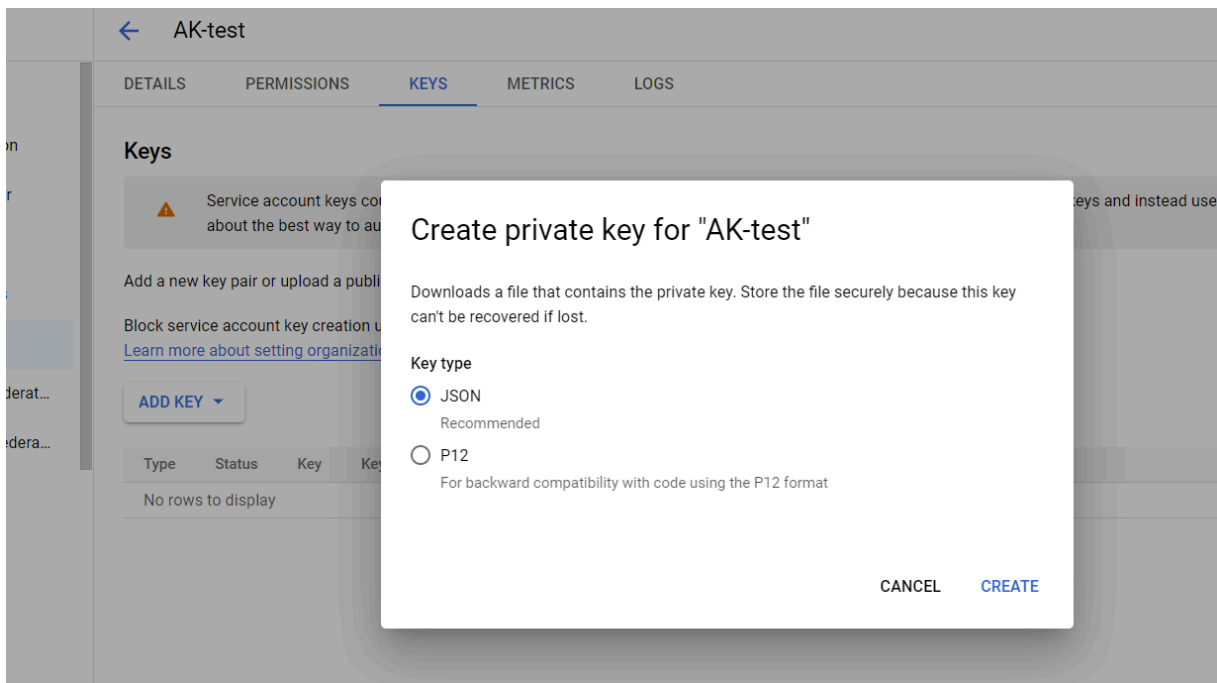
Add a new key pair or upload a public key certificate from an existing key pair.

Block service account key creation using [organization policies](#).  
[Learn more about setting organization policies for service accounts](#)

[ADD KEY](#)

Type	Status	Key	Key creation date	Key expiration date
No rows to display				

**Step 13:** Click the “Add key“ button and “Create new key “ . Chosen Key type should be JSON format.



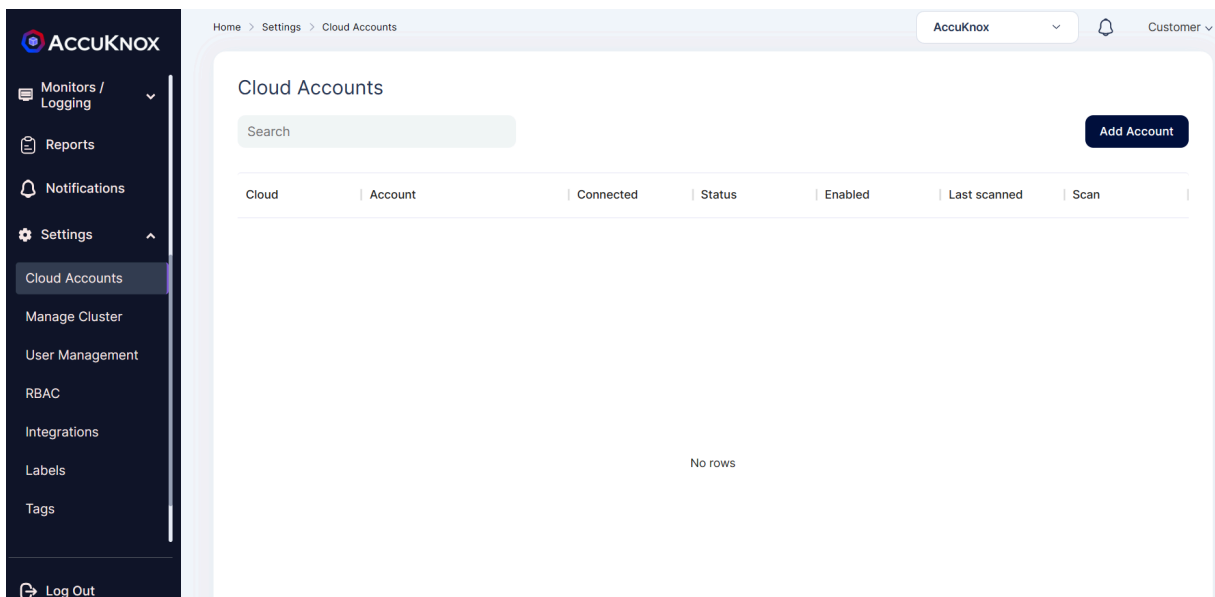
**Step 12:** Click the “Create“ button it will automatically download the JSON key.

## 3. Cloud Onboarding

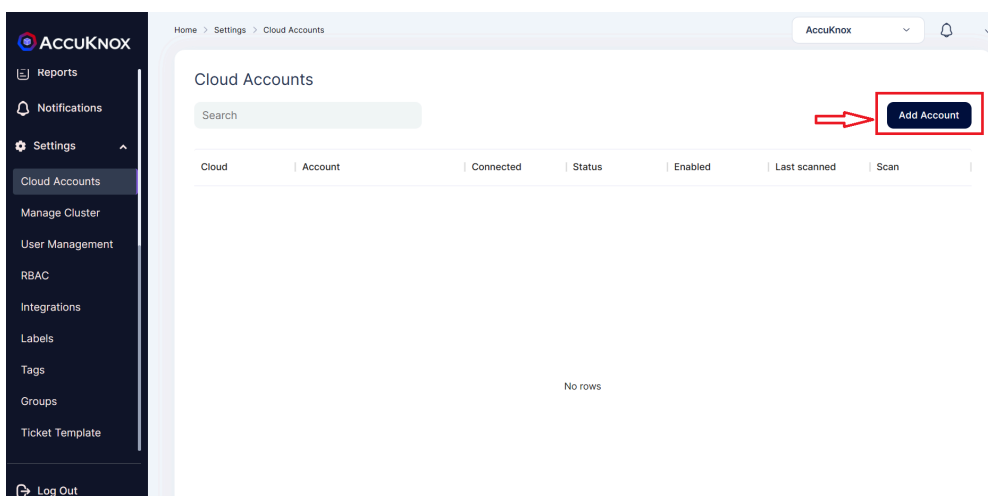
### 3.1 AWS Onboarding

In this example we are onboarding an AWS account using the Access Keys method.

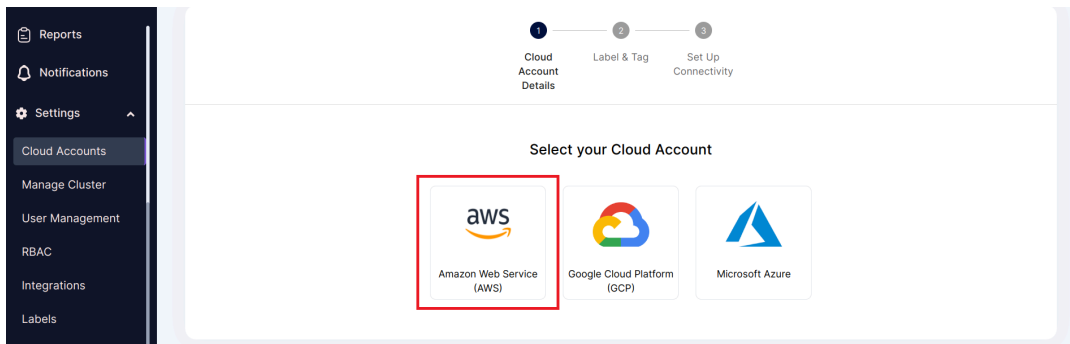
**Step 1:** To onboard Cloud Account Navigate to *Settings*→*cloud Accounts*



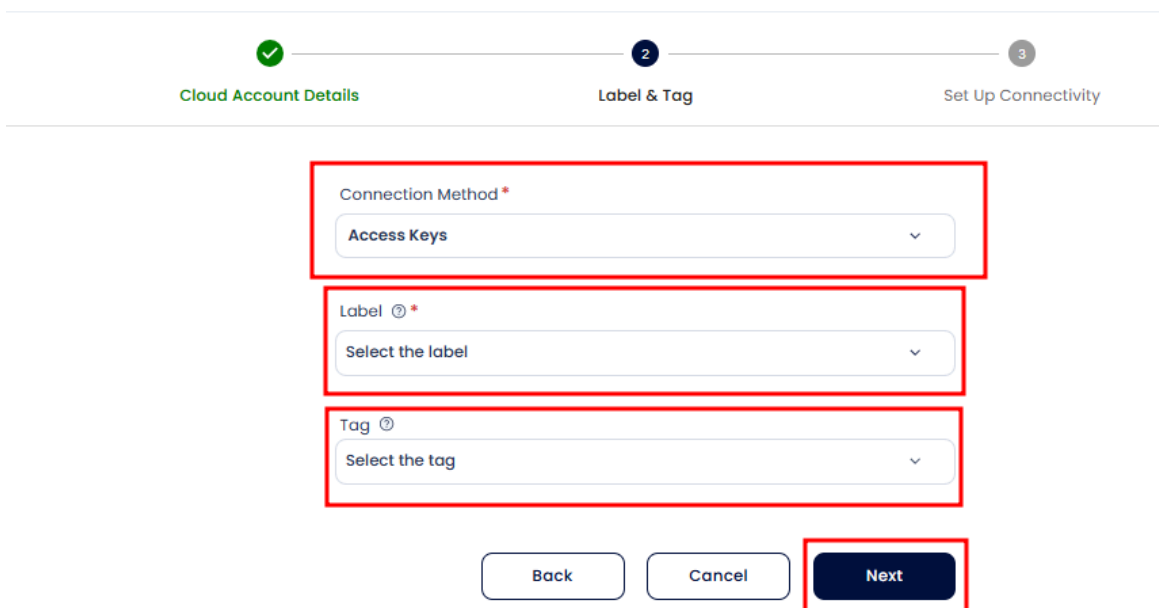
**Step 2:** In the Cloud Account Page select *Add Account* option



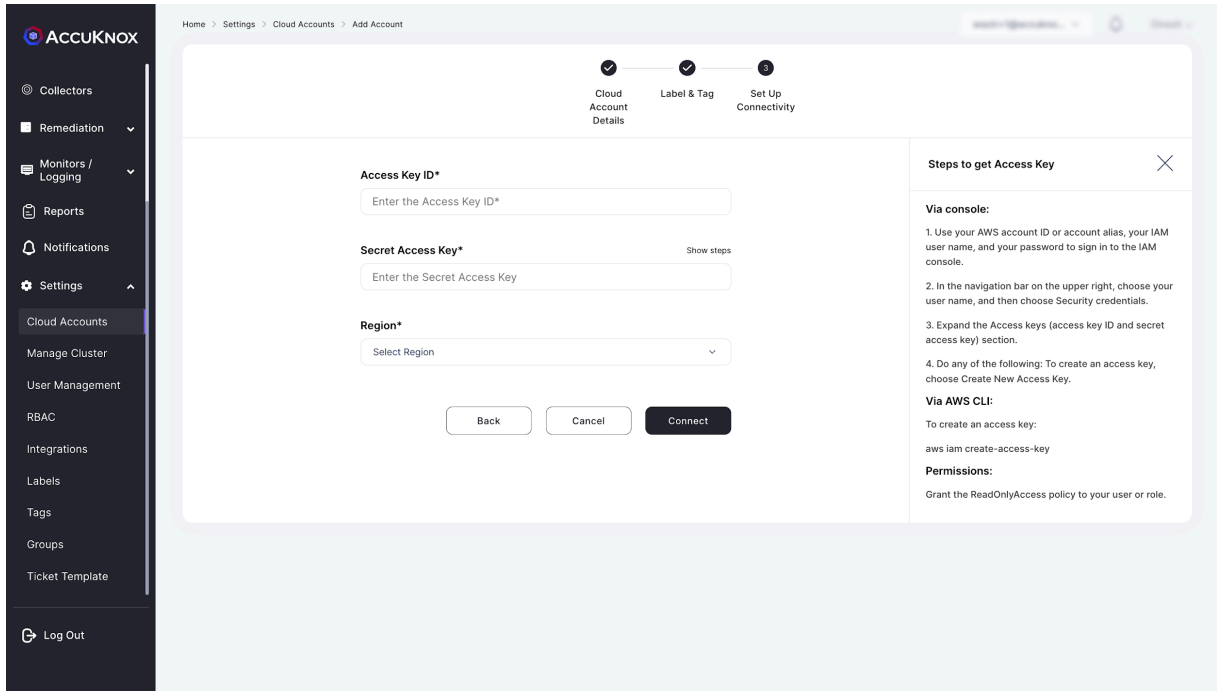
### Step 3: Select the AWS option



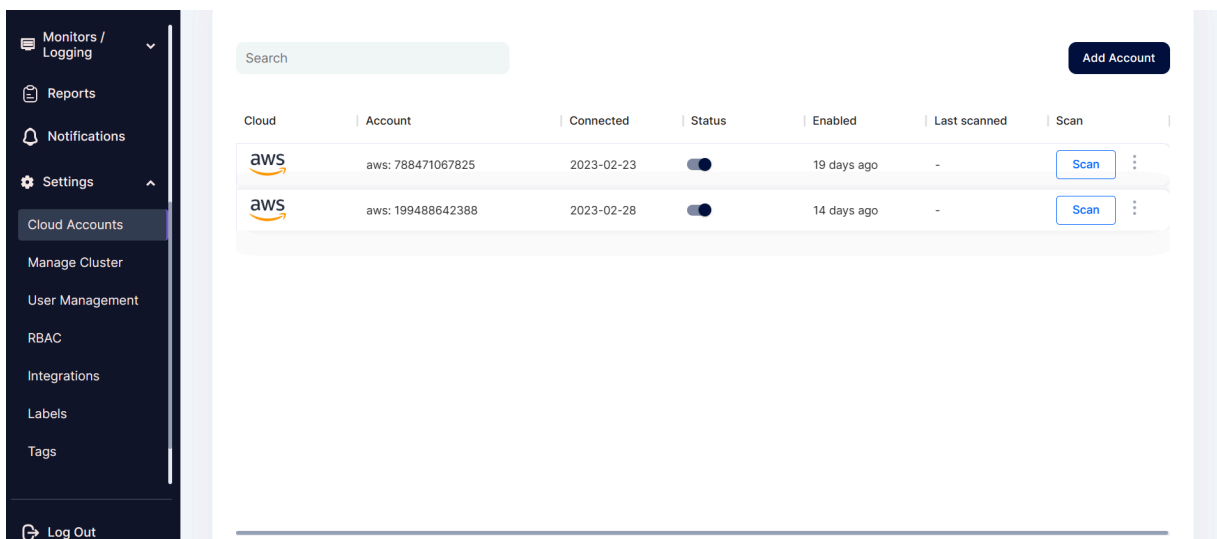
### Step 4: In the next Screen select the Connection method, labels and Tags field from the dropdown Menu.



### Step 5: After giving labels and Tag in the Next Screen Provide the AWS account's Access Key and Secret Access Key ID and Select the Region of the AWS account.



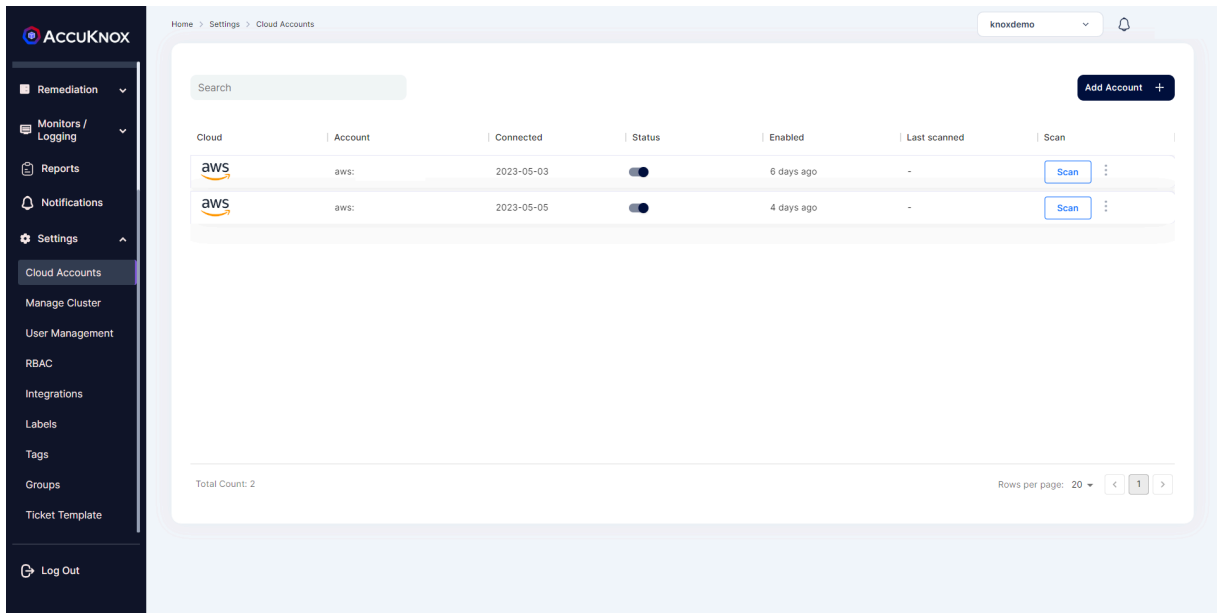
**Step 6:** AWS account is added to the AccuKnox using Access Key Method. We can see the onboarded cloud account by navigating to Settings→cloud Accounts option.



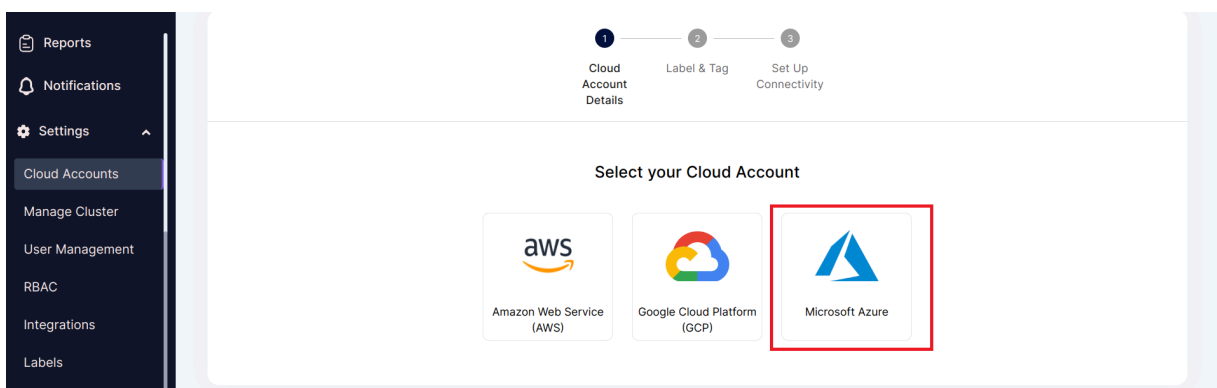
## 3.2 Azure Onboarding

In order to onboard the Azure cloud account onto AccuKnox SaaS Platform.

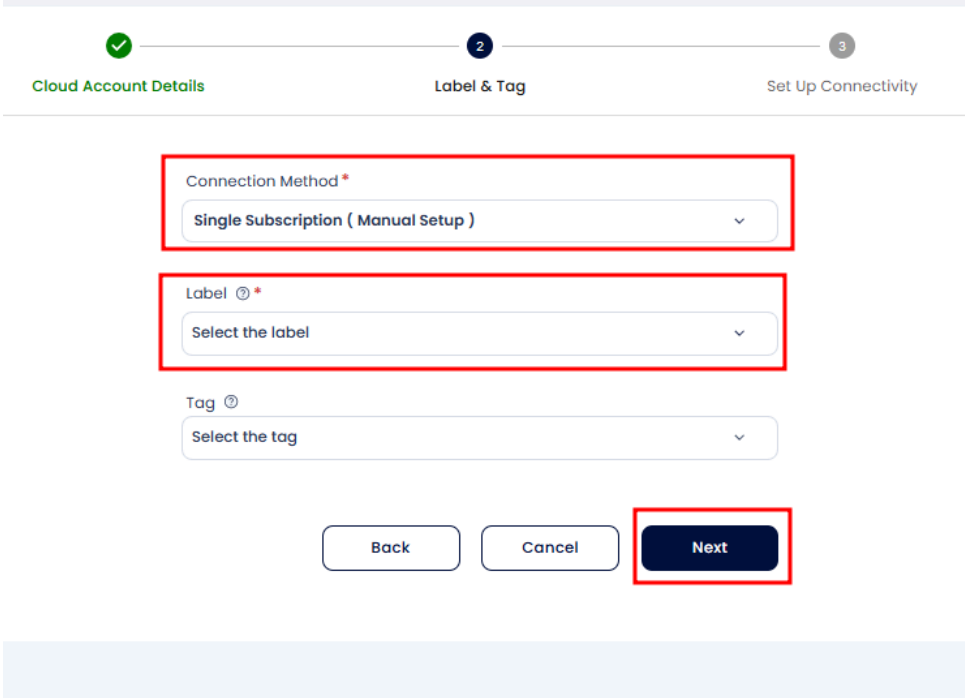
**Step 1:** Go to settings→ Cloud Account and click on Add Account



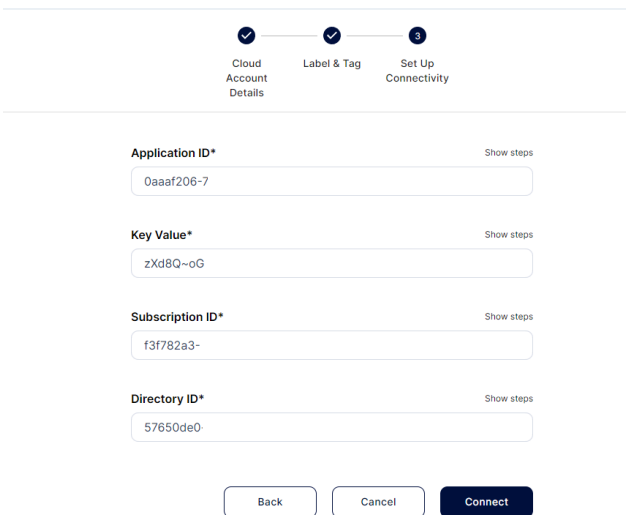
**Step 2:** Select Microsoft Azure as Cloud Account Type



**Step 3:** Select Connection Method, create label and Tags that will be associated with this Cloud Account



**Step 4:** Enter the details that we saved earlier during the steps for app registration and subscription id from subscriptions in azure portal and click on connect



**Step 5:** After successfully connecting your cloud account will show up in the list

Home > Settings > Cloud Accounts

knoxdemo

Search Add Account +

Cloud	Account	Connected	Status	Enabled	Last scanned	Scan
	aws:	2023-05-03	<input checked="" type="checkbox"/>	6 days ago	-	<span>Scan</span> <span>⋮</span>
	aws:	2023-05-05	<input checked="" type="checkbox"/>	4 days ago	-	<span>Scan</span> <span>⋮</span>
	azure:	2023-05-09	<input checked="" type="checkbox"/>	a minute ago	-	<span>Scan</span> <span>⋮</span>

Total Count: 3

Rows per page: 20 < 1 >

Remediation

Monitors / Logging

Reports

Notifications

Settings

Cloud Accounts

Manage Cluster

User Management

RBAC

Integrations

Labels

Tags

Groups

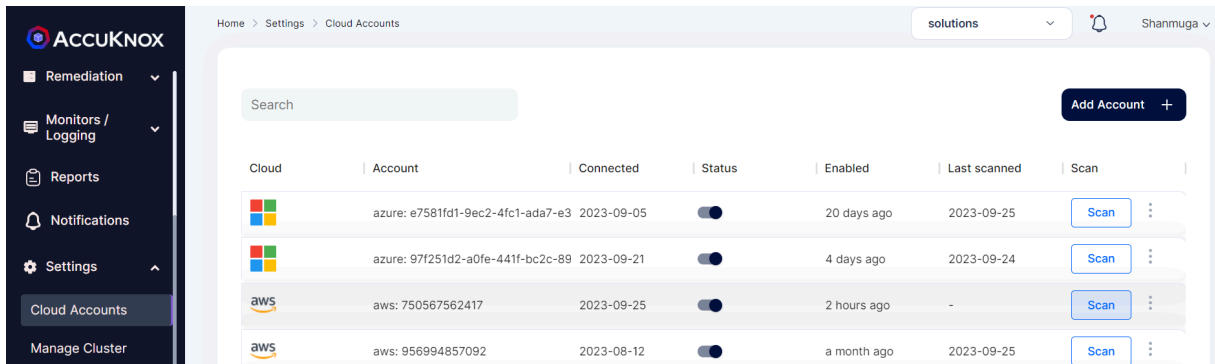
Ticket Template

Log Out

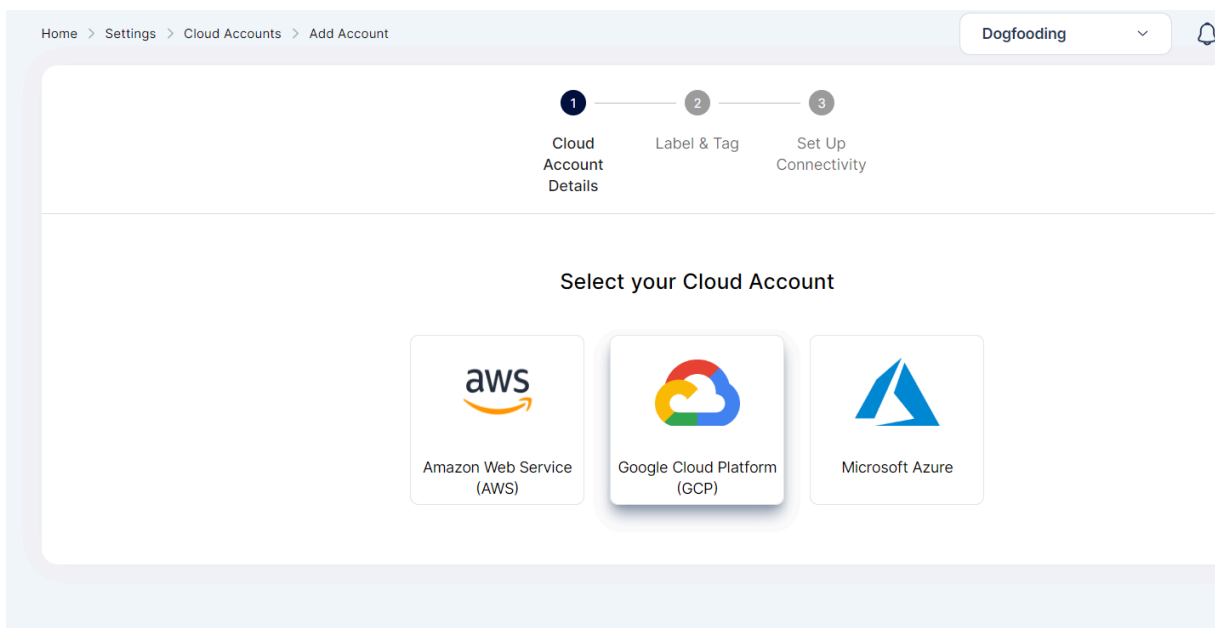


### 3.3 GCP Onboarding

**Step 1:** Go to the AccuKnox SaaS. Navigate to the “Settings” → “Cloud Accounts” then “Add Account”.



**Step 2:** Click the “GCP Platform”



**Step 3:** Select a Connection method, Create New Label and Add the Label for identifying the assets inside this account and add a Tag optionally.

Connection Method \*  
 Drag and Drop

Label ⓘ \*  
 Select the label

Tag ⓘ  
 Select the tag

Back
Cancel
Next

**Step 4:** Enter the “Project ID“, “Client Email”(The Service Account mail ID) and “Private Key” from the downloaded File. Copy paste the entire downloaded file into the ”Private Key” field . Then Click “Connect“

✓ ✓ 3  
 Cloud Account Details    Label & Tag    Set Up Connectivity

[Show steps](#)

Project ID

Client Email

Private Key  

```
study-396808.iam.gserviceaccount.com",
"universe_domain": "googleapis.com"
}
```



Back
Cancel
Connect

The cloud account has been onboarded successfully

Search

✓ Account Connected Successfully ×

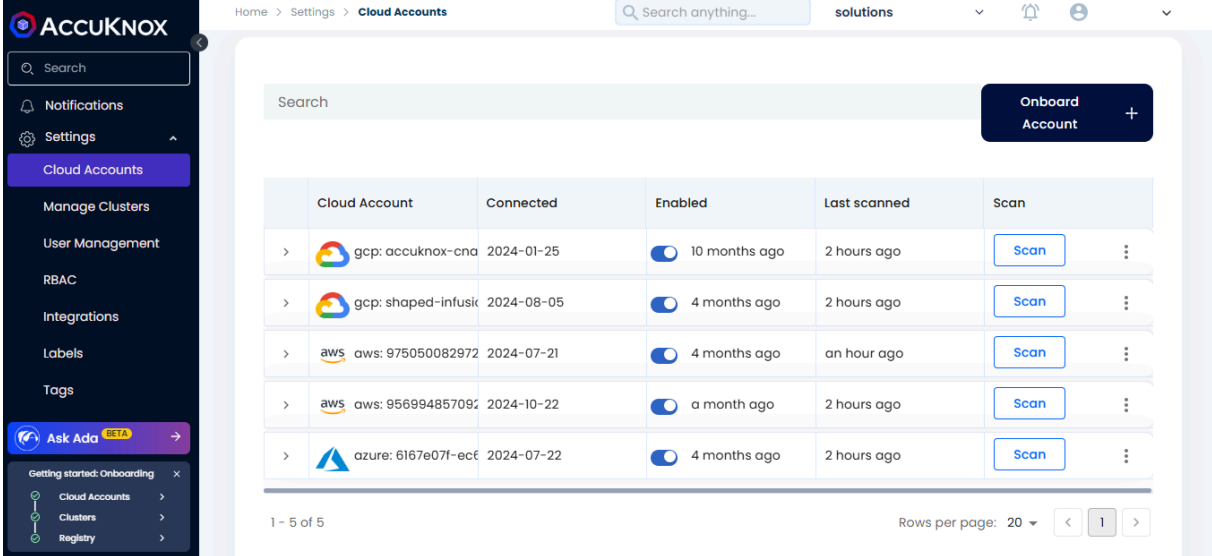
Add Account +

Cloud	Account	Connected	Status	Enabled	Last scanned	Scan
	aws: 956994857092	2023-09-21	<input checked="" type="checkbox"/>	5 days ago	2023-09-25	<span style="border: 1px solid #006633; padding: 2px 5px;">Scan</span> <span style="font-size: 0.8em; vertical-align: middle;">⋮</span>
	gcp: centering-study-396808	2023-09-26	<input checked="" type="checkbox"/>	a few seconds ago	-	<span style="border: 1px solid #006633; padding: 2px 5px;">Scan</span> <span style="font-size: 0.8em; vertical-align: middle;">⋮</span>

## 3.4 Cloud Account Deboarding

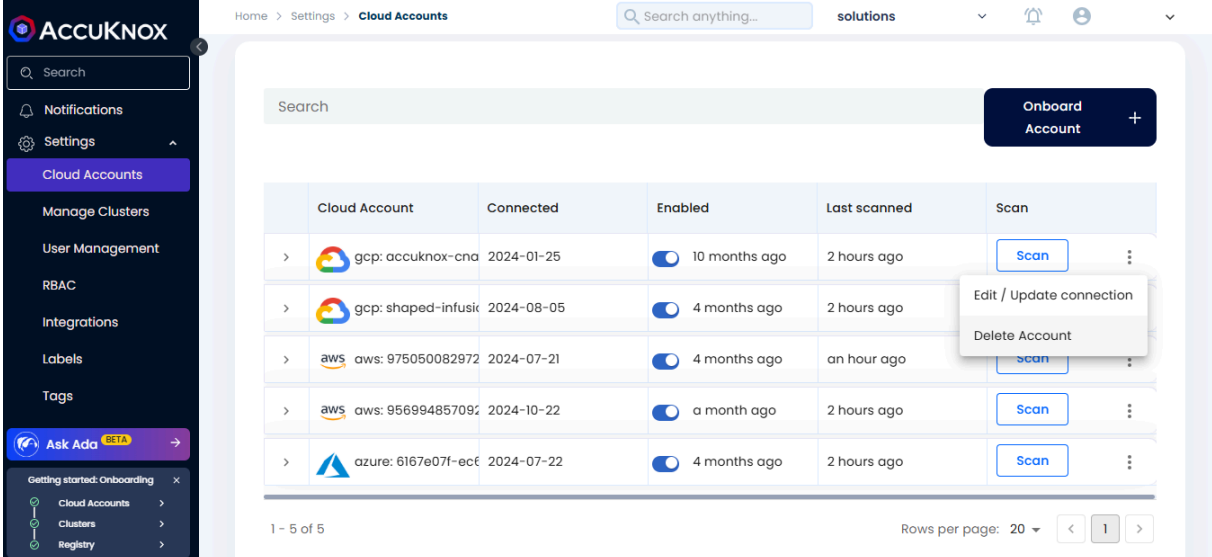
This guide outlines the steps for offboarding a cloud account from AccuKnox SaaS.

**Step 1:** Login to AccuKnox SaaS and Go to Cloud Accounts under Settings.



Cloud Account	Connected	Enabled	Last scanned	Scan
gcp: accuknox-cna	2024-01-25	<input checked="" type="checkbox"/> 10 months ago	2 hours ago	<a href="#">Scan</a>
gcp: shaped-infusik	2024-08-05	<input checked="" type="checkbox"/> 4 months ago	2 hours ago	<a href="#">Scan</a>
aws: 975050082972	2024-07-21	<input checked="" type="checkbox"/> 4 months ago	an hour ago	<a href="#">Scan</a>
aws: 956994857092	2024-10-22	<input checked="" type="checkbox"/> a month ago	2 hours ago	<a href="#">Scan</a>
azure: 6167e07f-ecf	2024-07-22	<input checked="" type="checkbox"/> 4 months ago	2 hours ago	<a href="#">Scan</a>

**Step 2:** Select the cloud account and click “Delete” to delete the account from SaaS.



Cloud Account	Connected	Enabled	Last scanned	Scan
gcp: accuknox-cna	2024-01-25	<input checked="" type="checkbox"/> 10 months ago	2 hours ago	<a href="#">Scan</a>
gcp: shaped-infusik	2024-08-05	<input checked="" type="checkbox"/> 4 months ago	2 hours ago	<a href="#">Scan</a>
aws: 975050082972	2024-07-21	<input checked="" type="checkbox"/> 4 months ago	an hour ago	<a href="#">Scan</a>
aws: 956994857092	2024-10-22	<input checked="" type="checkbox"/> a month ago	2 hours ago	<a href="#">Scan</a>
azure: 6167e07f-ecf	2024-07-22	<input checked="" type="checkbox"/> 4 months ago	2 hours ago	<a href="#">Scan</a>

This will delete the cloud account from AccuKnox SaaS.

## 4. CWPP Prerequisites

### 4.1 Minimum Resource required

Deployments	Resource Usage	Ports	Connection Type	AccuKnox Endpoint
KubeArmor	CPU: 200 m, Memory: 200 Mi	-	-	-
Agents Operator	CPU: 50 m, Memory: 50 Mi	8081, 9090	Outbound	*.accuknox.com:8081 --> SPIRE Access *.accuknox.com:9090 --> SPIRE Health Check
Discovery Engine	CPU: 200 m, Memory: 200 Mi	-	-	-
Shared Informer Agent	CPU: 20 m, Memory: 50 Mi	3000	Outbound	*.accuknox.com:3000 --> Knox-gateway
Feeder Service	CPU: 50 m, Memory: 100 Mi	3000	Outbound	*.accuknox.com:3000 --> Knox-gateway
Policy Enforcement	CPU: 10 m, Memory: 20 Mi	443	Outbound	*.accuknox.com:443 --> Policy Provider Service

These ports need to be allowed through the firewall.

### 4.2 AccuKnox Agents

We have the agent-based model for CWPP. This offers a balanced approach providing non-intrusive scanning for cloud accounts – not to mention the deep visibility for workloads using eBPF-based agents.

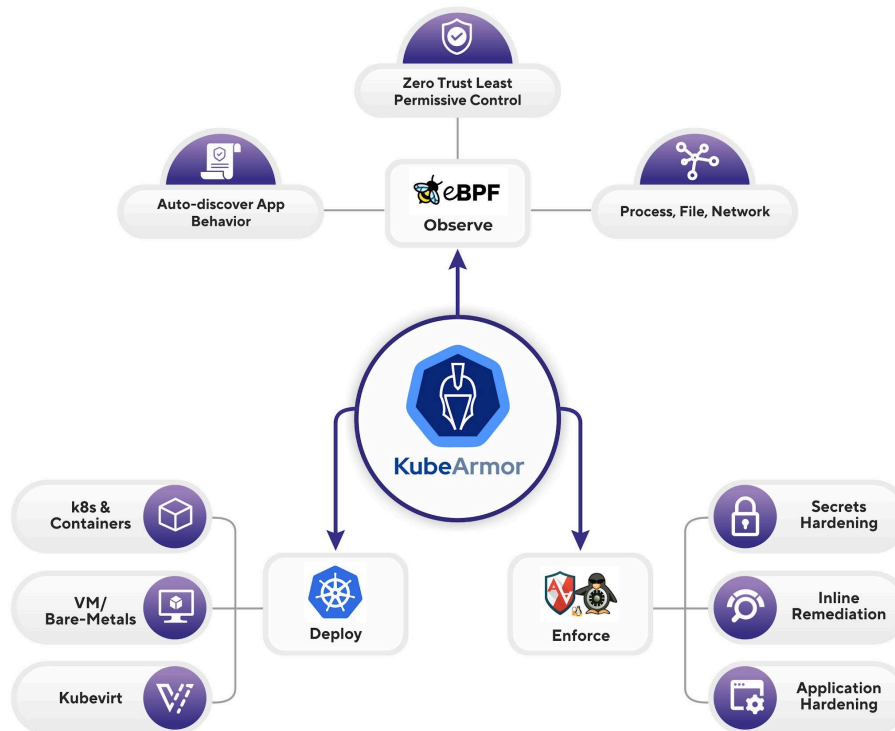
<b>CWPP (Requires Agent)</b>	<b>Protects cloud workloads against Zero-Day Attacks and safeguards against runtime exploits.</b>
Runtime Anomaly Detection	Detects anomalies in application behavior, compliance drift, and attack analysis with detailed context.
Container Forensics Analysis	Helps analyze containers for deep packet-level inspection and understanding of security controls.
Mitigation of Runtime Vulnerabilities	Mitigates exploitable vulnerabilities by applying least permissive security posture.
Protection from Cloud Native Attacks	Essential for safeguarding against sophisticated cloud native attacks that can evade agentless detection.

Note that we also offer an agentless model for CSPM. This is a lightweight, non-intrusive approach that provides deep visibility into cloud accounts without the need for agents. AccuKnox’s hybrid approach optimizes cloud security for diverse organizational needs.

**Listed below are the various agents that are part of the AccuKnox solution.**

## **1. KubeArmor**

KubeArmor is a cloud-native runtime security enforcement system that restricts the behavior (such as process execution, file access, and networking operation) of containers and nodes at the system level. It operates with Linux security modules LSMs, meaning that it can work on top of any Linux platforms (such as Alpine, Ubuntu, and Container-optimized OS from Google) if Linux security modules (e.g., AppArmor, SELinux, or BPF-LSM) are enabled in the Linux Kernel. KubeArmor will use the appropriate LSMs to enforce the required policies.



KubeArmor allows operators to define security policies and apply them to Kubernetes. Then, KubeArmor will automatically detect the changes in security policies from Kubernetes and enforce them on the corresponding containers and nodes. If there are any violations against security policies, KubeArmor immediately generates alerts with container identities. If operators have any logging systems, it automatically sends the alerts to their systems as well.

## 2. Feeder Service

The feeder service sends information from the Client Cluster to the AccuKnox SaaS Control Plane. Feeder Service is an agent which runs on every node, collects telemetry/alert events from source systems & messages, and emits them to Messaging Cluster for Storage & Analysis. Ways in which the Feeder service communicates to the central control plane:

- Directly posting messages to Kafka Topic
- List of topics (Each component has a separate topic name) on where the feeder service publishes feeds.
- Posting via a GRPC or REST API Service

All communication between Feeder and Control plane (Kafka. etc) is encrypted using TLS. Feeder Service uses a secret key from Kubernetes secrets to be applied to it when connecting to the control plane. This secret key allows the feeder to talk to the control plane and exchange data for a particular tenant-id/workspace-id. This is an API key that is generated as part of the cluster onboarding. The feeder service will self-assess some metrics and logs and send that information to the Control plane for its own health assessment for one or more components including its own (running on nodes).The Feeder Service makes it simpler to monitor the detailed communication between each entity.

### **3. Shared Informer Agent**

Shared Informer Agent watches all the changes occurring in Kubernetes entities such as Pods, Nodes, Namespaces, Endpoints, and Services.

- Any changes to an entity can be easily tracked by the Shared Informer Agent such as the Creation of an entity, the update of an entity, and if any entity has been deleted and as soon as the changes occur to the entities, the Shared Informer Agent pushes the information to the backend.
- The Shared Informant Agent makes it simpler to track and manage all of the entities that are present in Kubernetes as well as see changes in entities as they occur in real-time.

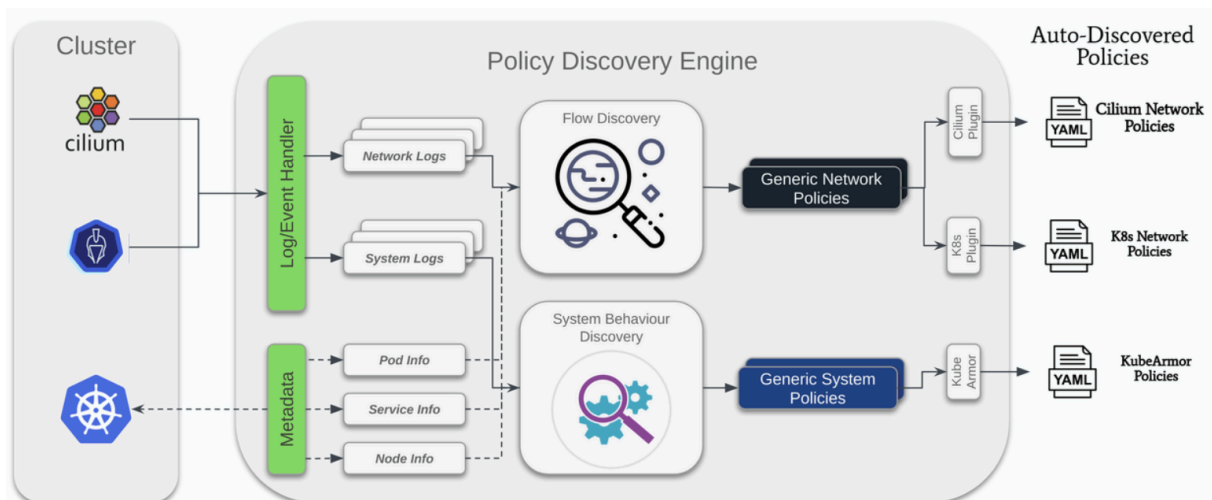
### **4. Policy Enforcement Agent**

AccuKnox's Policy Enforcement Agent enforces the policies by leveraging KubeArmor and Cillium. Policy Enforcement Agent not only keeps the track of the policies but is capable of doing tasks such as applying policies, denying policies, updating policies, and deleting the policies.

- The policy enforcement agent encrypts and decrypts the policies while handing them to and from the policy provider service. It reads the specification of the policies and provides back to the policy provider service.
- All of the changes done to the policy can be tracked granularly with the help of the Policy Enforcement Agent and Policy Gitops Flow which helps with version control and robust management of the security policies.

## 5. Discovery Engine

AccuKnox policy enforcement engine based on KubeArmor is very flexible and powerful. However, these policy engines must be fed with policies. With 10s or 100s of pods and workloads running in a cluster, it is insanely difficult to handcraft such policies. AccuKnox policy auto-discovery engine leverages the pod visibility provided by KubeArmor to auto-generate network and system policies.



AccuKnox's Runtime security solution is able to provide full visibility into all of these application interactions with the host kernel and provide the ability to filter or restrict specific actions at runtime.

With AccuKnox you can automatically discover the application interaction and network interaction (as described below) in the form of policy as code subsequently these policies can be audited or enforced at runtime giving you the ability to restrict specific behaviors of the application.

For example, you could have a policy that states the following:

- Pod A cannot access the/etc/bin folder
- Pod B cannot initiate ptrace i.e. trace the execution of other processes.
- Pod C cannot communicate to a remote TCP server running on port 5000.

This list can be as exhaustive as you like, and these policies are enforced within the kernel using kernel primitives and technologies as listed below:



## **Network Security using eBPF**

- Network runtime protection in the form of L3, L4, and L7 rules using identity (x509 certificates or K8s labels) for your K8s workloads. In K8s policies, this is implemented as a native K8s networkpolicy object.
- For Virtual Machine workloads, labels are used to provide host-level network policies for L3, L4, and L7.

## **Application security using Linux Security Modules (LSM) / KubeArmor**

- The Linux Security Module (LSM) framework provides a mechanism for various security checks to be hooked by new kernel extensions. It denies access to essential kernel objects, such as files, inodes, task structures, credentials, and inter-process communication objects.
- AccuKnox supports AppArmor, SELinux and BPF LSM as of today for its enforcement engine at runtime.

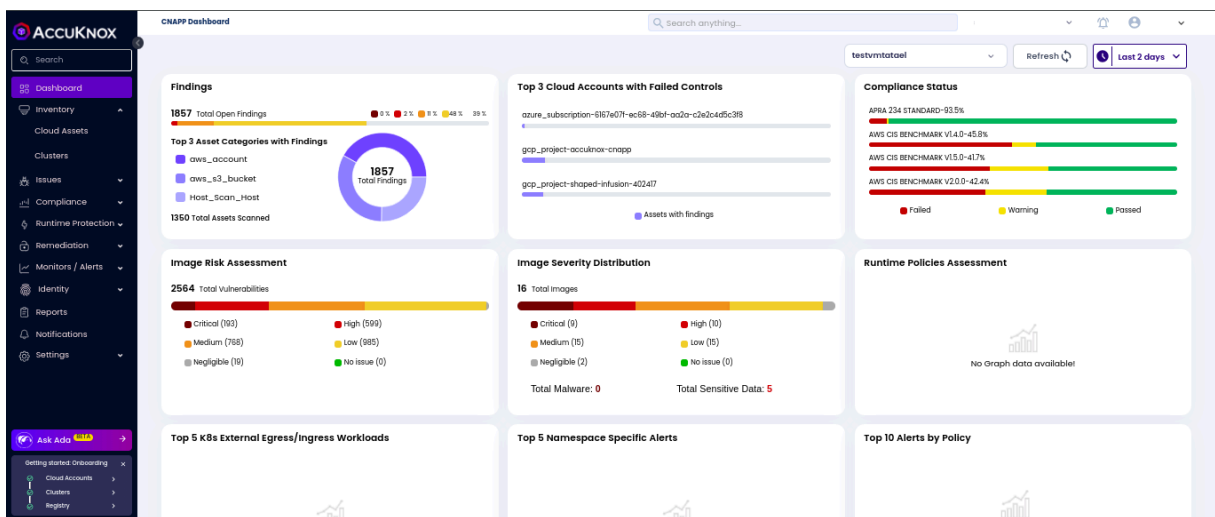
## 5. Cluster Onboarding

The cluster onboarding steps are the same for both managed and unmanaged clusters as follows:

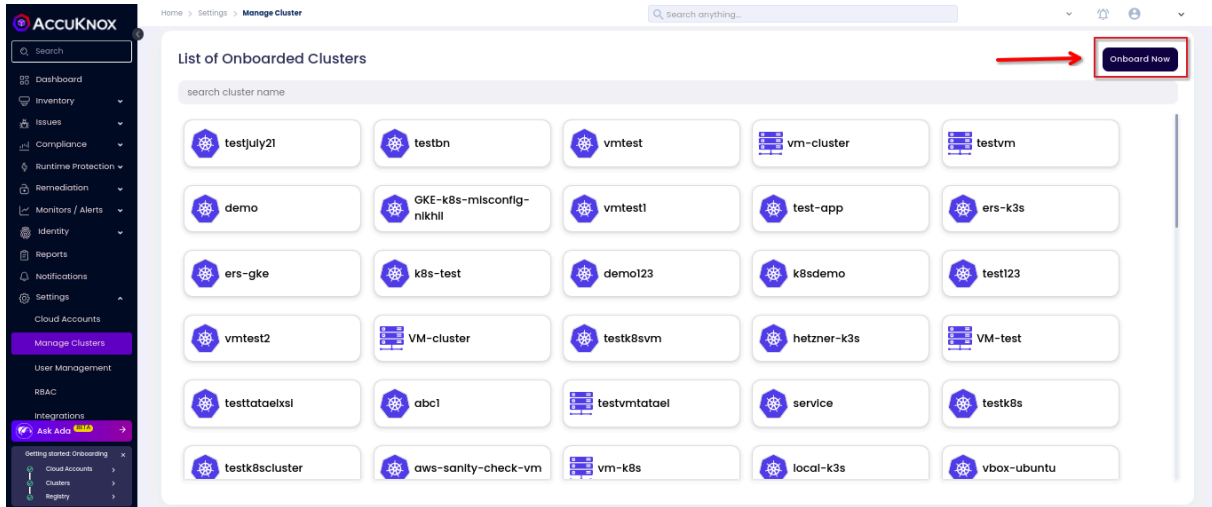
Below shown image is from an k3s cluster running in a local machine with Kali Linux Operating System. We can onboard this cluster by following the steps shown below

```
(Accuknox@kali) - [~]
└─$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
nginx-demo    1/1     Running   0           22s
redis-demo    1/1     Running   0           14s
```

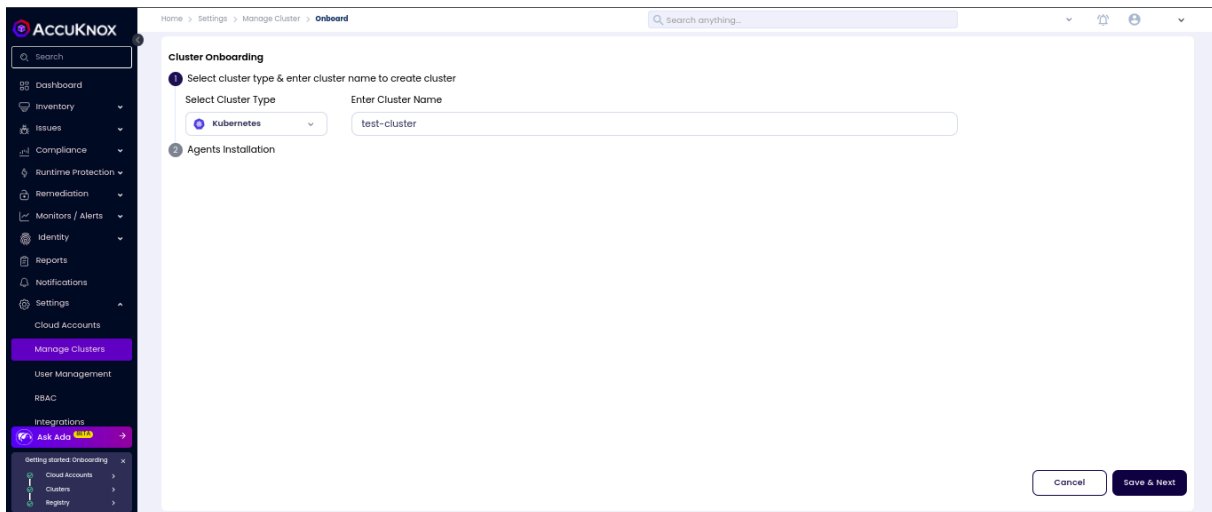
**Step 1:** As a first time user, the management console will show up the CNAPP dashboard without any data mentioned in widgets, since the cloud account and cluster onboarding is not done.



**Step 2:** Navigate to Manage Cluster from Settings Tab: From this page we can onboard the clusters running in various cloud platforms like GCP,AWS and Azure. We can onboard locally setup clusters using an cloud option. To onboard cluster select onboard now option



**Step 3:** In this screen, give any name to the cluster that you are going to onboard now.



**Step 4:** Installing KubeArmor and AccuKnox agents

We are going to install KubeArmor and AccuKnox-agents to connect to the AccuKnox SaaS application. For the agent installation selection click on the Runtime Visibility & Protection.

## Step 4.1 KubeArmor Installation

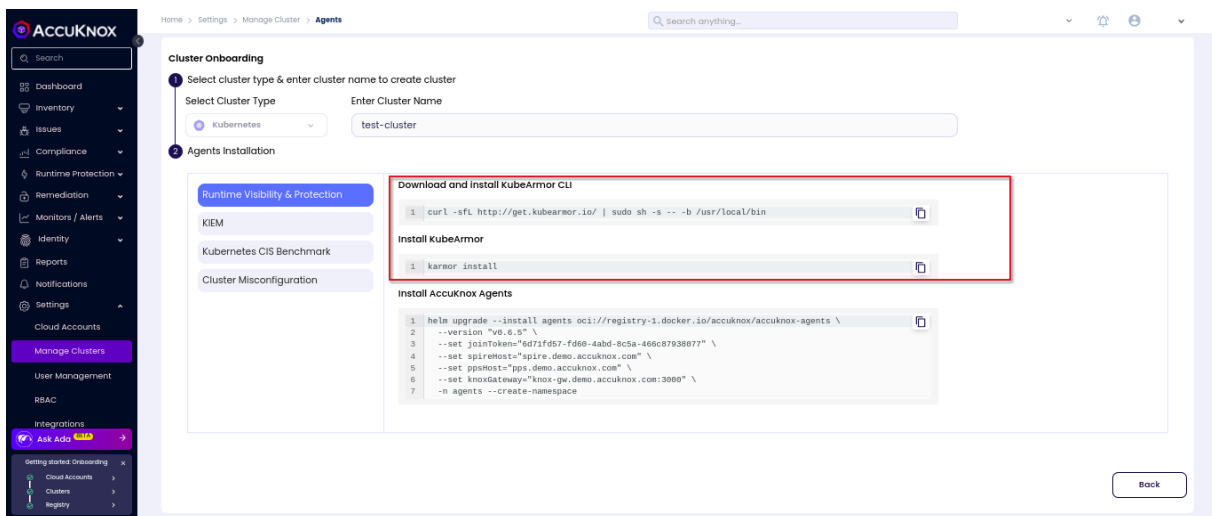
# KubeArmor

KubeArmor is a cloud-native runtime security enforcement system that restricts the behavior (such as process execution, file access, and networking operation) of containers and nodes at the system level.

With KubeArmor, a user can:

- Restrict file system access for certain processes
- Restrict what processes can be spawned within the pod
- Restrict the capabilities that can be used by the processes within the pod

KubeArmor differs from seccomp-based profiles, wherein KubeArmor allows to dynamically set the restrictions on the pod. With seccomp, the restrictions must be placed during the pod startup and cannot be changed later. KubeArmor leverages Linux Security Modules (LSMs) to enforce policies at runtime.



KubeArmor is installed using the following commands:

```
>> curl -sL http://get.kubearmor.io/ | sudo sh -s -- -b /usr/local/bin
>> karmor install
```

Sample Output:

```
(Accuknox@kali)-[~]
└─$ curl -sL http://get.kubearmor.io/ | sudo sh -s -- -b /usr/local/bin
kubearmor/kubearmor-client info checking GitHub for latest tag
kubearmor/kubearmor-client info found version: 1.2.3 for
v1.2.3/linux/amd64
kubearmor/kubearmor-client info installed /usr/local/bin/karmor
kubearmor/kubearmor-client info karmor is installed in /usr/local/bin
kubearmor/kubearmor-client info invoke /usr/local/bin/karmor or move
karmor to your desired PATH

(Accuknox@kali)-[~]
└─$ karmor install
🛡️ Installed helm release : kubearmor-operator
😊 KubeArmorConfig updated
🕒 This may take a couple of minutes
🎉 KubeArmor Snitch Deployed!
🎉 KubeArmor Daemonset Deployed!
😊 Done Checking , ALL Services are running!
🕒 Execution Time : 1m22.006691427s
🔧 Verifying KubeArmor functionality (this may take upto a minute)
|.
🛡️ Your Cluster is Armored Up!

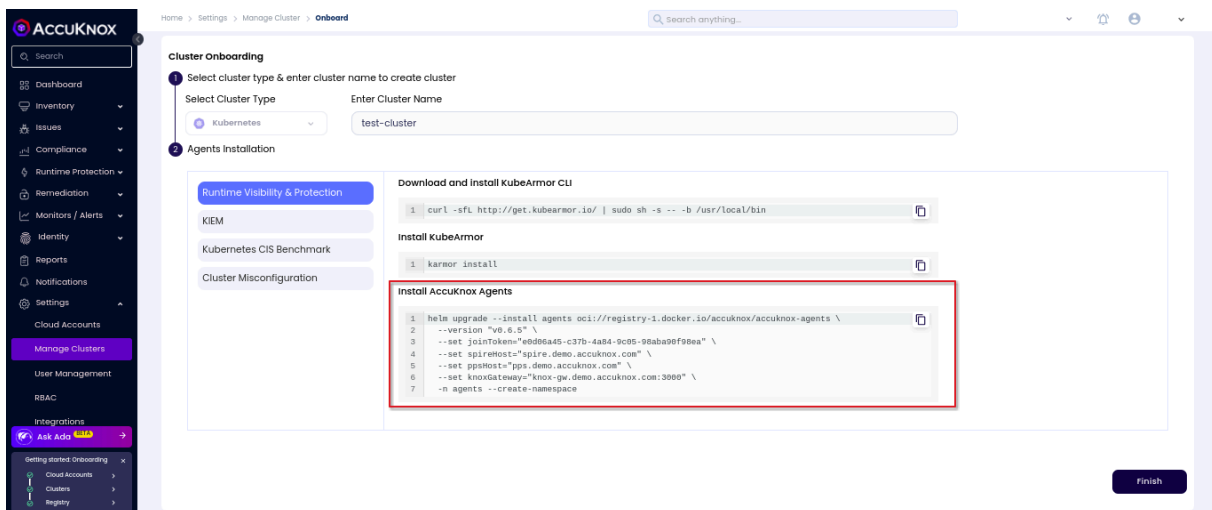
(Accuknox@kali)-[~]
└─$
```

## Step 4.2: AccuKnox-Agents installation

After installing KubeArmor we are going to install AccuKnox Agents in the cluster.

## AccuKnox Agents

1. **KubeArmor:** KubeArmor is a cloud-native runtime security enforcement system that restricts the behavior (such as process execution, file access, and networking operation) of containers and nodes at the system level. KubeArmor dynamically set the restrictions on the pod. KubeArmor leverages Linux Security Modules (LSMs) to enforce policies at runtime.
2. **Feeder Service:** It collects the feeds from kubeArmor and relays to the app.
3. **Shared Informer Agent:** It collects information about the cluster like pods, nodes, namespaces etc.,
4. **Policy Discovery Engine:** It discovers the policies using the workload and cluster information that is relayed by a shared informer Agent.



AccuKnox Agents can be installed using the following command:

```
helm upgrade --install agents
oci://registry-1.docker.io/accuknox/accuknox-agents
--version "v0.6.5"
--set joinToken="*****_*****_*****"
--set spireHost="spire.demo.accuknox.com"
--set ppsHost="pps.demo.accuknox.com"
--set knoxGateway="knox-gw.demo.accuknox.com:3000"
-n agents --create-namespace
```

Sample Output:

```
WARNING: Kubernetes configuration file is group-readable. This is
insecure. Location: /etc/rancher/k3s/k3s.yaml

WARNING: Kubernetes configuration file is world-readable. This is
insecure. Location: /etc/rancher/k3s/k3s.yaml
Release "agents" does not exist. Installing it now.
Pulled: registry-1.docker.io/accuknox/accuknox-agents:v0.6.5
Digest:
sha256:420a4dae8225ce1eb201b5468c588eeb71bbf532f9d9f1eafac2281760f61e11
NAME: agents
LAST DEPLOYED: Fri Jul 26 15:23:37 2024
NAMESPACE: agents
STATUS: deployed
REVISION: 1
TEST SUITE: None

(Accuknox@kali)-[~]
└─$
```

Note: In the above command joinToken is specific to this example and it will vary based on the cluster

### Step 5: Onboarded Cluster

After installing all the AccuKnox agents the cluster is onboarded successfully into the SaaS application. We can see the workload details of the onboarded cluster by Navigating to Inventory→cloud Workloads option. There all the onboarded clusters will be listed out and all the inactive ones would be grayed out. By double clicking on the active cluster user can get a more detailed view of the cluster.

Home > Inventory > Clusters

Search anything...

Clusters

Last 15 min | off | Onboard Cluster

LIST | GRAPH

test-cluster, testtabotest, VM-test, VM-sysid, k8s-demo, testk8s, testvmtoat, VM-cluster, alright, vm-k8s, vm-cluster, testk8s, OKE-k8s-miscconfig-rikshl, testk8s, era-k8s, testvm, testk8s, tensorflow-light, tensorflowwattack, test-demo, test, July-demo, detu-hardening-engine-te-wt-2, tensorflow-cluster, pytorch-test, demo-june, tensorflow-cluster, demo-28-june, tensorflow-cluster, DO-demo-cluster, service

+ | -

Getting started Onboarding x

- Cloud Accounts
- Clusters
- Registry

Ask Ads

Home > Inventory > Clusters

Search anything...

Clusters > test-cluster > Workloads

Last 15 min | off | Onboard Cluster

LIST | GRAPH

agents

- agents-operator
- shared-informer-agent
- policy-enforcement-agent
- discovery-engine
- feeder-service

default

- nginx-demo
- redis-demo

kube-system

- evictor-torath-220622ru
- local-path-provisioner
- troubk
- coredns
- metrics-server

+ | -

Getting started Onboarding x

- Cloud Accounts
- Clusters
- Registry

Ask Ads



## 6. Cluster Offboarding

This guide outlines the steps for offboarding a cluster from AccuKnox SaaS. The process involves uninstalling the agents from the cluster and deleting the cluster from AccuKnox SaaS.

Below, you will find detailed instructions for agent uninstallation from your cluster CLI and deleting the cluster from AccuKnox SaaS. These steps apply to all clusters.

### 1. Agents Uninstallation

Uninstall AccuKnox agents using the following commands:

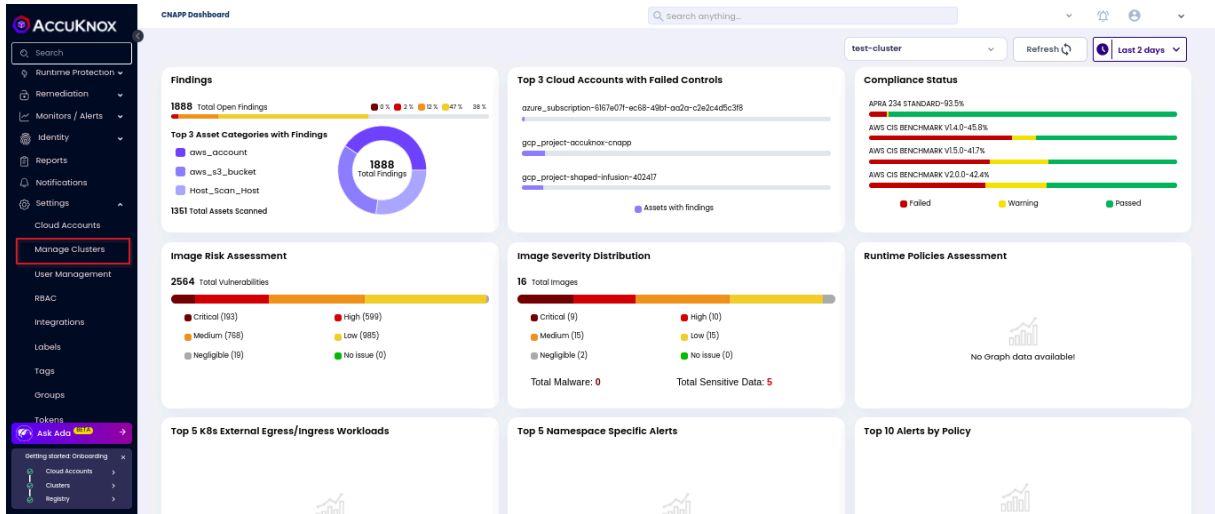
```
helm uninstall agents -n agents && kubectl delete ns agents;  
helm uninstall cis-k8s-job;  
helm uninstall kiem-job;  
helm uninstall k8s-risk-assessment-job
```

### 2. Sample for Uninstalling Runtime Visibility & Protection agents

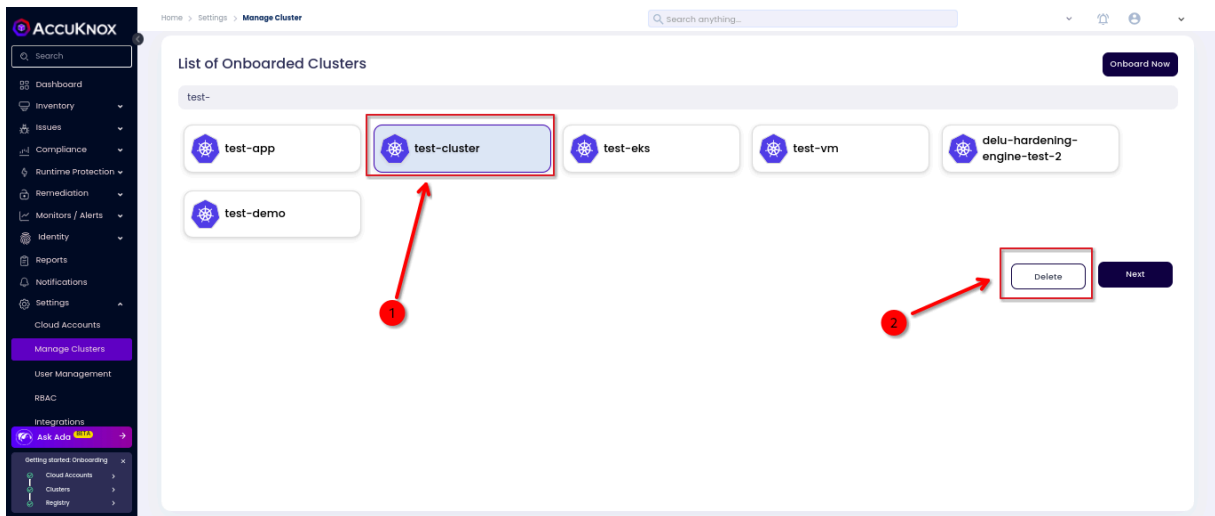
```
(Accuknox@kali)-[~]  
└─$ helm uninstall agents -n agents && kubectl delete ns agents  
  
WARNING: Kubernetes configuration file is group-readable. This is insecure.  
Location: /etc/rancher/k3s/k3s.yaml  
WARNING: Kubernetes configuration file is world-readable. This is insecure.  
Location: /etc/rancher/k3s/k3s.yaml  
release "agents" uninstalled  
namespace "agents" deleted
```

### 3. Cluster Deletion

**Step 1:** Login to AccuKnox SaaS and Go to Manage Cluster under Settings



**Step 2:** Select the cluster and click Delete to delete the cluster from SaaS.



This will delete the cluster from AccuKnox SaaS

## 7. VM Onboarding with Systemd/Docker Mode

### 7.1 Systemd

**Systemd** is a core component of modern Linux systems responsible for managing services and processes. It ensures that essential services start automatically during boot, remain running, and restart if they fail. In simple terms, systemd acts like a **controller** that organizes and oversees everything needed to keep the system stable and functional.

Currently, **root/sudo** permissions are needed for onboarding systemd. This is because KubeArmor requires privileges to protect the host and systemd services, packages are currently installed on the root directory.

Only in case of the control plane node, a working RabbitMQ server is required. This can be installed using Docker.

```
# Latest RabbitMQ 3.13
docker run -it --rm --name rabbitmq -p 5672:5672 -p 15672:15672
rabbitmq:3.13-management
```

Alternatively, you can install RabbitMQ using a package manager:

- **Linux, BSD, UNIX:** [Debian, Ubuntu](#) | [RHEL, CentOS Stream, Fedora](#) | [Generic binary build](#) | [Solaris](#)
- **Windows:** [Chocolatey package](#) | [Windows Installer](#) | [Binary build](#)
- **MacOS:** [Homebrew](#) | [Generic binary build](#)
- [Erlang/OTP for RabbitMQ](#)

BTF support is needed. Any kernel version which has this should work. Check if BTF info is present with the script below:

```
if [ ! -e "/sys/kernel/btf/vmlinux" ]; then
  echo "BTF info not present"
else
  echo "BTF info present"
fi
```

If the script returns "BTF info not present," [BTF support is not available](#), and you should run the script below to build the required files on your system:

```
# Download KubeArmor
git clone https://github.com/kubearmor/KubeArmor/
cd KubeArmor/KubeArmor/packaging
./post-install.sh
```

For detailed instructions specific to SystemD Based Non-BTF Environments, please refer to this [guide](#).

### 7.1.1 Container Protection Requirements (Optional)

If container protection is needed, a Linux Kernel with **BPF LSM** is desired. Generally, it is present in v5.8+. Here's a guide on enabling BPF LSM: [KubeArmor Getting Started FAQ](#).

If BPF LSM is not available, AppArmor should still work out of the box for host policy application. However, follow the guide [Support for non orchestrated containers](#) for each container.

### 7.1.2 Resource Requirements

Node Type	CPU	Memory	Disk
Control plane node	2vCPU	4 GB	1 GB

Worker node	2vCPU	2 GB	500 MB
-------------	-------	------	--------

### 7.1.3 Network Requirements

Connectivity between control plane node and worker nodes is a must. They should either be:

- Part of the same private network (**recommended & secure**)
- Control plane has a public IP (not recommended)

Component	Port	Endpoint	Purpose
Knox-Gateway	3000	knox-gw.<env>.accuknox.com:3000	For Knox-Gateway service
PPS	443	pps.<env>.accuknox.com	For PPS (Policy Provisioning Service)
Spire-Server	8081, 9090	spire.<env>.accuknox.com	For Spire-Server communication
KubeArmor Relay Server	32768	-	For Kubearmor relay server on control plane
Shared Informer Agent	32769	-	For Shared Informer agent on control plane

Policy Enforcement Agent (PEA)	32770	-	For Policy Enforcement Agent on control plane
Hardening Module	32771	-	For Discovery Engine Hardening Module on control plane
VM Worker Nodes	32768-32771	-	For VM worker nodes to connect to the control plane

Check the CWPP documentation for more details on the [network requirements](#).

You can check the connectivity between nodes using curl. Upon a successful connection, the message returned by curl will be:

```
$ curl <control-plane-addr>:32770
curl: (1) Received HTTP/0.9 when not allowed
```

#### 7.1.4 Onboarding

Navigate to the onboarding page (Settings → Manage Cluster → Onboard Now) and choose the "VM" option on the instructions page. Then, provide a name for your cluster. You will be presented with instructions to download `accuknox-cli` and onboard your cluster.

The following agents will be installed:

1. **Feeder-service** which collects KubeArmor feeds.

2. **Shared-informer-agent** authenticates with your VMs and collects information regarding entities like hosts, containers, and namespaces.
3. **Policy-enforcement-agent** authenticates with your VMs and enforces labels and policies.

## Install knoxctl/accuknox-cli

```
curl -sL https://knoxctl.accuknox.com/install.sh | sudo sh -s --  
-b /usr/bin
```

### 7.1.5 Onboarding Control Plane

The command may look something like this:

```
$ knoxctl onboard vm cp-node \  
--version "v0.2.10" \  
--join-token="843ef458-cecc-4fb9-b5c7-9f1bf7c34567" \  
--spire-host="spire.dev.accuknox.com" \  
--pps-host="pps.dev.accuknox.com" \  
--knox-gateway="knox-gw.dev.accuknox.com:3000"
```

By default, if Docker is not found, systemd mode of installation would be used. If you want to explicitly onboard using systemd services, add the `--vm-mode=systemd` flag to the above command.

The above command will emit the command to onboard worker nodes. You may also use the `--cp-node-addr` flag to specify the address that other nodes will use to connect with your cluster.

By default, the network created by onboarding commands reserves the subnet `172.20.32.0/27` for the `accuknox-net` Docker network. If you want to change it for your environment, you can use the `--network-cidr` flag.

## 7.1.6 Onboarding Worker Nodes

The second command will be for onboarding worker nodes. It may look something like this:

```
knoxcctl onboard vm node --cp-node-addr=<control-plane-addr>
```

Example:

```
$ knoxcctl onboard vm node --cp-node-addr=192.168.56.106
Pulling kubearmor-init      ... done
Pulling kubearmor          ... done
Pulling kubearmor-vm-adapter ... done
Creating network "accuknox-config_accuknox-net" with the default
driver
Creating kubearmor-init ... done
Creating kubearmor      ... done
Creating kubearmor-vm-adapter ... done
onboard-vm-node.go:41: VM successfully joined with control-plane!
```

If you encounter any issues while onboarding, use the commands below to debug:

```
docker logs spire-agent -f

docker logs shared-informer-agent -f

docker logs kubearmor-init -f

docker logs kubearmor -f
```



## 7.1.7 Deboarding

Deboard the cluster from SaaS first.

To deboard the worker-vm/Node:

```
knoxctl deboard vm node
```

To deboard the Control-Plane VM:

```
knoxctl deboard vm cp-node
```

Sample Output:

```
$ knoxctl deboard vm cp-node
[+] Running 10/10
✓ Container shared-informer-agent      Removed
✓ Container feeder-service             Removed
✓ Container policy-enforcement-agent   Removed
✓ Container wait-for-it                Removed
✓ Container kubearmor-vm-adapter        Removed
✓ Container kubearmor-relay-server      Removed
✓ Container spire-agent                 Removed
✓ Container kubearmor                   Removed
✓ Container kubearmor-init              Removed
✓ Network accuknox-config_accuknox-net Removed
```

Please remove any remaining resources at

```
/home/user/.accuknox-config
```

Control plane node deboarded successfully.

After that cleanup the ~/.accuknox-config directory

```
sudo rm -rf ~/.accuknox-config
```

## 7.2 Docker

Docker v19.0.3 and Docker Compose v1.27.0+ are required. Follow the latest [Install Docker Engine](#) for downloading. Ensure you also add your user to the docker user group: [Linux post-installation steps for Docker Engine](#).

Linux Kernel v5.8+ with BPF LSM support is needed. See how to [enable BPF LSM](#).

If Linux v5.8+ or BPF LSM is not supported in the given environment, host enforcement will still work out of the box. For protecting containers, new containers will have to be created with special options. See [Support for non orchestrated containers](#) for the same.

### 7.2.1 Resource Requirements

Node Type	CPU	Memory	Disk
Control plane node	2vCPU	4 GB	24 GB
Worker node	2vCPU	2 GB	12 GB

### 7.2.2 Network Requirements

Connectivity between control plane node and worker nodes is a must. They should either be:

- Part of the same private network (**recommended & secure**)
- Control plane has a public IP (not recommended)

Ports required on the control plane VM:

Component	Type	Ports	Endpoint	Purpose
Knox-Gateway	Outbound to SaaS	3000	knox-gw.<env>.accuknox.com:3000	For Knox-Gateway service
PPS	Outbound to SaaS	443	pps.<env>.accuknox.com	For PPS (Policy Provisioning Service)
Spire-Server	Outbound to SaaS	8081, 9090	spire.<env>.accuknox.com	For Spire-Server communication
KubeArmor Relay Server	Inbound in CP	32768	-	For KubeArmor relay server on control plane
Shared Informer Agent	Inbound in CP	32769	-	For Shared Informer agent on control plane

Policy Enforcement Agent (PEA)	Inbound in CP	32770	-	For Policy Enforcement Agent on control plane
Hardening Module	Inbound in CP	32771	-	For Discovery Engine Hardening Module on control plane
VM Worker Nodes	Outbound from worker node to CP	32768-32771	-	For VM worker nodes to connect to the control plane

By default, the network created by onboarding commands reserves the subnet 172.20.32.0/27. If you want to change it for your environment, you can use the `--network-cidr` flag.

You can check the connectivity between nodes using curl. Upon a successful connection, the message returned by curl will be:

```
$ curl <control-plane-addr>:32770
curl: (1) Received HTTP/0.9 when not allowed
```

## 7.3 Onboarding

Navigate to the onboarding page (Settings → Manage Cluster → Onboard Now) and choose the "VM" option on the instructions page. Then, provide a name for your cluster. You will be presented with instructions to download `accuknox-cli` and onboard your cluster.

The following agents are installed:

1. **Feeder-service** which collects KubeArmor feeds.
2. **Shared-informer-agent** authenticates with your VMs and collects information regarding entities like hosts, containers, and namespaces.
3. **Policy-enforcement-agent** authenticates with your VMs and enforces labels and policies.

### 7.3.1 Install `knoxctl/accuknox-cli`

```
curl -sfL https://knoxctl.accuknox.com/install.sh | sudo sh -s -- -b /usr/bin
```

### 7.3.2 Onboarding Control Plane

The command may look something like this:

```
$ knoxctl onboard vm cp-node \  
  --version "v0.2.10" \  
  --join-token="843ef458-cecc-4fb9-b5c7-9f1bf7c34567" \  
  --spire-host="spire.dev.accuknox.com" \  
  --pps-host="pps.dev.accuknox.com" \  
  --knox-gateway="knox-gw.dev.accuknox.com:3000"
```

The above command will emit the command to onboard worker nodes. You may also use the `--cp-node-addr` flag to specify the address that other nodes will use to connect with your cluster.

By default, the network created by onboarding commands reserves the subnet 172.20.32.0/27 for the `accuknox-net` Docker network. If you want to change it for your environment, you can use the `--network-cidr` flag.

### 7.3.3 Onboarding Worker Nodes

The second command will be for onboarding worker nodes. It may look something like this:

```
knoxcctl onboard vm node --cp-node-addr=<control-plane-addr>
```

Example:

```
$ knoxcctl onboard vm node --cp-node-addr=192.168.56.106
Pulling kubearmor-init      ... done
Pulling kubearmor          ... done
Pulling kubearmor-vm-adapter ... done
Creating network "accuknox-config_accuknox-net" with the default
driver
Creating kubearmor-init ... done
Creating kubearmor      ... done
Creating kubearmor-vm-adapter ... done
onboard-vm-node.go:41: VM successfully joined with control-plane!
```

If you encounter any issues while onboarding, use the commands below to debug:

```
docker logs spire-agent -f
docker logs shared-informer-agent -f
docker logs kubearmor-init -f
docker logs kubearmor -f
```

## 7.4 Deboarding

Deboard the cluster from SaaS first.

To deboard the worker-vm/Node:

```
knoxctl deboard vm node
```

To deboard the Control-Plane VM:

```
knoxctl deboard vm cp-node
```

Sample Output:

```
$ knoxctl deboard vm cp-node
[+] Running 10/10
✓ Container shared-informer-agent      Removed
✓ Container feeder-service             Removed
✓ Container policy-enforcement-agent    Removed
✓ Container wait-for-it                 Removed
✓ Container kubearmor-vm-adapter        Removed
✓ Container kubearmor-relay-server      Removed
✓ Container spire-agent                 Removed
✓ Container kubearmor                   Removed
✓ Container kubearmor-init              Removed
✓ Network accuknox-config_accuknox-net Removed
Please remove any remaining resources at
/home/user/.accuknox-config
Control plane node deboarded successfully.
```

After that cleanup the ~/.accuknox-config directory

```
sudo rm -rf ~/.accuknox-config
```

## 8. Registry Onboarding

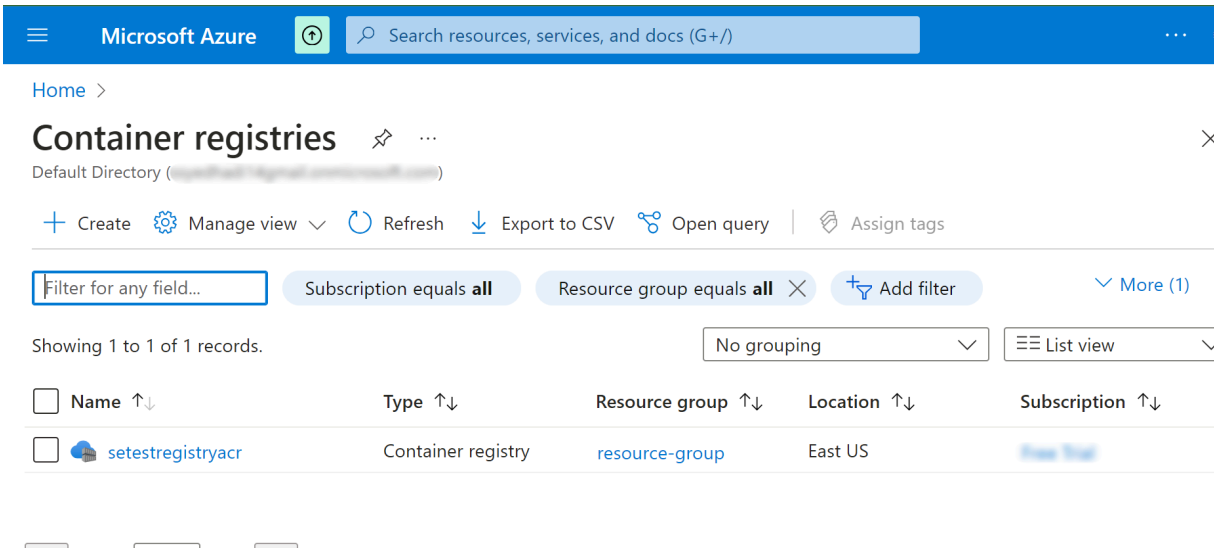
### 8.1 Azure Container Registry

#### ACR Onboarding

AccuKnox CSPM security tool scans images that are present in the onboarded Azure Container Registry and has the capability to find the risks and vulnerabilities associated with these images. The risks are identified and shown in the scan results. Users will be getting a comprehensive view of these risks and vulnerabilities in the dashboard along with their remediation.

#### 8.1.1 Steps to generate credentials for onboarding ACR

**Step 1:** Open the Azure Management Console and sign in with your Azure account credentials. Search for the **Container Registry** service in the search bar.

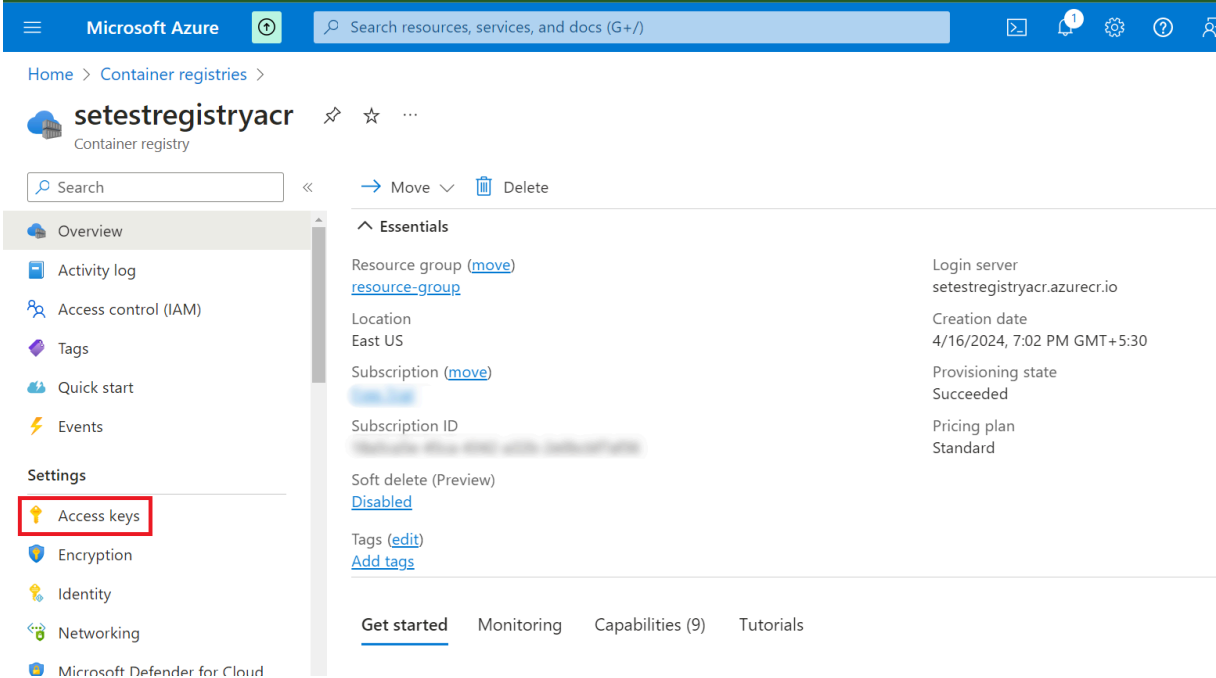


The screenshot shows the Azure Management Console interface. At the top, there is a search bar with the text "Search resources, services, and docs (G+/)". Below the search bar, the page title is "Container registries". There are several action buttons: "+ Create", "Manage view", "Refresh", "Export to CSV", "Open query", and "Assign tags". A filter bar is visible with the text "Filter for any field..." and "Subscription equals all". Below the filter bar, it says "Showing 1 to 1 of 1 records." and "No grouping". The table below has columns: Name, Type, Resource group, Location, and Subscription. The table contains one row with the following data:

Name	Type	Resource group	Location	Subscription
setestregistryacr	Container registry	resource-group	East US	

**Step 2:** Click on the name of the registry to be onboarded. In the navigation menu for the container registry, click on **Access Keys** under the Settings section.

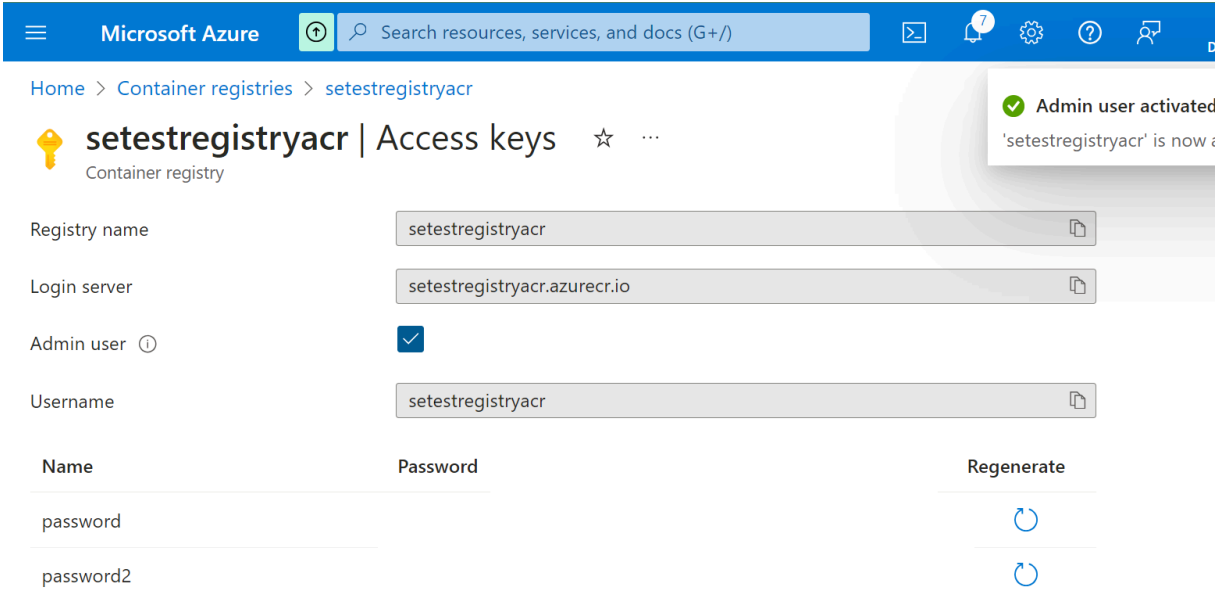




The screenshot shows the Microsoft Azure portal interface. At the top, there is a navigation bar with the Microsoft Azure logo and a search bar. Below the navigation bar, the breadcrumb trail reads 'Home > Container registries > setestregistryacr'. The main content area is divided into two sections: 'Essentials' and 'Settings'. The 'Settings' section is expanded, and the 'Access keys' option is highlighted with a red box. The 'Access keys' page displays the following information:

- Resource group:** [\(move\) resource-group](#)
- Location:** East US
- Subscription:** [\(move\)](#)
- Subscription ID:** [Redacted]
- Soft delete (Preview):** [Disabled](#)
- Tags:** [\(edit\)](#) [Add tags](#)
- Login server:** setestregistryacr.azurecr.io
- Creation date:** 4/16/2024, 7:02 PM GMT+5:30
- Provisioning state:** Succeeded
- Pricing plan:** Standard

**Step 3:** Click on the **Admin User** checkbox to activate Admin access.



The screenshot shows the Microsoft Azure portal interface. At the top, there is a navigation bar with the Microsoft Azure logo and a search bar. Below the navigation bar, the breadcrumb trail reads 'Home > Container registries > setestregistryacr'. The main content area is divided into two sections: 'Essentials' and 'Settings'. The 'Settings' section is expanded, and the 'Access keys' option is highlighted. The 'Access keys' page displays the following information:

- Registry name:** setestregistryacr
- Login server:** setestregistryacr.azurecr.io
- Admin user:**
- Username:** setestregistryacr

Below the 'Admin user' checkbox, there is a table with the following columns: Name, Password, and Regenerate.

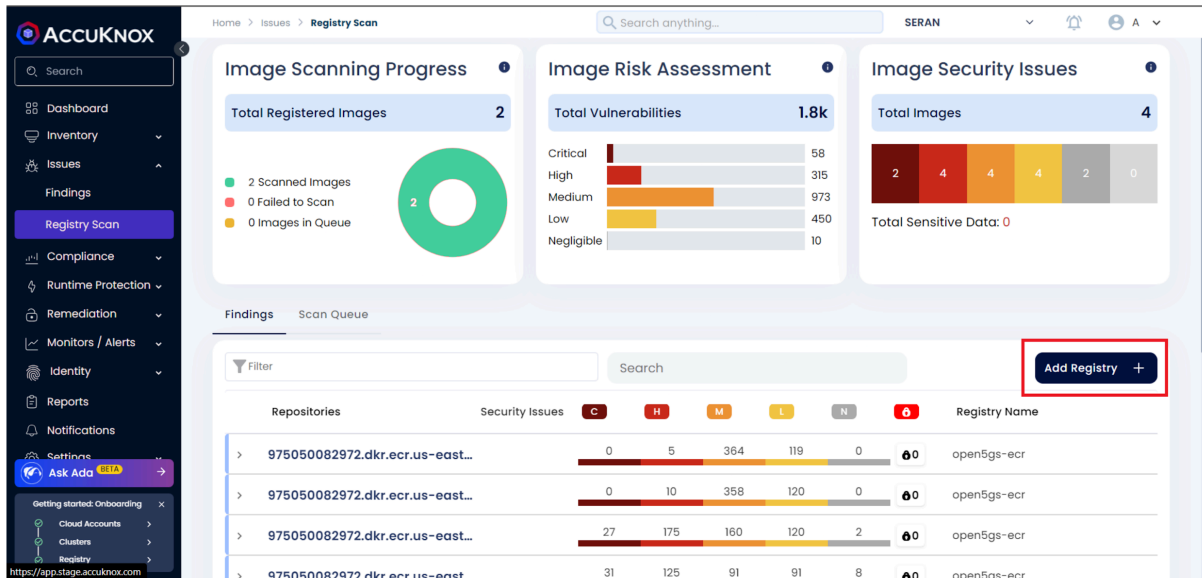
Name	Password	Regenerate
password	[Redacted]	<a href="#">Regenerate</a>
password2	[Redacted]	<a href="#">Regenerate</a>

A notification banner at the top right of the page reads: 'Admin user activated' and 'setestregistryacr' is now available for use.

Copy the generated **Login Server**, **Username** and **Password** for onboarding on AccuKnox SaaS.

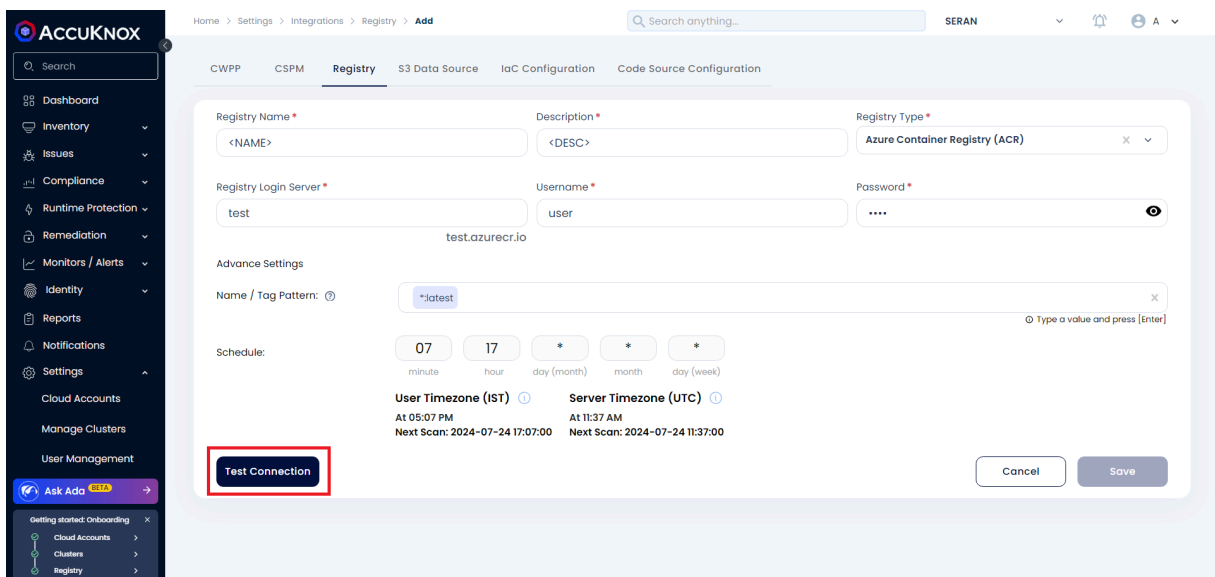
## 8.1.2 Steps to onboard the registry on AccuKnox SaaS

**Step 1:** Login to the AccuKnox SaaS and Navigate to Issues → Registry Scan. Click on **Add Registry**

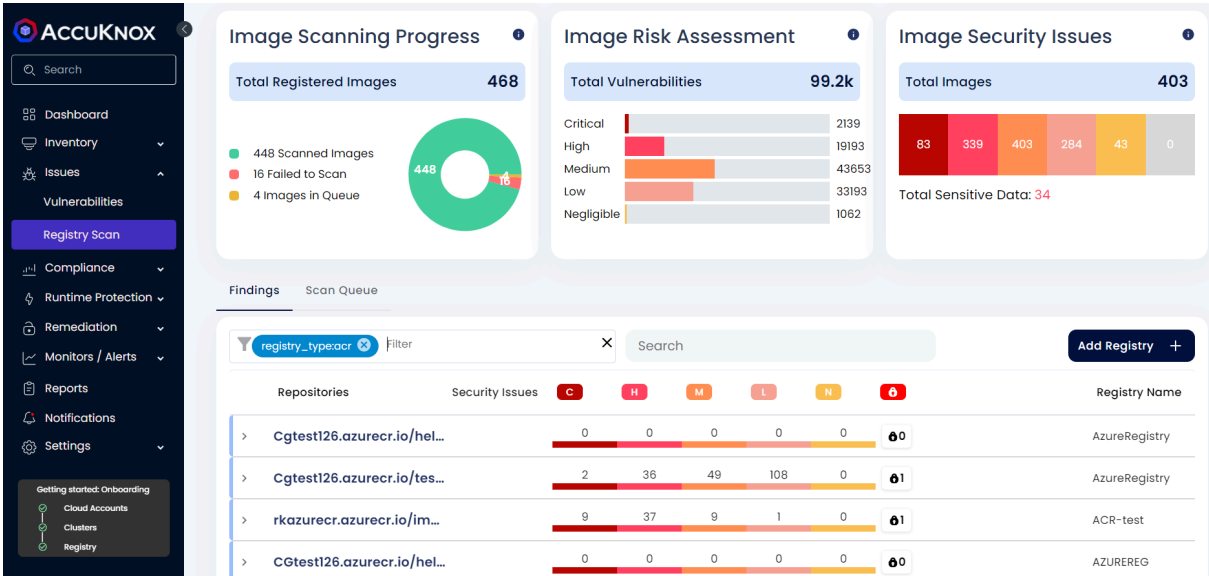


**Step 2:** Enter any Registry Name and Description. Select Registry Type as ACR and paste the Login Server, Username and Password that was copied.

Click on **Test Connection** and then click on the enabled **Save** button



**Step 3:** A popup appears that the registry is added on successful onboarding. Navigate to Issues → Registry Scan to view the scan results. The status of the scan can be checked from the **Scan Queue** tab



The screenshot displays the AccuKnox dashboard with the following sections:

- Image Scanning Progress:** Total Registered Images: 468. A donut chart shows 448 Scanned Images (green), 16 Failed to Scan (red), and 4 Images in Queue (yellow).
- Image Risk Assessment:** Total Vulnerabilities: 99.2k. A horizontal bar chart shows the distribution: Critical (2139), High (19193), Medium (43653), Low (33193), and Negligible (1062).
- Image Security Issues:** Total Images: 403. A bar chart shows counts for different severity levels: 83 (Critical), 339 (High), 403 (Medium), 284 (Low), 43 (Negligible), and 0 (Other). Total Sensitive Data: 34.

The **Findings** tab is active, showing a table of scan results:

Repositories	Security Issues	Registry Name
	C H M L N O	
> Cgtest126.azurecr.io/hel...	0 0 0 0 0 0	AzureRegistry
> Cgtest126.azurecr.io/tes...	2 36 49 108 0 0	AzureRegistry
> rkazurecr.azurecr.io/im...	9 37 9 1 0 0	ACR-test
> CGtest126.azurecr.io/hel...	0 0 0 0 0 0	AZUREREG

## 8.2 Harbor Registry

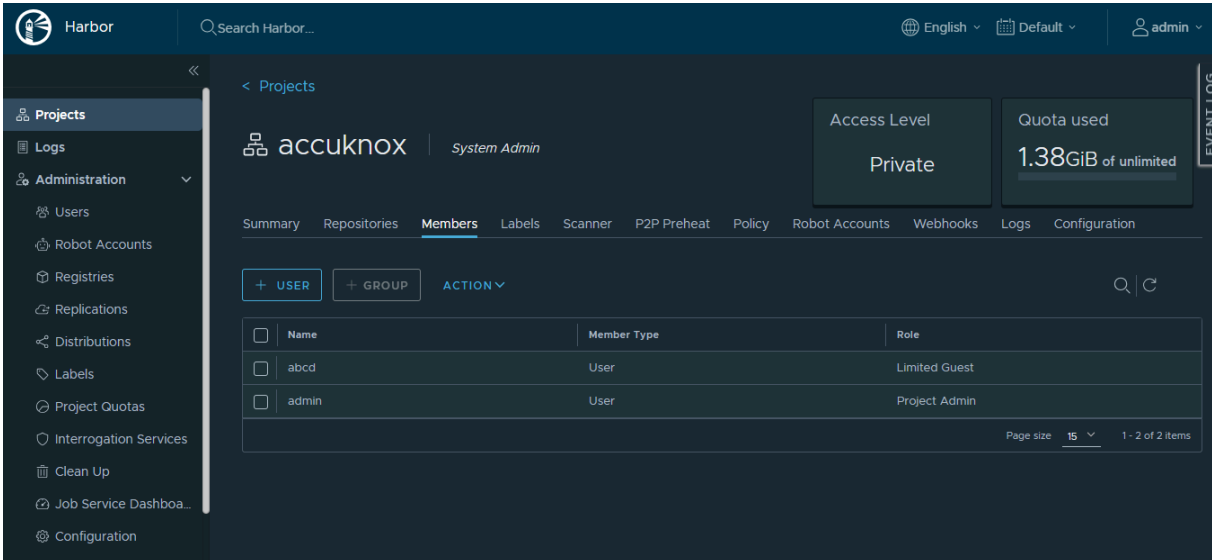
Harbor is an open source registry that secures artifacts with policies and role-based access control, ensures images are scanned and free from vulnerabilities, and signs images as trusted.

### 8.2.1 Prerequisites for Harbor Registry Onboarding in Accuknox:

In Harbor, users and groups are created by the admin.  
If you have the admin access, then login through those credentials.

To create a new user in Harbor, you can follow the steps mentioned here:  
<https://goharbor.io/docs/administration/managing-users/create-users-db/>

After creating the user, we need to add this user as a member in the Project.  
Click on “Projects”, select the Project in which you have to add the user.

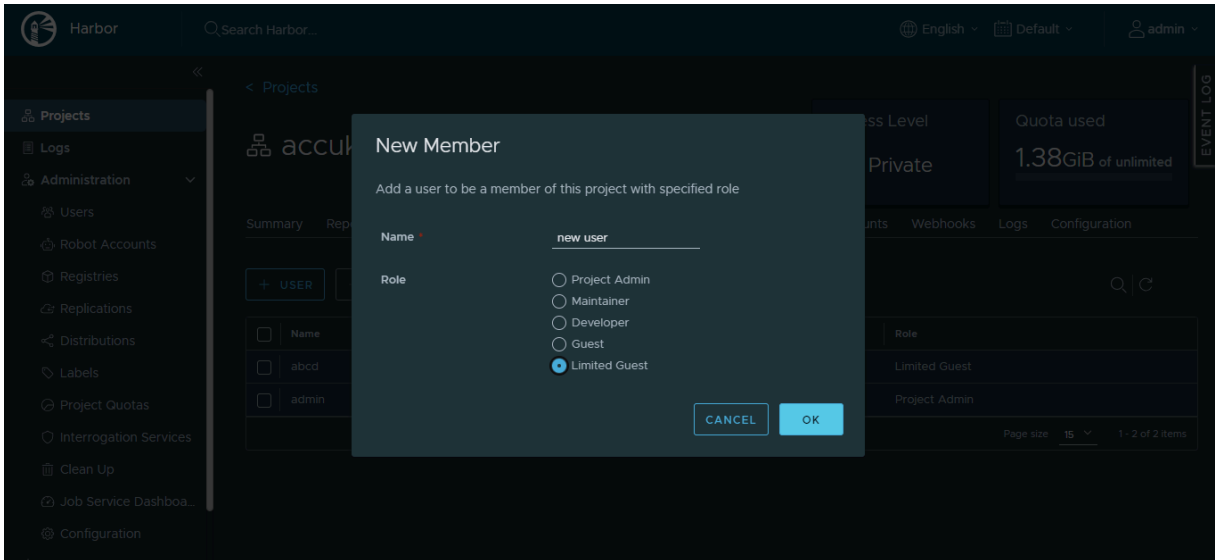


The screenshot shows the Harbor Registry interface. The top navigation bar includes the Harbor logo, a search bar, language settings (English), default settings, and the user 'admin'. The left sidebar lists various management options. The main content area is titled 'Projects' and shows the 'accuknox' project. The 'Members' tab is active, displaying a table with the following data:

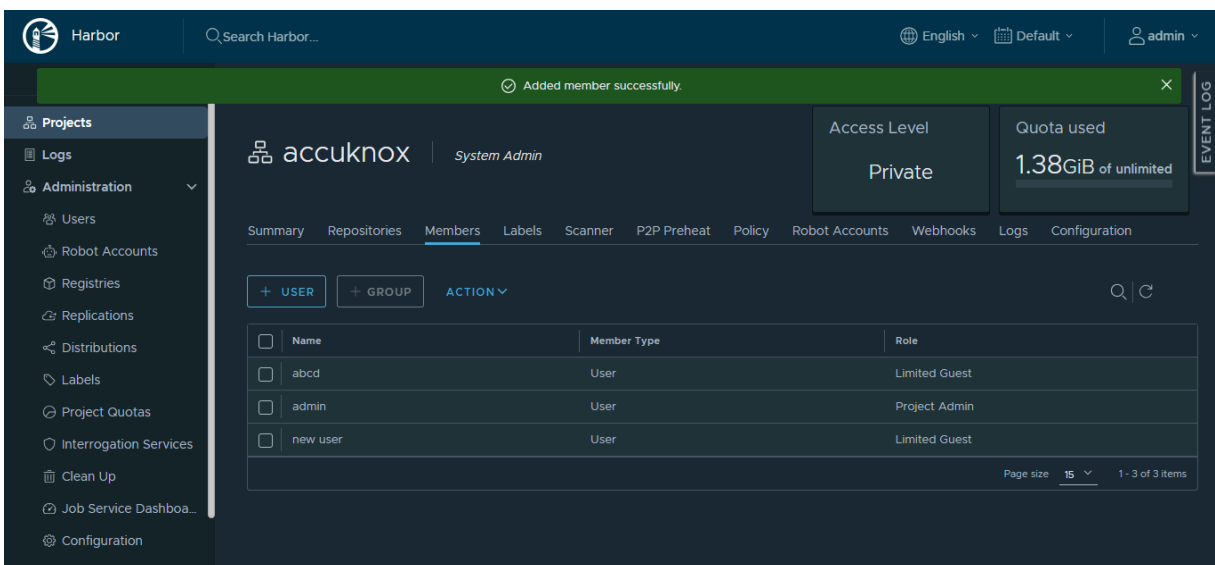
Name	Member Type	Role
abcd	User	Limited Guest
admin	User	Project Admin

Buttons for '+ USER', '+ GROUP', and 'ACTION' are visible above the table. The page size is set to 15, and there are 1 - 2 of 2 items.

Click on the “Members” tab -> +User -> give the user name and select “Limited Guest” role -> ok



We have now added the member (user) in the Project.



Now we can onboard this user in the Accuknox dashboard.

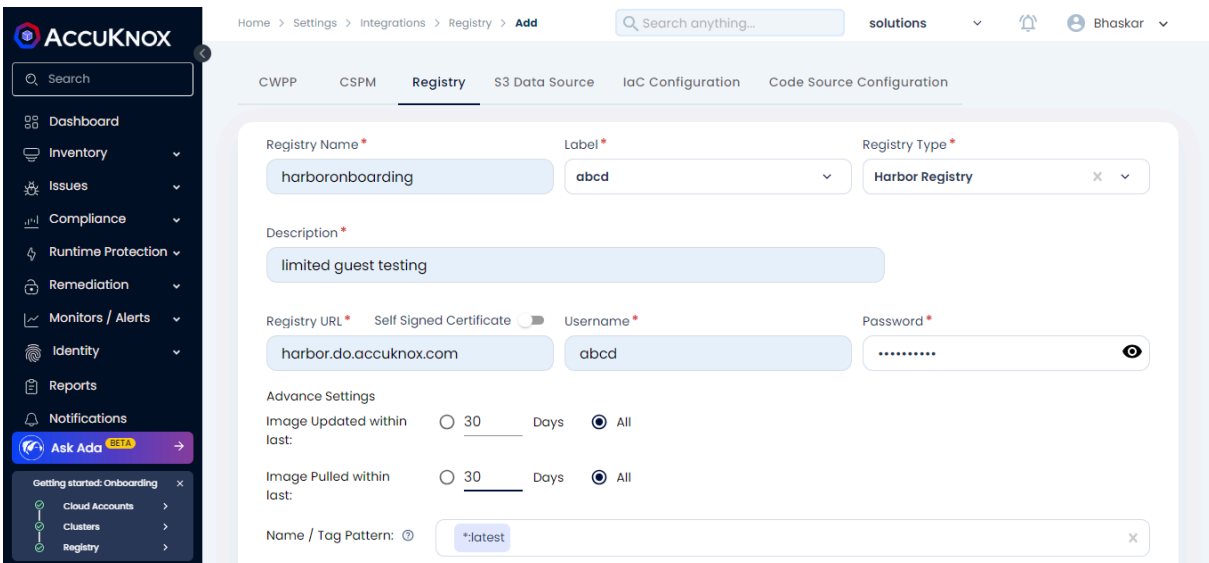
## 8.2.2 Steps to Onboard Harbor Registry on Accuknox:

In Accuknox dashboard, under Issues, click on “Registry Scan”  
Now, click on “Add Registry”



Give the registry name, select Label and select “Harbor Registry” from the Registry type dropdown.

Then, paste the Registry URL and provide the user credentials.

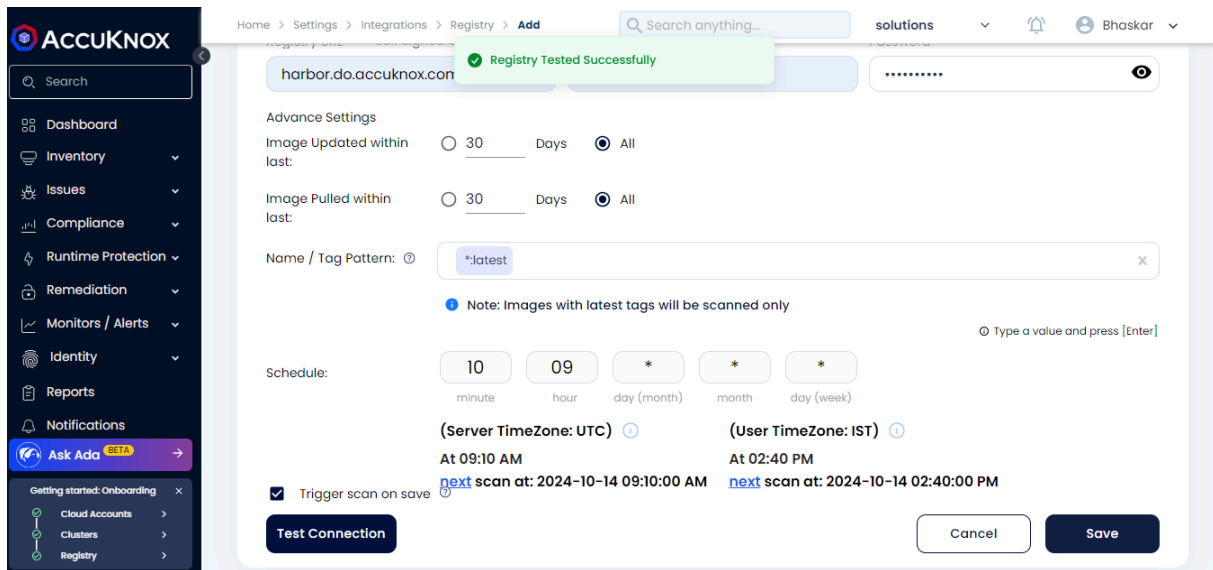


Provide the Tag pattern, and schedule time for the scanning.

If you need to trigger the scan after saving then click on the “Trigger scan on save” checkbox.

After providing all the information, click on “Test Connection”, it should show “Registry Tested Successfully”.

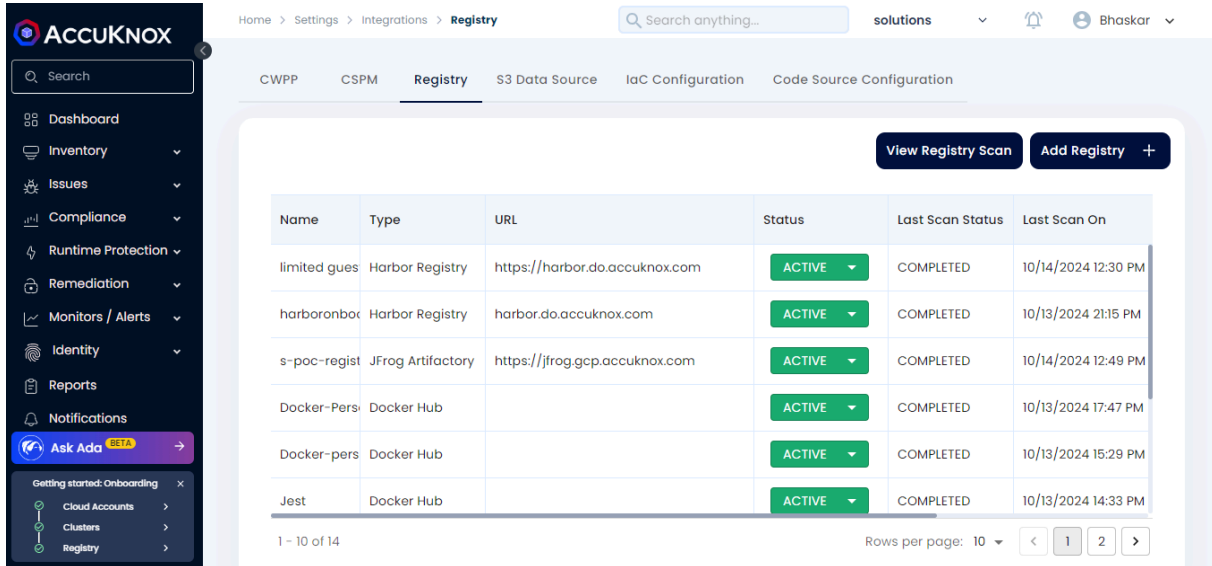
Now, click on Save.



After saving the registry, the scan will start based on the scheduled time.

To see that the scanning is completed or not, go to Settings -> Integrations -> Registry

Here, we can see the list of onboarded registries and their details.



Once the scanning is completed, we can see the scan results in the Issues -> Registry Scan

-> Under “Image Scanning Progress” pie chart, select your registry to view the progress.



To view the details of your registry, you can use the filter such as “registry\_type”, then select the “harbor” registry or you can also use the filter “registry\_name” and provide the name of your registry.



Home > Issues > Registry Scan

Search anything...

Findings Scan Queue

registry\_name:harboronboarding

Repositories	Security Issues	C	H	M	L	N	0	Registry Name
> harbor.do.accuknox...	8	20	127	0	0	0	0	harboronboarding
> harbor.do.accuknox...	10	25	66	89	1	0	0	harboronboarding
> harbor.do.accuknox...	2	11	40	11	0	0	0	harboronboarding
> harbor.do.accuknox...	18	366	1745	289	0	0	1	harboronboarding
> harbor.do.accuknox...	8	42	74	80	0	0	0	harboronboarding
> harbor.do.accuknox...	0	0	20	2	0	0	0	harboronboarding
> harbor.do.accuknox...	1	6	34	32	0	0	0	harboronboarding

By clicking on the repositories, we can get more details about the scan results.

Home > Issues > Registry Scan > Image Details

Search anything...

harbor.do.accuknox.com/accuknox/analyser:multiarch

Overview Vulnerabilities Resources Sensitive Data Scan History Layers

Architecture: amd64

Content Digest: sha256:6c87603de17d61bf71bb05bb145c0451ef30db46c083facc04e755140cb d2b58

Created: 08/04/2024 03:31 PM

Docker Digest: harbor.do.accuknox.com/accuknox/analyser@sha256:1cddb35b963f2a4278 7d308261ac610cdf3bc7a6d0be5a20c8f7926ec6a946d

Docker ID: sha256:6c87603de17d61bf71bb05bb145c0451ef30db46c083facc04e755140cb d2b58

Docker Labels: No Data Available

Docker Version: No Data Available

Environment: PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin

Operating System: linux (alpine)

Vulnerability Scan Details

harbor.do.accuknox.com/accuknox/analyser:multiarch created 72 day(s) ago

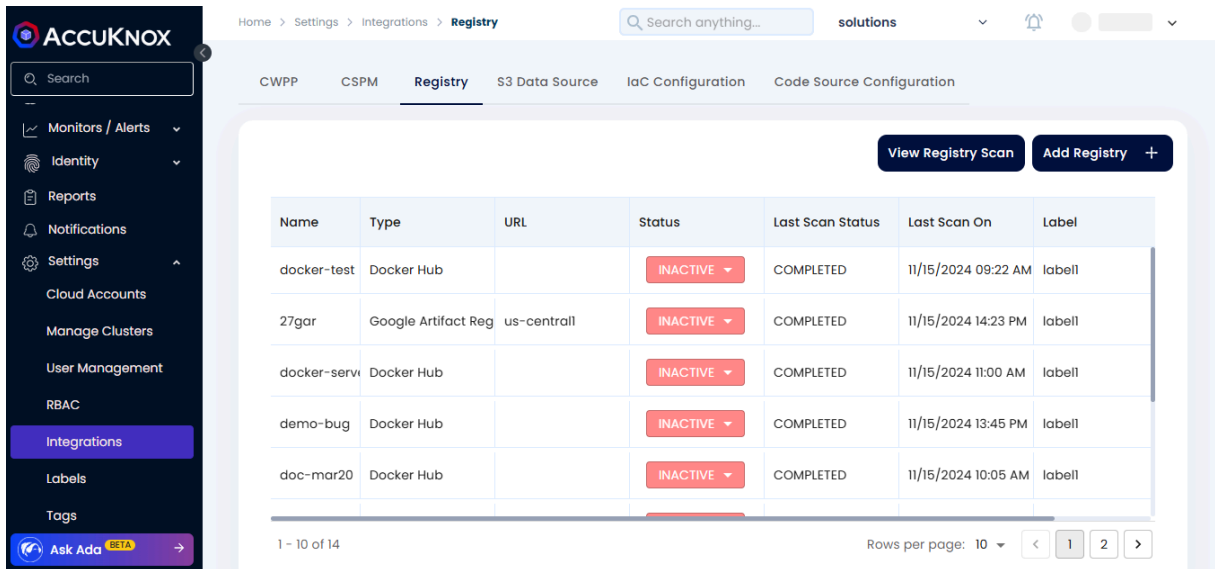
Total 38

2 5 31 0 0

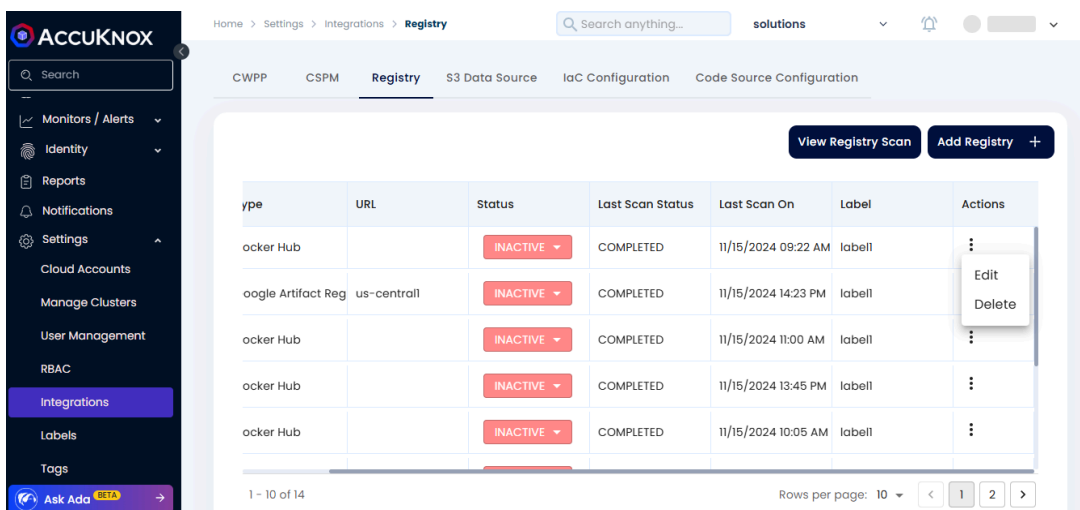
## 8.3 Deboarding a Registry

This guide outlines the steps for offboarding a registry from AccuKnox SaaS.

**Step 1:** Login to AccuKnox SaaS and Go to Settings -> Integrations -> Registry.



**Step 2:** Scroll to the left and click on the 3 dots of the registry you want to delete and click "Delete".



This will delete the Registry from AccuKnox SaaS.

## 9. AccuKnox CNAPP Dashboard Widgets

### 9.1 CWPP Widgets

In CWPP, Accuknox has 32 widgets to visualize the findings. Some of them are shown below.

#### 1. Top 5 cluster findings Widget



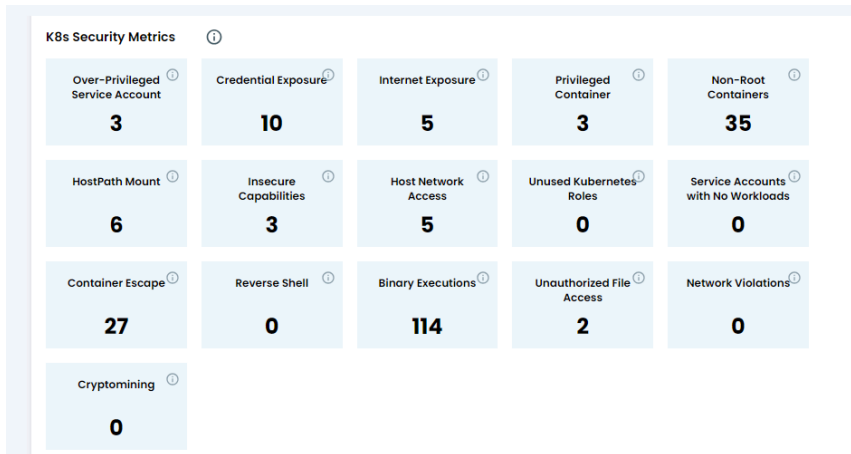
This widget provides an overview of the top findings and the number of affected resources.

#### 2. Findings by Asset Categories Widget




This widget categorizes failed findings by asset type and includes severity details, aiding users in pinpointing which asset categories have severe issues that need immediate attention.

### 3. K8S Security Metrics Widgets




This widget highlights key security metrics related to misconfigurations and vulnerabilities within your Kubernetes clusters, helping to identify and mitigate potential security risks.

### 4. Workload Alerts Widgets



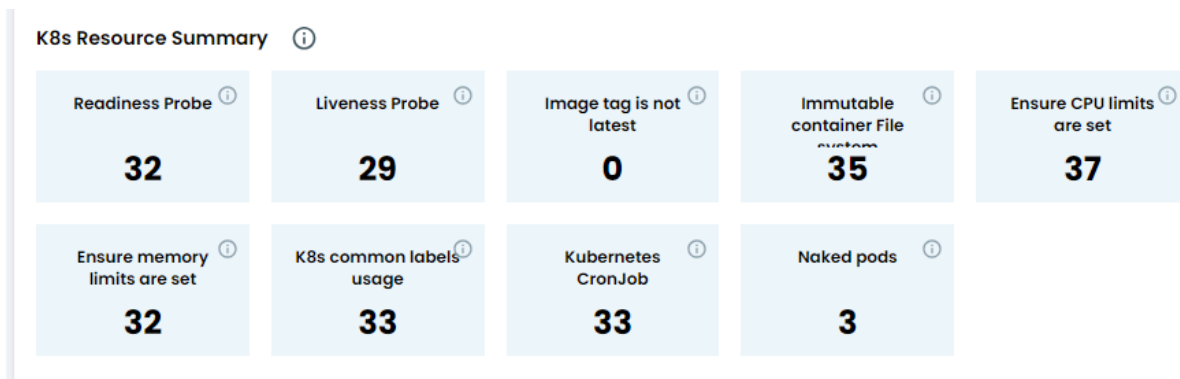
Workload Alerts shows us the alerts generated by each container or VM.

## 5. Workloads without any Policy Applied Widget



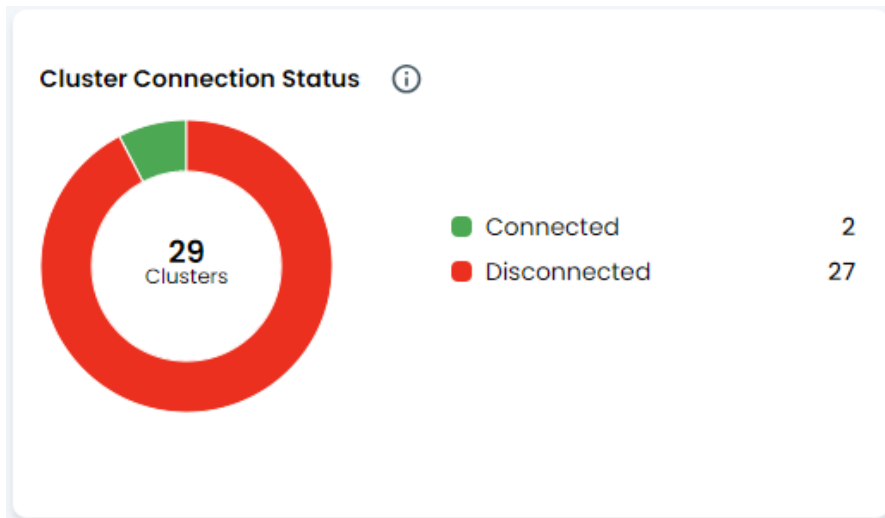
This widget shows us the total number of workloads with policies and the number of workloads policies which do not have a policy applied. The widget allows filtering based on clusters.

## 6. K8s Resource SummaryWidget



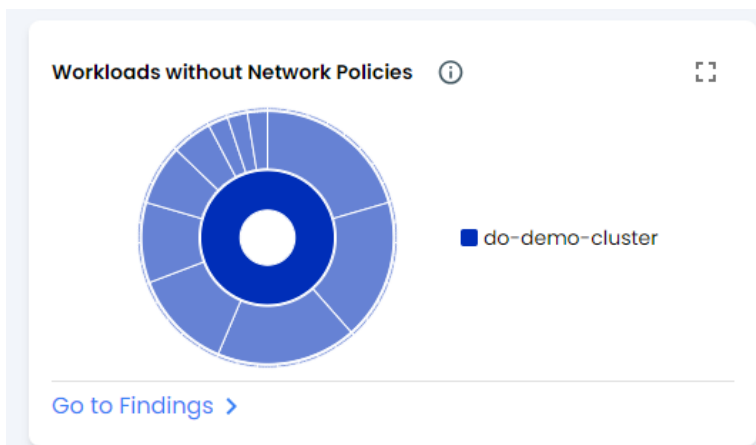
This widget displays key metrics related to resource limits, label usage, health checks, and best practices in your Kubernetes clusters.

## 7. Cluster Connection Status Widget



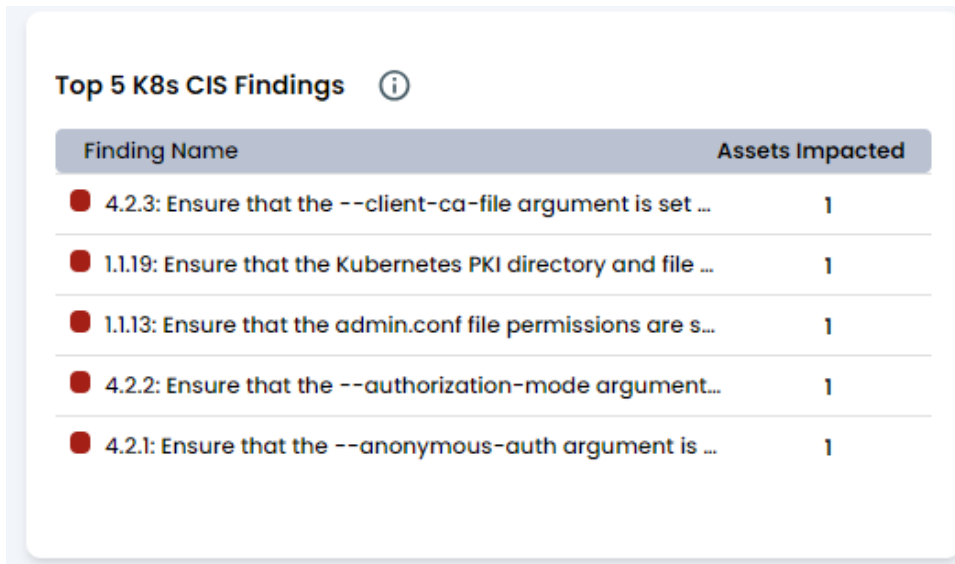
This widget will show us the connection status of Clusters which are onboarded.

## 8. Workloads without Network Policies Widget



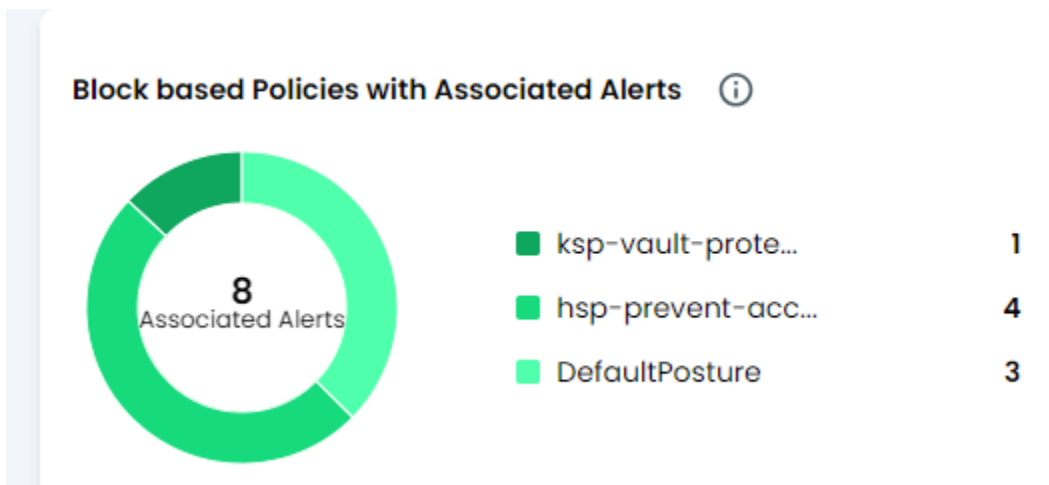
This widget displays the number of workloads that lack network policies. The aim is to help users quickly identify potential security gaps where network policies are not enforced.

## 9. Top 5 K8s CIS Findings Widget



This widget highlights the top 5 CIS benchmark related findings in your Kubernetes clusters, sorted by criticality and affected assets. It helps prioritize remediation to improve cluster security and compliance.

## 10. Block based Policies with Associated Alerts Widget

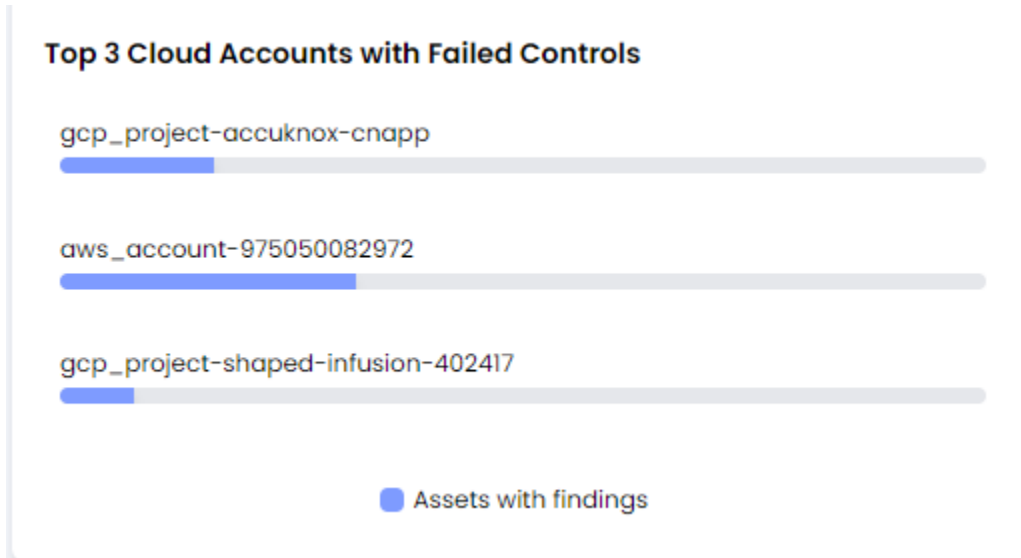


This widget shows all the block based policies which are of high severity and have alerts associated.

## 9.2 CSPM

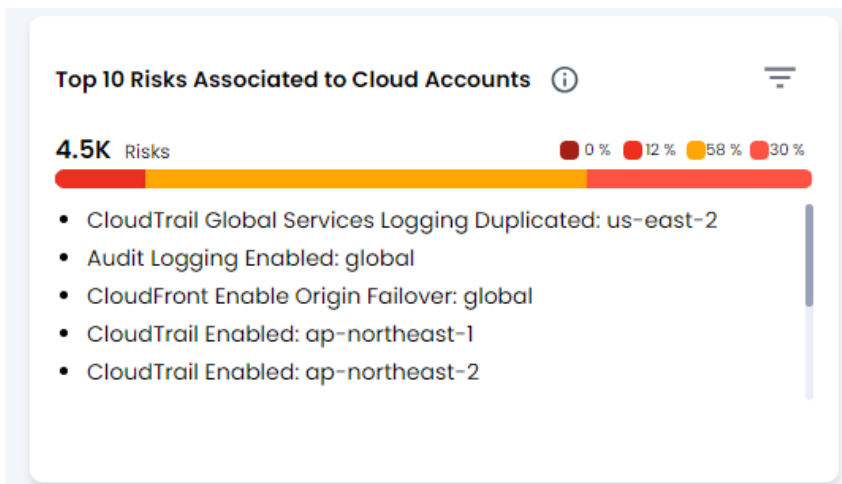
There are 5 widgets under the CSPM section

### 1. Top 3 cloud accounts with failed controls Widget



This widget will show the top 3 cloud accounts based on the highest number of failed controls.


### 2. Top 10 risk associated to cloud accounts Widget





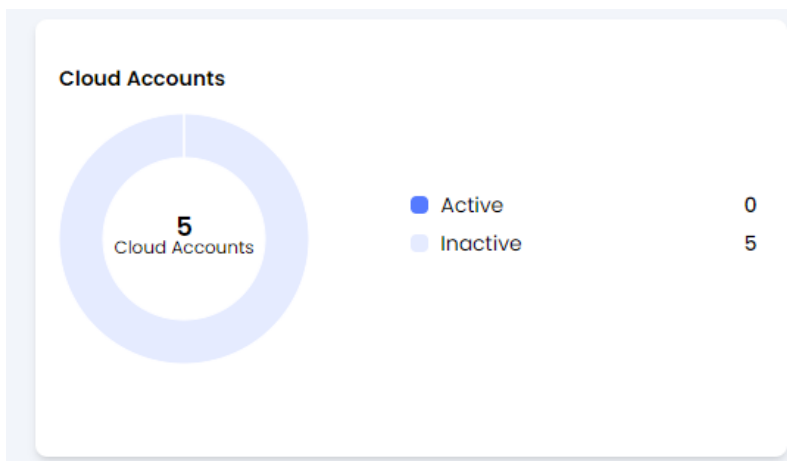
This widget assesses and prioritizes risks associated with IAM policies, S3 bucket, security groups, load balancers, etc... across your cloud accounts.

### 3. Findings Trends Widget



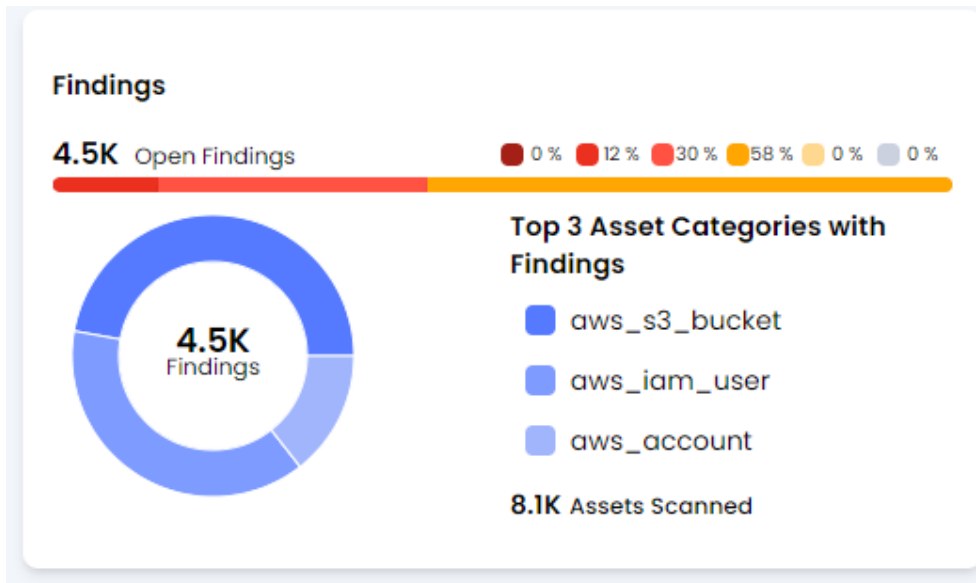
Trend analysis showing the status of findings and their changes over time in the environment.

### 4. Cloud Accounts Widget



This widget shows the number of cloud accounts on boarded on the AccuKnox platform and the status of their connection, i.e: Active/Inactive.

## 5. Findings Widget



This widget shows us the total number of Findings along with top 3 asset categories that have the highest number of Findings associated with them.

## 9.3 KIEM


The KIEM section consists of 3 widgets

### 1. Kiem Risk Assessment Widget



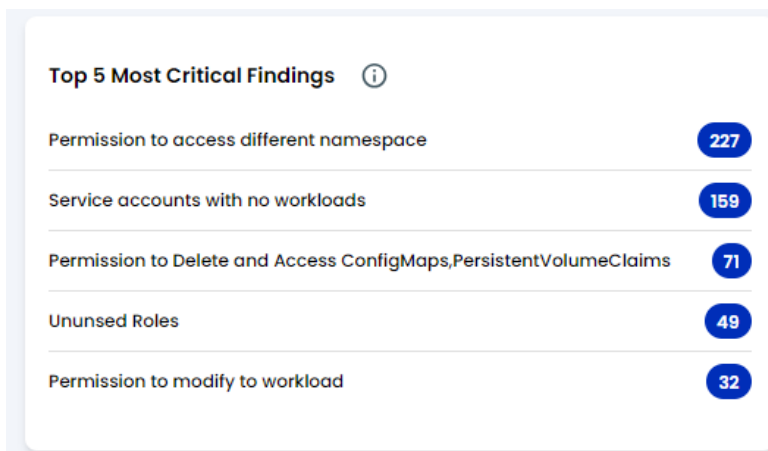
This widget shows us the distribution of KIEM findings by criticality.

## 2. KIEM Findings by Asset type Widget



This widget shows the distribution of KIEM findings by the type of assets they were identified in.

## 3. Top 5 most critical findings Widget




This widget shows the most critical findings with the highest number of occurrences or assets affected for prioritization.

## 9.4 Cloud Misconfiguration Widget

This section currently contains the following widgets

### Cloud Account Risk Assessment Widget




This widget shows the total number of checks that were performed and their result in a pie chart. This can be further filtered to include only the checks for specific cloud accounts.

## 9.5 Container images Widgets


This section consists of two widgets

### 1. Image Severity Distribution Widget



This widget shows the total number of vulnerable ./.images along with the severity level of the vulnerability identified in them. Eg. In the above image, there are 138 ./.images identified to contain a critical vulnerability

## 2. Image Risk Assessment Widget




This widget shows the total number of vulnerabilities identified in all the container images along with the severity levels.

## 9.6 Tickets Widgets

Here, the following widget exists for visualization.

### Tickets by status Widget



This widget shows us the total number of tickets generated for the findings along with their current status.

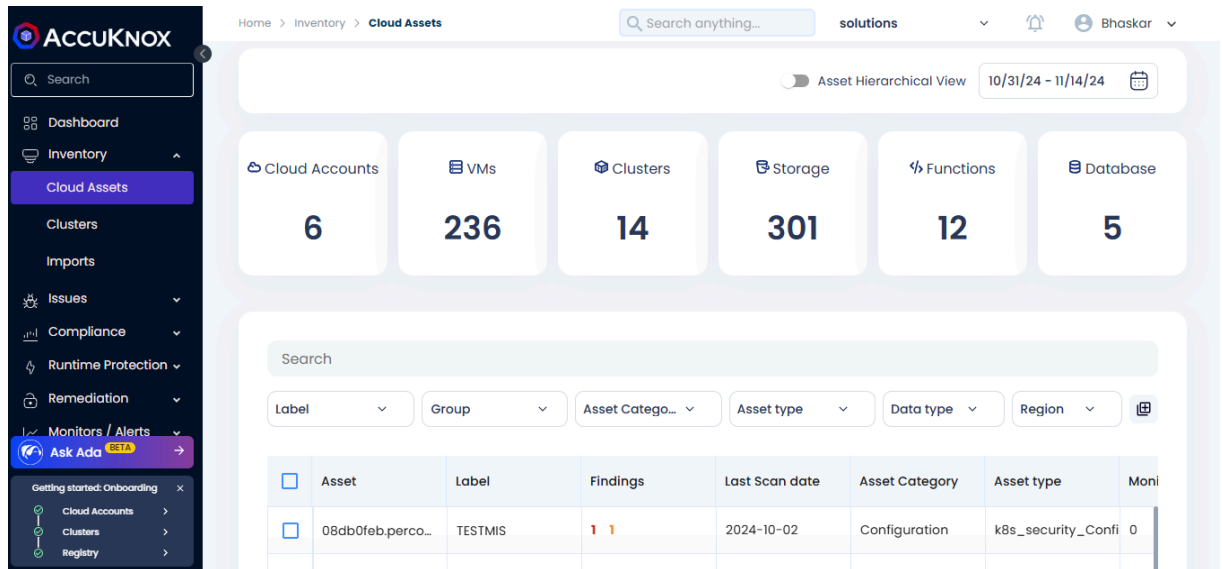
# 10. CSPM (Cloud Security Posture Management)

## 10.1 Asset Inventory

### Cloud Assets

#### 10.1.1 How to find a particular asset

- First navigate to the Cloud Assets screen under Inventory:



The screenshot shows the AccuKnox Cloud Assets interface. The left sidebar contains a navigation menu with 'Cloud Assets' highlighted. The main content area features a dashboard with six asset categories: Cloud Accounts (6), VMs (236), Clusters (14), Storage (301), Functions (12), and Database (5). Below the dashboard is a search bar and a table of assets. The table has columns for Asset, Label, Findings, Last Scan date, Asset Category, Asset type, and Moni. A single asset is visible in the table with the label 'TESTMIS' and a finding count of 1.

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
08db0feb.perco...	TESTMIS	1	2024-10-02	Configuration	kBs_security_Confi	0

- If the name of the Asset is not known but the Asset type is known, the Filter by “Asset type” can be used to filter the Assets list. The search functionality can also be used on the filtered result:

Home > Inventory > Cloud Assets

Search anything...

solutions

Bhaskar

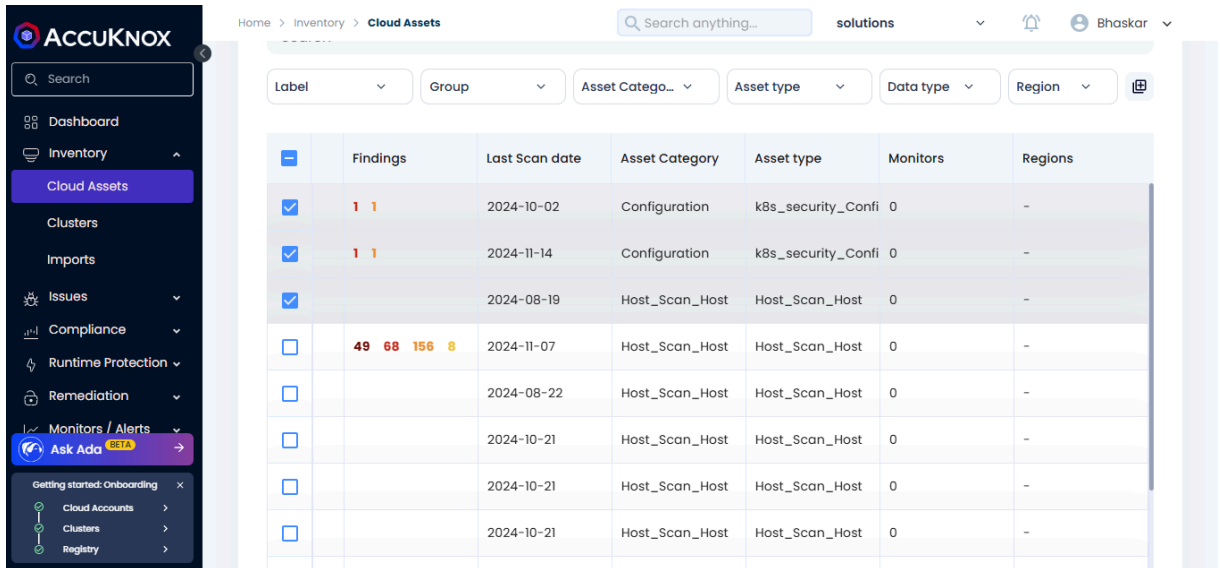
Search

Label Group Asset Catego... Asset type Data type Region

Asset	Label	Findings	Asset type	Data type	Region	Moni
<input type="checkbox"/>	08ab0feb.perco...	TESTMIS 1 1	aws_account			
<input type="checkbox"/>	08ab0feb.perco...	CMDO 1 1	nul_parent			
<input type="checkbox"/>	10.0.0.13(VAGRAN...	NessusTest	aws_cloudformation_stack			Conf 0
<input type="checkbox"/>	10.0.0.167(agent-...	WINSERVERS 49 68 156 8	aws_ebs_volume			Conf 0
<input type="checkbox"/>	10.21.0.5(PIBSRVN...	NessusTest	aws_ec2_instance			ost 0
<input type="checkbox"/>	173.201.177.205(ip...	WINSERVERS	aws_ec2_launch_template			ost 0
<input type="checkbox"/>	173.201.185.104(ip...	WINSERVERS	aws_ec2_load_balancer_listener			ost 0
<input type="checkbox"/>	173.231.229.252(v...	WINSERVERS	aws_ec2_network_interface			ost 0
			aws_ec2_network_load_balancer			
			2024-08-22	Host_Scan_Host	Host_Scan_Host	0
			2024-10-21	Host_Scan_Host	Host_Scan_Host	0
			2024-10-21	Host_Scan_Host	Host_Scan_Host	0
			2024-10-21	Host_Scan_Host	Host_Scan_Host	0

## 10.1.2 How to group assets

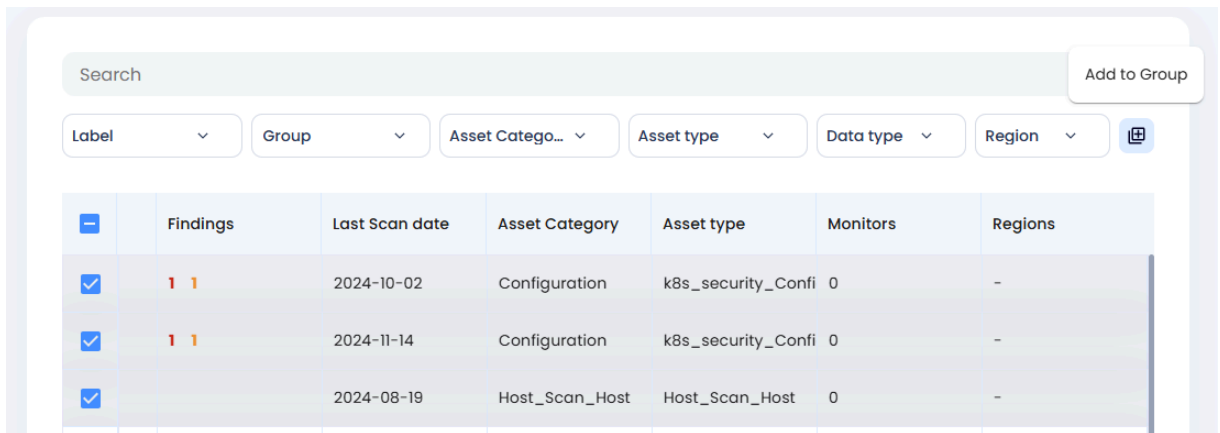
- Select the assets to be grouped in the Assets screen:



The screenshot shows the AccuKnox interface with the 'Cloud Assets' page selected in the sidebar. The main content area displays a table of assets with the following columns: Findings, Last Scan date, Asset Category, Asset type, Monitors, and Regions. Three rows are selected with blue checkmarks in the first column.

	Findings	Last Scan date	Asset Category	Asset type	Monitors	Regions
<input checked="" type="checkbox"/>	1 1	2024-10-02	Configuration	k8s_security_Confi	0	-
<input checked="" type="checkbox"/>	1 1	2024-11-14	Configuration	k8s_security_Confi	0	-
<input checked="" type="checkbox"/>		2024-08-19	Host_Scan_Host	Host_Scan_Host	0	-
<input type="checkbox"/>	49 68 156 8	2024-11-07	Host_Scan_Host	Host_Scan_Host	0	-
<input type="checkbox"/>		2024-08-22	Host_Scan_Host	Host_Scan_Host	0	-
<input type="checkbox"/>		2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-
<input type="checkbox"/>		2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-
<input type="checkbox"/>		2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-

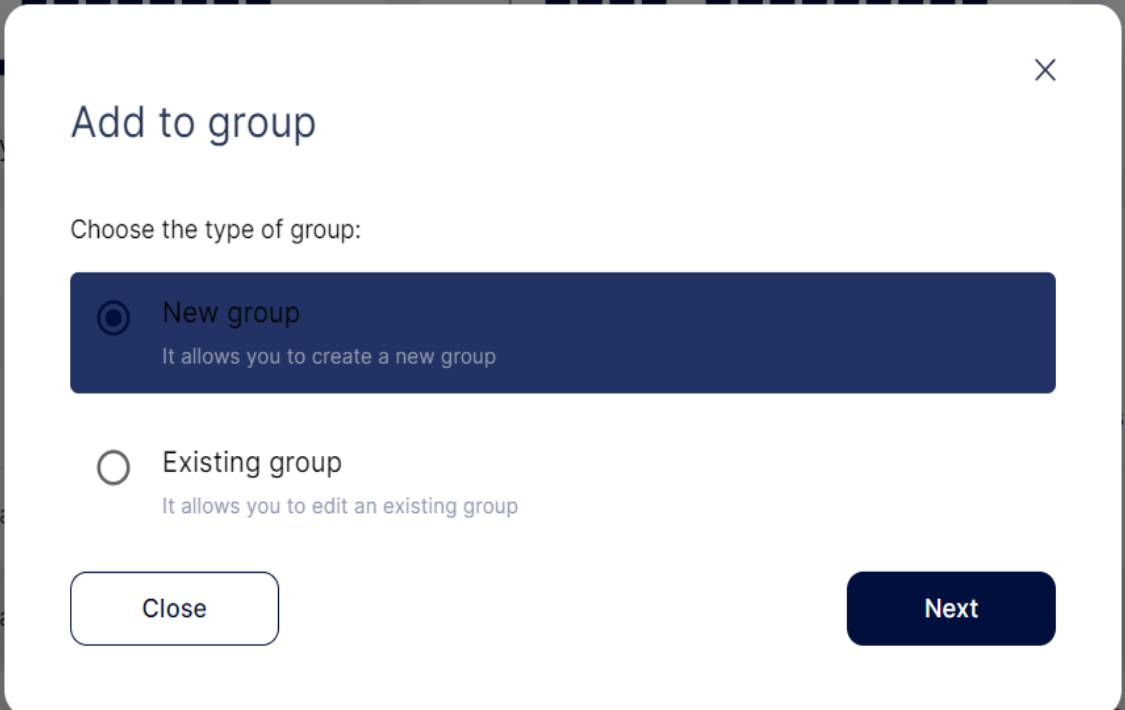
- Click on the Add to group button on the top right:



The screenshot shows the same table as above, but with the 'Add to Group' button highlighted in a light blue box in the top right corner of the table area.

- In the pop-up that follows, create a new group or add to an existing group:





×

## Add to group

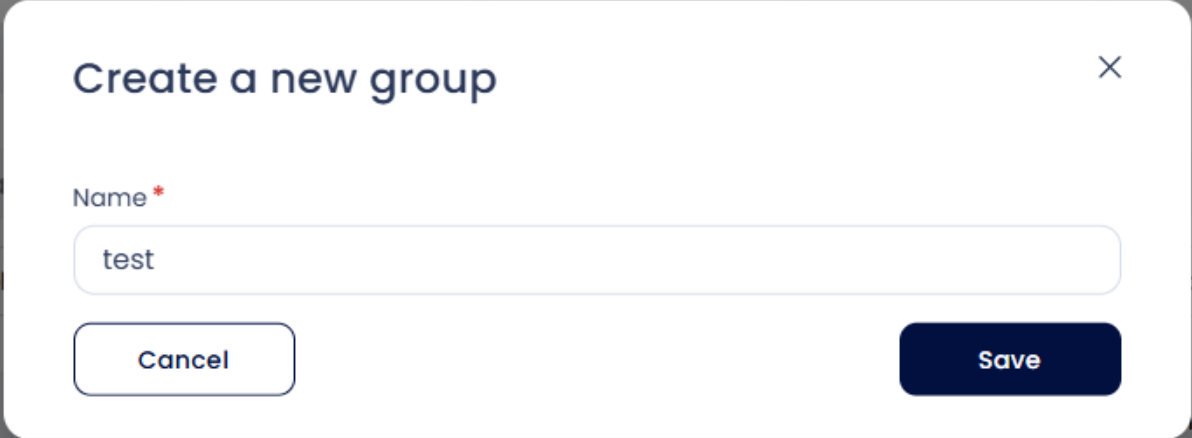
Choose the type of group:

New group  
It allows you to create a new group

Existing group  
It allows you to edit an existing group

Close Next

- After entering a name for the group or selecting an existing group, click on Save to finish adding the assets to a group:



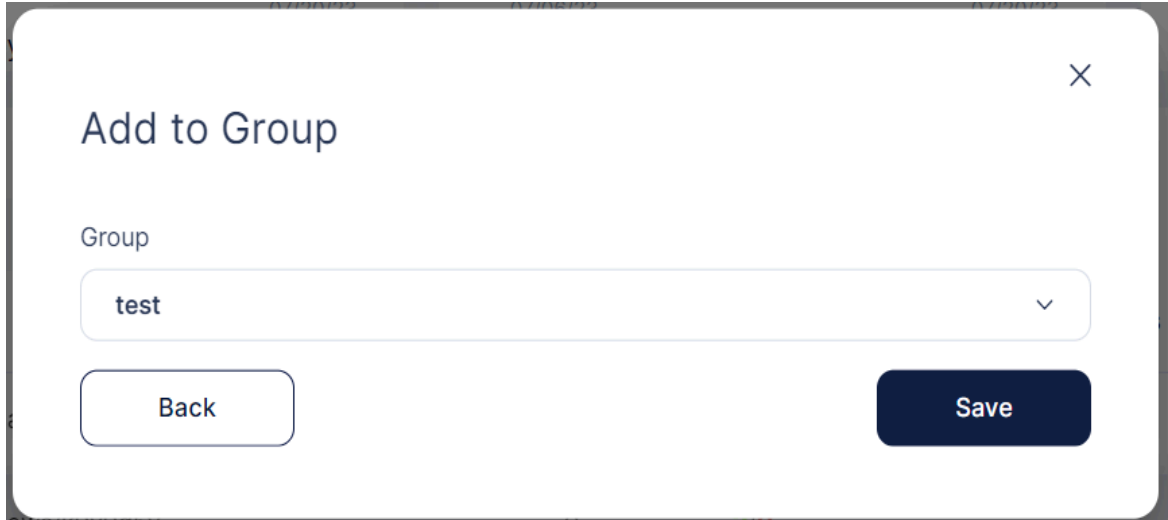
×

## Create a new group

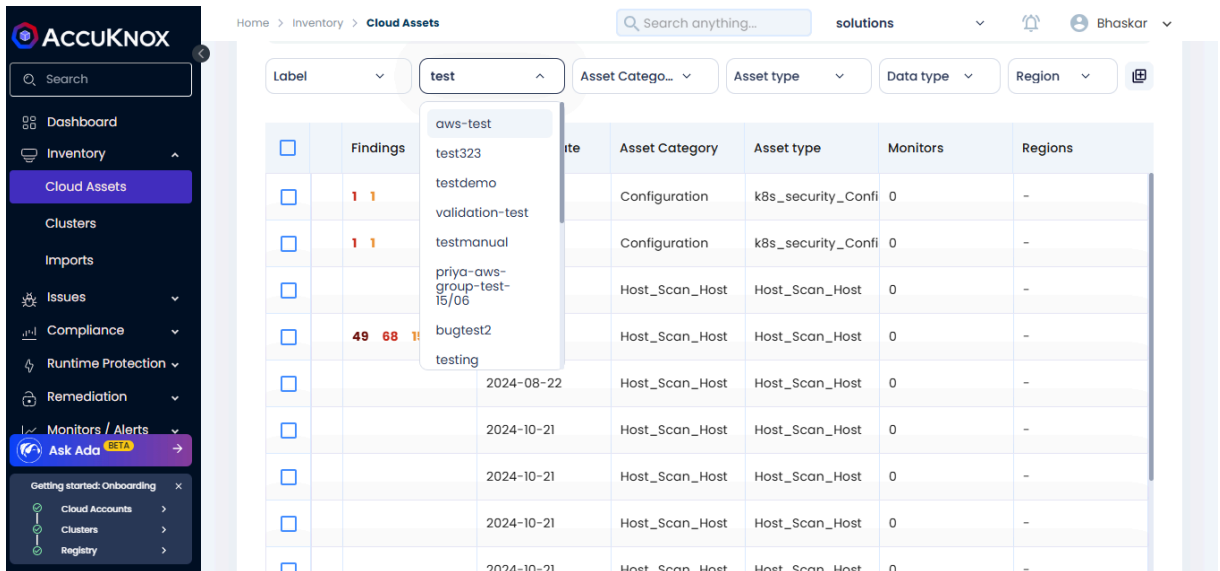
Name \*

test

Cancel Save

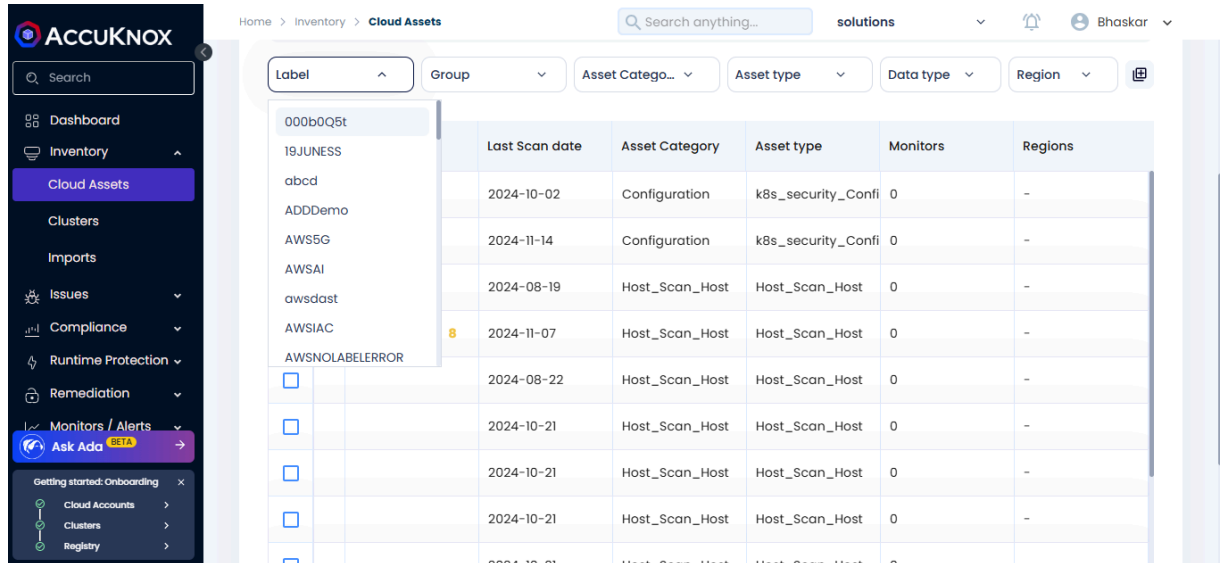


- Now, filtering by group allows us to see only the assets that were added to the group:



### 10.1.3 How to search asset by label

- To find all the assets that have a particular label, select the label from the Filter by Label drop down in the Assets screen:



The screenshot shows the AccuKnox interface for the 'Cloud Assets' section. A dropdown menu is open under the 'Label' filter, listing various labels such as '000b0Q5t', '19JUNESS', 'abcd', 'ADDDemo', 'AWS5G', 'AWSAI', 'awsdast', 'AWSIAC', and 'AWSNOLABELERROR'. The table below displays the assets filtered by the selected label.

Label	Last Scan date	Asset Category	Asset type	Monitors	Regions
abcd	2024-10-02	Configuration	k8s_security_Confi	0	-
AWS5G	2024-11-14	Configuration	k8s_security_Confi	0	-
AWSAI	2024-08-19	Host_Scan_Host	Host_Scan_Host	0	-
awsdast	2024-08-19	Host_Scan_Host	Host_Scan_Host	0	-
AWSIAC	2024-11-07	Host_Scan_Host	Host_Scan_Host	0	-
AWSNOLABELERROR	2024-08-22	Host_Scan_Host	Host_Scan_Host	0	-
	2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-
	2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-
	2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-
	2024-10-21	Host_Scan_Host	Host_Scan_Host	0	-

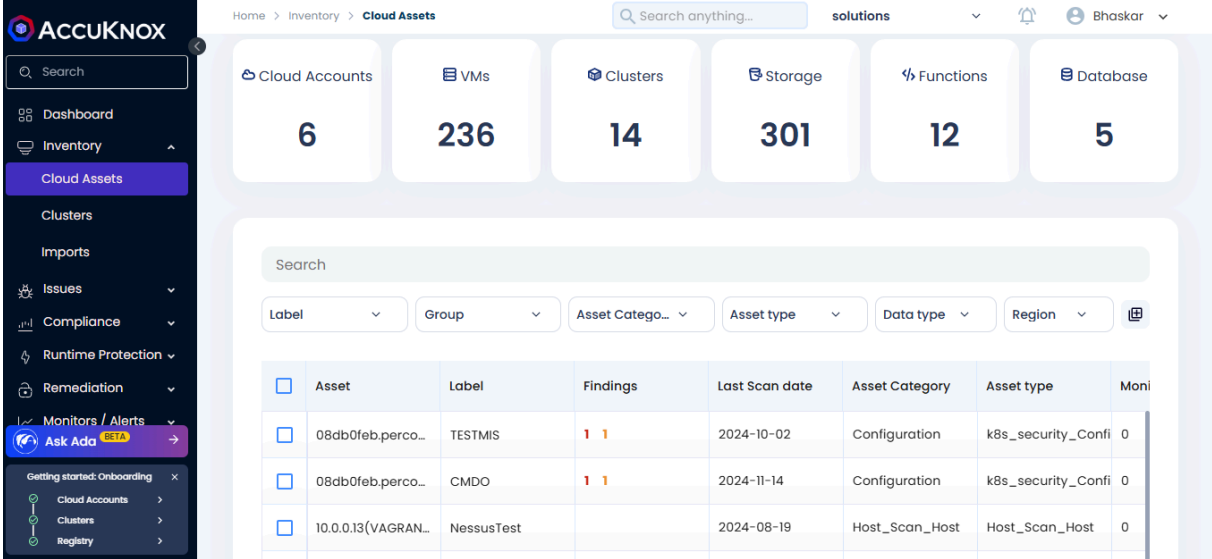
- To further refine the results, we can use the search bar or add additional filters such as Assets.

## 10.2 Misconfigurations

### 10.2.1 Where to find misconfigurations

- **Cloud Assets Page**

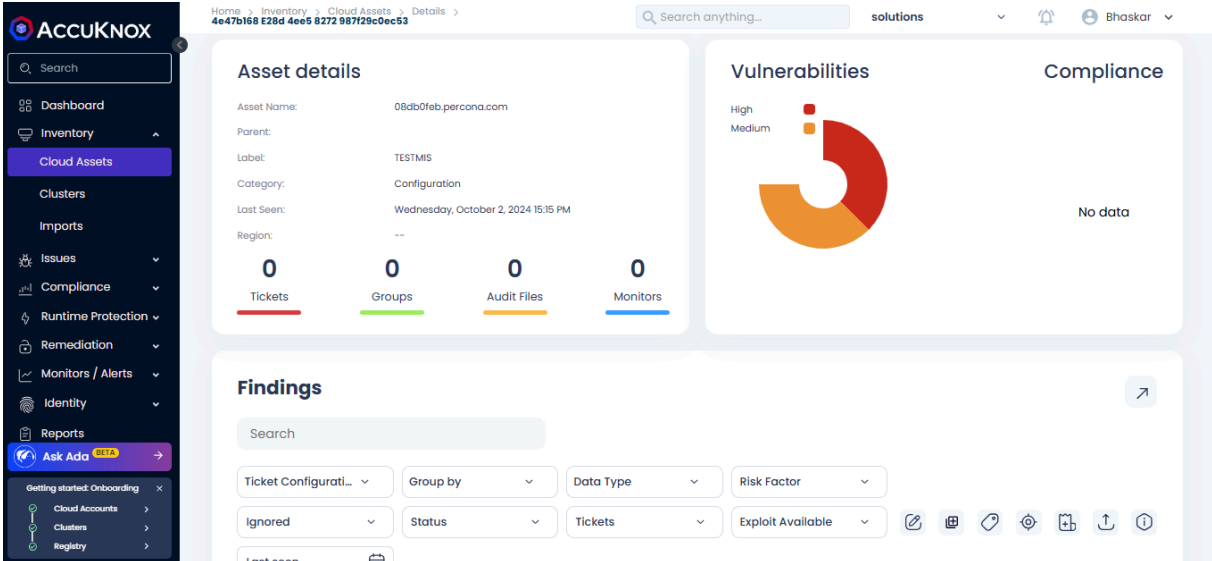
Once we have onboard the Cloud Account, we can navigate to the Inventory → Cloud Assets, here we can see the list of Assets with vulnerabilities.



The screenshot shows the AccuKnox interface for the 'Cloud Assets' page. At the top, there are navigation links for 'Home', 'Inventory', and 'Cloud Assets'. A search bar is present. Below the navigation, there are six summary cards for different asset types: Cloud Accounts (6), VMs (236), Clusters (14), Storage (301), Functions (12), and Database (5). Below these cards is a search bar and several filter dropdowns: Label, Group, Asset Category, Asset type, Data type, and Region. A table lists the assets with columns for Asset, Label, Findings, Last Scan date, Asset Category, Asset type, and Monitors.

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Monitors
08db0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0
08db0feb.perco...	CMDO	1 1	2024-11-14	Configuration	k8s_security_Confi	0
10.0.0.13(VAGRAN...	NessusTest		2024-08-19	Host_Scan_Host	Host_Scan_Host	0

From the Asset listing click any Asset for the Asset Details.



The screenshot shows the 'Asset details' page for a specific asset. The asset name is '08db0feb.percona.com'. Other details include Parent, Label (TESTMIS), Category (Configuration), Last Seen (Wednesday, October 2, 2024 15:15 PM), and Region. Below this information are four summary cards: Tickets (0), Groups (0), Audit Files (0), and Monitors (0). To the right, there is a 'Vulnerabilities' section with a donut chart showing High and Medium risk levels, and a 'Compliance' section with 'No data'. At the bottom, there is a 'Findings' section with a search bar and several filter dropdowns: Ticket Configurati..., Group by, Data Type, Risk Factor, Ignored, Status, Tickets, and Exploit Available. There are also icons for editing, deleting, and other actions.

Scroll down for the Findings, here you can see the Risk Factor for the particular Findings.

The screenshot shows the ACCUKNOX interface with the 'Findings' section active. The left sidebar contains navigation options like Dashboard, Inventory, Cloud Assets, Clusters, Imports, Issues, Compliance, Runtime Protection, Remediation, Monitors / Alerts, Identity, Reports, and Ask Ada. The main content area displays a table of findings with the following data:

Last seen	Risk Factor	Finding	Status	Ignored	Exploit Available	Tickets
2024-11-07	High	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2	Active	False	False	0
2024-11-07	Medium	OpenSSL 1.0.x < 1.0.2r Information Disclo	Active	False	False	0
2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False	0
2024-11-07	Critical	Ubuntu 20.04 LTS / 22.04 LTS / 23.04 / 23.1	Active	False	False	0
2024-11-07	Medium	OpenSSL 1.0.1 < 1.0.1k Multiple Vulnerabilit	Active	False	False	0

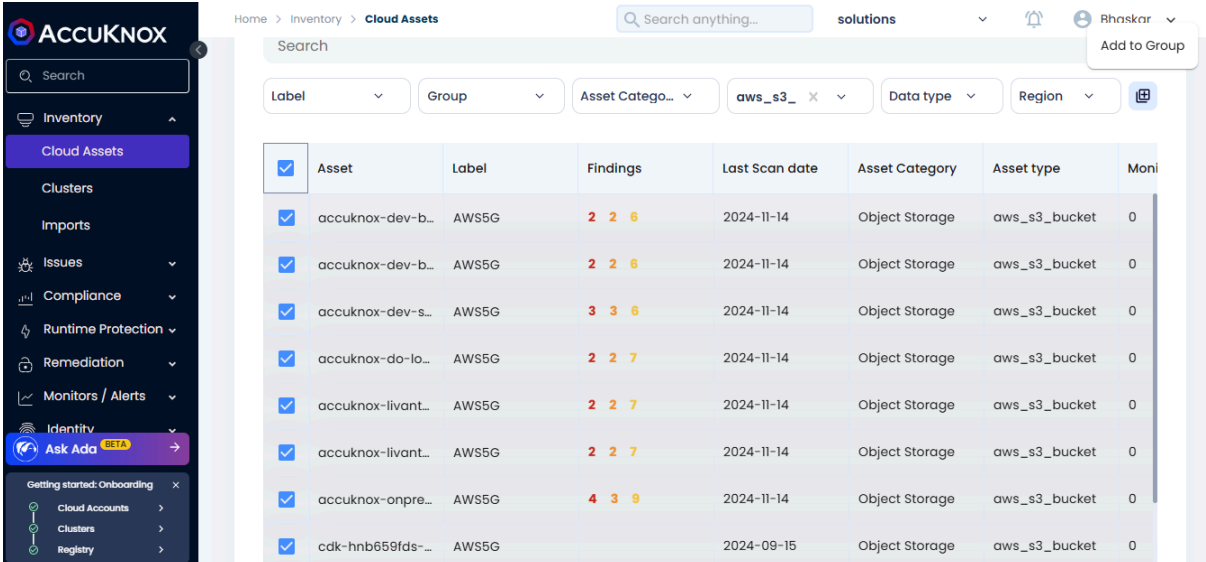
## 10.2.2 How to group by Asset and find misconfiguration

**Step1 :** In the Assets screen under Inventory, filter by Assets to view only the particular Asset type (for example s3 bucket)

The screenshot shows the ACCUKNOX interface with the 'Assets' section active. The left sidebar is the same as in the previous screenshot. The main content area displays a table of assets with the following data:

Asset	Label	Findings	Date	Category	Asset type	Monitors
08db0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0
08db0feb.perco...	CMDO	1 1	2024-11-14	Configuration	k8s_security_Confi	0
10.0.0.13(VAGRAN...	NessusTest		2024-08-19	Host_Scan_Host	Host_Scan_Host	0
10.0.0.167(agent-...	WINSERVERS	49 68 156 8	2024-11-07	Host_Scan_Host	Host_Scan_Host	0
10.21.0.5(PIBSRVN...	NessusTest		2024-08-22	Host_Scan_Host	Host_Scan_Host	0
173.201.177.205(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.201.185.104(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.231.229.252(v...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0

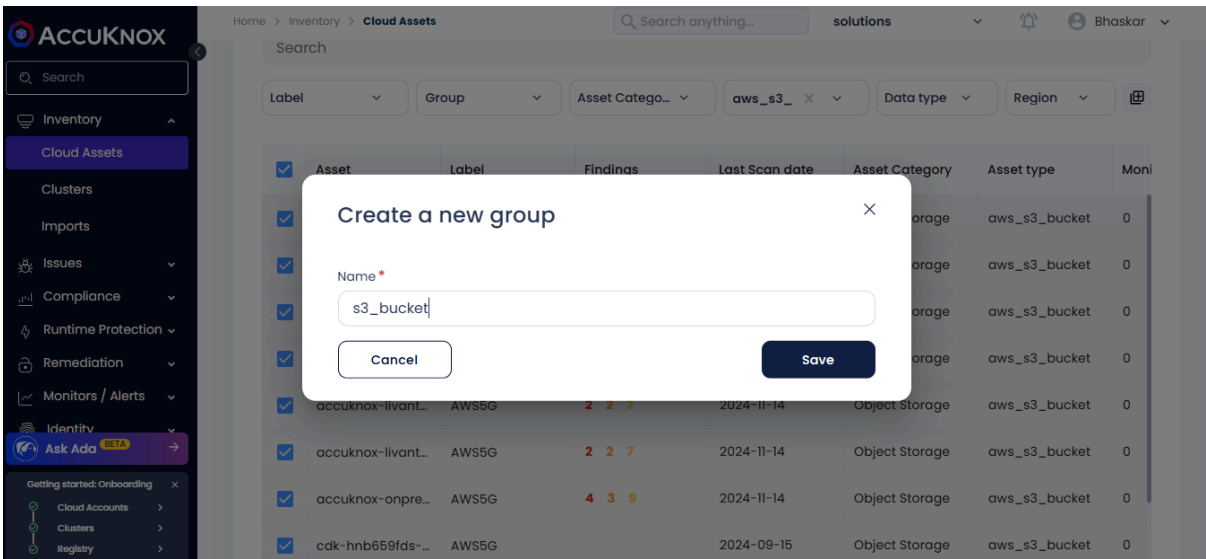
**Step2 :** Select all and Add to a group by clicking the Add to group button:



The screenshot shows the Accuknox interface with the 'Cloud Assets' page selected. A table of assets is displayed with the following columns: Asset, Label, Findings, Last Scan date, Asset Category, Asset type, and Moni. All rows in the table have their checkboxes selected. The 'Add to Group' button is visible in the top right corner of the table area.

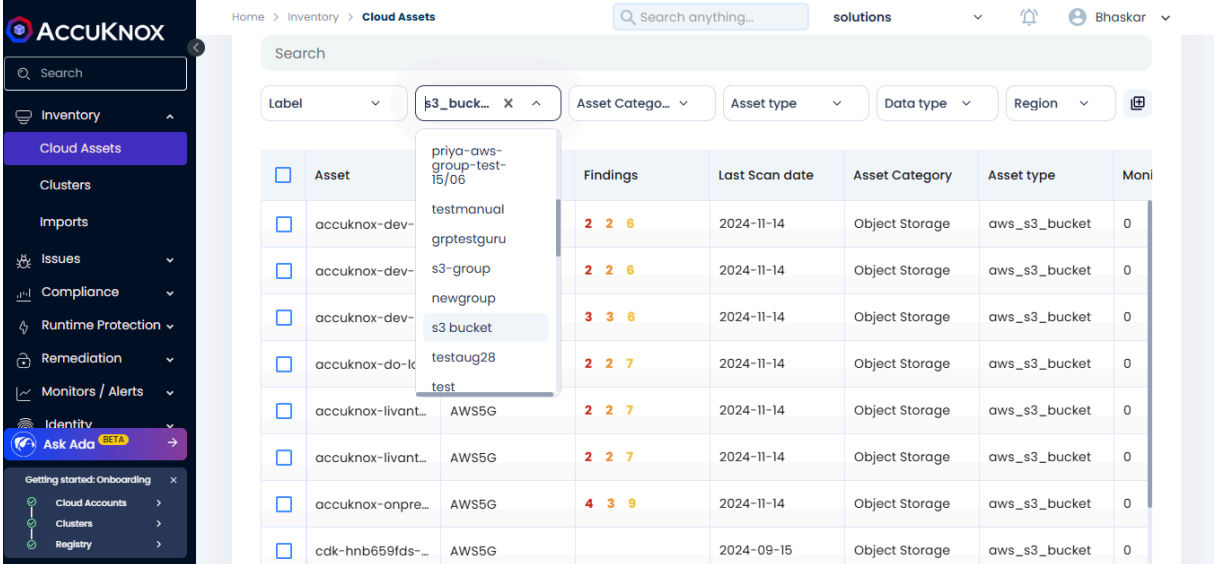
Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
accuknox-dev-b...	AWS5G	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-b...	AWS5G	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-s...	AWS5G	3 3 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-do-lo...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livan...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livan...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-onpre...	AWS5G	4 3 9	2024-11-14	Object Storage	aws_s3_bucket	0
cdk-hnb659fds...	AWS5G		2024-09-15	Object Storage	aws_s3_bucket	0

**Step3:** Click on Save



The screenshot shows the same Accuknox interface as Step 2, but with a 'Create a new group' dialog box open. The dialog box has a 'Name' field with the text 's3\_bucket' entered. There are 'Cancel' and 'Save' buttons at the bottom of the dialog box.

**Step 4 :** To view the Grouped S3 bucket details, click on Issues -> Cloud Assets, select the group that was created from the drop down:



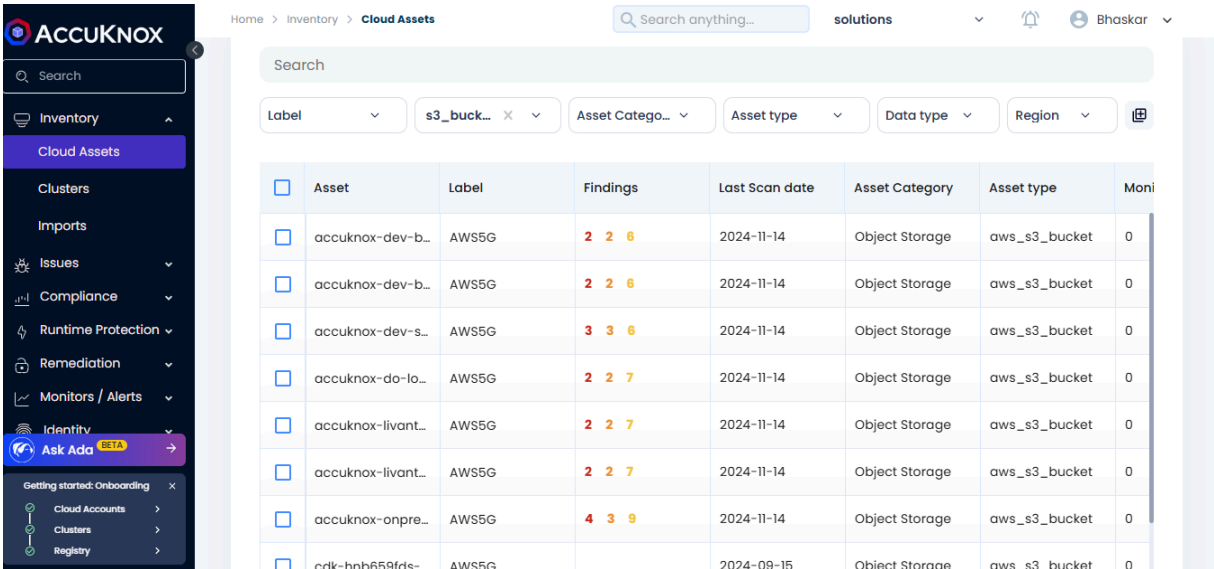
Home > Inventory > Cloud Assets

Search anything...

Label: s3\_buck... X

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
accuknox-dev-	testmanual	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-	grptestguru	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-	s3-group	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-	newgroup	3 3 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-do-l	s3 bucket	3 3 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-do-l	testaug28	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-do-l	test	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livant...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livant...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-onpre...	AWS5G	4 3 9	2024-11-14	Object Storage	aws_s3_bucket	0
cdk-hnb659fds-...	AWS5G		2024-09-15	Object Storage	aws_s3_bucket	0

**Step 5:** Select the Group, the list of s3 buckets with any misconfigurations associated with them can be seen



Home > Inventory > Cloud Assets

Search anything...

Label: s3\_buck... X

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
accuknox-dev-b...	AWS5G	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-b...	AWS5G	2 2 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-dev-s...	AWS5G	3 3 6	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-do-lo...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livant...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-livant...	AWS5G	2 2 7	2024-11-14	Object Storage	aws_s3_bucket	0
accuknox-onpre...	AWS5G	4 3 9	2024-11-14	Object Storage	aws_s3_bucket	0
cdk-hnb659fds-...	AWS5G		2024-09-15	Object Storage	aws_s3_bucket	0

**Step 7:** Click on any of them to get more details

Home > Inventory > Cloud Assets > Details > 5498ee521c8a4fb280790fdb7b88a7f4

Search anything... solutions

**Asset details**

Asset Name: accuknox-dev-back-up-consul  
 Parent: 975060082972  
 Label: AWS5G  
 Category: Object Storage  
 Last Seen: Thursday, November 14, 2024 08:57 AM  
 Region: us-east-1

**4** Tickets    **1** Groups    **0** Audit Files    **0** Monitors

**Vulnerabilities**    **Compliance**

High  
Low  
Medium

No data

**Findings**

Search

Ticket Confli...    Group by    Data Type    Risk Factor

Similarly, we can use only the group by option to view all the misconfigurations grouped together for each Asset.

### 10.2.3 How to group by Findings

1. Go to **Issues** tab, click on **Findings** section

Home > Issues > Findings

Search anything... solutions

**Findings**    Rule Engine

Cloud Findings    Group by    Insights    Saved Filters

Search

<input type="checkbox"/>	Last seen	Assetname	Name	Message	Risk factor
<input type="checkbox"/>	2024-11-18 09:08:30	weaviate-backup-user	Users MFA Enabled: glo...	User: weaviate-backup...	Medium
<input type="checkbox"/>	2024-11-18 09:08:30	ses-smtp-user.2023021...	Access Keys Rotated: gl...	User access key 1 was L...	Medium
<input type="checkbox"/>	2024-11-18 09:08:30	velero-prom-backup	Users Password And Ke...	User has console acces...	Low
<input type="checkbox"/>	2024-11-18 09:08:30	vishnu@accuknox.com	Access Keys Last Used: ...	User access key 1: was L...	Low
<input type="checkbox"/>	2024-11-18 09:08:30	knoxorg3test172794436...	No User IAM Policies: glo...	User is using attached ...	Low

1 - 20 of 6022    Rows per page: 20    1 2 3 4 5 ... 302



## 2. Navigate to **Group by** filter and choose Findings

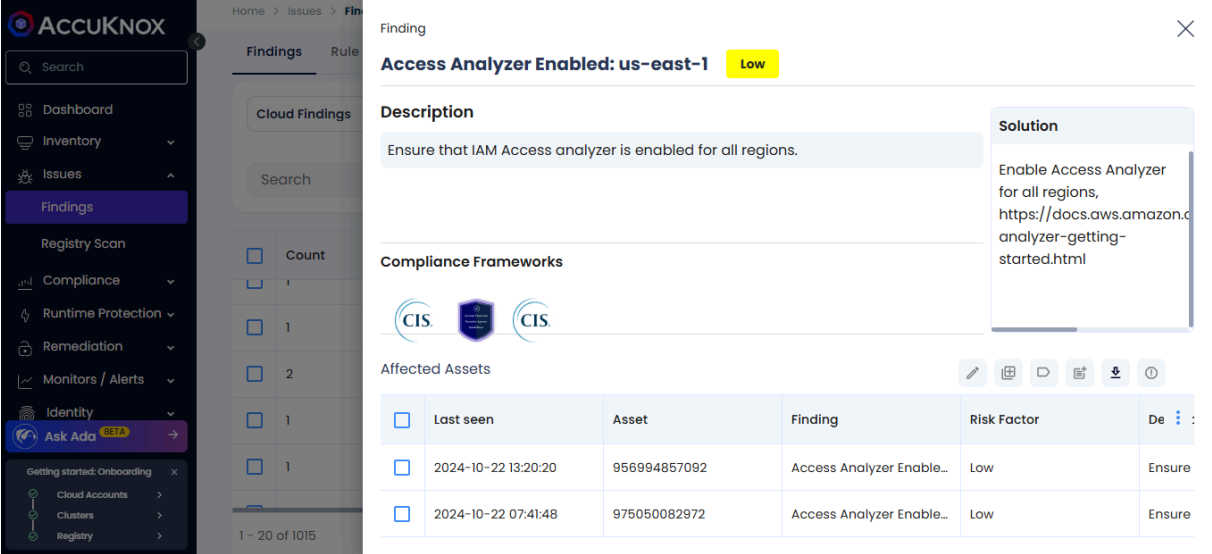
The screenshot shows the ACCUKNOX interface with the 'Findings' section selected. The 'Group by' dropdown menu is open, showing 'Finding' and 'Asset' options. The table below displays a list of findings with columns for 'Last seen', 'Assetname', 'Name', 'Message', and 'Risk factor'.

<input type="checkbox"/>	Last seen	Assetname	Name	Message	Risk factor
<input type="checkbox"/>	2024-11-18 09:08:30	weaviate-backup-user	Users MFA Enabled: glo...	User: weaviate-backup...	Medium
<input type="checkbox"/>	2024-11-18 09:08:30	ses-smtp-user.2023021...	Access Keys Rotated: gl...	User access key I was l...	Medium
<input type="checkbox"/>	2024-11-18 09:08:30	velero-prom-backup	Users Password And Ke...	User has console acces...	Low
<input type="checkbox"/>	2024-11-18 09:08:30	vishnu@accuknox.com	Access Keys Last Used: ...	User access key I: was l...	Low
<input type="checkbox"/>	2024-11-18 09:08:30	knoxorg3test172794436...	No User IAM Policies: glo...	User is using attached ...	Low

The screenshot shows the ACCUKNOX interface with the 'Findings' section selected. The 'Group by' dropdown is now set to 'Finding'. The table below displays a list of findings grouped by 'Finding', with columns for 'Count', 'Last seen', 'Assetname', 'Name', and 'Message'.

<input type="checkbox"/>	Count	Last seen	Assetname	Name	Message
<input type="checkbox"/>	1	2024-10-22 08:53:12	975050082972	Access Analyzer Enable...	
<input type="checkbox"/>	1	2024-10-22 08:03:44	975050082972	Access Analyzer Enable...	
<input type="checkbox"/>	1	2024-10-22 08:08:42	975050082972	Access Analyzer Enable...	
<input type="checkbox"/>	1	2024-10-22 07:43:29	975050082972	Access Analyzer Enable...	
<input type="checkbox"/>	1	2024-10-22 07:55:00	975050082972	Access Analyzer Enable...	

Now, you can view that similar findings are grouped. On clicking the arrow button in the findings list, you will be able to view all the assets it is found in it.



**Access Analyzer Enabled: us-east-1** Low

**Description**  
Ensure that IAM Access analyzer is enabled for all regions.

**Solution**  
Enable Access Analyzer for all regions, <https://docs.aws.amazon.com/iam/latest/userguide/access-analyzer-getting-started.html>

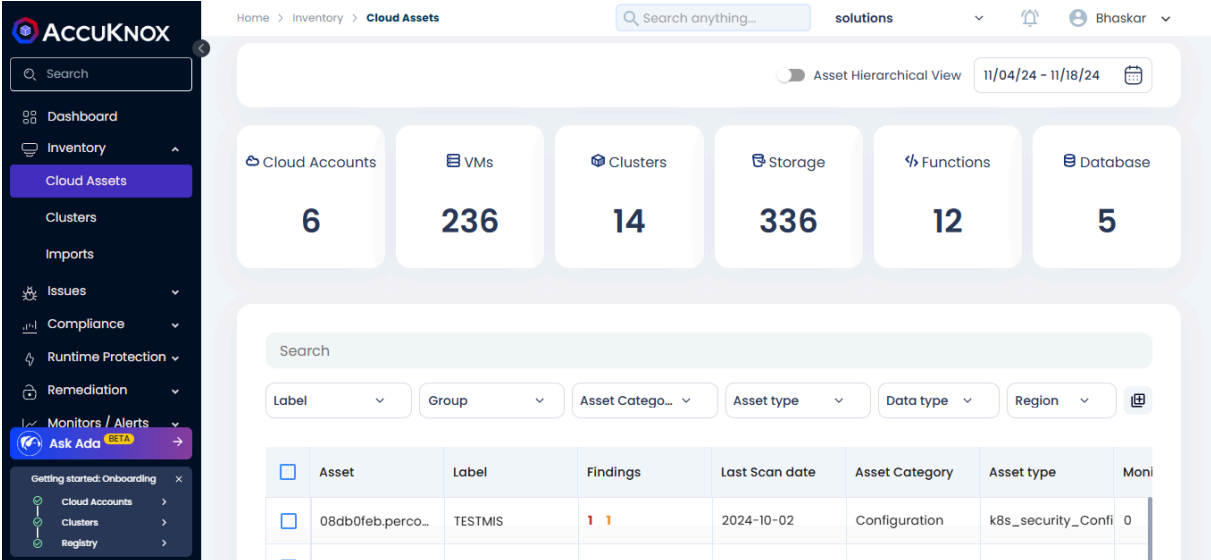
**Compliance Frameworks**  
CIS, CIS

**Affected Assets**

<input type="checkbox"/>	Last seen	Asset	Finding	Risk Factor	De
<input type="checkbox"/>	2024-10-22 13:20:20	956994857092	Access Analyzer Enable...	Low	Ensure
<input type="checkbox"/>	2024-10-22 07:41:48	975050082972	Access Analyzer Enable...	Low	Ensure

## 10.2.4 How to group by criticality and Status

### 1. Goto Inventory tab, click on **Cloud Assets** section



Home > Inventory > **Cloud Assets** Search anything... solutions Bhaskar

Asset Hierarchical View 11/04/24 - 11/18/24

Cloud Accounts	VMs	Clusters	Storage	Functions	Database
6	236	14	336	12	5

Search

Label Group Asset Catego... Asset type Data type Region

<input type="checkbox"/>	Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
<input type="checkbox"/>	08db0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0

2. Scroll down and click on the particular asset for which misconfiguration need to be viewed

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
08db0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0
08db0feb.perco...	CMDO	1 1	2024-11-18	Configuration	k8s_security_Confi	0
10.0.0.13(VAGRAN...	NessusTest		2024-08-19	Host_Scan_Host	Host_Scan_Host	0
10.0.0.167(agent-...	WINSERVERS	49 68 156 8	2024-11-07	Host_Scan_Host	Host_Scan_Host	0
10.21.0.5(PIBSRVN...	NessusTest		2024-08-22	Host_Scan_Host	Host_Scan_Host	0
173.201.177.205(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.201.185.104(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.231.229.252(v...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
1R.161.229.43(serv...	WINSFRVRS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0

3. You will land on the page as shown below. Scroll down and navigate to **Findings** sections.

**Asset details**

Asset Name: 10.0.0.167(agent-name)  
 Parent: F644504a35cb4f84911e61906c13e20  
 Label: WINSERVERS  
 Category: Host\_Scan\_Host  
 Last Seen: Thursday, November 7, 2024 18:24 PM  
 Region: --

**Vulnerabilities**

Critical  
High  
Low  
Medium

**Compliance**

No data

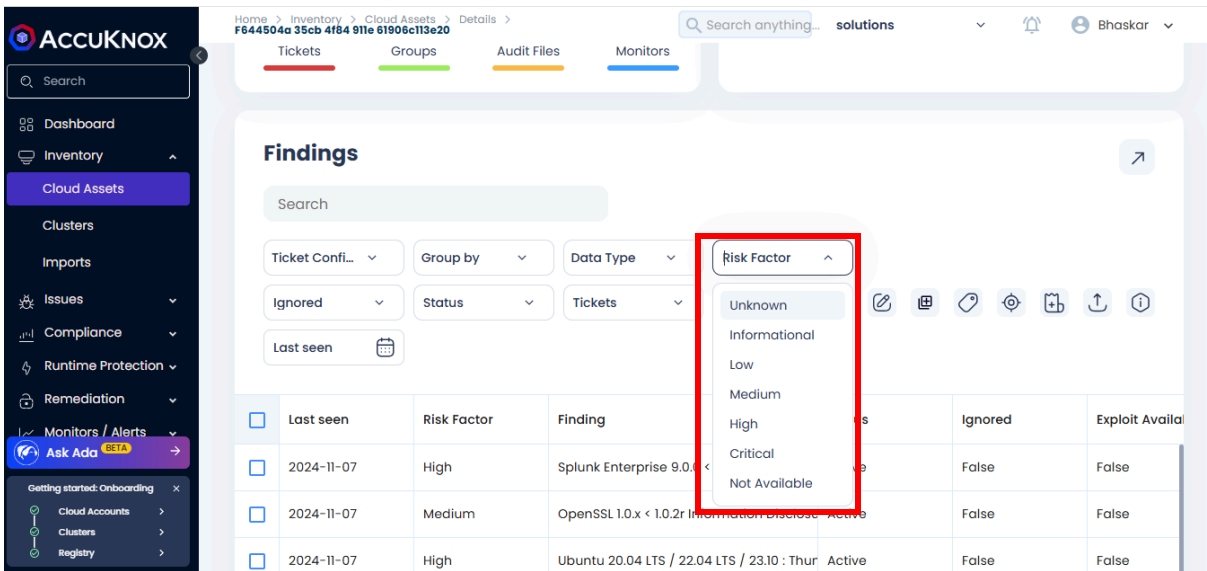
**Findings**

0 Tickets   0 Groups   0 Audit Files   0 Monitors

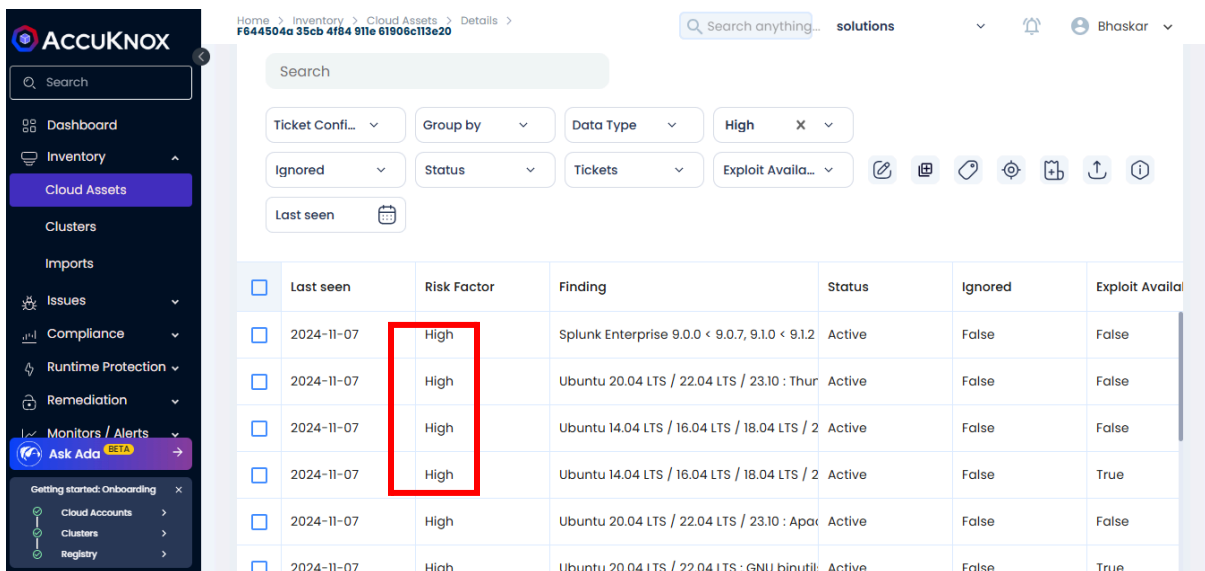
Search

Ticket Confi...   Group by   Data Type   Risk Factor

4. Navigate to the **Risk Factor** filter, and choose the severity level.



5. Now, you can find the findings as per the criticality level as shown below



6. Navigate to the **Group by** filter, and choose **Status**.

Home > Inventory > Cloud Assets > Details > F644504a35cb4f84911e61906c113e20

Search anything... solutions

Search

Ticket Conf... Ignored Last seen

Group by: Status

Data Type: High

Tickets: Exploit Availa...

	Last seen	Risk Factor	Finding	Status	Ignored	Exploit Availa
<input type="checkbox"/>	2024-11-07	High	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2	Active	False	False
<input type="checkbox"/>	2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False
<input type="checkbox"/>	2024-11-07	High	Ubuntu 14.04 LTS / 16.04 LTS / 18.04 LTS / 2	Active	False	False
<input type="checkbox"/>	2024-11-07	High	Ubuntu 14.04 LTS / 16.04 LTS / 18.04 LTS / 2	Active	False	True
<input type="checkbox"/>	2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Apar	Active	False	False
<input type="checkbox"/>	2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS : GNU binutil:	Active	False	True

Now, you can view the findings grouped by the status, such as active and accepted risk

Home > Inventory > Cloud Assets > Details > F644504a35cb4f84911e61906c113e20

Search anything... solutions

Search

Ticket Conf... Ignored Last seen

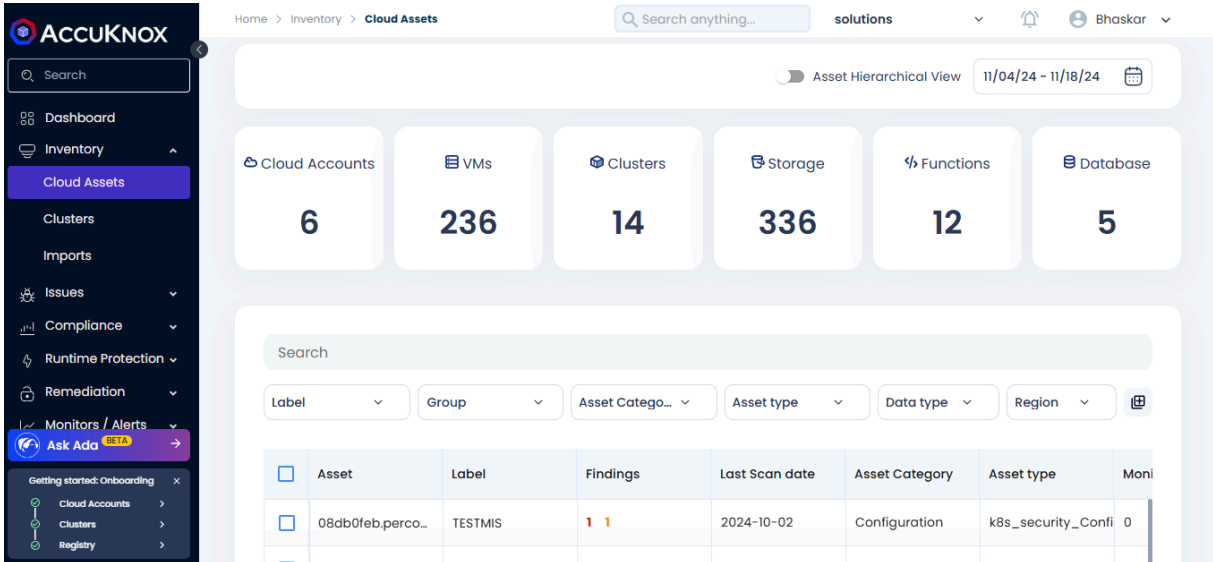
Status: High

Status: Tickets: Exploit Availa...

	Last seen	Risk Factor	Finding	Status	Ignored	
<input type="checkbox"/>	64 ▶	2024-11-07	High	Ubuntu 16.04 ESM / 18.04 ESM / 20.04 LTS	Active	False
<input type="checkbox"/>	4 ▶	2024-10-21	High	Ubuntu 22.04 LTS / 23.10 / 24.04 LTS : libar	Waiting for Verificat	False

## 10.2.5 How to create a Ticket

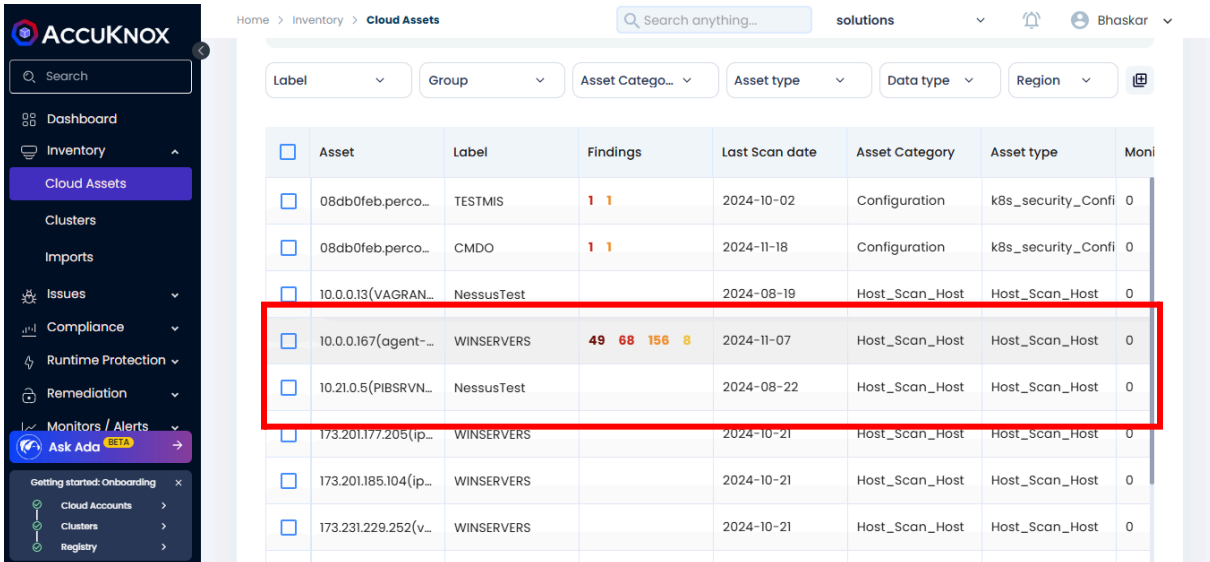
1. Go to Inventory tab, click on Cloud Assets section



The screenshot shows the AccuKnox interface with the 'Cloud Assets' section selected. The dashboard displays a summary of assets across six categories: Cloud Accounts (6), VMs (236), Clusters (14), Storage (336), Functions (12), and Database (5). Below the summary is a search bar and a table of assets.

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
08ab0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0

- a. Scroll down and click on the particular asset for which misconfiguration need to be viewed



The screenshot shows the same AccuKnox interface, but with the table of assets expanded. A red box highlights the row for the asset with ID 10.0.0.167, which has 49 findings in total (68 Critical, 156 High, 8 Medium).

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
08ab0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0
08ab0feb.perco...	CMDO	1 1	2024-11-18	Configuration	k8s_security_Confi	0
10.0.0.13(VAGRAN...	NessusTest		2024-08-19	Host_Scan_Host	Host_Scan_Host	0
10.0.0.167(agent-...	WINSERVERS	49 68 156 8	2024-11-07	Host_Scan_Host	Host_Scan_Host	0
10.21.0.5(PIBSRVN...	NessusTest		2024-08-22	Host_Scan_Host	Host_Scan_Host	0
173.201.177.205(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.201.185.104(ip...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0
173.231.229.252(v...	WINSERVERS		2024-10-21	Host_Scan_Host	Host_Scan_Host	0

- You will land on the page as shown below. Scroll down and navigate to **Findings** sections.

The screenshot shows the ACCUKNOX dashboard. The left sidebar contains navigation options: Dashboard, Inventory, Cloud Assets (highlighted), Clusters, Imports, Issues, Compliance, Runtime Protection, Remediation, Monitors / Alerts, and Ask Ada. The main content area is titled 'Asset details' for asset ID 'F844504a35cb4f84911e61906c113e20'. It displays fields for Asset Name (10.0.167(agent-name)), Parent, Label (WINSERVERS), Category (Host\_Scan\_Host), Last Seen (Thursday, November 7, 2024 18:24 PM), and Region. Below these are four metrics: Tickets (0), Groups (0), Audit Files (0), and Monitors (0). To the right, there are sections for 'Vulnerabilities' (with a donut chart) and 'Compliance' (with 'No data'). The 'Findings' section is highlighted with a red box and contains a search bar and filter options: Ticket Confi..., Group by, Data Type, and Risk Factor.

- Select the check mark behind the **Findings** for which ticket needs to be created.

The screenshot shows the 'Findings' section of the ACCUKNOX dashboard. It features a search bar and filter options: Ticket Confi..., Group by, Data Type, Risk Factor, Ignored, Status, Tickets, and Exploit Availa... Below the filters is a table of findings. The table has columns for Last seen, Risk Factor, Finding, Status, Ignored, and Exploit Availa... Two rows are highlighted with red boxes, indicating the findings for which tickets need to be created.

Last seen	Risk Factor	Finding	Status	Ignored	Exploit Availa...
<input checked="" type="checkbox"/> 2024-11-07	High	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2	Active	False	False
<input checked="" type="checkbox"/> 2024-11-07	Medium	OpenSSL 1.0.x < 1.0.2r Information Disclosu	Active	False	False
<input type="checkbox"/> 2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False
<input type="checkbox"/> 2024-11-07	Critical	Ubuntu 20.04 LTS / 22.04 LTS / 23.04 / 23.1	Active	False	False
<input type="checkbox"/> 2024-11-07	Medium	OpenSSL 1.0.1 < 1.0.1k Multiple Vulnerabiliti	Active	False	False

- Select the desired ticket configuration by which ticket will be created (Create a ticket configuration if it doesn't exist already) and click on the ticket icon.

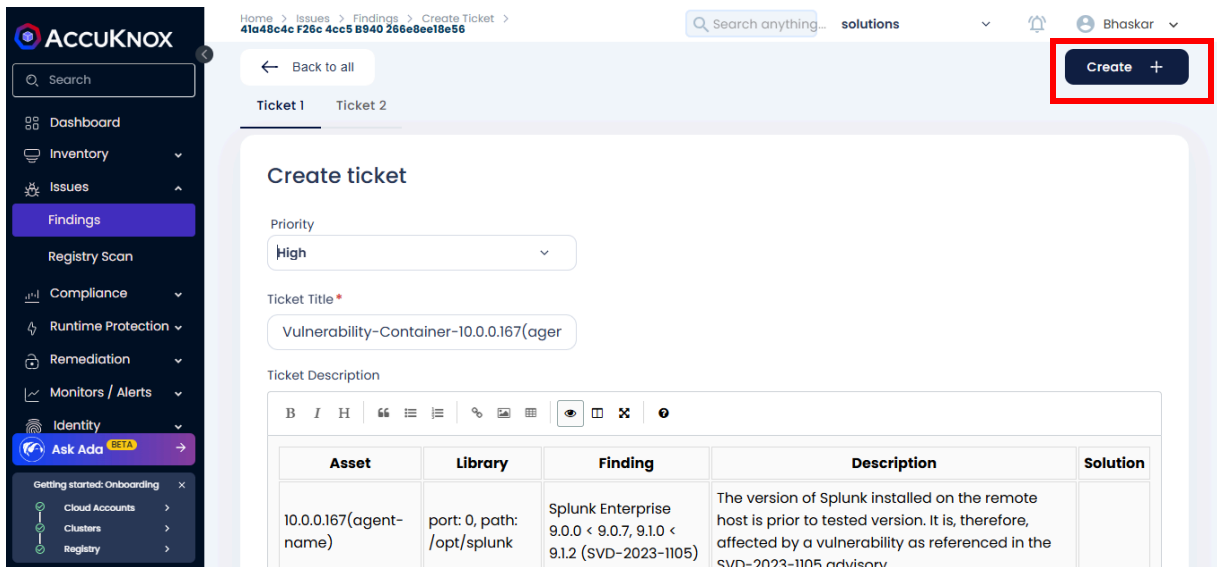
Risk Factor	Finding	Status	Ignored	Exploit Availa
High	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2	Active	False	False
Medium	OpenSSL 1.0.x < 1.0.2r Information Disclosu	Active	False	False
High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False
Critical	Ubuntu 20.04 LTS / 22.04 LTS / 23.04 / 23.1	Active	False	False
Medium	OpenSSL 1.0.1 < 1.0.1k Multiple Vulnerabiliti	Active	False	False

- Choose the **Priority** from the dropdown.

Asset	Library	Finding	Description	Solution
10.0.0.167(agent-name)	port: 0, path: /opt/splunk	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2 (SVD-2023-1105)	The version of Splunk installed on the remote host is prior to tested version. It is, therefore, affected by a vulnerability as referenced in the SVD-2023-1105 advisory.	

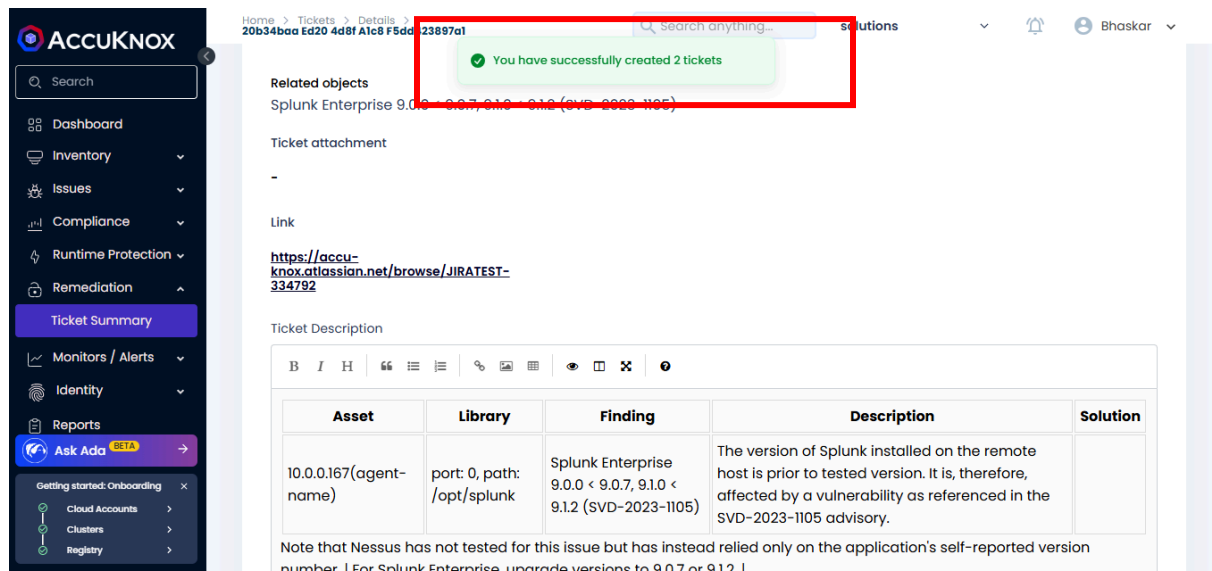
- Edit the **Ticket Title** and **Ticket Description**, as required.





7. Click on the **Create** button at the top right corner.

You can see the tickets were created successfully.



You can manage the created tickets in the **Ticket Summary** section, under the **Remediation** tab.

Home > Tickets Search anything... solutions Bhaskar

ACCUKNOX

Search

- Dashboard
- Inventory
- Issues
- Compliance
- Runtime Protection
- Remediation**
- Ticket Summary
- Monitors / Alerts
- Identity
- Reports
- Ask Ada BETA
- Getting started: Onboarding
- Cloud Accounts
- Clusters
- Registry

### Tickets by status

Total Tickets **1.4k**

opened: 1399
ongoing: 3
closed: 42
cancelled: 4

### Open Tickets by Priority

Open Tickets by Priority **1.4k**

Priority	Count	Percentage
Highest	1.2k	87%
High	173	12%
Medium	7	1%
Low	8	1%
Lowest	0	0%

### Top 5 Tickets by Age

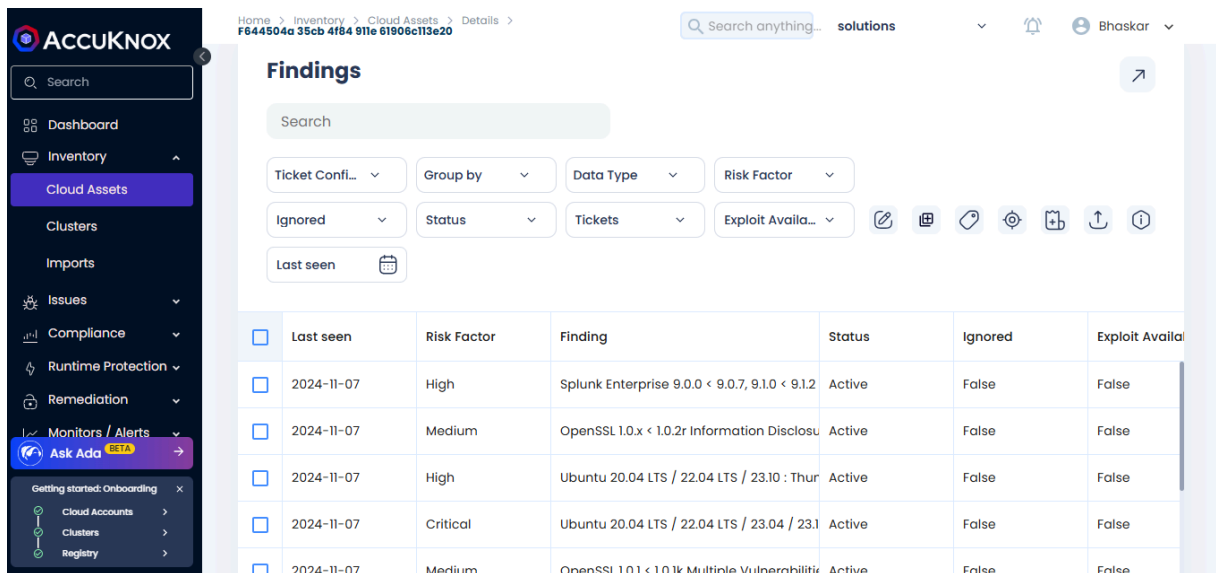
Ticket Number	Summary	Priority	Age
JIRATEST-225	Integer overflow in defineAttribute in xmlparse.c: (expat@2.4.1-r0)	High	514 D
JIRATEST-226	Local users can trigger security-relevant memory corruption via maL...	High	514 D
JIRATEST-195	systemd: buffer overrun in format_timespan() function: (libsystemd...	Medium	549 D
JIRATEST-203	kernel: null-ptr-deref caused by x25_disconnect: (linux-libc-dev@5...	Medium	537 D
JIRATEST-192	openssl: RSA authentication weakness: (libssl1.1@1.1.1n-0+deb11u4)	Low	550 D

## 10.3 Issues/Findings

### 10.3.1 Group findings by source and severity

AccuKnox automatically scans assets with the help of various open-source tools. It uses tools like Clair, Trivy, CLOC, Fortify, Snyk, SonarQube, Cloudsploit, Kube Bench, and various other open-source tools for Scanning.

Findings can be grouped according to the tools that were used to do the scan by selecting the “Data Type” option from the “Group By” drop down in the Vulnerabilities screen.



The screenshot shows the AccuKnox interface with a sidebar on the left containing navigation options like Dashboard, Inventory, Cloud Assets, Clusters, Imports, Issues, Compliance, Runtime Protection, Remediation, and Monitors / Alerts. The main content area is titled "Findings" and includes a search bar and several filter dropdowns: Ticket Confi..., Group by, Data Type, Risk Factor, Ignored, Status, Tickets, and Exploit Availa... There is also a "Last seen" filter with a calendar icon. Below the filters is a table of findings with the following columns: Last seen, Risk Factor, Finding, Status, Ignored, and Exploit Availa... The table contains five rows of data:

Last seen	Risk Factor	Finding	Status	Ignored	Exploit Availa...
2024-11-07	High	Splunk Enterprise 9.0.0 < 9.0.7, 9.1.0 < 9.1.2	Active	False	False
2024-11-07	Medium	OpenSSL 1.0.x < 1.0.2r Information Disclosu	Active	False	False
2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False
2024-11-07	Critical	Ubuntu 20.04 LTS / 22.04 LTS / 23.04 / 23.1	Active	False	False
2024-11-07	Medium	OpenSSL 1.0.1 < 1.0.1k Multiple Vulnerabiliti	Active	False	False

Users can further filter the findings with respect to their Risk factor so that they can have a view of most critical findings from each tool being used.

The screenshot shows the ACCUKNOX interface with the 'Findings' section active. A dropdown menu for 'Risk Factor' is open, showing options: Unknown, Informational, Low, Medium, High, Critical, and Not Available. The table below shows findings with columns: Last seen, Risk Factor, Finding, Status, Ignored, and Exploit Available.

Last seen	Risk Factor	Finding	Status	Ignored	Exploit Available
2024-11-07	High	Splunk Enterprise 9.0.0 <	Active	False	False
2024-11-07	Medium	OpenSSL 1.0.x < 1.0.2r Information Disclosu	Active	False	False
2024-11-07	High	Ubuntu 20.04 LTS / 22.04 LTS / 23.10 : Thur	Active	False	False
2024-11-07	Critical	Ubuntu 20.04 LTS / 22.04 LTS / 23.04 / 23.1	Active	False	False
2024-11-07	Medium	OpenSSL 1.0.1 < 1.0.1k Multiple Vulnerabiliti	Active	False	False

### 10.3.2 How to group by Findings and severity

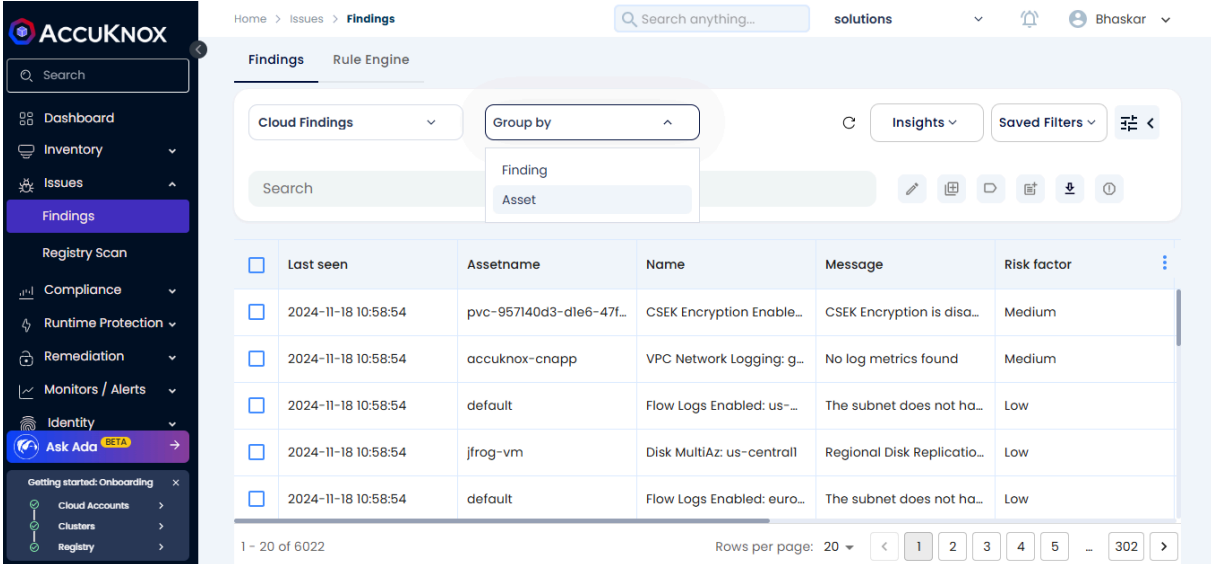
When resolving and patching vulnerabilities it is important to tackle the findings that are most abundant and most severe first. Users can use the Group by Findings feature to look for the vulnerabilities or misconfiguration that exist in large no. of assets and prioritize them accordingly.

The screenshot shows the ACCUKNOX interface with the 'Findings' section active. A dropdown menu for 'Group by' is open, showing options: Last seen, Risk Factor, Data Type, and Finding. The table below shows findings with columns: Last seen, Risk Factor, Finding, Status, Ignored, and Exploit Available.

Last seen	Risk Factor	Finding	Status	Ignored	Exploit Available
2024-11-07	High	Apache 2.4.x < 2.4.55 Multiple Vulnerabiliti	Active	False	False
2024-11-07	Critical	Apache 2.4.x < 2.4.56 Multiple Vulnerabiliti	Active	False	True
2024-11-07	High	Apache 2.4.x < 2.4.58 Multiple Vulnerabiliti	Active	False	False
2024-11-07	High	Apache 2.4.x < 2.4.58 Out-of-Bounds Rea	Active	False	False
2024-11-07	Medium	Apache 2.4.x < 2.4.59 Multiple Vulnerabiliti	Active	False	False

### 10.3.3 How to group by Asset and severity

Users can have an Asset wise view of the findings. Grouping by assets, groups the vulnerabilities or misconfigurations together with respect to the asset that they are associated with.



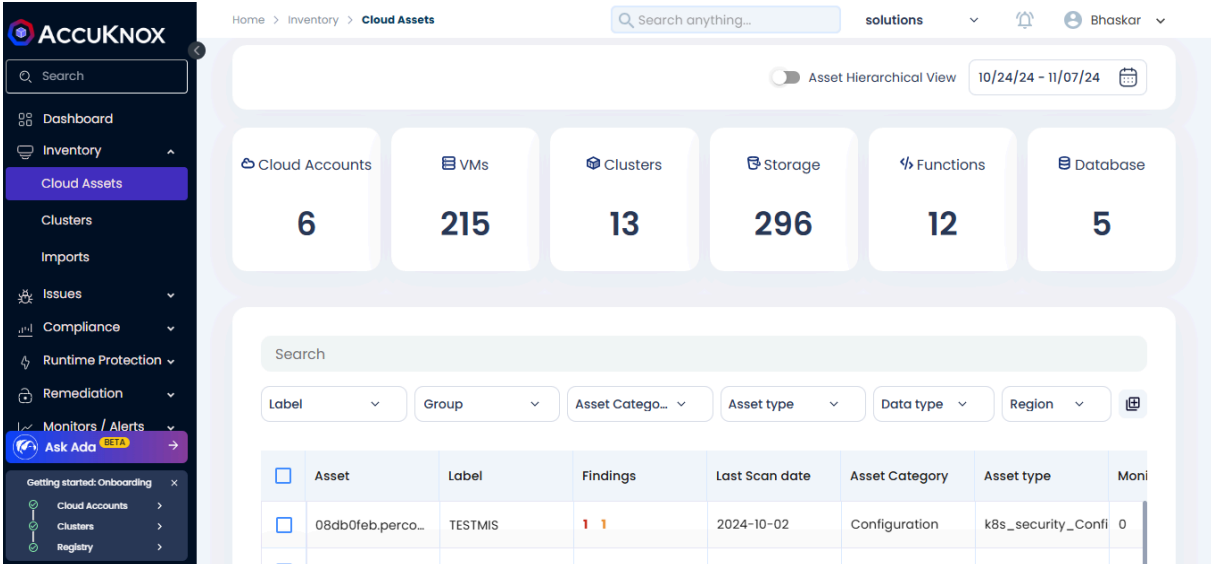
The screenshot shows the Accuknox Findings page. The left sidebar contains navigation options like Dashboard, Inventory, Issues, Findings, Registry Scan, Compliance, Runtime Protection, Remediation, Monitors / Alerts, Identity, and Ask Ada. The main content area shows a table of findings with columns: Last seen, Assetname, Name, Message, and Risk factor. A 'Group by' dropdown menu is open, showing options for 'Finding' and 'Asset'. The table contains several rows of findings, all dated 2024-11-18 10:58:54, with various asset names and risk factors ranging from Low to Medium.

Last seen	Assetname	Name	Message	Risk factor
2024-11-18 10:58:54	pvc-957140d3-d1e6-47f...	CSEK Encryption Enable...	CSEK Encryption is disa...	Medium
2024-11-18 10:58:54	accuknox-cnapp	VPC Network Logging: g...	No log metrics found	Medium
2024-11-18 10:58:54	default	Flow Logs Enabled: us...	The subnet does not ha...	Low
2024-11-18 10:58:54	jfrog-vm	Disk MultiAz: us-centra...	Regional Disk Replicatio...	Low
2024-11-18 10:58:54	default	Flow Logs Enabled: euro...	The subnet does not ha...	Low

### 9.4 Asset Hierarchical View

In Accuknox Dashboard, under Inventory -> Cloud Assets, Asset Hierarchical View is present.

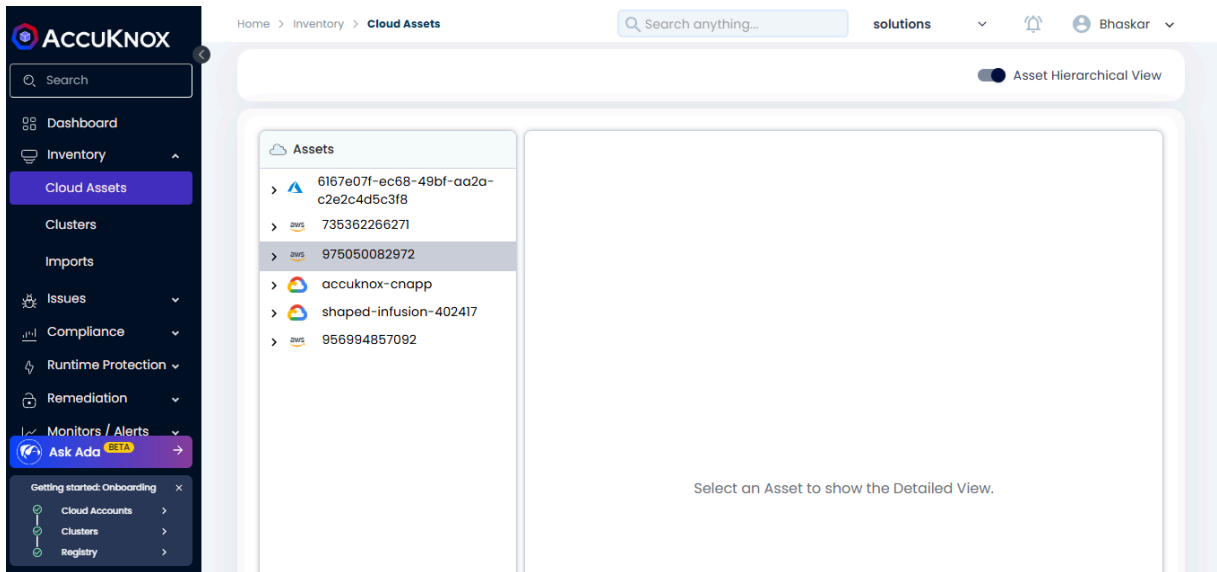
It shows the cloud assets in order.



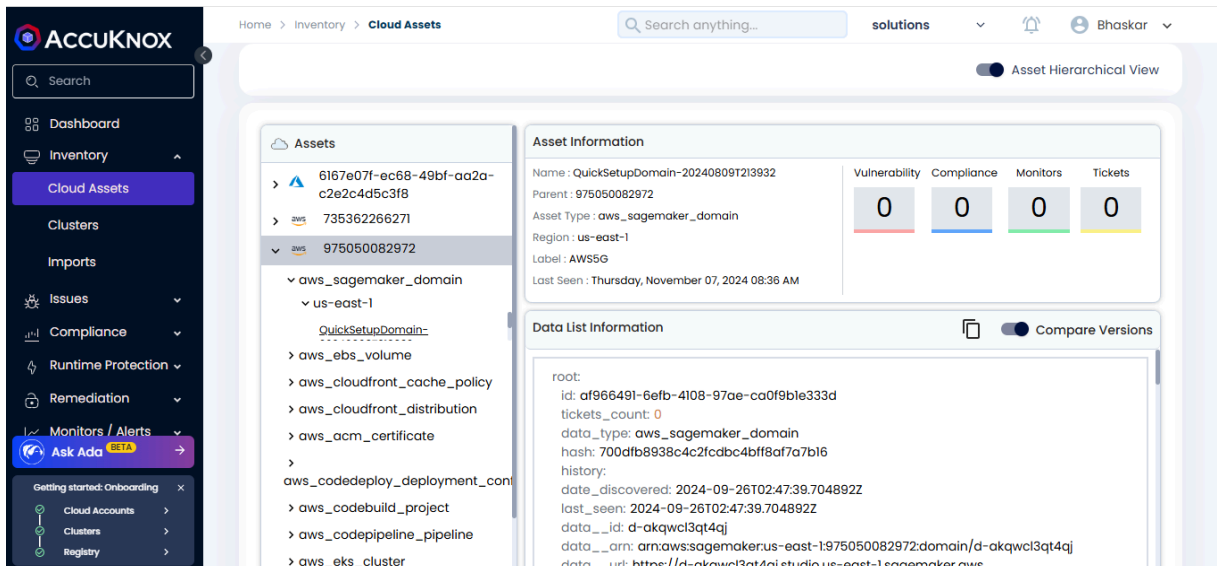
The screenshot shows the Accuknox Cloud Assets page. The left sidebar is the same as in the previous screenshot. The main content area shows a summary of cloud assets with a toggle for 'Asset Hierarchical View' and a date range of 10/24/24 - 11/07/24. Below the summary are six cards representing different asset types: Cloud Accounts (6), VMs (215), Clusters (13), Storage (296), Functions (12), and Database (5). Below these cards is a search bar and a table of assets with columns: Asset, Label, Findings, Last Scan date, Asset Category, Asset type, and Moni.

Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Moni
08db0feb.perco...	TESTMIS	1 1	2024-10-02	Configuration	k8s_security_Confi	0

Toggle Asset Hierarchical View by clicking on the trigger.



You can click on any of the cloud accounts to get more information about it.



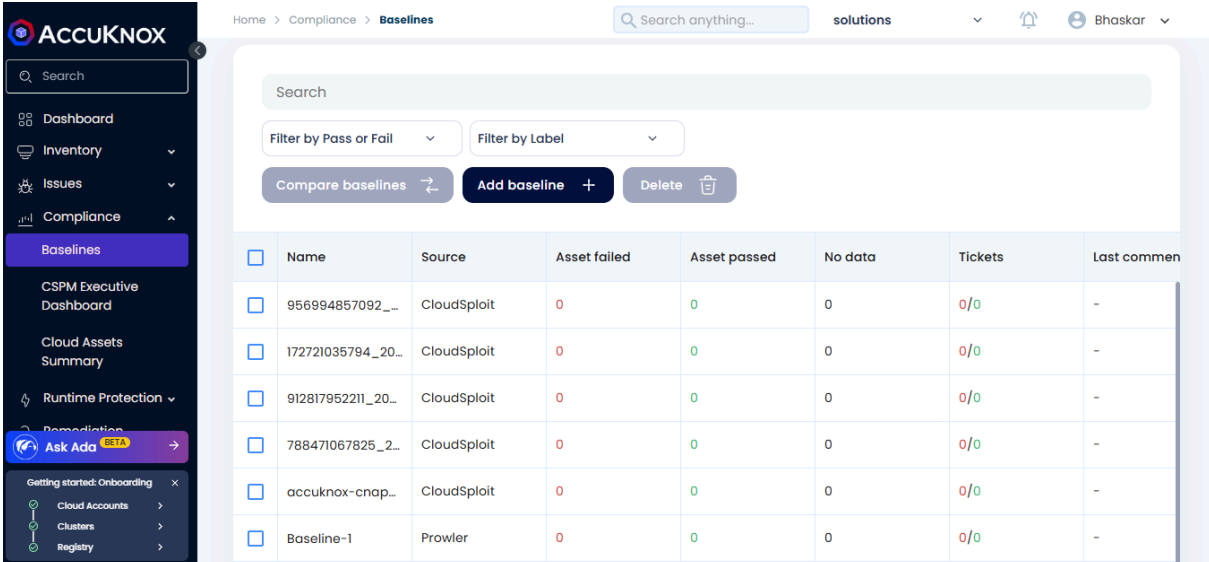
## 10.4 Baselines

### 10.4.1 How to create a Baseline out of a data source

AccuKnox's Baseline is an approach to detect drift in configuration from the conformance suite from multiple 'data sources' that can be associated with a specific 'asset' or 'group' of assets. It is a golden benchmark that is used to detect any change in compliance behavior proactively.

To create a baselines follow these steps:

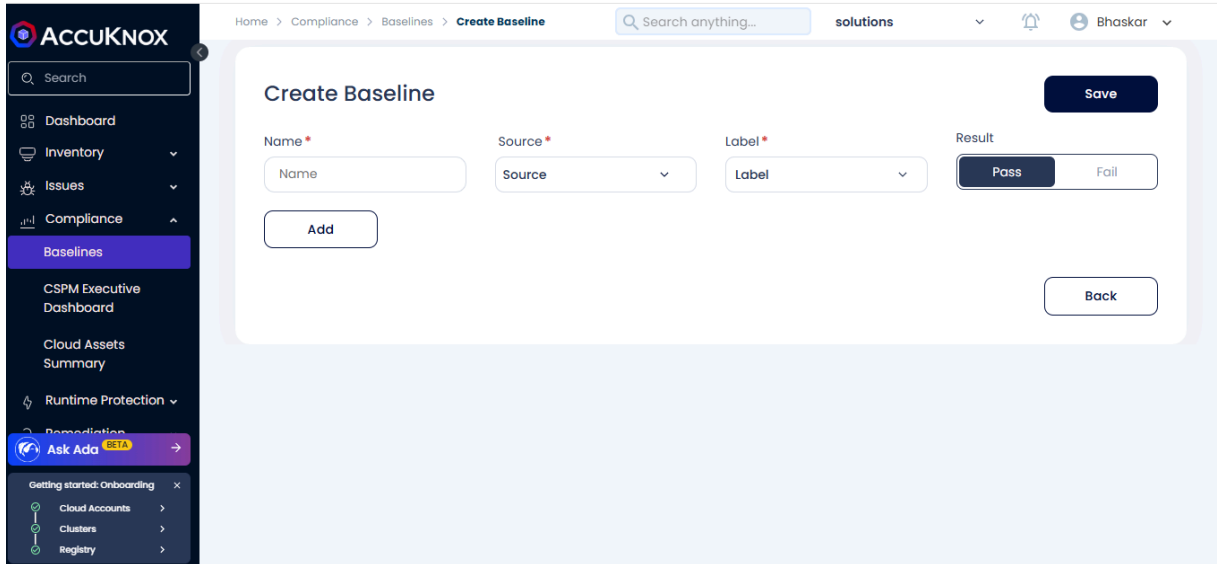
**Step 1:** Head to the Baselines page under the Compliance section and click on "Add baseline".



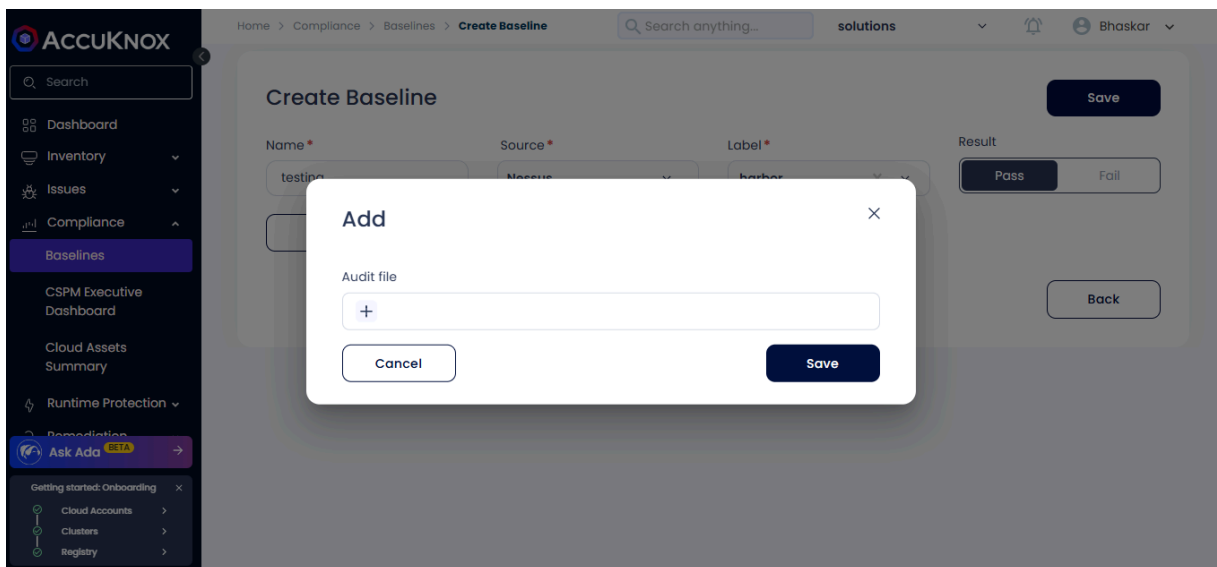
The screenshot shows the AccuKnox interface for the 'Baselines' page. The breadcrumb trail is 'Home > Compliance > Baselines'. There is a search bar and filters for 'Filter by Pass or Fail' and 'Filter by Label'. Action buttons include 'Compare baselines', 'Add baseline', and 'Delete'. Below is a table of existing baselines:

<input type="checkbox"/>	Name	Source	Asset failed	Asset passed	No data	Tickets	Last commen
<input type="checkbox"/>	956994857092_...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	172721035794_20...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	912817952211_20...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	788471067825_2...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	accuknox-cnap...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	Baseline-1	Prowler	0	0	0	0/0	-

**Step 2:** Provide a name, select the source and select the Result for your baseline and add a label for your baseline.



**Step 3:** Finally add the audit files by clicking on add, these files contain the compliance analysis from different cloud accounts.



Now you can see the compliance analysis by clicking on the baseline that you created.



Home > Compliance > Baselines > Baseline Configuration > 890ff57e 458d 47e4 908e 8d6aad537dd5

Search anything solutions

AC-25-07

Result: Result, Expected Value: Expected Value, Ticket(s) Created: Ticket(s) Created, Annotation: Annotation

Edit Fields, Edit Sources

Data, Assigned Hosts, Assigned Groups, History

Search

Ticket Configuration, Group by, Filter by Compliance

<input type="checkbox"/>	Control	Result	Tags	Description	Info	Last Comment	Tickets
<input type="checkbox"/>	Access Approval...	PASSED	False	Ensure that Acce...	Category: Securi...	-	
<input type="checkbox"/>	Alias IP Ranges E...	PASSED	False	Ensures all Kuber...	Category: Kuber...	-	

## 10.4.2 How to compare baselines

Once you have created a baseline for your cloud infrastructure, to ensure continuous compliance you can create another baseline and compare them to see if there is any drift in the configuration between your past baseline and your current baseline.

To compare your baselines, select multiple baselines and click on compare baselines to see the comparison.

Home > Compliance > Baselines

Search anything...

Filter by Pass or Fail, Filter by Label

Compare baselines, Add baseline, Delete

<input type="checkbox"/>	Name	Source	Asset failed	Asset passed	No data	Tickets	Last commen
<input checked="" type="checkbox"/>	956994857092_...	CloudSploit	0	0	0	0/0	-
<input checked="" type="checkbox"/>	172721035794_20...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	912817952211_20...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	788471067825_2...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	accuknox-cnapp...	CloudSploit	0	0	0	0/0	-
<input type="checkbox"/>	Baseline-1	Prowler	0	0	0	0/0	-
<input type="checkbox"/>	multiHost	Security Hub	0	0	0	0/0	-

The comparison will look like following,

Compare			
Finding	baseline-aws100723	Baseline-1	multiHost
Ensure, a log metric filter and alarm exist for Manage...	✓	✗	✗
A, log metric filter and alarm should exist for usage o...	✓	✗	✗
Ensure, a log metric filter and alarm exist for change...	✓	✗	✗
Ensure, a log metric filter and alarm exist for AWS Co...	✓	✗	✗
Avoid, the use of the root user, Multi region CloudTra...	✓	✗	✗
Ensure, a log metric filter and alarm exist for route ta...	✓	✗	✗
Ensure, a log metric filter and alarm exist for IAM poli...	✓	✗	✗
Ensure, a log metric filter and alarm exist for VPC ch...	✓	✗	✗
Ensure, a log metric filter and alarm exist for disablin...	✓	✗	✗
Ensure, a log metric filter and alarm exist for change...	✓	✗	✗

Total Count: undefined

## 10.5 Compliance

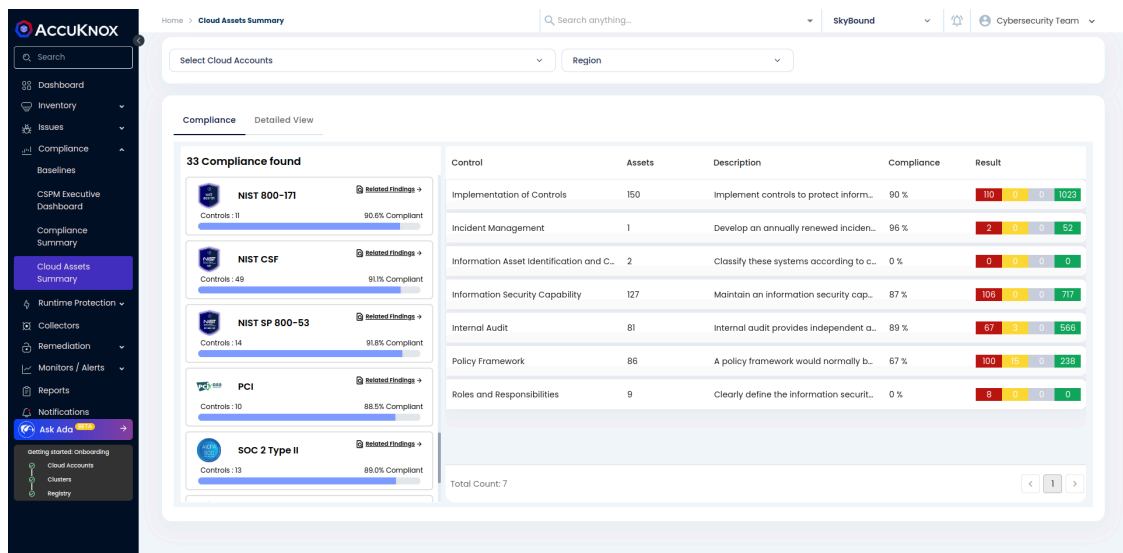
AccuKnox helps you to review your cloud infrastructure health and compliance posture. AccuKnox also helps you to generate reports that contain summary and detailed assessment of vulnerability/findings and compliance risks in your cloud infrastructure or in applications.

### 10.5.1 How to get Compliance for Cloud Assets

- In order to check for compliance Navigate to **Compliance > Cloud Asset Summary**

Users can click on any Compliance Program or their Sub-control which will navigate to the list of misconfiguration. Further user can filter based on Cloud Account, Region, Severity, Checks, and many more on the **Detailed View** Tab.

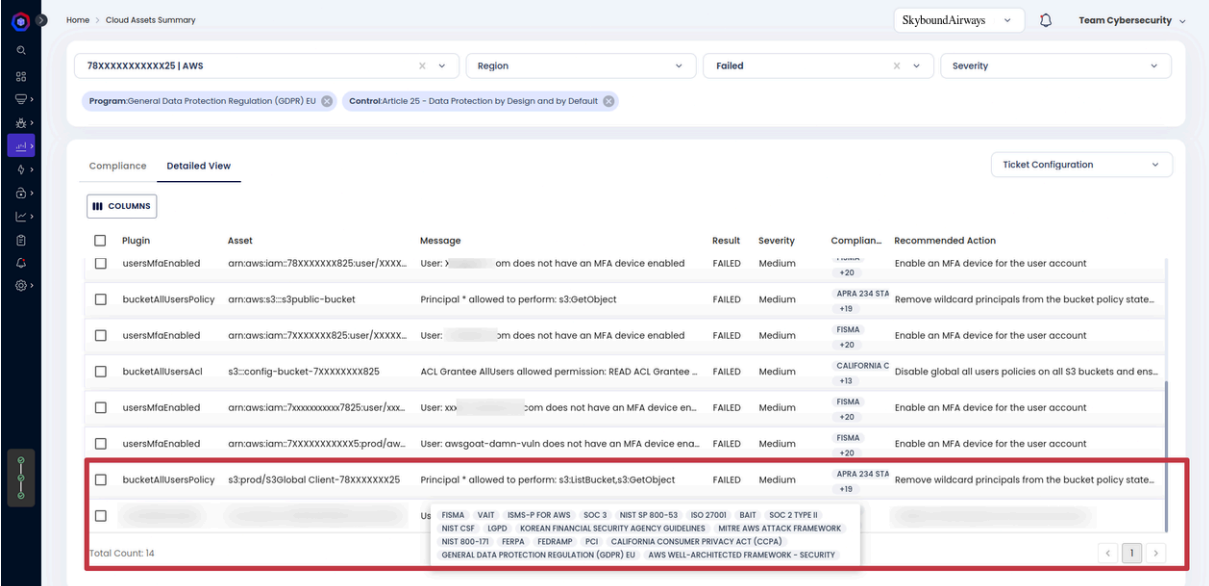
- **Compliance:** A detailed report that gives you insight into how you score against a framework's requirements and rules.



The screenshot displays the AccuKnox interface for the 'Cloud Assets Summary' page. The left sidebar contains navigation options such as Dashboard, Inventory, Issues, Compliance, Baselines, and various security tools. The main content area shows a 'Compliance' section with a 'Detailed View' tab. It lists 33 compliance findings across several frameworks:

Control	Assets	Description	Compliance	Result
<b>NIST 800-171</b> Controls: 11 90.6% Compliant	Implementation of Controls	150	Implement controls to protect inform...	90% 10 9 0 1023
<b>NIST CSF</b> Controls: 49 91.1% Compliant	Incident Management	1	Develop an annually renewed incident...	96% 2 0 0 52
<b>NIST SP 800-53</b> Controls: 14 91.8% Compliant	Information Asset Identification and C...	2	Classify these systems according to c...	0% 0 0 0 0
<b>PCI</b> Controls: 10 88.5% Compliant	Information Security Capability	127	Maintain an information security cap...	87% 106 0 0 717
<b>SOC 2 Type II</b> Controls: 13 89.0% Compliant	Internal Audit	81	Internal audit provides independent a...	89% 67 3 0 566
	Policy Framework	86	A policy framework would normally b...	67% 100 16 0 238
	Roles and Responsibilities	9	Clearly define the information securit...	0% 8 0 0 0
Total Count: 7				< 1 >

- Detailed View:** A filtered view of the **Misconfigurations** page that shows resources with misconfigurations for the selected Compliance Program.



Home > Cloud Assets Summary SkyboundAirways Team Cybersecurity

78XXXXXXXXXX25 | AWS Region Failed Severity

Program: General Data Protection Regulation (GDPR) EU Control: Article 25 - Data Protection by Design and by Default

Compliance **Detailed View** Ticket Configuration

**COLUMNS**

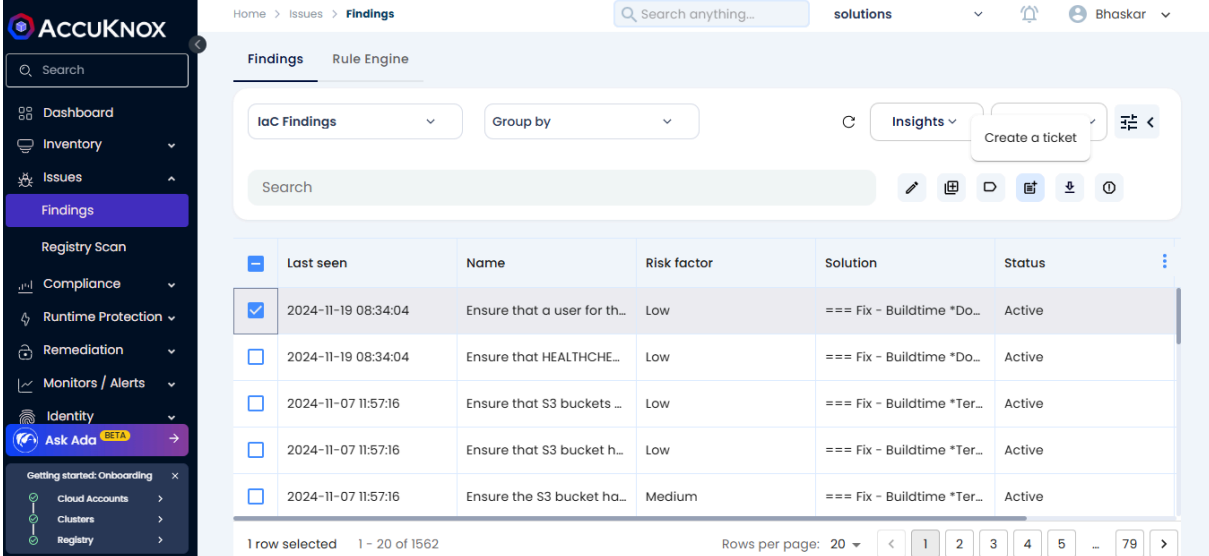
Plugin	Asset	Message	Result	Severity	Compliance	Recommended Action
<input type="checkbox"/>	usersMfaEnabled	arn:aws:iam:78XXXXXXXX825:user/XXXX... User: >om does not have an MFA device enabled	FAILED	Medium	FISMA +20	Enable an MFA device for the user account
<input type="checkbox"/>	bucketAllUsersPolicy	arn:aws:s3::s3public-bucket Principal * allowed to perform: s3:GetObject	FAILED	Medium	APRA 234 STA +19	Remove wildcard principals from the bucket policy state...
<input type="checkbox"/>	usersMfaEnabled	arn:aws:iam:7XXXXXXXX825:user/XXXX... User: >om does not have an MFA device enabled	FAILED	Medium	FISMA +20	Enable an MFA device for the user account
<input type="checkbox"/>	bucketAllUsersAcl	s3::config-bucket-7XXXXXXXX825 ACL Grantee AllUsers allowed permission: READ ACL Grantee ...	FAILED	Medium	CALIFORNIA C +13	Disable global all users policies on all S3 buckets and ens...
<input type="checkbox"/>	usersMfaEnabled	arn:aws:iam:7XXXXXXXX825:user/xxx... User: xxx>om does not have an MFA device en...	FAILED	Medium	FISMA +20	Enable an MFA device for the user account
<input type="checkbox"/>	usersMfaEnabled	arn:aws:iam:7XXXXXXXXX5:prod/aw... User: awsgoat-damn-vuln does not have an MFA device ena...	FAILED	Medium	FISMA +20	Enable an MFA device for the user account
<input type="checkbox"/>	bucketAllUsersPolicy	s3prod/s3Global Client-78XXXXXXXX25 Principal * allowed to perform: s3:ListBuckets3.GetObject	FAILED	Medium	APRA 234 STA +19	Remove wildcard principals from the bucket policy state...

Total Count: 14

FISMA | VAIT | ISMS-P FOR AWS | SOC 3 | NIST SP 800-53 | ISO 27001 | BAIT | SOC 2 TYPE II  
 NIST CSF | LOGD | KOREAN FINANCIAL SECURITY AGENCY GUIDELINES | MITRE AWS ATTACK FRAMEWORK  
 NIST 800-171 | FERPA | FEDRAMP | PCI | CALIFORNIA CONSUMER PRIVACY ACT (CCPA)  
 GENERAL DATA PROTECTION REGULATION (GDPR) EU | AWS WELL-ARCHITECTED FRAMEWORK - SECURITY

## 10.6 Remediation - Fix Problems/Create Tickets

To remediate any findings, users will need to select the finding or group of findings From the Issues→ Findings page and click Create Ticket as shown in the below screenshot.

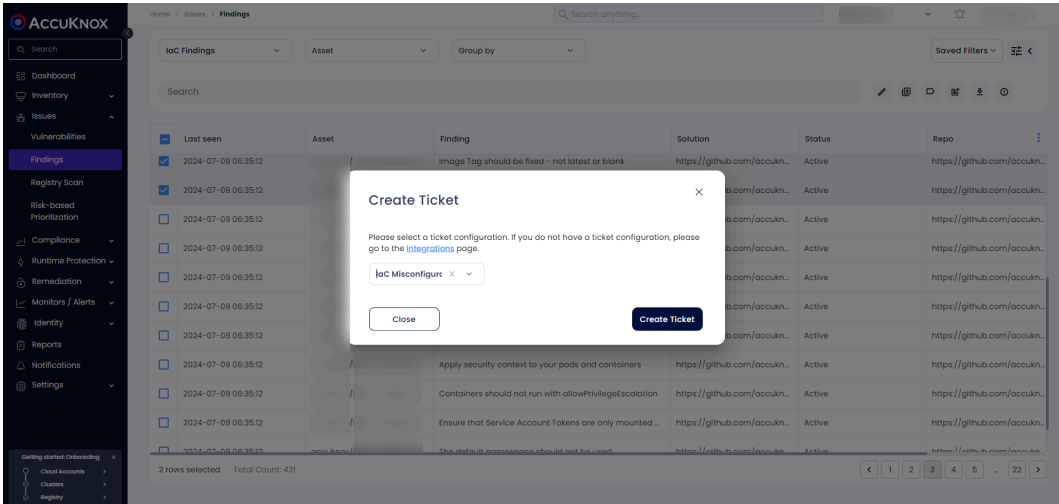


The screenshot shows the AccuKnox interface with the 'Findings' page selected. A table of findings is displayed with columns: Last seen, Name, Risk factor, Solution, and Status. The first row is selected, and a 'Create a ticket' button is visible in the top right corner of the table area.

Last seen	Name	Risk factor	Solution	Status
2024-11-19 08:34:04	Ensure that a user for th...	Low	=== Fix - Buildtime *Do...	Active
2024-11-19 08:34:04	Ensure that HEALTHCHE...	Low	=== Fix - Buildtime *Do...	Active
2024-11-07 11:57:16	Ensure that S3 buckets ...	Low	=== Fix - Buildtime *Ter...	Active
2024-11-07 11:57:16	Ensure that S3 bucket h...	Low	=== Fix - Buildtime *Ter...	Active
2024-11-07 11:57:16	Ensure the S3 bucket ha...	Medium	=== Fix - Buildtime *Ter...	Active

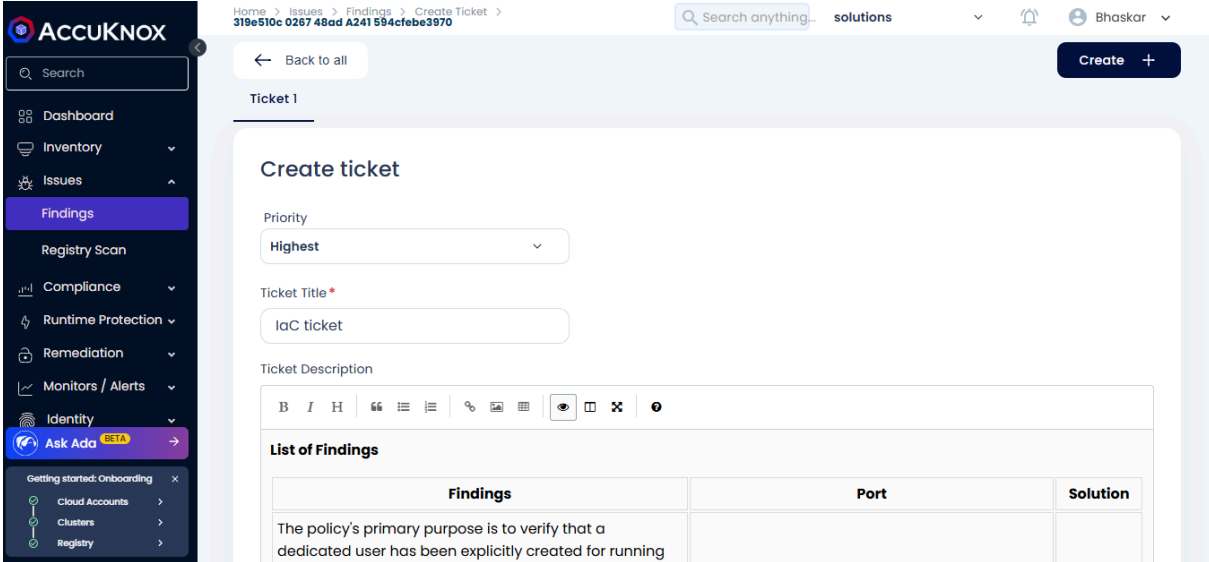
**NOTE: Before this users must have integrated their Ticketing backend like Jira Servicenow or connects or Freshservice under Integrations → CSPM section. (Refer to ticketing integrations in Integrations section)**

After clicking on the create ticket Icon the next page will popup.



The screenshot shows the AccuKnox interface with the 'Findings' page selected. A 'Create Ticket' modal dialog box is open, prompting the user to select a ticket configuration. The dialog box contains the text: 'Please select a ticket configuration. If you do not have a ticket configuration, please go to the integrations page.' and a dropdown menu showing 'IaC Misconfigur'.

Once the user clicks on Create Ticket new page with all the information related to the IaC findings and with a predefined Priority based on the Risk Factor. The user has to click on Create to confirm the ticket creation.



Home > Issues > Findings > Create Ticket >  
 319e510c 0287 48ad A241 594cfebe3970

Search anything... solutions

← Back to all

**Create +**

### Create ticket

Priority  
 Highest

Ticket Title \*  
 IaC ticket

Ticket Description

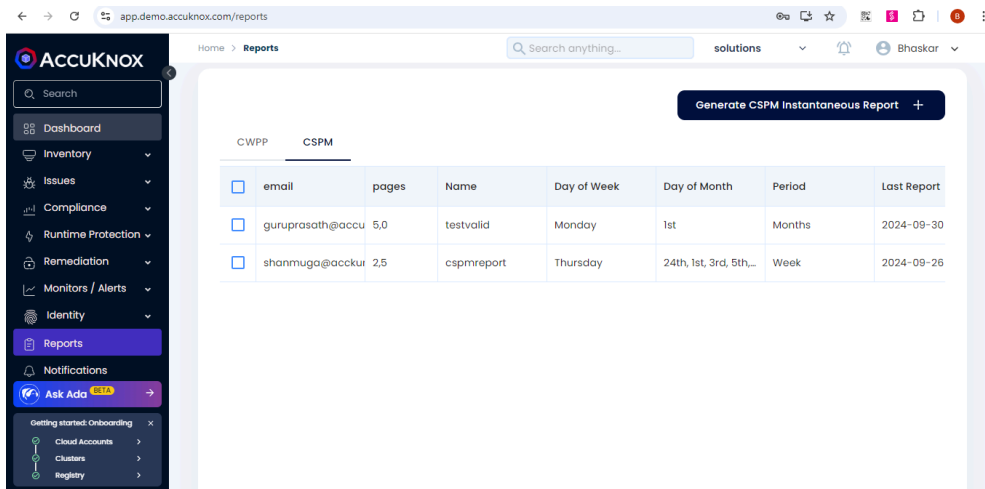
B I H | 🗨️ | 🗨️ | 🔗 | 📎 | 🗑️ | 🗑️ | 🗑️

#### List of Findings

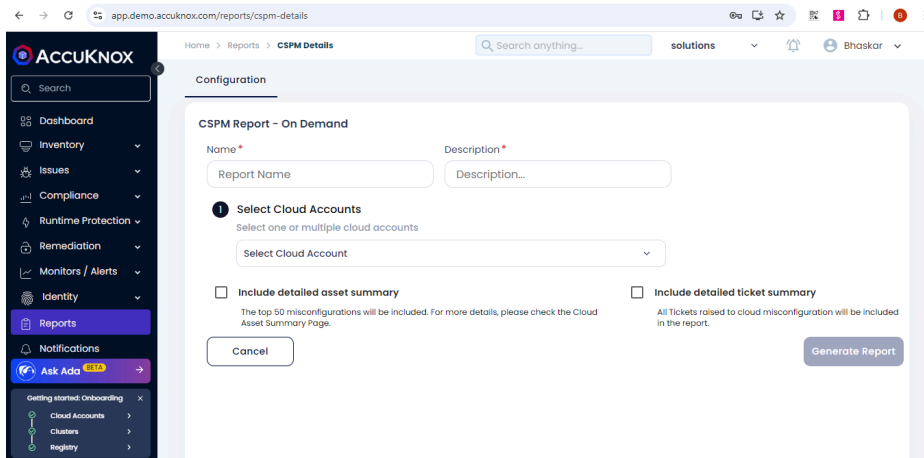
Findings	Port	Solution
The policy's primary purpose is to verify that a dedicated user has been explicitly created for running		

## 10.7 CSPM Reports

In the Accuknox dashboard, under Settings, go to Reports section, click on the “CSPM” tab, then click on “Generate CSPM Instantaneous Report”.



Now, give the Name, Description and select the Cloud account. Under the Cloud Account section, we can select multiple cloud accounts.



After selecting the cloud account, under the Compliance Program, select one of the two:

1. Compliance Report
2. Cloud Account Misconfiguration Report.

### Compliance Report:

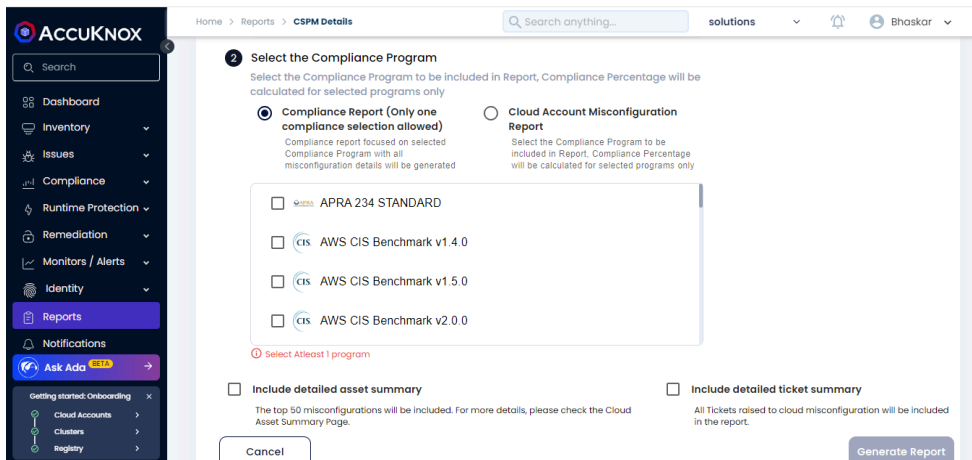
Allows selecting only a single compliance program. The report generated will be focused on the selected compliance program, it will show us all the misconfiguration details.

### Cloud Account Misconfiguration Report:

Allows us in selecting more than one compliance program. In this report, the Compliance Percentage will be calculated for selected programs.

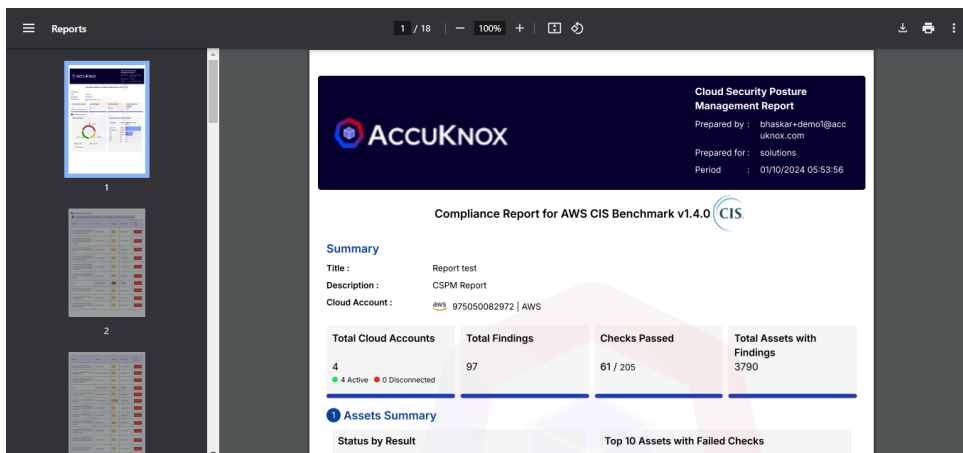
After filling the required details, click on “Generate Report”.

You can also select the checkboxes to include the detailed asset summary and detailed ticket summary.



Now we can see that the CSPM Report has been generated.

You can click on the Download icon to download it.





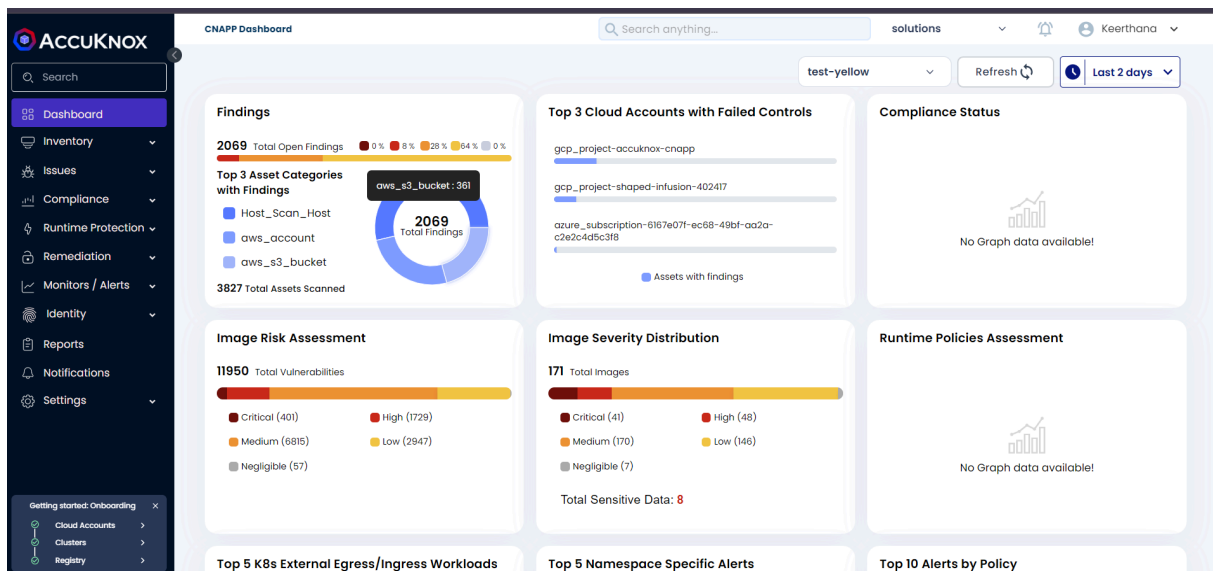
## 10.8 Rules Engine

The Rules Engine allows users to customize and automate ticket creation by selecting the data type, defining the criticality, and configuring specific ticket settings. This ensures that tickets are created based on the selected criteria, providing more control over the ticketing process.

### Automated Ticket Creation using Rules Engine

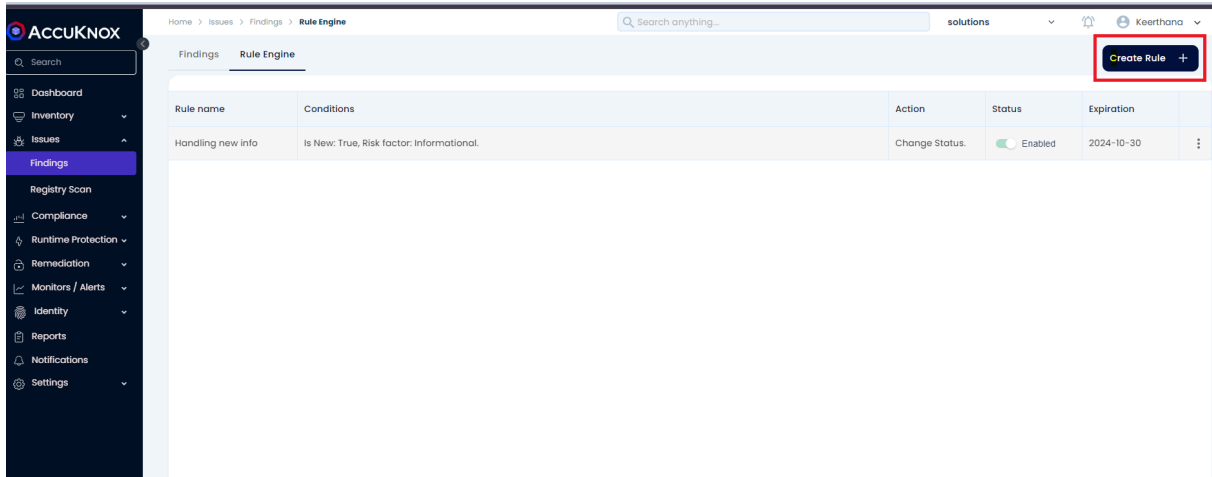
In this section we can find the steps to create a ticket using Rule Engine in the AccuKnox SaaS platform:

**Step 1:** Log in to [app.demo.accuknox.com](http://app.demo.accuknox.com) and navigate to the CNAPP dashboard.



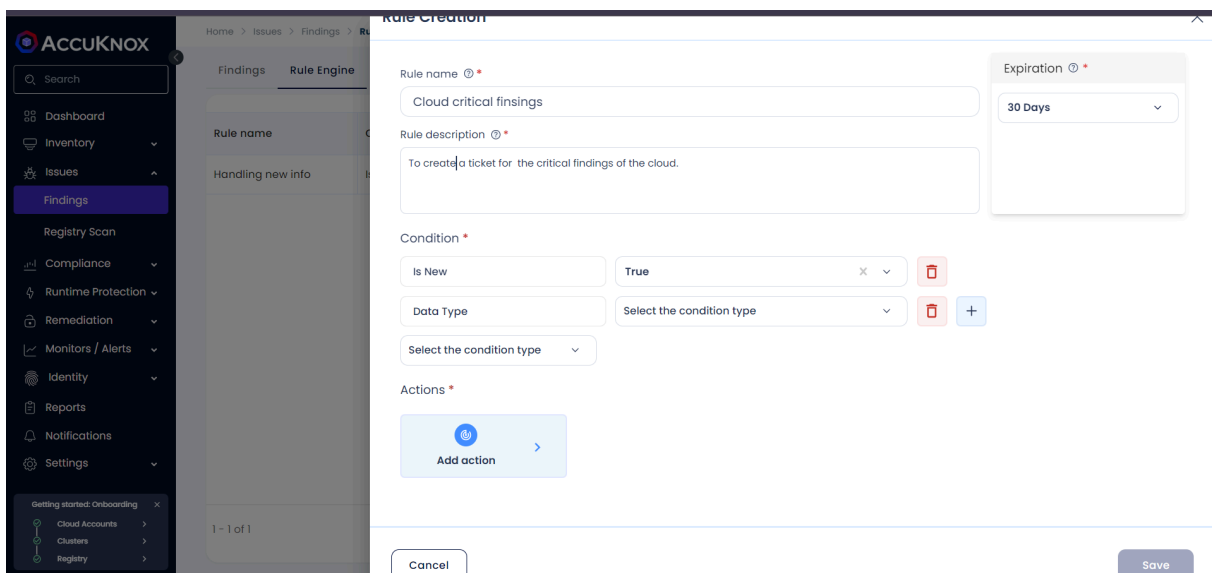
**Step 2:** Hover over to **Rule Engine** in Findings (Issues>Findings>Rule Engine)

**Step 3:** Click on **Create Rule** to create an automated rule.

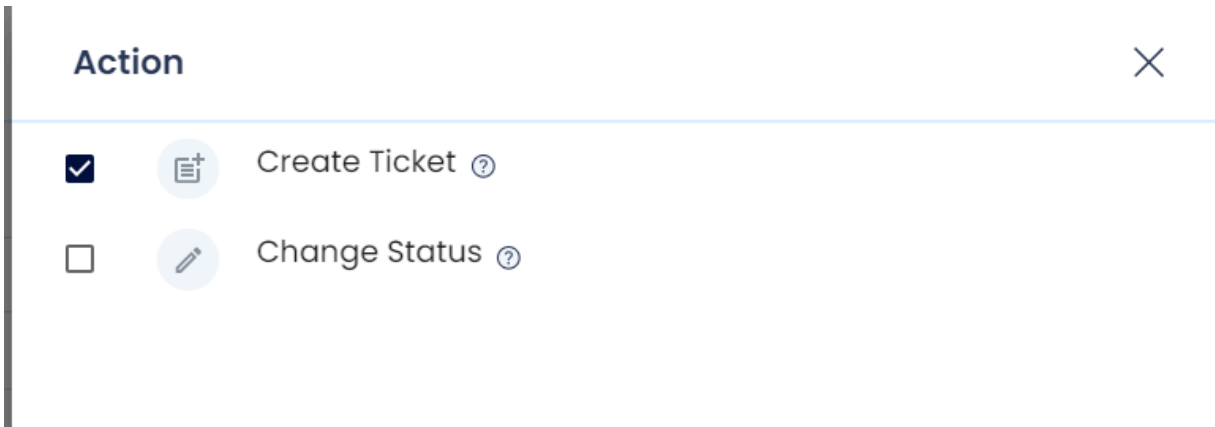


**Step 4:** Provide the necessary details, including the rule name, rule description, condition type (true or false), and click on "Action" to add the specific action.

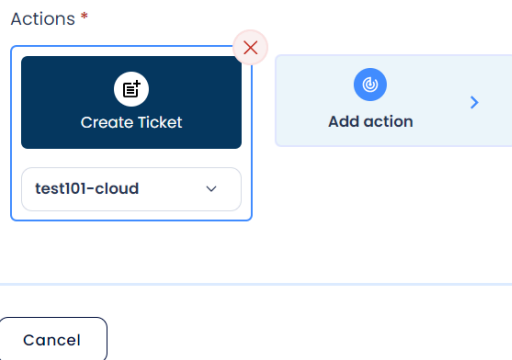
Clicking on the condition type will trigger the action that is required, it is crucial to select the condition type as per the rule created and the description of it.



**Step 5:** Click on **Create Ticket** to initiate creation of ticket when a finding with a matching the is found.



**Step 6:** After finalizing the condition, select the ticketing template from the drop down and save it to execute.



**Step 7:** Users need to create the ticketing configuration via Fresh Service Integration, which helps in automating the process of generating Freshservice “Problem alerts” with the existing security workflow.

# 11. ASPM (Application Security Posture Management)

This section makes use of Gitlab as an example to demonstrate the ASPM integrations with the CI/CD pipelines. In case a different platform is in use, please refer to the [help docs](#) for customized steps.

## 11.1 SAST

### 11.1.1 Integrating SonarQube SAST with AccuKnox in a GitLab CI/CD Pipeline

This guide demonstrates how to incorporate AccuKnox into a CI/CD pipeline using GitLab to enhance security. We'll use SonarQube SAST scanning to identify code vulnerabilities and send the results to AccuKnox for further analysis and remediation.

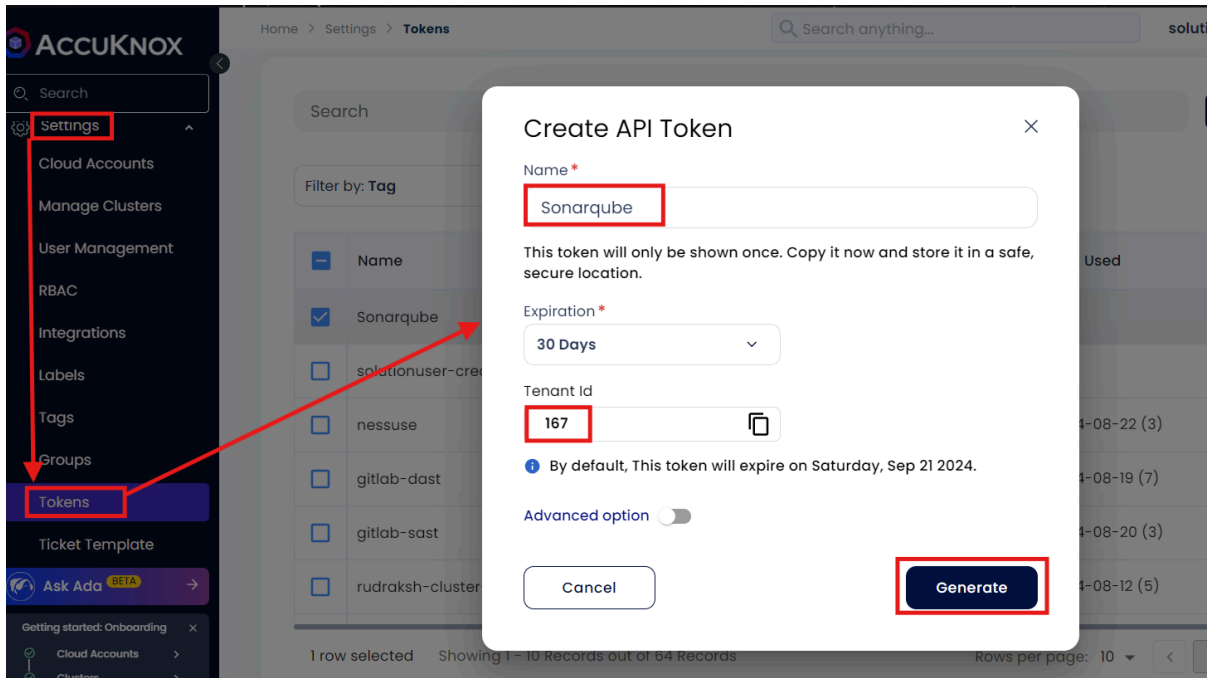
#### 11.1.2 Pre-requisites

- **GitLab Access**
- **AccuKnox UI Access**
- **SonarQube Access**

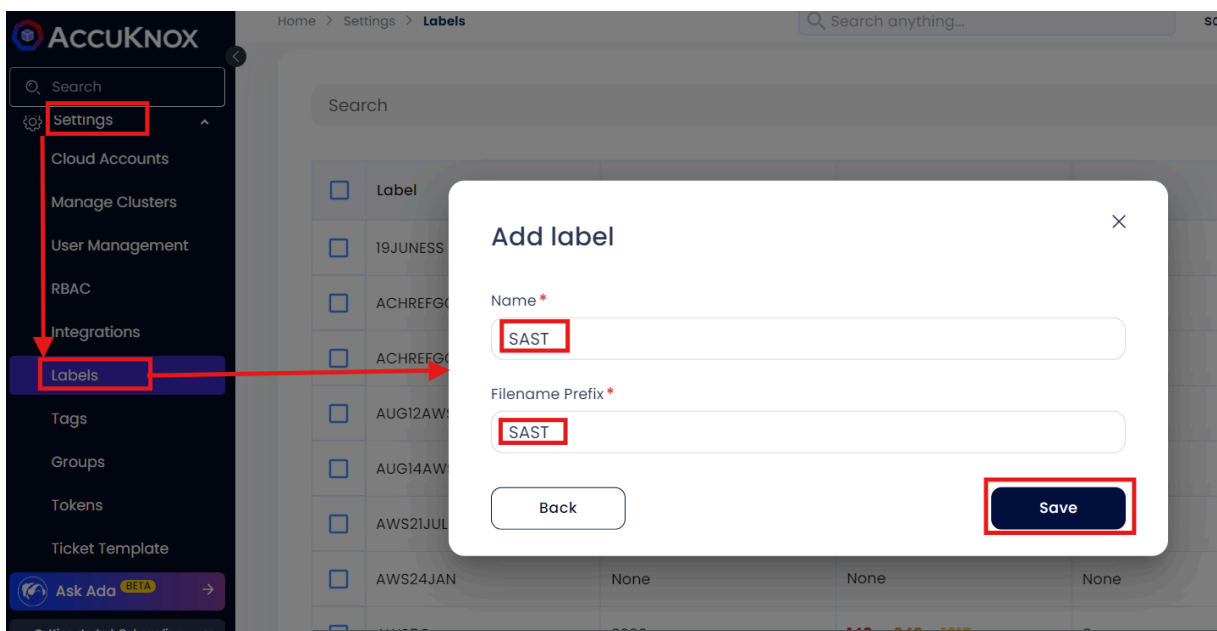
#### 11.1.3 Steps for Integration

Step 1: Log in to AccuKnox

- Navigate to **Settings** and select **Tokens** to create an AccuKnox token for forwarding scan results to SaaS.

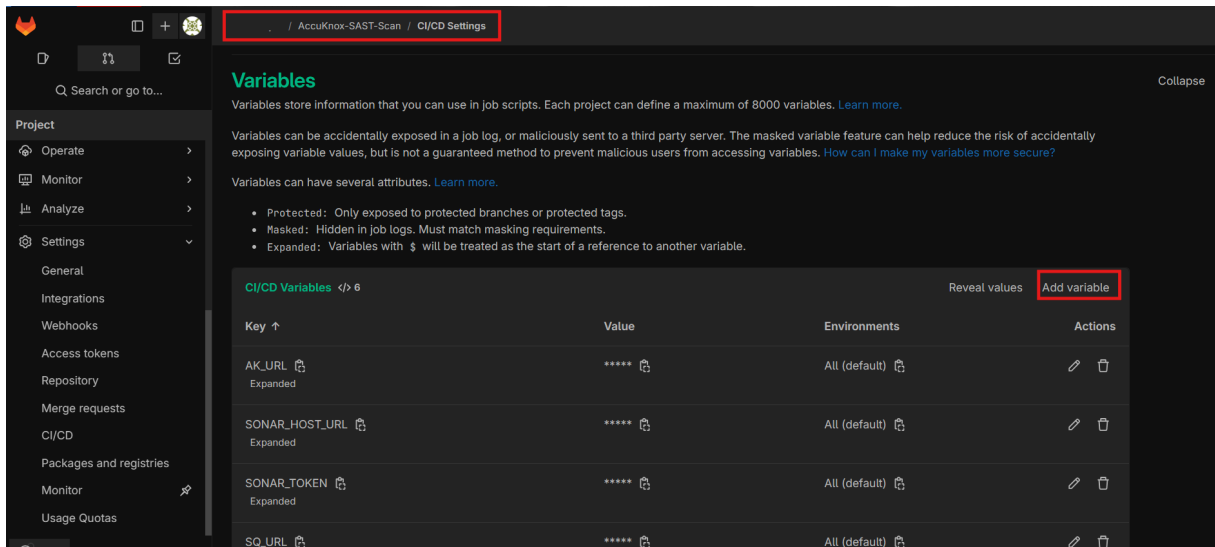


- Go to **AccuKnox > Settings > Labels** and create a label. This label will be used in the GitLab pipeline YAML file.



## Step 2: Create GitLab CI/CD Variables

- Copy the AccuKnox token and create a GitLab CI/CD masked variable for it.
- Additionally, create variables for the **tenant ID**, **AccuKnox URL**, **SonarQube token**, and the **SonarQube project URL**.



## Step 3: Set Up GitLab CI/CD Pipeline

Create a new pipeline in your GitLab project with the following YAML configuration:

```

stages:
  - sonarqube-check
  - fetch-report
  - upload-report

sonarqube-check:
  stage: sonarqube-check
  image:
    name: sonarsource/sonar-scanner-cli:latest
    entrypoint: [""]
  variables:
    SONAR_USER_HOME: "${CI_PROJECT_DIR}/.sonar" # Defines the location of the
analysis task cache
    GIT_DEPTH: "0" # Ensures all branches are fetched, required by the analysis

```

```
task
  cache:
    key: "${CI_JOB_NAME}"
    paths:
      - .sonar/cache
    script:
      - sonar-scanner -Dsonar.qualitygate.wait=true || true
    allow_failure: true
    rules:
      - if: $CI_COMMIT_REF_NAME == 'main' || $CI_PIPELINE_SOURCE ==
'merge_request_event'

  fetch-report:
    stage: fetch-report
    image: docker:latest
    services:
      - docker:dind
    dependencies:
      - sonarqube-check
    script:
      - |
        docker run --rm \
          -e SQ_URL=$SQ_URL \
          -e SQ_AUTH_TOKEN=$SONAR_TOKEN \
          -e REPORT_PATH=/app/data/ \
          -e SQ_PROJECTS="^gitlab-sast-testing$" \
          -v $PWD:/app/data/ \
          accuknox/sastjob:latest
    artifacts:
      paths:
        - SQ-*.json
      expire_in: 1 hour # Optional: Set expiration time for artifacts

  upload-report:
    stage: upload-report
    image: curlimages/curl:latest
    dependencies:
      - fetch-report
    script:
      - |
        for file in `ls -1 SQ-*.json`; do
          curl --location --request POST
            "<https://$AK_URL/api/v1/artifact/?tenant_id=$TENANT_ID&data_type=SQ&save_to_s3=
            false>" \
            --header "Tenant-Id: $TENANT_ID" \
```

```
--header "Authorization: Bearer $TOKEN" \  
--form "file=@\"$file\""  
done
```

### 11.1.4 Initial CI/CD Pipeline Without AccuKnox Scan

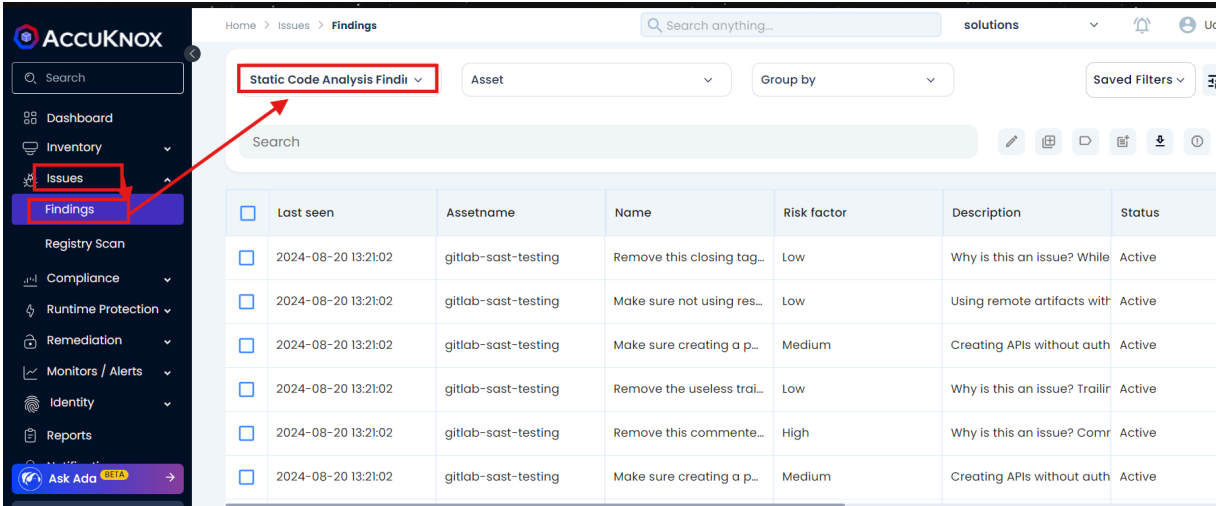
Initially, the CI/CD pipeline does not include the AccuKnox scan. Vulnerabilities in the code could go unnoticed without security checks.

### 11.1.5 CI/CD Pipeline After AccuKnox Integration

After integrating AccuKnox into the pipeline, pushing changes triggers the SonarQube scan, and results are sent to AccuKnox. AccuKnox helps identify potential code vulnerabilities.

### 11.1.6 View Results in AccuKnox SaaS

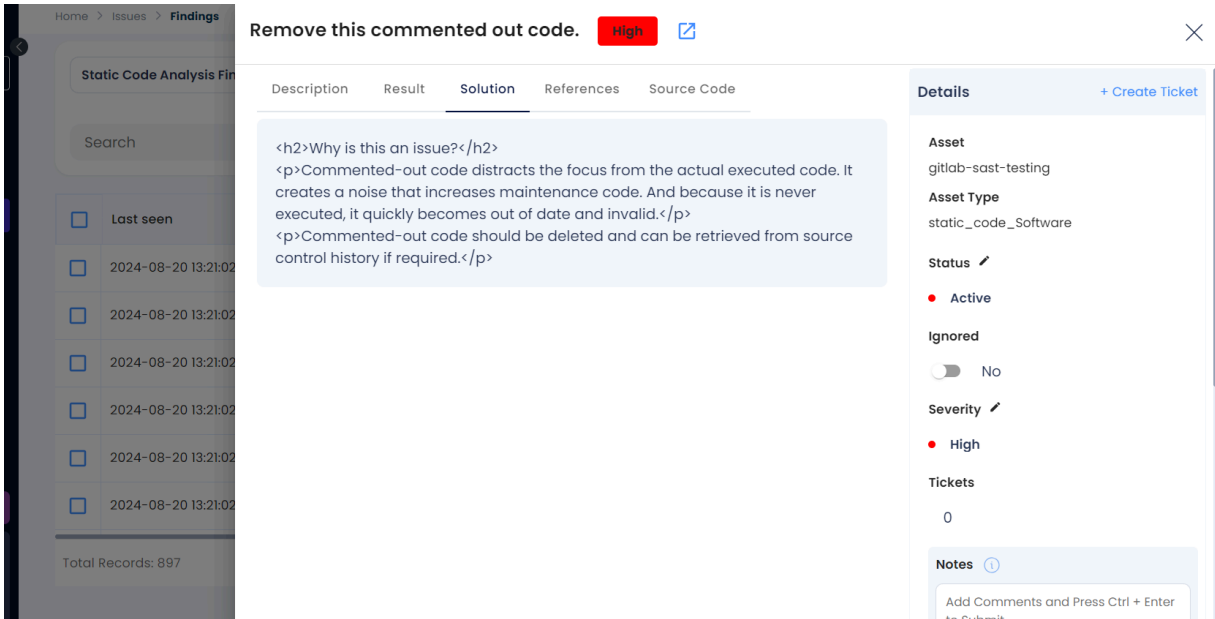
1. **Access the Dashboard:** After the pipeline completes, navigate to the AccuKnox SaaS dashboard.
2. **View Findings:** Go to **Issues > Findings** and select **SAST Findings** to see identified vulnerabilities



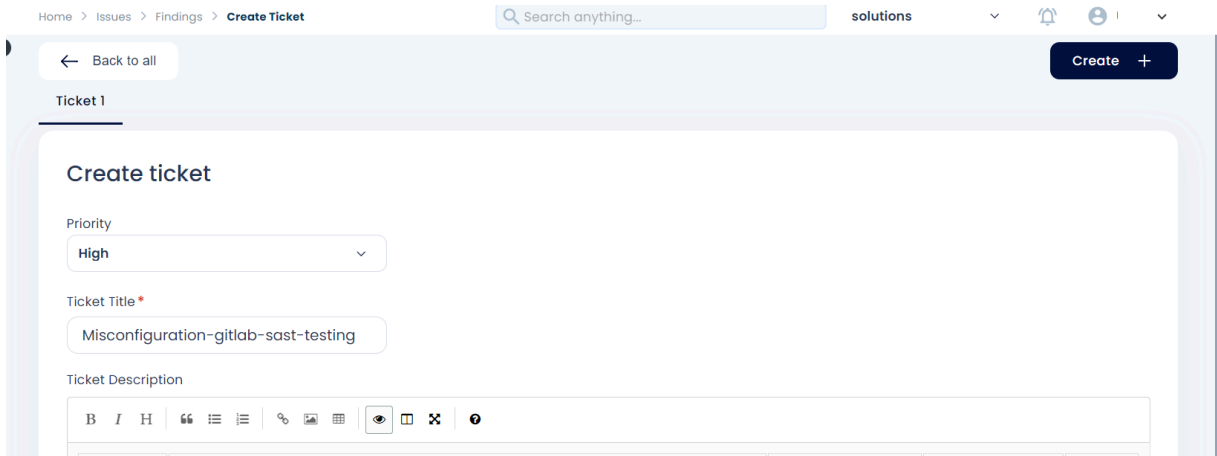
<input type="checkbox"/>	Last seen	Assetname	Name	Risk factor	Description	Status
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Remove this closing tag...	Low	Why is this an issue? While	Active
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Make sure not using res...	Low	Using remote artifacts with	Active
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Make sure creating a p...	Medium	Creating APIs without auth	Active
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Remove the useless trai...	Low	Why is this an issue? Trallir	Active
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Remove this commente...	High	Why is this an issue? Comr	Active
<input type="checkbox"/>	2024-08-20 13:21:02	gitlab-sast-testing	Make sure creating a p...	Medium	Creating APIs without auth	Active

3. **Analyze and Fix Vulnerabilities:** Click on a vulnerability to view more details and follow the instructions in the **Solutions** tab.





4. **Create a Ticket:** For unresolved vulnerabilities, create a ticket in your issue tracking system.



5. **Re-run the Pipeline:** After fixing the vulnerabilities, rerun the GitLab CI/CD pipeline and verify that the issues have been resolved in the AccuKnox dashboard.

## 11.2 DAST

### 11.2.1 Gitlab DAST Scan

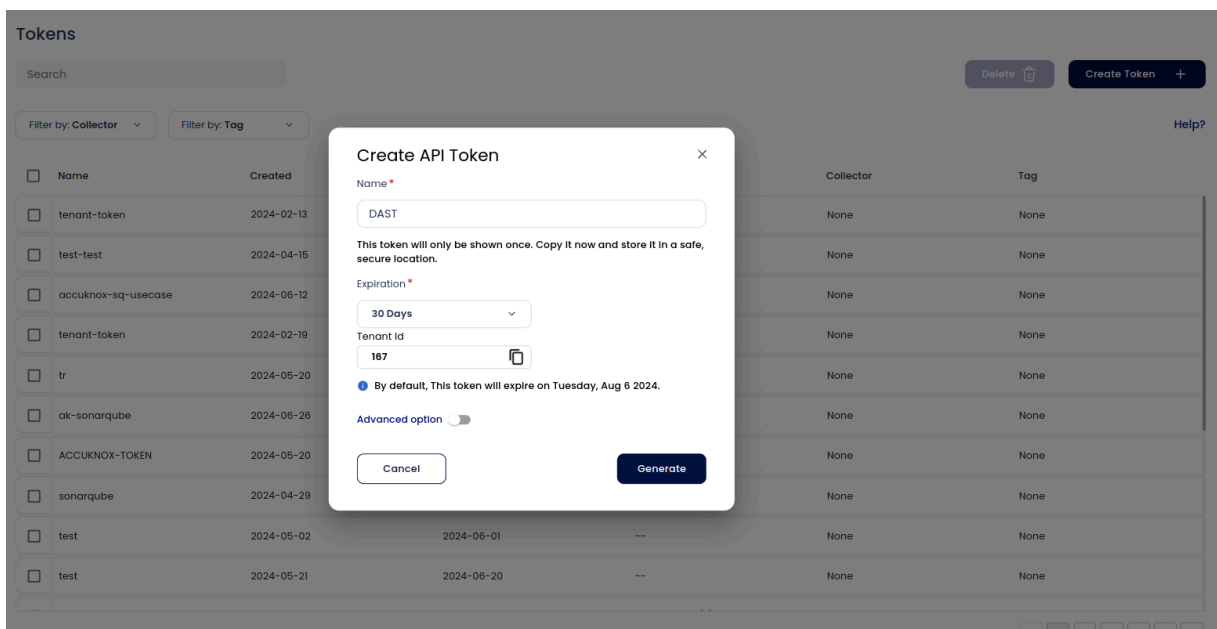
To demonstrate the benefits of incorporating AccuKnox into a CI/CD pipeline using GitLab to enhance security, consider a specific scenario involving a domain with known vulnerabilities. By integrating AccuKnox scanning into the pipeline, we can identify and resolve these security issues.

### 11.2.2 Pre-requisites

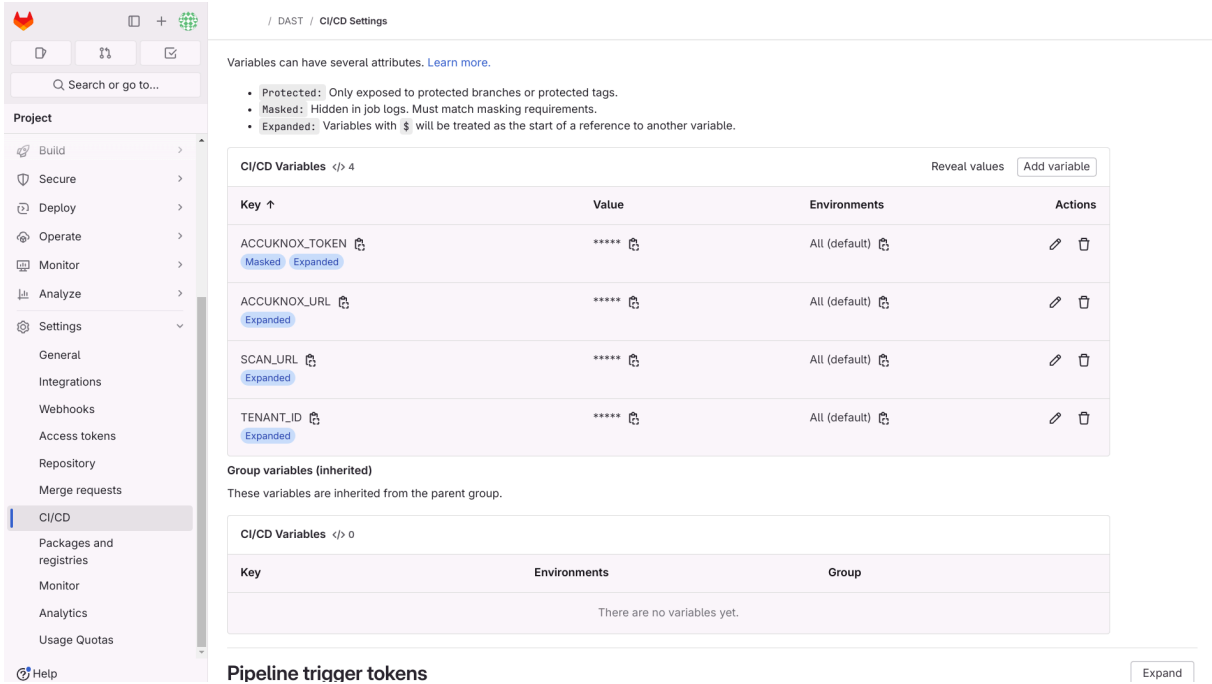
- GitLab Access
- AccuKnox UI access

### 11.2.3 Steps for Integration

**Step 1:** Log in to AccuKnox Navigate to Settings and select Tokens to create an AccuKnox token for forwarding scan results to SaaS.



**Step 2:** Copy the token and create a GitLab CI/CD masked variable for the token to be used in the pipeline. Also, create variables for the tenant id, AccuKnox URL (cspm.accuknox.com or cspm.demo.accuknox.com), and the target URL that you want to use for DAST.

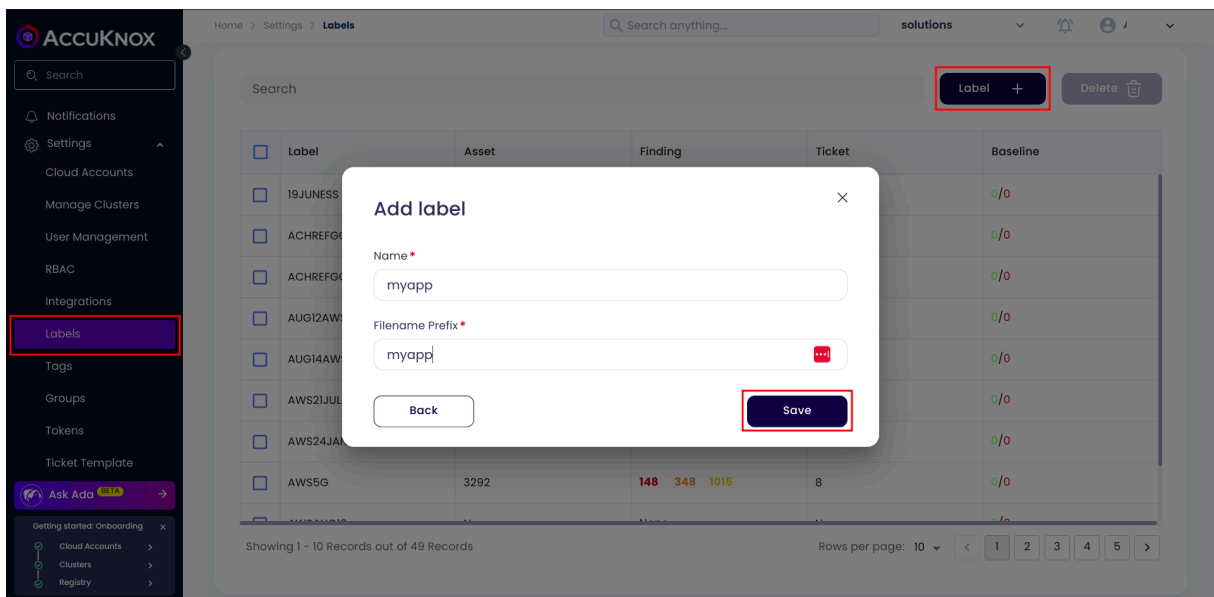


The screenshot shows the 'CI/CD Settings' page in AccuKnox. On the left is a navigation sidebar with 'CI/CD' selected. The main content area has a heading 'Variables can have several attributes. Learn more.' followed by a list of attributes: Protected, Masked, and Expanded. Below this is a table of 'CI/CD Variables' with 4 entries:

Key ↑	Value	Environments	Actions
ACCUKNOX_TOKEN <small>Masked Expanded</small>	*****	All (default)	[Edit] [Delete]
ACCUKNOX_URL <small>Expanded</small>	*****	All (default)	[Edit] [Delete]
SCAN_URL <small>Expanded</small>	*****	All (default)	[Edit] [Delete]
TENANT_ID <small>Expanded</small>	*****	All (default)	[Edit] [Delete]

Below the table is a section for 'Group variables (inherited)' which is currently empty, showing 'There are no variables yet.'

**Step 3:** To create a label, navigate to AccuKnox > Settings > Labels, assign a name to your label, click the save button, and then configure it as a GitLab CI/CD variable



The screenshot shows the 'Labels' configuration page in AccuKnox. A modal window titled 'Add label' is open in the center. The modal contains the following fields and buttons:

- Name \***: Input field containing 'myapp'
- Filename Prefix \***: Input field containing 'myapp'
- Buttons**: 'Back' and 'Save' buttons.

The background shows a table of existing labels with columns for Label, Asset, Finding, Ticket, and Baseline. A 'Label +' button is visible in the top right of the table area.



This is a close-up of a single variable entry from the CI/CD Variables table. It shows:

- Key:** LABEL
- Value:** \*\*\*\*\*
- Environments:** All (default)
- Actions:** [Edit] [Delete]
- Status:** Expanded

## Step 4: Set Up GitLab CI/CD Pipeline

Create a new pipeline in your GitLab project with the following YAML configuration:

```
stages:
  - DAST
  - upload-report
DAST:
  stage: DAST
  image: docker:latest
  services:
    - docker:dind
  script:
    - docker run --rm -v $(pwd):/zap/wrk -t zaproxy/zap-stable zap-full-scan.py
    -t $SCAN_URL -J report.json -I
  artifacts:
    paths:
      - report.json

upload-report-to-accuknox:
  stage: upload-report
  image: curlimages/curl:latest
  dependencies:
    - DAST
  script:
    - |
      curl --location --request POST
      "https://$ACCUKNOX_URL/api/v1/artifact/?tenant_id=$TENANT_ID&label_id=$LABEL_ID&data_type=ZAP&save_to_s3=true" \
        --header "Tenant-Id: $TENANT_ID" \
        --header "Authorization: Bearer $ACCUKNOX_TOKEN" \
        --form "file=@report.json"
    - echo "Checking for critical vulnerabilities..."
    - |
      if grep -q -i -E "HIGH|CRITICAL" report.json; then
        echo "AccuKnox Scan has halted the deployment because it detected
        high/critical vulnerabilities"
        exit 1
      else
        exit 0
      fi
```

## 11.2.4 Initial CI/CD Pipeline Without AccuKnox Scan

Initially, the CI/CD pipeline does not include the AccuKnox scan. When you push changes to the repository, no security checks are performed, potentially allowing security issues in the application.

## 11.2.5 CI/CD Pipeline After AccuKnox Scan Integration

After integrating AccuKnox into your CI/CD pipeline, the next push triggers the CI/CD pipeline. The AccuKnox scan identifies potential vulnerabilities in the application.

```

7  Preparing environment
8  Running on runner-jhcjxvh-project-60898060-concurrent-0 via runner-jhcjxvh-s-l-s-amd64-1724328658-fcf2315
8...
9  Getting source from Git repository
10 Fetching changes with git depth set to 20...
11 Initialized empty Git repository in /builds/affan22/DAST/.git/
12 Created fresh repository.
13 Checking out 594f5322 as detached HEAD (ref is main)...
14 Skipping Git submodules setup
15 $ git remote set-url origin "${CI_REPOSITORY_URL}"
16 Downloading artifacts
17 Downloading artifacts for DAST (7642139528)...
18 Downloading artifacts from coordinator... ok      host=storage.googleapis.com id=7642139528 responseStat
us=200 OK token=glcibt-66
19 Executing "step_script" stage of the job script
20 Using docker image sha256:65019fbb78d5aa95b9ce0fff1fedffebefbe33dfbd886c8aeaabd49be8c909407 for curlimages/
curl:latest with digest curlimages/curl@sha256:8addc281f0ea517409209f76832b6ddc2cab3264feb1ebbec2a2521ffa
d24e4 ...
21 $ curl --location --request POST "https://$ACCUKNOX_URL/api/v1/artifact/?tenant_id=$TENANT_ID&label_id=$LA
BEL&data_type=ZAP&save_to_s3=false" \ # collapsed multi-line command
22   % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
23             Dload  Upload  Total   Spent    Left   Speed
24 100 32333 100    39 100 32294    83  68741  --:--:--  --:--:--  --:--:--  68793
25 {"detail":"File received successfully"}$ echo "Checking for critical vulnerabilities..."
26 Checking for critical vulnerabilities...
27 $ if grep -q -i -E "HIGH|CRITICAL" report.json; then # collapsed multi-line command
28 AccuKnox Scan has halted the deployment because it detected high/critical vulnerabilities
29 Cleaning up project directory and file based variables
30 ERROR: Job failed: exit code 1

```

## 11.2.6 View Results in AccuKnox SaaS

**Step 1:** After the workflow completes, navigate to the AccuKnox SaaS dashboard.

**Step 2:** Go to **Issues > Findings** and select **DAST Findings** to see identified vulnerabilities.

Last seen	Name	Assetname	Description	Risk factor	Location
2024-08-19 11:51:51	Storable and Cacheabl...	https://juice-shop.hero...	<p>The response conte...	Informational	https://juice-shop.hero
2024-08-19 11:51:51	Re-examine Cache-co...	https://juice-shop.hero...	<p>The cache-control h...	Informational	https://juice-shop.hero
2024-08-19 11:51:51	Storable but Non-Cach...	https://juice-shop.hero...	<p>The response conte...	Informational	https://juice-shop.hero
2024-08-19 11:51:51	Cross-Domain Misconfi...	https://juice-shop.hero...	<p>Web browser data L...	Medium	https://juice-shop.hero
2024-08-19 11:51:51	Storable but Non-Cach...	https://juice-shop.hero...	<p>The response conte...	Informational	https://juice-shop.hero
2024-08-19 11:51:51	Timestamp Disclosure -...	https://juice-shop.hero...	<p>A timestamp was di...	Low	https://juice-shop.hero
2024-08-19 11:51:51	Cross-Domain Misconfi...	https://juice-shop.hero...	<p>Web browser data L...	Medium	https://juice-shop.hero
2024-08-19 11:51:51	Deprecated Feature Pol...	https://juice-shop.hero...	<p>The header has now...	Low	https://juice-shop.hero
2024-08-19 11:51:51	Re-examine Cache-co...	https://juice-shop.hero...	<p>The cache-control h...	Informational	https://juice-shop.hero

**Step 3:** Click on a vulnerability to view more details.

**Cross-Domain Misconfiguration** Medium

**Description** | Result | Solution | References | Source Code

<p>Web browser data loading may be possible, due to a Cross Origin Resource Sharing (CORS) misconfiguration on the web server.</p>

Finding for in resource [WebApp | https://juice-shop.herokuapp.com](https://juice-shop.herokuapp.com)

Failing since on 19/08/2024

Last detected on 19/08/2024

**Compliance Frameworks**  
Coming Soon...

**Asset Information**  
{}

**Details** [+ Create Ticket](#)

**Asset**  
https://juice-shop.herokuapp.com

**Asset Type**  
WebApp

**Status**

**Ignored**  
 No

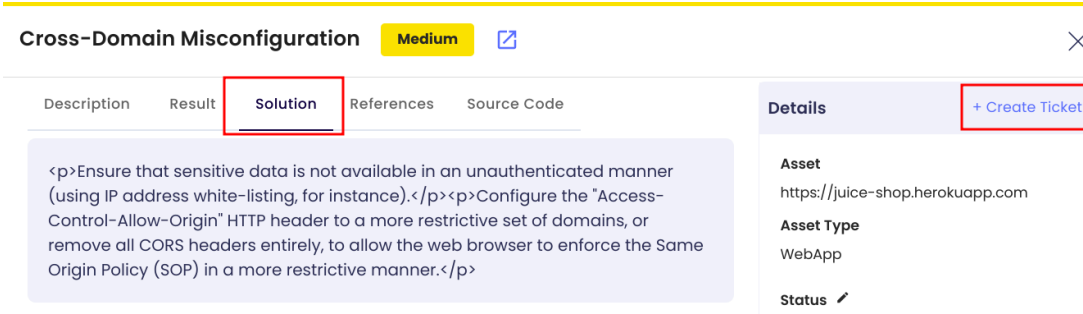
**Severity**

**Tickets**  
0

**Notes**  
Add Comments and Press Ctrl + Enter to Submit

## Step 4: Fix the Vulnerability

Follow the instructions in the Solutions tab to fix the vulnerability (e.g., Cross-Domain Misconfiguration).



**Cross-Domain Misconfiguration** Medium

Description Result **Solution** References Source Code

**Details** + Create Ticket

**Asset**  
https://juice-shop.herokuapp.com

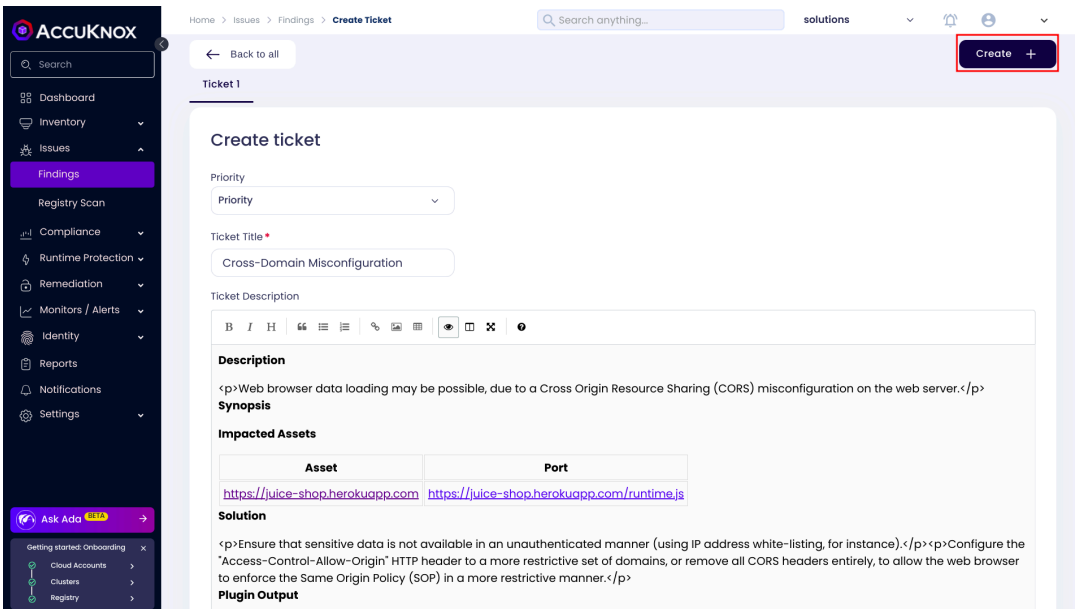
**Asset Type**  
WebApp

**Status** ✓

Ensure that sensitive data is not available in an unauthenticated manner (using IP address white-listing, for instance). Configure the "Access-Control-Allow-Origin" HTTP header to a more restrictive set of domains, or remove all CORS headers entirely, to allow the web browser to enforce the Same Origin Policy (SOP) in a more restrictive manner.

## Step 5: Create a Ticket for Fixing the Vulnerability

Create a ticket in your issue tracking system to address the identified vulnerability.



Home > Issues > Findings > Create Ticket

Search anything...

Back to all Create +

Ticket 1

**Create ticket**

Priority  
Priority

Ticket Title  
Cross-Domain Misconfiguration

Ticket Description

**Description**  
<p>Web browser data loading may be possible, due to a Cross Origin Resource Sharing (CORS) misconfiguration on the web server.</p>

**Synopsis**

**Impacted Assets**

Asset	Port
<a href="https://juice-shop.herokuapp.com">https://juice-shop.herokuapp.com</a>	<a href="https://juice-shop.herokuapp.com/runtime.js">https://juice-shop.herokuapp.com/runtime.js</a>

**Solution**  
<p>Ensure that sensitive data is not available in an unauthenticated manner (using IP address white-listing, for instance).</p><p>Configure the "Access-Control-Allow-Origin" HTTP header to a more restrictive set of domains, or remove all CORS headers entirely, to allow the web browser to enforce the Same Origin Policy (SOP) in a more restrictive manner.</p>

**Plugin Output**

## Step 6: Review Updated Results

- After fixing the vulnerability, rerun the GitLab CI/CD pipeline.
- Navigate to the AccuKnox SaaS dashboard and verify that the vulnerability has been resolved.

## 11.3 Iac GitLab Scan

### 11.3.1 Integrating IaC with AccuKnox in a GitLab CI/CD Pipeline

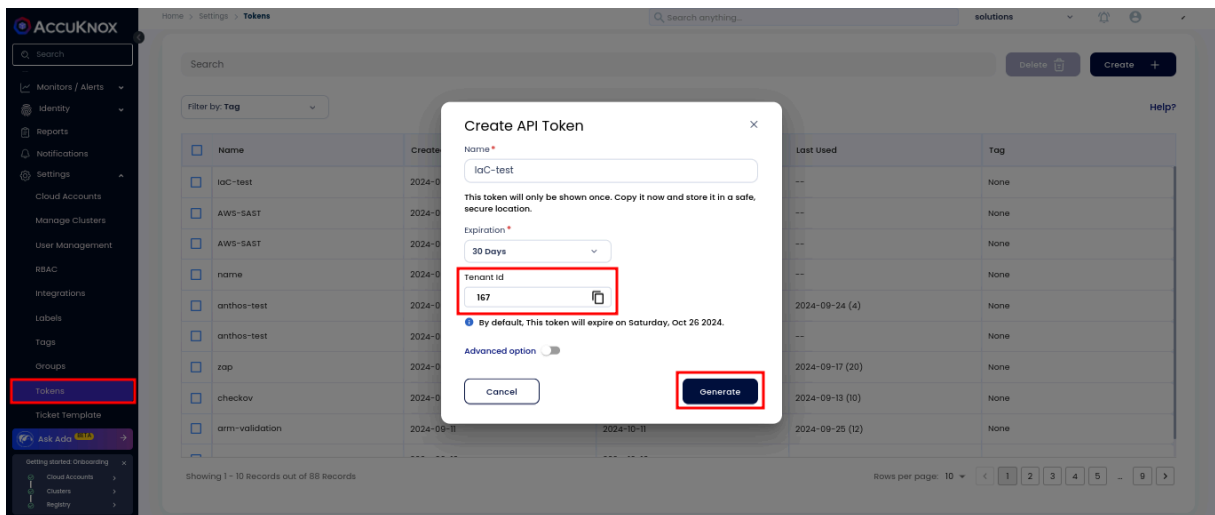
This guide demonstrates how to integrate Infrastructure as Code (IaC) security into a GitLab CI/CD pipeline using AccuKnox. We will implement automated checks to identify configuration vulnerabilities in your IaC templates and send the results to AccuKnox for thorough analysis and remediation. This approach ensures your infrastructure is resilient and aligns with security best practices, effectively minimizing deployment risks.

### 11.3.2 Pre-requisites

- GitLab Access
- AccuKnox UI Access

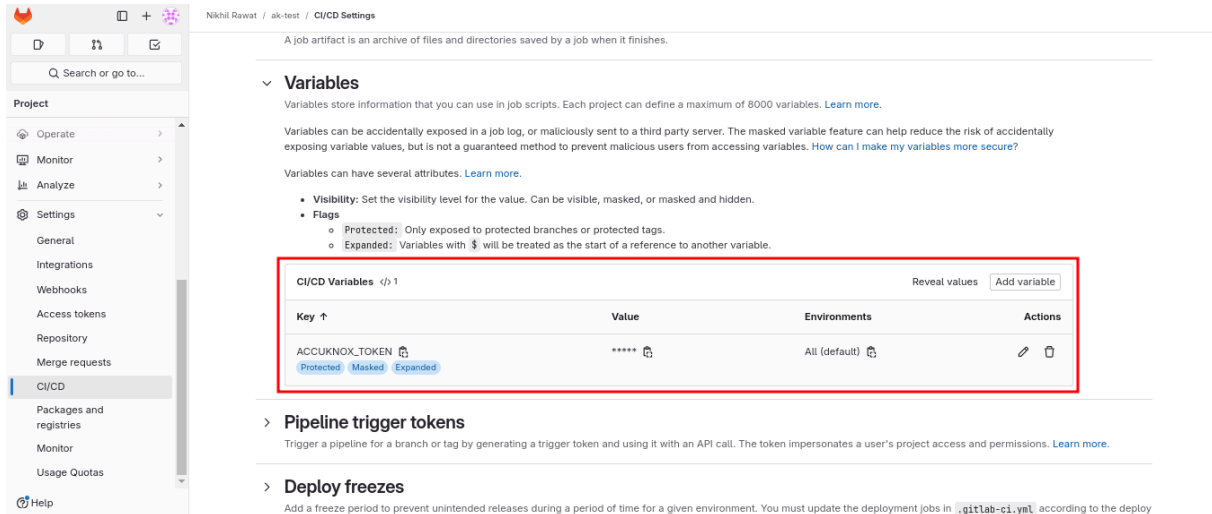
### 11.3.3 Steps for Integration

**Step 1:** Log in to Accuknox Navigate to Settings and select Tokens to create a token for forwarding scan results to Accuknox SaaS. Additionally tenant ID can also be found there which would be helpful for later use.

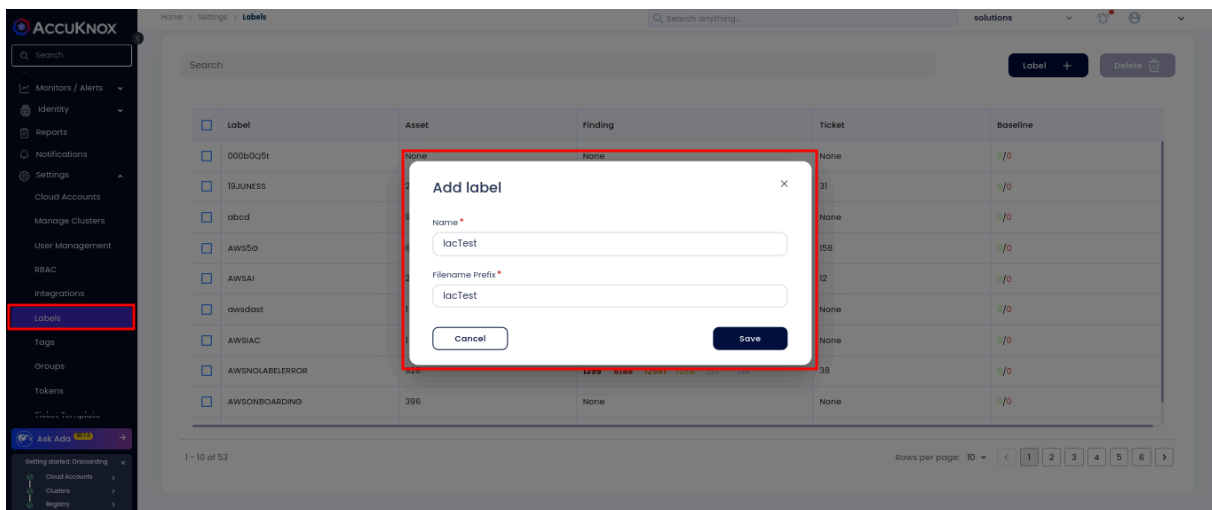


**Step 2:** Copy the token and create a GitLab CI/CD masked variable for the token to be used in the pipeline.





**Step 3:** In order to create a label Go to Accuknox > Settings > Labels and create a label, this label is needed to be configured in the pipeline configuration later.



**Step 4:** Set Up GitLab CI/CD Pipeline

Create a new pipeline in your GitLab project, and add the following YAML configuration. Update the variables and configurations to match your own project values.

```
stages:
  - scan
  - process
  - deploy
```

```
variables:
  GITLAB_SERVER_URL: 'https://gitlab.com'
  GITLAB_REPOSITORY: 'nikhil120/ak-test' # Update to your GitLab repository
  CSPM_URL: 'cspm.demo.accuknox.com' # Replace with your CSPM endpoint
  TENANT_ID: '000' # Set your unique tenant ID
  ACCUKNOX_API_TOKEN: $ACCUKNOX_API_TOKEN # Ensure this environment variable is
set

clone_repo:
  stage: scan
  script:
    - apt-get update
    - apt-get install -y python3 python3-venv python3-pip jq
    - python3 -m venv venv
    - source venv/bin/activate
    - pip install --upgrade pip
    - pip install checkov
    - git clone https://gitlab.com/${GITLAB_REPOSITORY}.git AccuKnox_Iac
    - checkov -d AccuKnox_Iac --output json > checkov_report.json || true
  artifacts:
    paths:
      - checkov_report.json

process_report:
  stage: process
  script:
    - echo "Setting up the environment"
    - apt-get update
    - apt-get install -y python3 python3-venv python3-pip jq
    - echo "Checkov scan complete."
    - ls -al
    - echo "Manipulating JSON report..."
    - |
      # Define repo and branch variables (set these appropriately)
      REPO_LINK="https://gitlab.com/${GITLAB_REPOSITORY}"
      BRANCH_NAME="main" # GitLab predefined variable for branch name

      # Check if the report is empty or not and manipulate JSON
      if [ -s checkov_report.json ]; then
        jq --arg repoLink "$REPO_LINK" --arg branch "$BRANCH_NAME" \
          '. += [{"details": {"repo": $repoLink, "branch": $branch}}]' \
          checkov_report.json > temp.json && \
          mv temp.json checkov_report.json
      else
```

```
    echo "[]" > checkov_report.json # Initialize an empty array if the file
is empty
    jq --arg repoLink "$REPO_LINK" --arg branch "$BRANCH_NAME" \
      '. += [{"details": {"repo": $repoLink, "branch": $branch}}]' \
      checkov_report.json > temp.json && \
      mv temp.json checkov_report.json
  fi
artifacts:
  paths:
    - checkov_report.json

push_report:
  stage: deploy
  image: curlimages/curl:latest
  script:
    - echo "Uploading checkov_report.json to CSPM endpoint..."
    - |
      curl --location --request POST
      "https://${CSPM_URL}/api/v1/artifact/?tenant_id=${TENANT_ID}&data_type=IAC&label
      _id=iactest&save_to_s3=false" \
        --header "Tenant-Id: ${TENANT_ID}" \
        --header "Authorization: Bearer ${ACCUKNOX_API_TOKEN}" \
        --form "file=@checkov_report.json"
  dependencies:
    - process_report
```

### Configuration Notes:

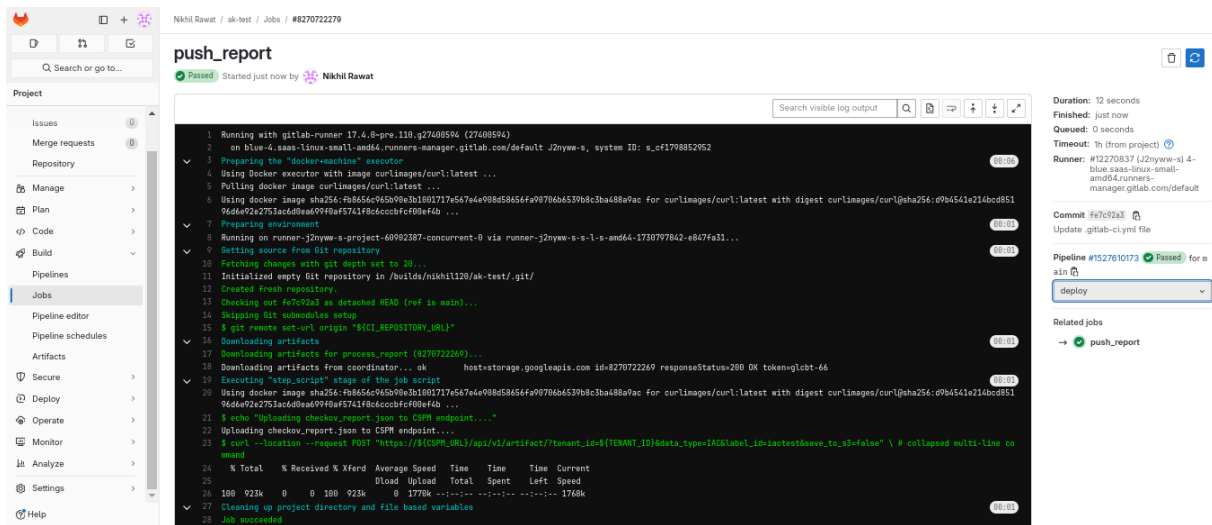
- **GITLAB\_REPOSITORY:** Replace this with the path to your GitLab repository in the format username/repository-name. Replace with your actual GitLab repository path.
- **CSPM\_URL:** Set this to your specific AccuKnox CSPM endpoint URL.
- **TENANT\_ID:** Set your tenant ID here.
- **ACCUKNOX\_API\_TOKEN:** Ensure this variable is stored securely in GitLab CI/CD settings.
- **LABEL\_ID:** Customize this label to identify the scan report (e.g. "iac-test"). This label helps categorize and retrieve reports within Accuknox CSPM.

### 11.3.4 Initial CI/CD Pipeline Without AccuKnox IaC Scan

Initially, the CI/CD pipeline does not include the AccuKnox IaC scan. When changes are pushed to the repository, no infrastructure security checks are performed, potentially allowing misconfigurations or vulnerabilities in the IaC code.

### 11.3.5 CI/CD Pipeline After AccuKnox IaC Scan Integration

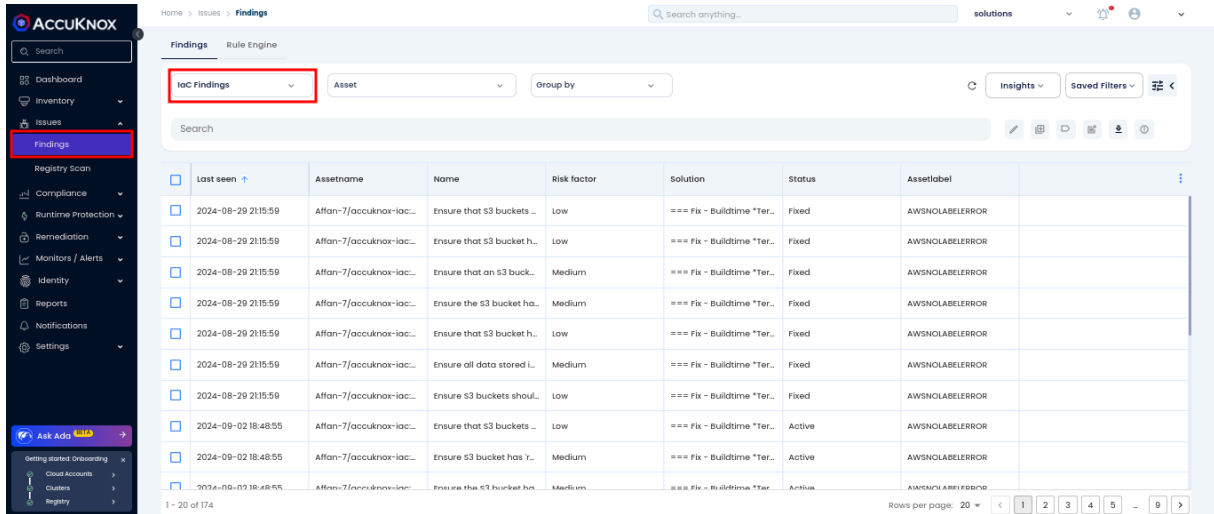
Once the AccuKnox IaC scan is integrated into the CI/CD pipeline, every push triggers an IaC security scan. This scan identifies potential security vulnerabilities or misconfigurations in the infrastructure code, enhancing security prior to deployment. The findings are then sent to the AccuKnox platform.



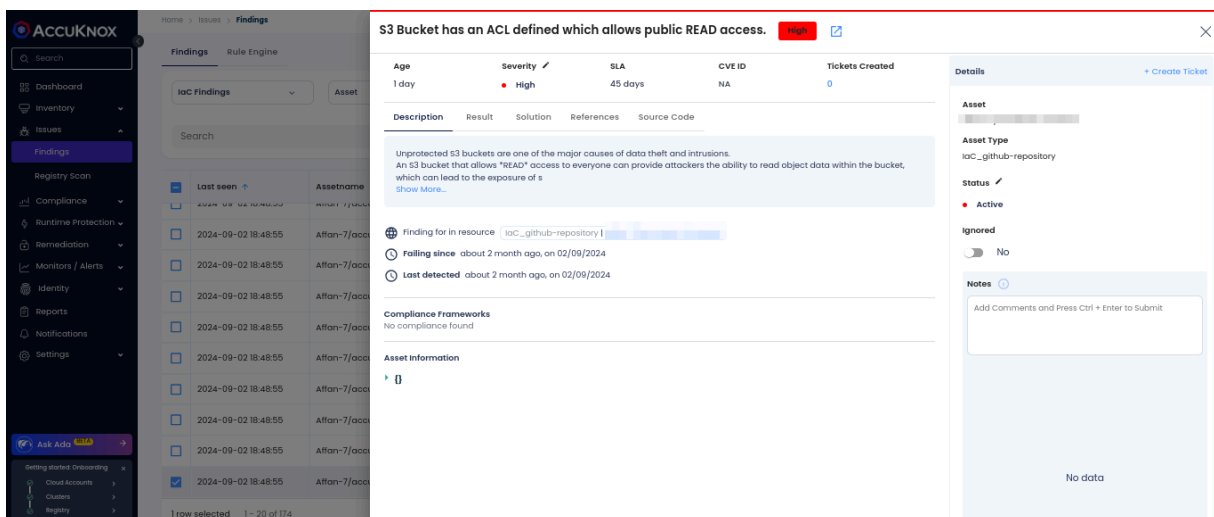
### 11.3.6 View Results in AccuKnox SaaS

**Step 1:** After the pipeline completes, navigate to the Accuknox SaaS dashboard.

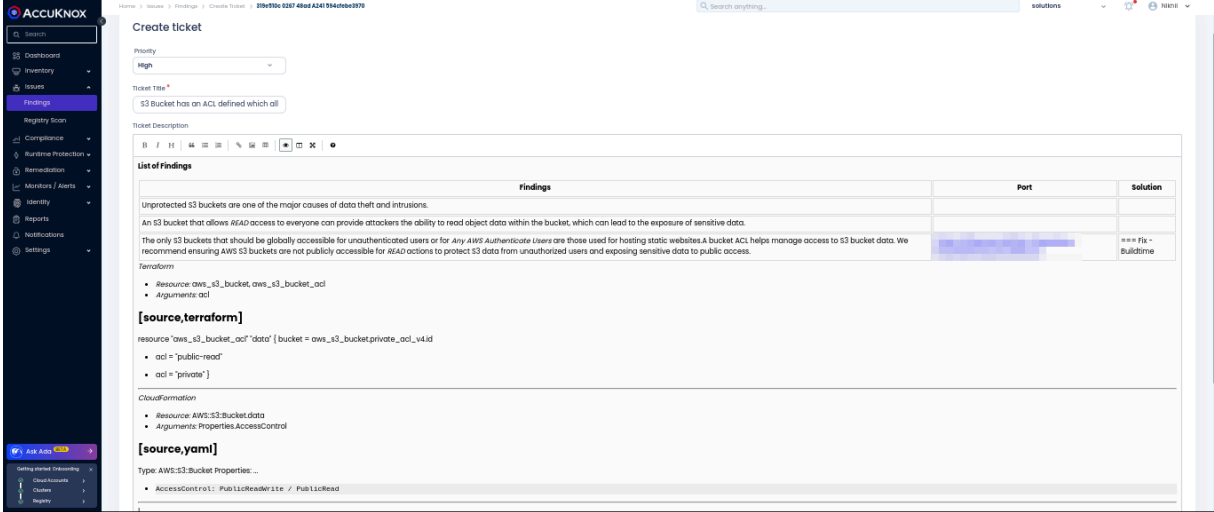
**Step 2:** Go to **Issues > Findings** and select **IaC Findings** to see identified vulnerabilities.



**Step 3:** Click on a vulnerability to view more details and follow the instructions in the Solutions tab.



**Step 4:** For unresolved vulnerabilities, create a ticket in your issue tracking system.



**Create ticket**

Priority: High

Ticket Title: S3 Bucket has an ACL defined which all

Ticket Description:

**List of Findings**

Findings	Port	Solution
Unprotected S3 buckets are one of the major causes of data theft and intrusions. An S3 bucket that allows ATAD access to everyone can provide attackers the ability to read object data within the bucket, which can lead to the exposure of sensitive data. The only S3 buckets that should be globally accessible for unauthenticated users or for Any AWS Authenticate Users are those used for hosting static websites. A bucket ACL helps manage access to S3 bucket data. We recommend ensuring AWS S3 buckets are not publicly accessible for ATAD actions to protect S3 data from unauthorized users and exposing sensitive data to public access.		*** Fix - Buildtime

**Terraform**

- Resource: aws\_s3\_bucket, aws\_s3\_bucket\_acl
- Arguments: acl

**[source,terraform]**

```
resource "aws_s3_bucket_acl" "data" { bucket = aws_s3_bucket.private_acl_v4.id
  acl = "public-read"
  acl = "private" }
```

**CloudFormation**

- Resource: AWS::S3::Bucket.data
- Arguments: Properties.AccessControl

**[source,yaml]**

Type: AWS::S3::Bucket Properties: ...

- AccessControl: PublicReadWrite / PublicRead

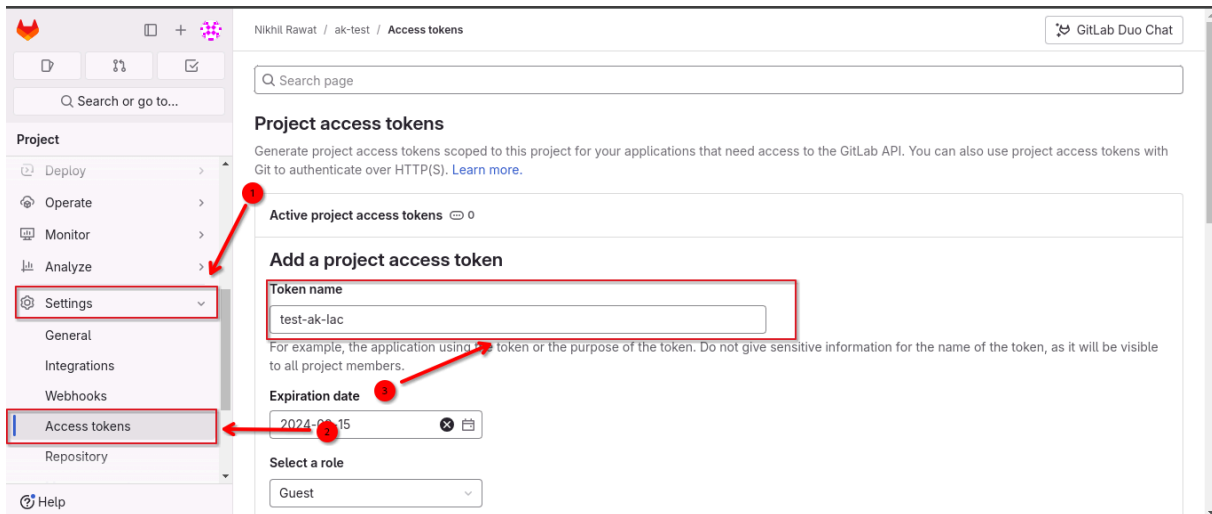
**Step 5:** After fixing the vulnerabilities, rerun the GitLab CI/CD pipeline and verify that the issues have been resolved in the AccuKnox dashboard.

## 11.4 GitLab IaC Scan via Accuknox

This guide demonstrates how to secure a CI/CD pipeline in GitLab using Accuknox to enhance security for Infrastructure as Code (IaC). We will identify code vulnerabilities and send the results to AccuKnox for analysis and remediation.

### 11.4.1 Prerequisites

1. **Public Repository:**
  - You only need the repository URL containing the IaC files.
2. **Private Repository:**
  - Go to your GitLab repository. Navigate to Settings > Access Tokens to get the token.



- Add a new token with `read_repository` as the scope and assign the role as Reporte

**Select a role**

Reporter

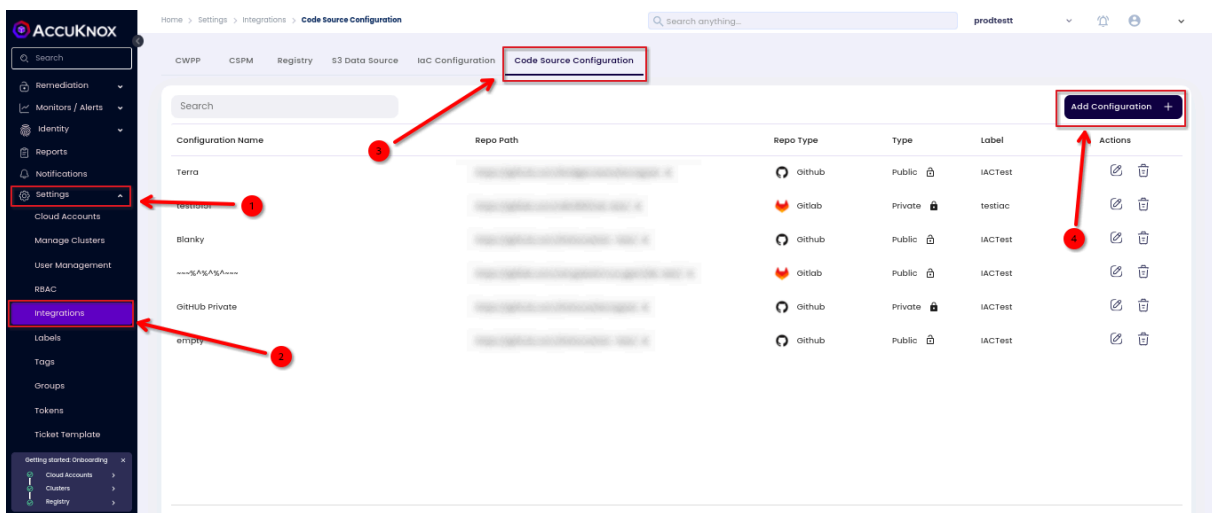
**Select scopes**

Scopes set the permission levels granted to the token. [Learn more.](#)

- api  
Grants complete read and write access to the scoped project API, including the container registry, the dependency proxy, and the package registry.
- read\_api  
Grants read access to the scoped project API, including the Package Registry.
- create\_runner  
Grants create access to the runners.
- manage\_runner  
Grants access to manage the runners.
- k8s\_proxy  
Grants permission to perform Kubernetes API calls using the agent for Kubernetes.
- read\_repository  
Grants read access (pull) to the repository.
- write\_repository  
Grants read and write access (pull and push) to the repository.
- read\_registry

## 11.4.2 Configuring Code Source in Accuknox

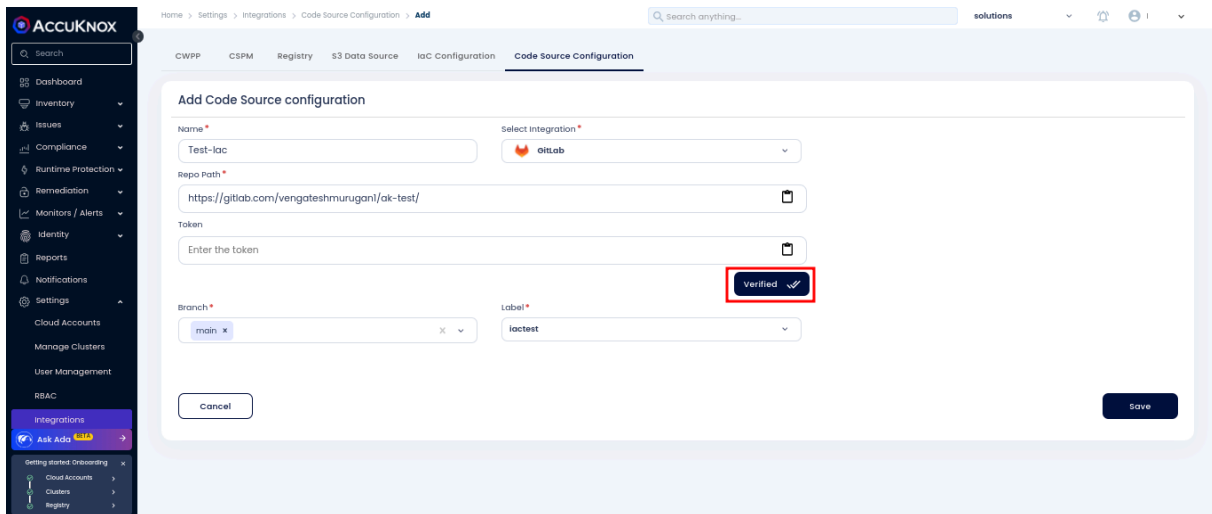
1. Go to Settings > Integration > Code Source Configuration on the Accuknox platform.



2. Enter the repository path:
  - **Public Repository:** No token is needed.
  - **Private Repository:** Enter the previously created access token.
3. Click on Test to verify the configuration and ensure there are no errors.
4. Select the branch type and label.

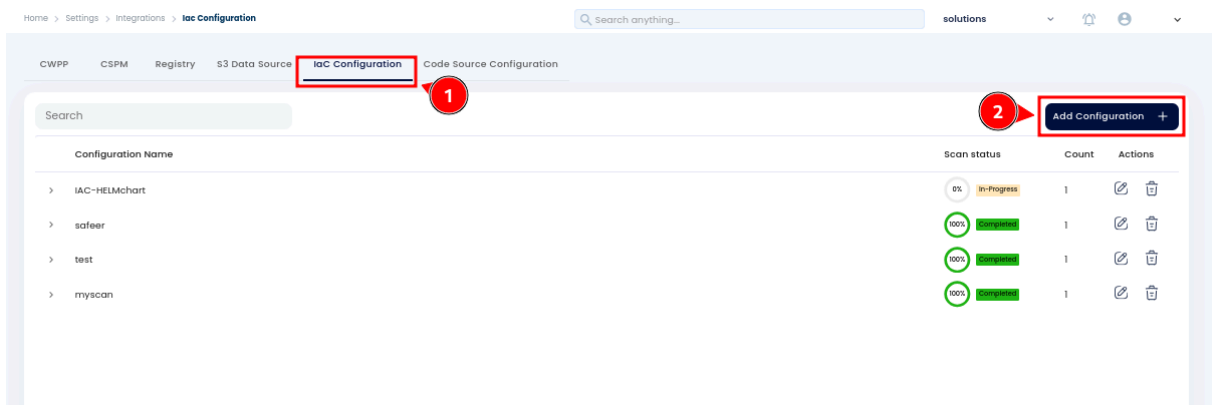


## 5. Save the configuration.

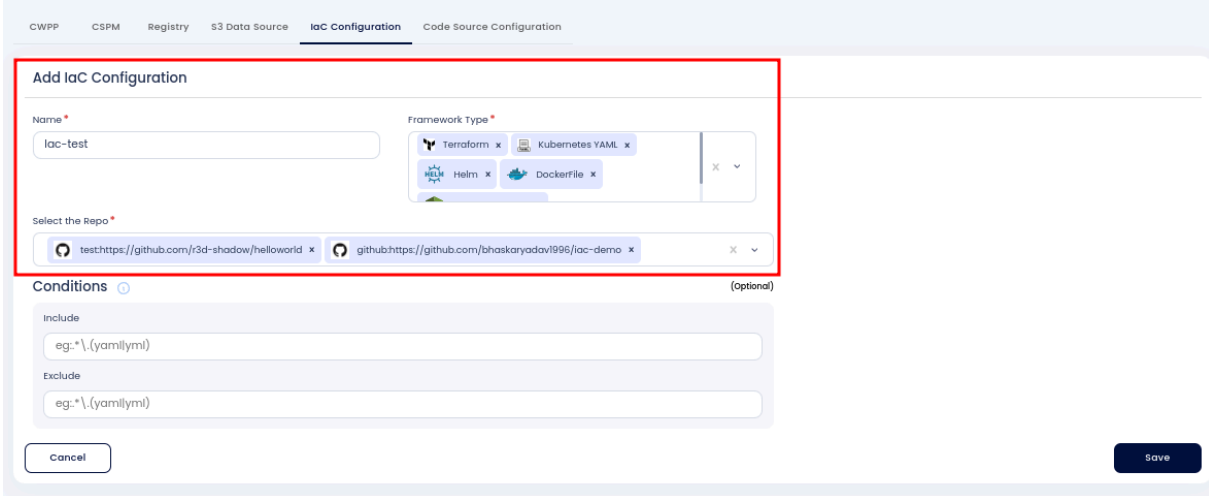


## 11.4.3 Setting Up IaC Configuration

1. Navigate to the IaC Configuration tab.
2. Click on Add Configuration.



1. Fill in the following details:
  - **Integration Name:** Provide a name for this integration.
  - **Framework Type:** Select the file types you want to scan in the repository (e.g., Terraform, Helm, Dockerfile).
2. Select the repository from the dropdown menu that you previously added.



**Add IaC Configuration**

Name \*  
iac-test

Framework Type \*  
Terraform x Kubernetes YAML x  
Helm x Dockerfile x

Select the Repo \*  
test: https://github.com/r3d-shadow/helloworld x  
github: https://github.com/bhaskaryadav1996/iac-demo x

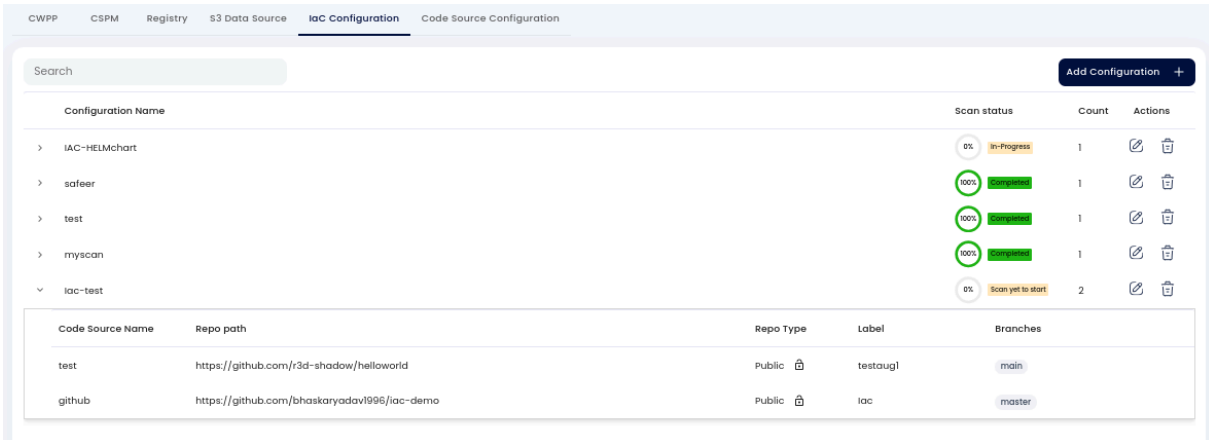
Conditions (optional)

Include  
eg: \*.yaml

Exclude  
eg: \*.yaml

Cancel Save

1. Under the conditions which is an Optional field, you can **include** or **exclude** specific files from the scan.
2. Save the configuration.

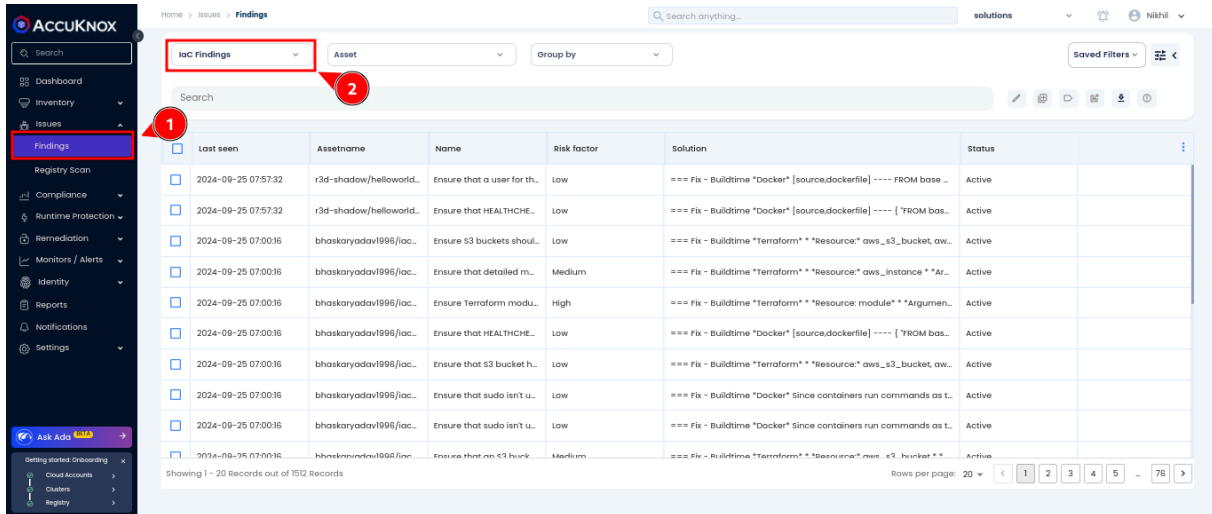


Configuration Name	Scan status	Count	Actions
IAC-HELMchart	0% In-Progress	1	[Edit] [Delete]
safer	100% Completed	1	[Edit] [Delete]
test	100% Completed	1	[Edit] [Delete]
myscan	100% Completed	1	[Edit] [Delete]
iac-test	0% Scan yet to start	2	[Edit] [Delete]

Code Source Name	Repo path	Repo Type	Label	Branches
test	https://github.com/r3d-shadow/helloworld	Public	testaug1	main
github	https://github.com/bhaskaryadav1996/iac-demo	Public	iac	master

## 11.4.4 Viewing and Managing IaC Findings on Accuknox

1. On the Accuknox platform, navigate to Issues > Findings.
2. Select the findings type as IaC Findings.
3. Add the appropriate labels to filter and view the specific IaC findings.



## 11.5 Container Scan Use Case

To show how incorporating AccuKnox into a CI/CD pipeline with Gitlab can improve security, let's look at a detailed example involving a Docker image that initially had known vulnerabilities. By running AccuKnox scanning in the pipeline, we can find and fix these vulnerabilities before deploying the image. The following narrative illustrates this process by comparing the situations before and after adding AccuKnox, as seen in the Gitlab jobs log.

### 11.5.1 Scenario Before Integrating AccuKnox

**Context:** We started with a Docker image built from a Dockerfile using an outdated base image (python:alpine) that contained many known security vulnerabilities. Using this old base image unintentionally introduced many security weaknesses to the Docker image.

#### Dockerfile Example

FROM python:alpine

#### Hypothetical GitLab jobs Log - Pre AccuKnox Scan:

```
Building Docker image...
Image built successfully: your-image:latest
Pushing your-image:latest to Docker Hub...
Image pushed successfully.
```

## 11.5.2 Scenario After Integrating AccuKnox

**Enhancing the GitLab Workflow:** We then added a step to our GitLab workflow to run the AccuKnox vulnerability scan on the newly built Docker image.

### Updated GitLab Workflow Snippet (Incorporating AccuKnox Scan):

```
build:
  stage: build
  script:
    - echo "Logging into Docker..."
    - echo "$DOCKER_LOGIN_PASSWORD" | docker login -u "$DOCKER_LOGIN_USER"
    --password-stdin
    - echo "Building Docker image..."
    - docker build . -t $IMAGE_NAME
    - docker images
    - echo "Running AccuKnox Container Scanner..."
    - docker run --rm -v /var/run/docker.sock:/var/run/docker.sock
    $SCAN_IMAGE_NAME image $IMAGE_NAME --format json >> report.json
  artifacts:
    paths:
      - report.json
    expire_in: 1 hour
```

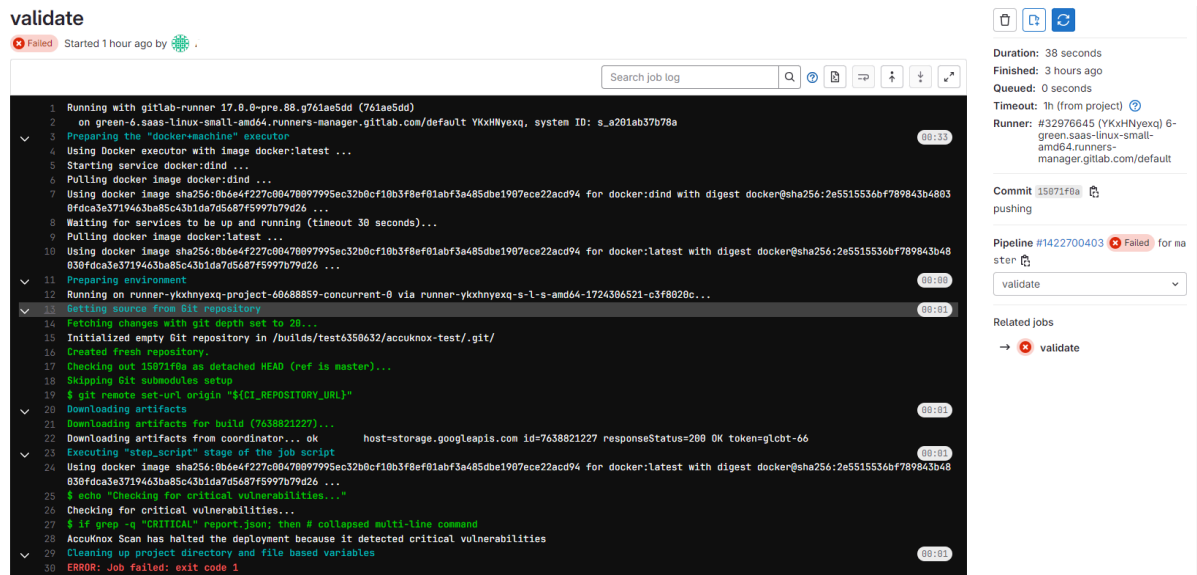
### GitLab Jobs Log - Post AccuKnox Integration:

```
Preparing environment
00:00
Running on runner-ykxhnyexq-project-60688859-concurrent-0 via
runner-ykxhnyexq-s-l-s-amd64-1724306521-c3f8020c...
Getting source from Git repository
00:01
Fetching changes with git depth set to 20...
Initialized empty Git repository in /builds/test6350632/accuknox-test/.git/
Created fresh repository.
Checking out 15071f0a as detached HEAD (ref is master)...
Skipping Git submodules setup
$ git remote set-url origin "${CI_REPOSITORY_URL}"
Downloading artifacts
00:01
Downloading artifacts for build (7638821227)...
```

```

Downloading artifacts from coordinator... ok          host=storage.googleapis.com
id=7638821227 responseStatus=200 OK token=glcbt-66
Executing "step_script" stage of the job script
00:01
Using docker image
sha256:0b6e4f227c00470097995ec32b0cf10b3f8ef01abf3a485dbe1907ece22acd94 for
docker:latest with digest
docker@sha256:2e5515536bf789843b4803fdca3e3719463ba85c43b1da7d5687f5997b79d26
...
$ echo "Checking for critical vulnerabilities..."
Checking for critical vulnerabilities...
$ if grep -q "CRITICAL" report.json; then # collapsed multi-line command
AccuKnox Scan has halted the deployment because it detected critical
vulnerabilities
Cleaning up project directory and file based variables
ERROR: Job failed: exit code 1

```



**validate**  
 Failed Started 1 hour ago

Search job log

```

1 Running with gitlab-runner 17.8.0-pre.88.g761ae5dd (761ae5dd)
2 on green-6.saas-linux-small-amd64.runners-manager.gitlab.com/default YKxHMyexq, system ID: s_a201ab37b78a
3 Preparing the "docker-machine" executor
4 Using Docker executor with image docker:latest ...
5 Starting service docker:dind ...
6 Pulling docker image docker:dind ...
7 Using docker image sha256:0b6e4f227c00470097995ec32b0cf10b3f8ef01abf3a485dbe1907ece22acd94 for docker:dind with digest docker@sha256:2e5515536bf789843b4803fdca3e3719463ba85c43b1da7d5687f5997b79d26 ...
8 Waiting for services to be up and running (timeout 30 seconds)...
9 Pulling docker image docker:latest ...
10 Using docker image sha256:0b6e4f227c00470097995ec32b0cf10b3f8ef01abf3a485dbe1907ece22acd94 for docker:latest with digest docker@sha256:2e5515536bf789843b4803fdca3e3719463ba85c43b1da7d5687f5997b79d26 ...
11 Preparing environment
12 Running on runner-ykxhnyexq-project-66888859-concurrent-0 via runner-ykxhnyexq-s-1-s-amd64-1724306521-c3f8020c...
13 Getting source from Git repository
14 Fetching changes with git depth set to 20...
15 Initialized empty Git repository in /builds/test6350632/accuknox-test/.git/
16 Created fresh repository.
17 Checking out 13071f0a as detached HEAD (ref is master)...
18 Skipping Git submodule setup
19 $ git remote set-url origin "${CI_REPOSITORY_URL}"
20 Downloading artifacts
21 Downloading artifacts for build (7638821227)...
22 Downloading artifacts from coordinator... ok          host=storage.googleapis.com id=7638821227 responseStatus=200 OK token=glcbt-66
23 Executing "step_script" stage of the job script
24 Using docker image sha256:0b6e4f227c00470097995ec32b0cf10b3f8ef01abf3a485dbe1907ece22acd94 for docker:latest with digest docker@sha256:2e5515536bf789843b4803fdca3e3719463ba85c43b1da7d5687f5997b79d26 ...
25 $ echo "Checking for critical vulnerabilities..."
26 Checking for critical vulnerabilities...
27 $ if grep -q "CRITICAL" report.json; then # collapsed multi-line command
28 AccuKnox Scan has halted the deployment because it detected critical vulnerabilities
29 Cleaning up project directory and file based variables
30 ERROR: Job failed: exit code 1

```

Duration: 38 seconds  
 Finished: 3 hours ago  
 Queued: 0 seconds  
 Timeout: 1h (from project)  
 Runner: #32976645 (YKxHMyexq) 6-green.saas-linux-small-amd64.runners-manager.gitlab.com/default

Commit 15071f0a pushing

Pipeline #1422700403 Failed for master  
 validate

Related jobs  
 → validate

AccuKnox carefully analyzed the image and found critical and high-severity vulnerabilities. Based on these findings, the workflow stopped and prevented the vulnerable image from being pushed to the Docker registry.

### 11.5.3 Remediation and Rescan

**Fortifying the Dockerfile:** After seeing the vulnerabilities, we updated the Dockerfile to use a newer, more secure base image (python:alpine instead, to fix the security issues.

#### **Dockerfile Post-Update:**

```
FROM python:alpine# Additional image enhancements and setup
```

#### **GitLab Jobs Log - After Remediation:**

Building Docker image...

Image built successfully: your-image:latest

Scanning your-image:latest with AccuKnox...

INF Scanning /path/to/your-image:latest

INF Number of language-specific files: 1

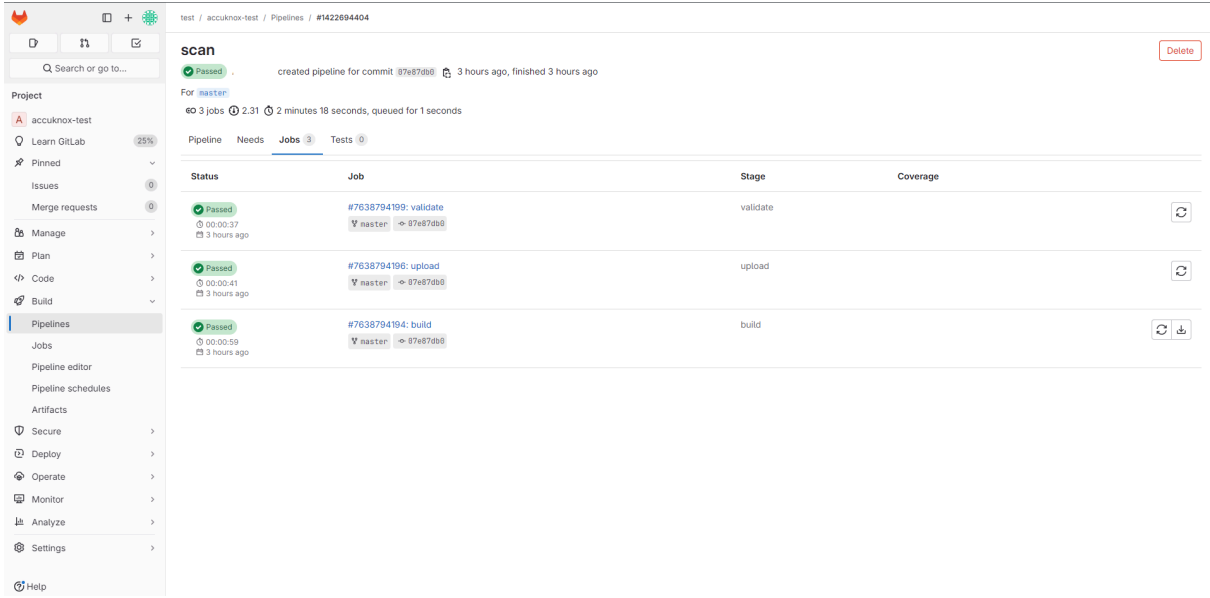
INF No critical vulnerabilities found.

Image scan passed successfully.

Pushing your-image:latest to Docker Hub...

Image pushed successfully.

Once the vulnerabilities were resolved, the AccuKnox scan approved the updated image, allowing it to be safely pushed to the registry. This example clearly shows how important it is to have vulnerability scanning in the pipeline - it prevents insecure images from being deployed to production, ensuring only secure images make it through.



The screenshot shows the GitLab CI/CD interface for a pipeline named 'scan'. The pipeline is in a 'Passed' state and was created for commit '07e870b0' 3 hours ago. It consists of three jobs: 'validate', 'upload', and 'build', all of which have completed successfully. The interface includes a sidebar with navigation options like Project, Pipelines, Jobs, and Settings, and a main area displaying the pipeline details and job status.

Status	Job	Stage	Coverage
Passed	#7638794109: validate	validate	
Passed	#7638794106: upload	upload	
Passed	#7638794104: build	build	

## 11.5.4 Steps needed to be taken for integration

**Step 1:** The user needs to create a GitLab workflow file inside their GitLab repository using the following workflow Template:

```
image: docker:latest # Docker image with Docker installed
```

services:

```
- docker:dind # Docker-in-Docker service for building Docker images
```

variables:

```
IMAGE_NAME: "[tag]/gitlab-pipeline:v1"
SCAN_IMAGE_NAME: "accuknox/accuknox-container-scan"
CSPM_URL: $ACCUKNOX_CSPM_URL
TENANT_ID: $TENANT_ID
DOCKER_LOGIN_USER: $DOCKER_LOGIN_USER
DOCKER_LOGIN_PASSWORD: $DOCKER_LOGIN_PASSWORD
ACCUKNOX_API_TOKEN: $ACCUKNOX_API_TOKEN
```

```
stages:
  - build
  - upload
  - validate
build:
  stage: build
  script:
    - echo "Logging into Docker..."
    - echo "$DOCKER_LOGIN_PASSWORD" | docker login -u "$DOCKER_LOGIN_USER"
--password-stdin
    - echo "Building Docker image..."
    - docker build . -t $IMAGE_NAME
    - docker images
    - echo "Running AccuKnox Container Scanner..."
    - docker run --rm -v /var/run/docker.sock:/var/run/docker.sock
$SCAN_IMAGE_NAME image $IMAGE_NAME --format json >> report.json
  artifacts:
    paths:
      - report.json
    expire_in: 1 hour

upload:
  stage: upload
  image: curlimages/curl:latest
  script:
    - echo "Uploading report.json to CSPM endpoint..."
    - |
      curl --location --request POST
      "https://${CSPM_URL}/api/v1/artifact/?tenant_id=${TENANT_ID}&data_type=TR&save_t
      o_s3=false" \
        --header "Tenant-Id: ${TENANT_ID}" \
        --header "Authorization: Bearer ${ACCUKNOX_API_TOKEN}" \
        --form "file=@\"report.json\""

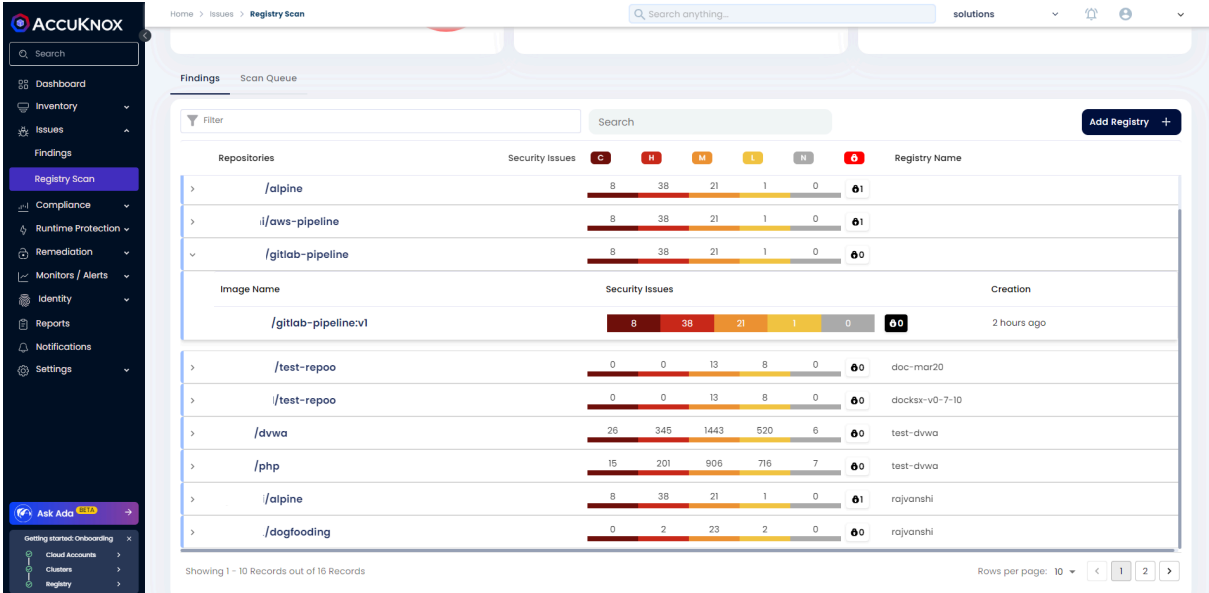
validate:
  stage: validate
  script:
    - echo "Checking for critical vulnerabilities..."
    - |
      if grep -q "CRITICAL" report.json; then
        echo "AccuKnox Scan has halted the deployment because it detected
critical vulnerabilities"
        exit 1
      else
        exit 0
      fi
```



**Note:** In the above template, the user needs to change some variables, including ACCUKNOX\_API\_TOKEN, CSPM\_URL(cspm.demo|stage|dev.accuknox.com), and TENANT\_ID. Values for these variables can be obtained from AccuKnox SaaS.

**Step 2:** Now, when a user attempts to make any changes to their repository, the workflow will be triggered, performing the necessary steps for scanning and posting the results to AccuKnox SaaS.

**Step 3:** Once the scan is complete, the user can go into the AccuKnox SaaS and navigate to Issues → RegistryScan where they can find their repository name and select it to see the associated findings



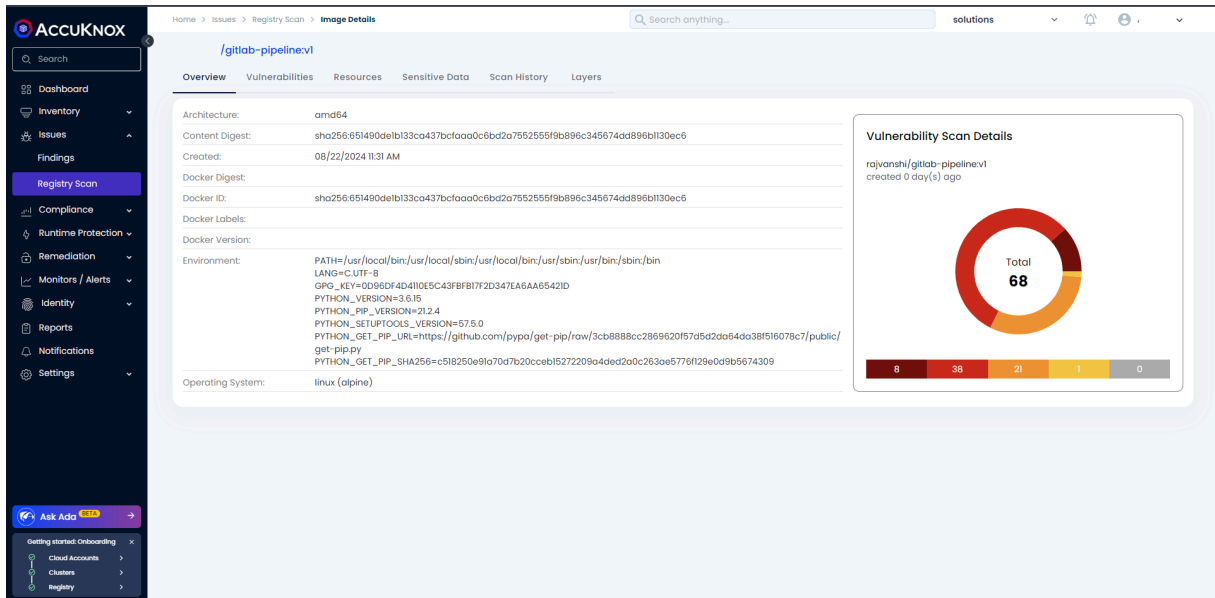
The screenshot shows the AccuKnox web interface. On the left is a dark sidebar with navigation options: Dashboard, Inventory, Issues, Findings, Registry Scan (highlighted), Compliance, Runtime Protection, Remediation, Monitors / Alerts, Identity, Reports, Notifications, and Settings. At the bottom of the sidebar are 'Ask Ada' and 'Getting started Onboarding' sections.

The main content area is titled 'Home > Issues > Registry Scan'. It features a search bar and a 'Filter' dropdown. Below this is a table with columns for 'Repositories', 'Security Issues', and 'Registry Name'. The 'Security Issues' column contains a bar chart and a score. The 'Registry Name' column contains the repository name and a 'Creation' timestamp.

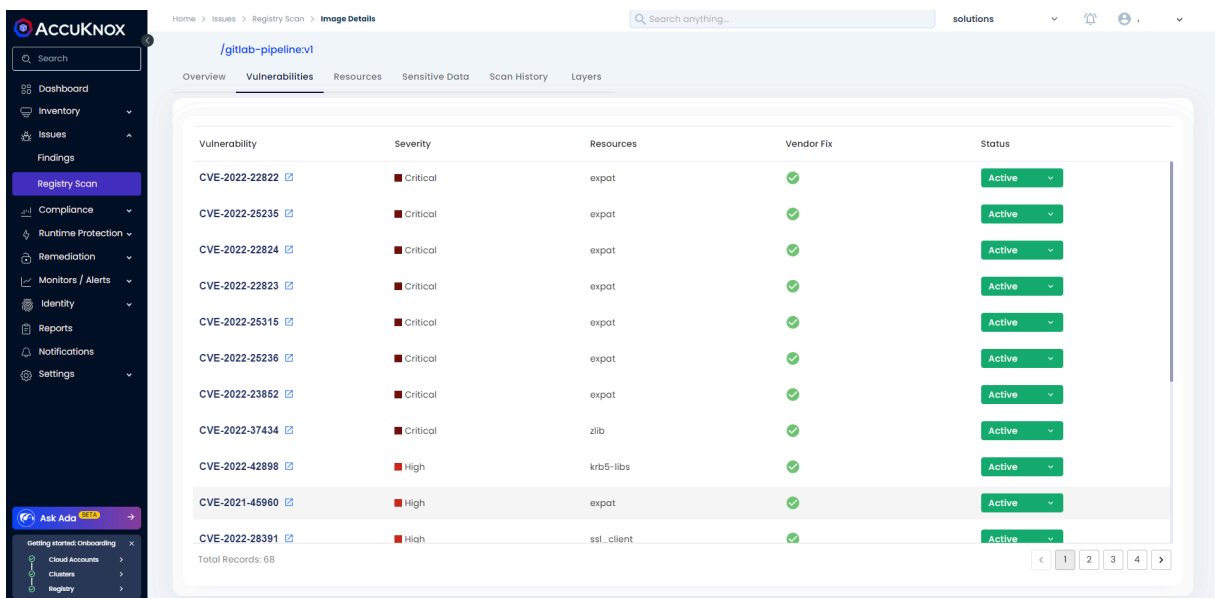
Repositories	Security Issues	Registry Name
> /alpine	8 38 21 1 0 01	
> i/aws-pipeline	8 38 21 1 0 01	
▼ /gitlab-pipeline	8 38 21 1 0 00	
<b>Image Name</b>		
/gitlab-pipeline.v1	8 38 21 1 0 00	2 hours ago
> /test-repoo	0 0 13 8 0 00	doc-mar20
> /test-repoo	0 0 13 8 0 00	docksr-v0-7-10
> /dvwa	26 345 1443 520 6 00	test-dvwa
> /php	15 201 806 716 7 00	test-dvwa
> /alpine	8 38 21 1 0 01	rajvanshi
> /dogfooding	0 2 23 2 0 00	rajvanshi

Showing 1 - 10 Records out of 16 Records. Rows per page: 10. Page 1 of 2.

**Step 4:** After clicking on the image name, the user will be able to see the metadata for the image that was built during the workflow execution.



**Step 5:** In the Vulnerabilities section, the user can see the image-specific vulnerabilities in a list manner that contains relevant information. These findings will also be available in the Issues → Vulnerabilities section where the user can manage these findings with others.



**Step 6:** The Resources section contains information about packages and modules that were used to build the code base into a container image.

Home > Issues > Registry Scan > Image Details

Search anything...

**/gitlab-pipeline.v1**

Overview Vulnerabilities **Resources** Sensitive Data Scan History Layers

Filter by Severity: Critical High Medium Low Negligible

Resource	Type	Version	Fix Version	Vulnerabilities
expat	package	2.4.1-r0	2.4.3-r0	7 9 1 0 0
libcrypto11	package	1.1.1-r7	1.1.1-r0	0 5 7 0 0
libssl1	package	1.1.1-r7	1.1.1-r0	0 5 7 0 0
libretls	package	3.3.4-r2	3.3.4-r3	0 1 0 0 0
setuptools	package	67.6.0	66.5.1	0 2 0 0 0
ssl_client	package	1.34.1-r3	1.34.1-r5	0 1 0 0 0
libcom_err	package	1.46.4-r0	1.46.6-r0	0 1 0 0 0
ncurses-libs	package	6.3_p20211120-r0	6.3_p20211120-r1	0 2 0 0 0
libtirpc-conf	package	1.3.2-r0	1.3.2-r1	0 1 0 0 0
libtirpc	package	1.3.2-r0	1.3.2-r1	0 1 0 0 0

Total Records: 21

**Step 7:** The user can see the scan history of every scan that happened while triggering the workflow.

Home > Issues > Registry Scan > Image Details

Search anything...

**/gitlab-pipeline.v1**

Overview Vulnerabilities Resources Sensitive Data **Scan History** Layers

Scan Date	Image ID	Security Status	Image Creation Date	Scan Results
08/22/2024 11:28 AM	sha256:8e213a4a36223e879af5b26c1d140b6f3eeb4b4f2644...	Passed	08/22/2024 11:27 AM	0 0 0 0 0
08/22/2024 11:24 AM	sha256:2799d144bdc9cb64ee8ce890daab432acc6f95fee7...	Passed	08/22/2024 11:23 AM	8 38 21 1 0
08/22/2024 11:19 AM	sha256:605ed707461b62af25d4351b36f38e38e9e8a0ec6e9e...	Passed	08/22/2024 11:18 AM	8 38 21 1 0
08/22/2024 11:32 AM	sha256:651490delb133ca437bcfaaa0c6bd2a7552555f9b89...	Passed	08/22/2024 11:31 AM	8 38 21 1 0

# 12. KSPM (Kubernetes Security Posture Management)

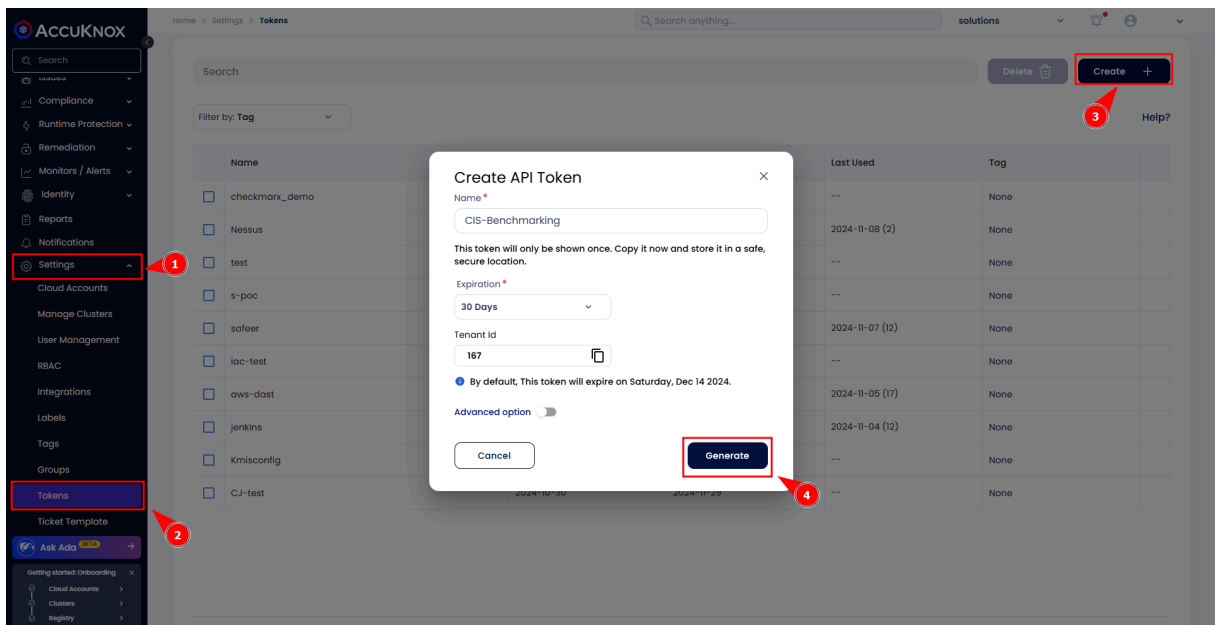
## 12.1 CIS Benchmarking Compliance Scan Onboarding

This guide details the steps to onboard a Kubernetes cluster to Accuknox SaaS for CIS Benchmarking compliance scanning, enabling you to monitor and improve cluster security in line with CIS standards.

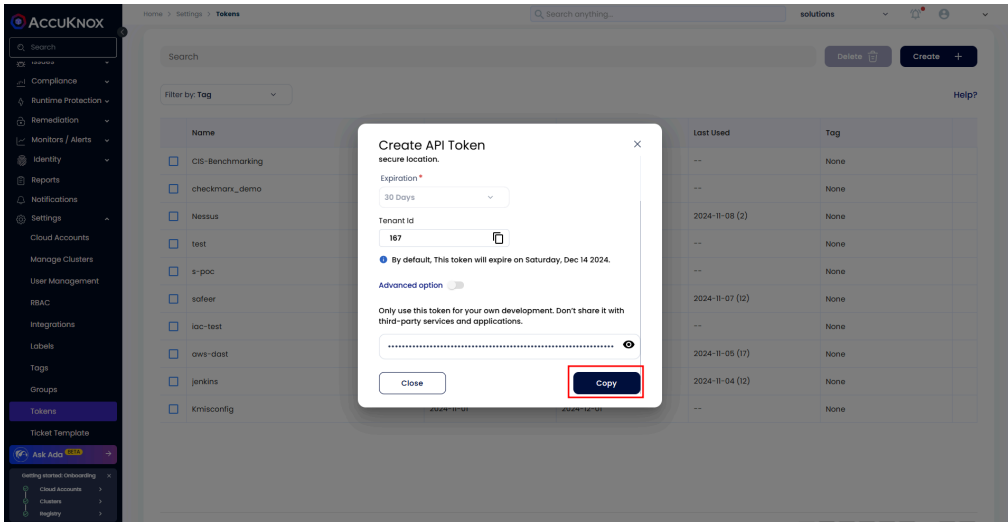
### Step 1: Generate an Access Token

To begin, create a token that will authenticate your cluster for scanning. Follow these steps:

1. Navigate to **Settings > Tokens** in the Accuknox platform and Click on the **Create** button, give your token a descriptive name (e.g., "CIS-Compliance-Token"), and click **Generate**.

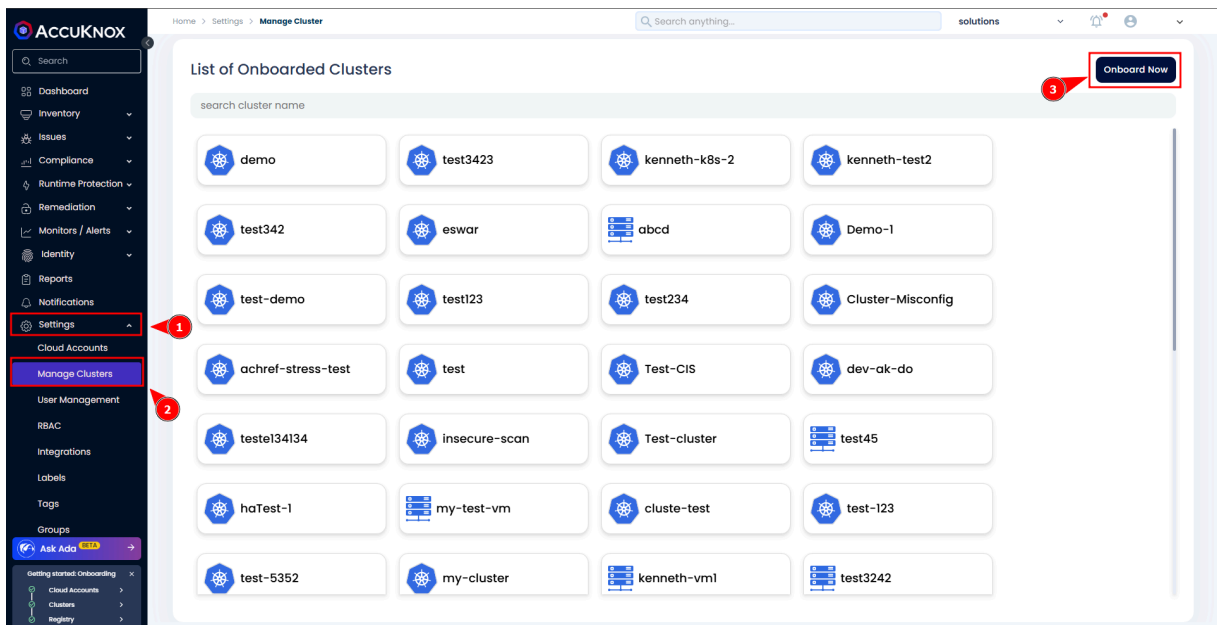


2. Once the token is generated, copy it and securely save it for later use.

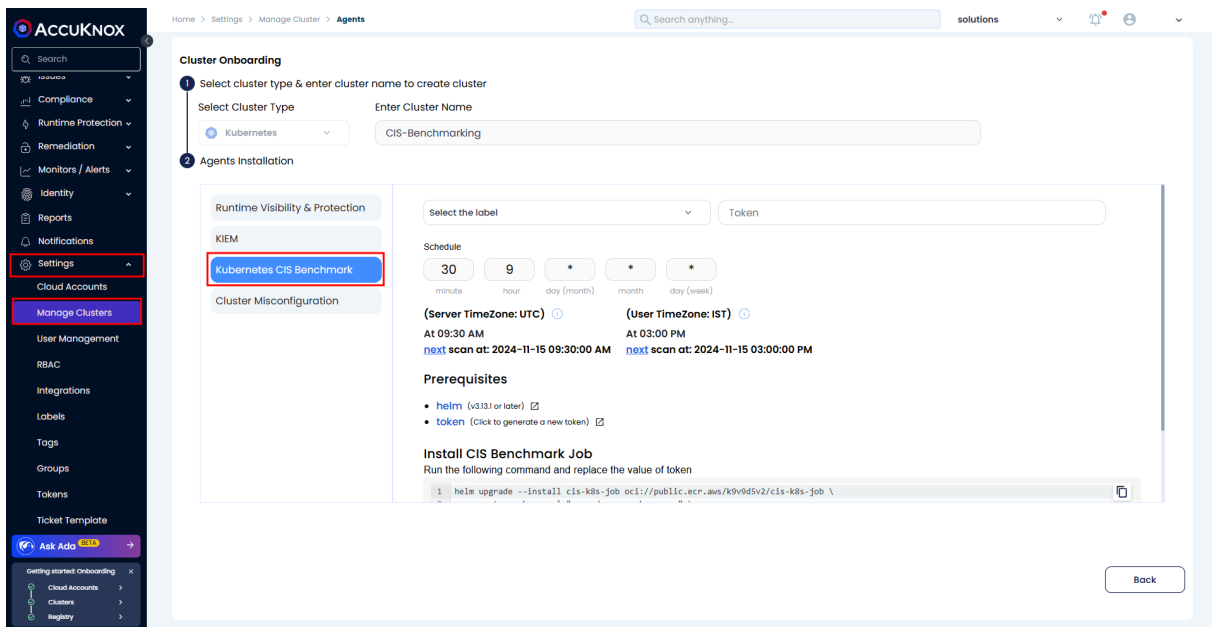


## Step 2: Onboard Your Cluster

1. Go to **Settings > Manage Clusters** and Click **Onboard Now** or select an existing cluster if you're updating a previously onboarded cluster.

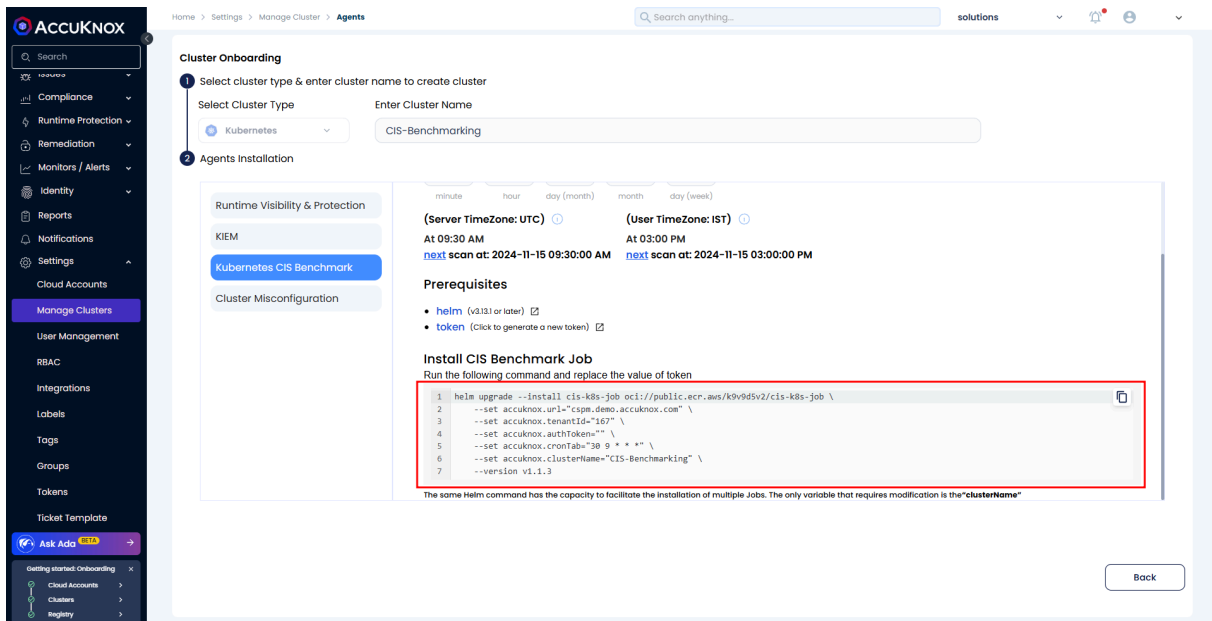


2. Enter a name for your cluster to identify it in Accuknox. From the scan type, choose **CIS Benchmarking**.
3. Select a label for easy identification and paste the token you generated in Step 1. Set a scan schedule based on your requirements. Accuknox will automatically run scans according to the selected schedule



## Step 3: Deploy the Scanner Using Helm

1. Scroll down to the **Helm Command** section and copy the provided command.



2. Run this command in your terminal on a machine that has access to your Kubernetes cluster. The command will schedule the scan for CIS Benchmarking compliance.
3. Once the Helm installation is complete, return to the Accuknox platform and click **Finish**.

## Step 4: View Compliance Findings

After the initial scan is completed, you can view the compliance results:

1. Go to **Issues > Findings** in Accuknox.
2. Use the **Findings** dropdown to filter and select CIS k8s Benchmarking finding results.

Home > Issues > Findings

Search anything...

Findings Rule Engine

CIS K8s Benchmark Findings Group by

Search

<input type="checkbox"/>	Assetname	Test number	Tool output	Cvss score	Description	Solution	Type	Group text
<input type="checkbox"/>	safer	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	Test-k3s	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	DO-demo-cluster	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	insecure-scan	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	DO-demo-cluster	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	rk-k8s-stage	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	Test-k3s	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	safer	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	insecure-scan	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...		
<input type="checkbox"/>	Test-k3s	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...		
<input type="checkbox"/>	safer	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...		
<input type="checkbox"/>	insecure-scan	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...		

1 - 20 of 571 Rows per page: 20

3. Each result will provide details on specific CIS controls and any non-compliant configurations detected.

Home > Issues > Findings

Search anything...

Findings Rule Engine

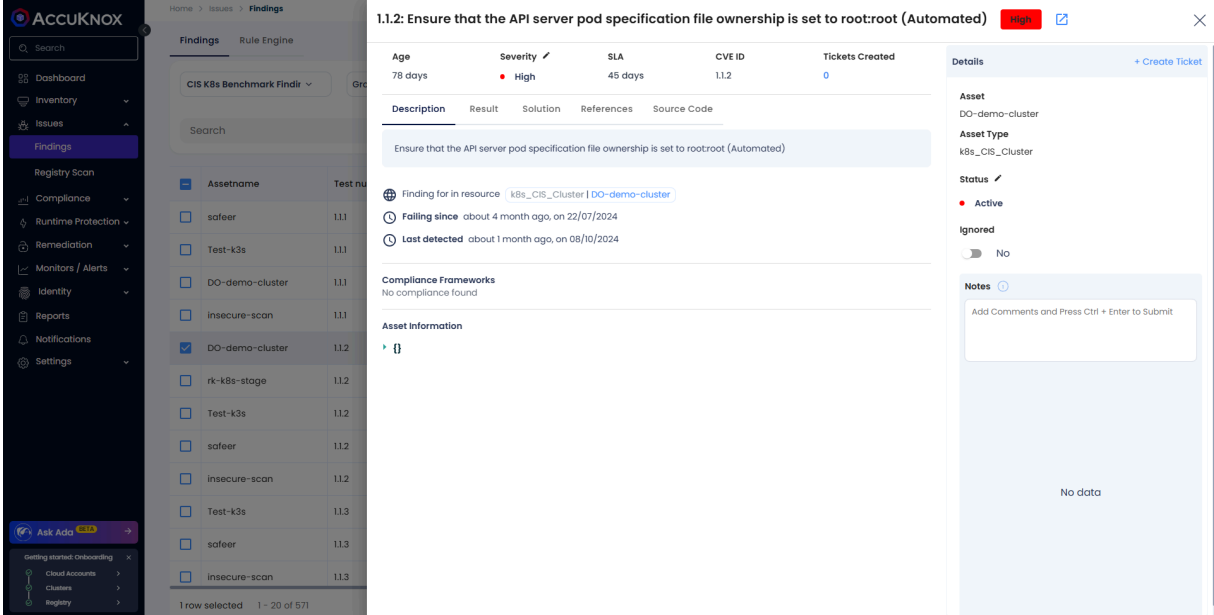
CIS K8s Benchmark Findings Group by

Search

<input type="checkbox"/>	Assetname	Test number	Tool output	Cvss score	Description	Solution	Expected result	Assetlabel
<input type="checkbox"/>	safer	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'permissions' is present	SAFEER
<input type="checkbox"/>	Test-k3s	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'permissions' is present	CISTEST
<input type="checkbox"/>	DO-demo-cluster	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'permissions' is present	KIEMDO
<input type="checkbox"/>	insecure-scan	1.1.1	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'permissions' is present	Test101
<input type="checkbox"/>	DO-demo-cluster	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'rootroot' is present	KIEMDO
<input type="checkbox"/>	rk-k8s-stage	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'rootroot' is present	nessus
<input type="checkbox"/>	Test-k3s	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'rootroot' is present	CISTEST
<input type="checkbox"/>	safer	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'rootroot' is present	SAFEER
<input type="checkbox"/>	insecure-scan	1.1.2	FAILED	0.0	Ensure that the API serv...	Run the below comman...	'rootroot' is present	Test101
<input type="checkbox"/>	Test-k3s	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...	'permissions' is present	CISTEST
<input type="checkbox"/>	safer	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...	'permissions' is present	SAFEER
<input type="checkbox"/>	insecure-scan	1.1.3	FAILED	0.0	Ensure that the controll...	Run the below comman...	'permissions' is present	Test101

1 - 20 of 571 Rows per page: 20





The screenshot displays the Accuknox Findings interface. On the left is a navigation sidebar with options like Dashboard, Inventory, Findings, Registry Scan, Compliance, Runtime Protection, Remediation, Monitors / Alerts, Identity, Reports, Notifications, and Settings. The main area shows a list of findings under the heading '1.1.2: Ensure that the API server pod specification file ownership is set to root:root (Automated)'. The selected finding is 'High' severity, with an age of 78 days and an SLA of 45 days. The description is 'Ensure that the API server pod specification file ownership is set to root:root (Automated)'. The finding is for resource 'k8s\_CIS\_Cluster | DO-demo-cluster'. It is marked as 'Falling since about 4 month ago, on 22/07/2024' and 'Last detected about 1 month ago, on 08/10/2024'. The status is 'Active'. The interface also shows a table of assets and a details panel on the right with a 'Notes' section.

Age	Severity	SLA	CVE ID	Tickets Created
78 days	High	45 days	1.1.2	0

Description	Result	Solution	References	Source Code
Ensure that the API server pod specification file ownership is set to root:root (Automated)				

Assetname	Test nu
safeer	1.1.1
Test-k3s	1.1.1
DO-demo-cluster	1.1.1
insecure-scan	1.1.1
DO-demo-cluster	1.1.2
rk-k8s-stage	1.1.2
Test-k3s	1.1.2
safeer	1.1.2
insecure-scan	1.1.2
Test-k3s	1.1.3
safeer	1.1.3
insecure-scan	1.1.3

Details panel:

- Asset: DO-demo-cluster
- Asset Type: k8s\_CIS\_Cluster
- Status: Active
- Ignored: No

Notes: Add Comments and Press Ctrl + Enter to Submit

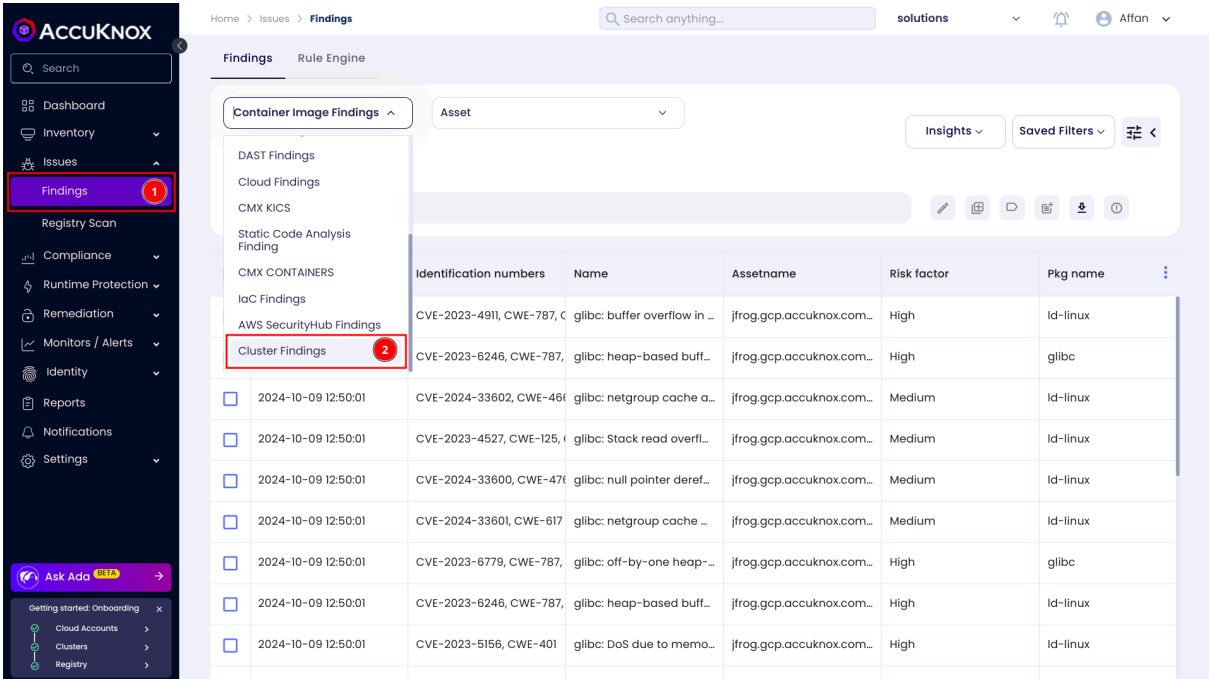
No data

This completes the onboarding process for CIS Benchmarking compliance scanning. You can review findings regularly to maintain and improve your cluster's CIS compliance.

## 12.2 Cluster Misconfiguration Scanning

Cyber attacks frequently occur due to security misconfigurations in applications and infrastructure. Preventing these vulnerabilities is crucial for maintaining a secure environment. AccuKnox empowers you to identify and remediate security misconfigurations within your Kubernetes clusters, ensuring that your applications and infrastructure are fully protected from potential threats.

In AccuKnox you can go to findings page and select the cluster findings to list all of the cluster misconfiguration findings.



The screenshot shows the AccuKnox interface with the 'Findings' page selected. The left sidebar contains a navigation menu with 'Findings' highlighted. The main content area shows a list of findings under the 'Container Image Findings' category. The 'Cluster Findings' option in the dropdown menu is highlighted.

Identification numbers	Name	Assetname	Risk factor	Pkg name
CVE-2023-4911, CWE-787, C	glibc: buffer overflow in ...	jfrog.gcp.accuknox.com...	High	ld-linux
CVE-2023-6246, CWE-787,	glibc: heap-based buff...	jfrog.gcp.accuknox.com...	High	glibc
2024-10-09 12:50:01 CVE-2024-33602, CWE-466	glibc: netgroup cache a...	jfrog.gcp.accuknox.com...	Medium	ld-linux
2024-10-09 12:50:01 CVE-2023-4527, CWE-125, C	glibc: Stack read overfl...	jfrog.gcp.accuknox.com...	Medium	ld-linux
2024-10-09 12:50:01 CVE-2024-33600, CWE-476	glibc: null pointer deref...	jfrog.gcp.accuknox.com...	Medium	ld-linux
2024-10-09 12:50:01 CVE-2024-33601, CWE-617	glibc: netgroup cache ...	jfrog.gcp.accuknox.com...	Medium	ld-linux
2024-10-09 12:50:01 CVE-2023-6779, CWE-787,	glibc: off-by-one heap-...	jfrog.gcp.accuknox.com...	High	glibc
2024-10-09 12:50:01 CVE-2023-6246, CWE-787,	glibc: heap-based buff...	jfrog.gcp.accuknox.com...	High	ld-linux
2024-10-09 12:50:01 CVE-2023-5156, CWE-401	glibc: DoS due to memo...	jfrog.gcp.accuknox.com...	High	ld-linux

You can click on a finding to see more details about it.

Home > Issues > Findings

Search anything...

Cluster Findings Asset

Group by

Search

<input type="checkbox"/>	Last seen	Name	Risk factor ↑	Assetname	Tool output
<input type="checkbox"/>	2024-10-09 10:55:45	Applications credentials in configuration files	High	mysql	FAILED
<input type="checkbox"/>	2024-10-09 08:40:16	Applications credentials in configuration files	High	cis-k8s-cronjob	FAILED
<input type="checkbox"/>	2024-09-30 17:06:29	Anonymous access enabled	High	kubeadm:bootstrap-sig...	FAILED
<input type="checkbox"/>	2024-10-02 15:52:54	Applications credentials in configuration files	High	reporter-config	FAILED
<input type="checkbox"/>	2024-10-02 15:15:22	Applications credentials in configuration files	High	mysql	FAILED
<input type="checkbox"/>	2024-09-30 17:06:29	Applications credentials in configuration files	High	k8s-risk-assessment-jo...	FAILED
<input type="checkbox"/>	2024-10-09 08:40:16	Anonymous access enabled	High	system:public-info-vie...	FAILED
<input type="checkbox"/>	2024-09-30 17:06:29	Anonymous access enabled	High	system:public-info-vie...	FAILED
<input type="checkbox"/>	2024-07-27 11:10:13	Anonymous access enabled	High	system:public-info-vie...	FAILED

1 - 20 of 11950 Rows per page: 20

Home > Issues > Findings

Cluster Findings

Group by

Search

1 row selected 1 - 20 of

### Applications credentials in configuration files High

Age: 7 days    Severity: ● High    SLA: 45 days    Tickets Created: 3

**Description**    Result    Solution    References    Source Code

Attackers who have access to configuration files can steal the stored secrets and use them. This control checks if ConfigMaps or pod specifications have sensitive information in their configuration.

Finding for in resource: [k8s\\_security\\_Deployment | mysql](#)

Failing since: about 6 day ago, on 03/10/2024

Last detected: on 09/10/2024

**Compliance Frameworks**  
No compliance found

**Asset Information**

```

{
  "id": "3bea491c-7049-4645-b292-2b6e8299f20c"
  "tickets_count": 0
  "data_type": "cluster-misconfiguration"
  "hash": "89d735e8eae43e55c68578df3edb0d57"
  "history": [
    {
      "changed": {}
      "scan_id": "0ef547b5-e688-40ec-bbb4-32460505d01f"
      "timestamp": "2024-10-03 03:53:49.195581+00:00"
    }
  ]
}

```

**Details** [+ Create Ticket](#)

Asset: mysql

Asset Type: k8s\_security\_Deployment

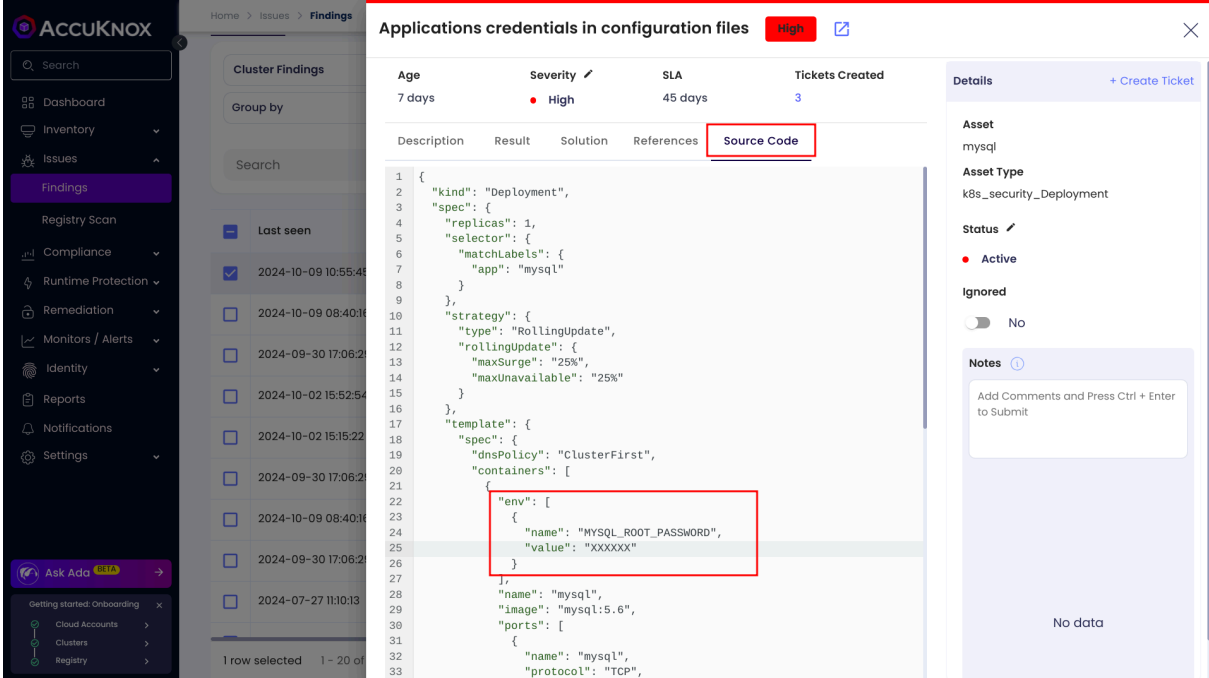
Status: ● Active

Ignored:  No

Notes:

No data

Here AccuKnox detected the application credentials leaked in the Kubernetes configuration. By clicking on the source code tab you can see that there is a hard coded password in a deployment manifest.

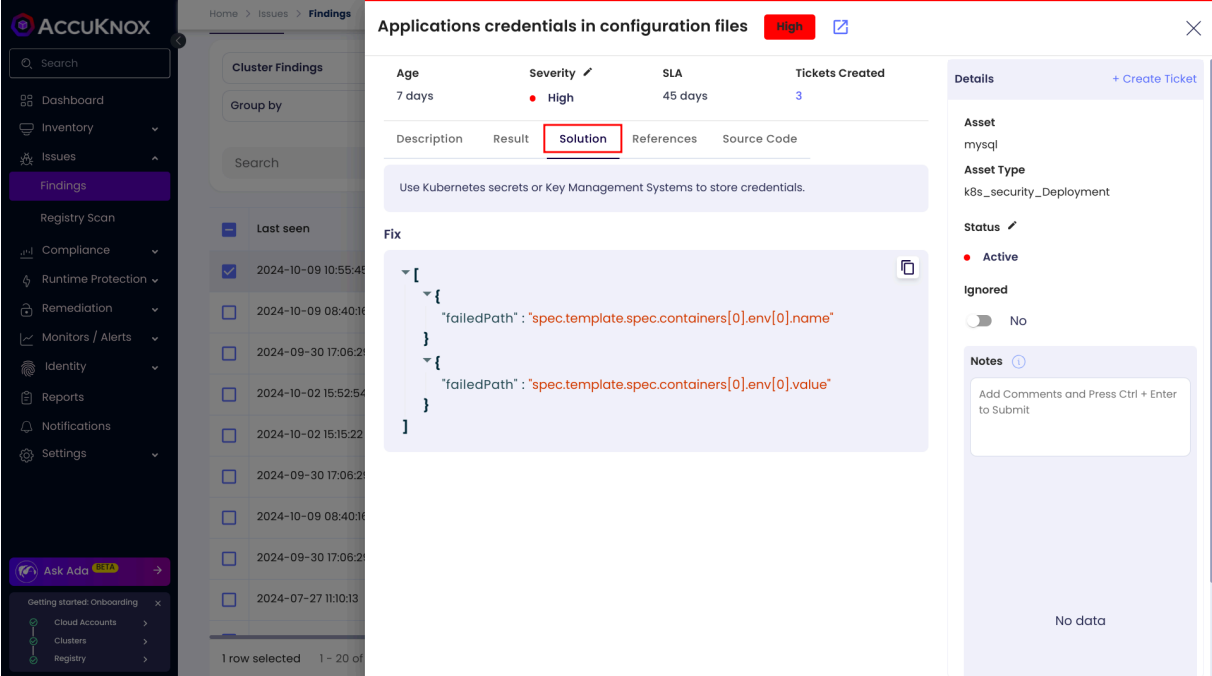


The screenshot displays the AccuKnox interface. On the left is a navigation sidebar with options like Dashboard, Inventory, Issues, Findings, Registry Scan, Compliance, Runtime Protection, Remediation, Monitors / Alerts, Identity, Reports, Notifications, and Settings. The main area shows a finding titled "Applications credentials in configuration files" with a severity of "High". Below the title, there are filters for Age (7 days), Severity (High), SLA (45 days), and Tickets Created (3). A table lists findings with columns for Description, Result, Solution, References, and Source Code. The Source Code tab is active, showing a Kubernetes deployment configuration snippet. A red box highlights the environment variable configuration: `"env": [{"name": "MYSQL_ROOT_PASSWORD", "value": "XXXXXX"}]`. On the right, a "Details" panel shows asset information (mysql), asset type (k8s\_security\_Deployment), status (Active), and a "Notes" section with a text input field.

An attacker can use these credentials and access your database. These sort of Kubernetes misconfigurations might get unnoticed by developers or DevOps engineers. By leveraging AccuKnox a user can detect vulnerabilities in time.

### 12.2.1 Remediation

AccuKnox provides you assistive remediation. Click on the solution tab and you will see what action can be preformed to remediate this issue.



The screenshot shows the AccuKnox interface with a finding titled "Applications credentials in configuration files" of High severity. The "Solution" tab is active, showing a JSON snippet for fixing the issue. The right sidebar shows details for the asset "mysql" and its type "k8s\_security\_Deployment".

Age	Severity	SLA	Tickets Created
7 days	High	45 days	3

```

[
  {
    "failedPath": "spec.template.spec.containers[0].env[0].name"
  },
  {
    "failedPath": "spec.template.spec.containers[0].env[0].value"
  }
]

```

Details:

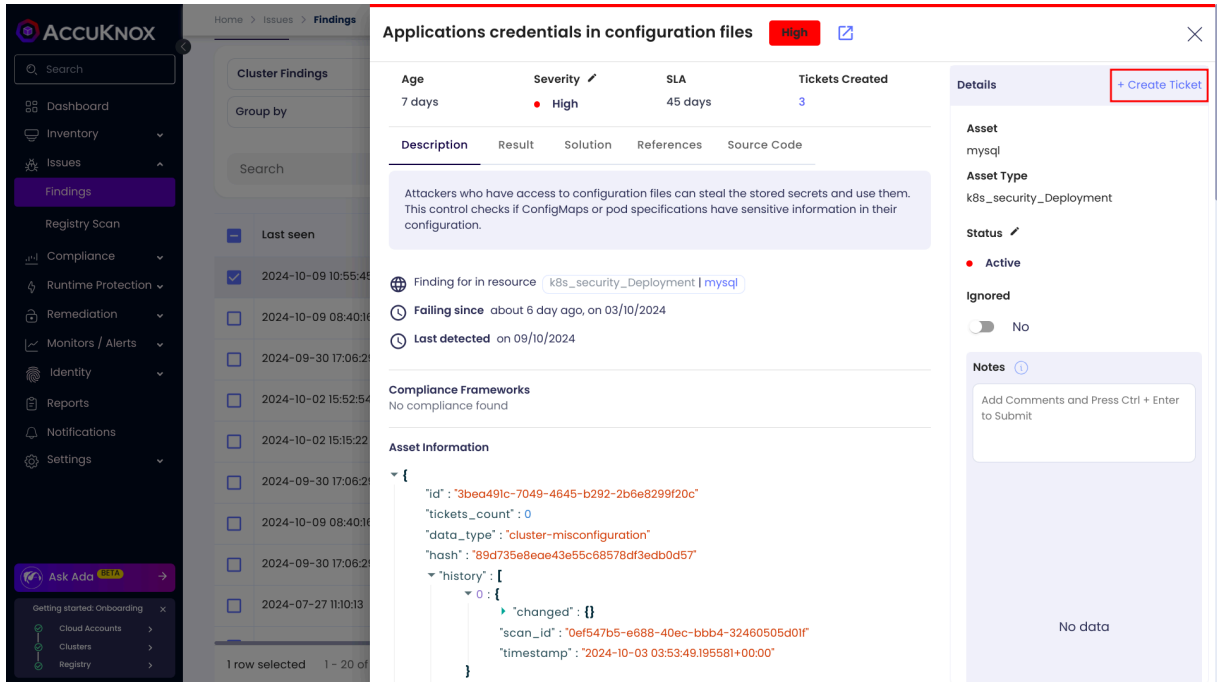
- Asset: mysql
- Asset Type: k8s\_security\_Deployment
- Status: Active
- Ignored: No

## 12.2.2 Vulnerability Management Lifecycle

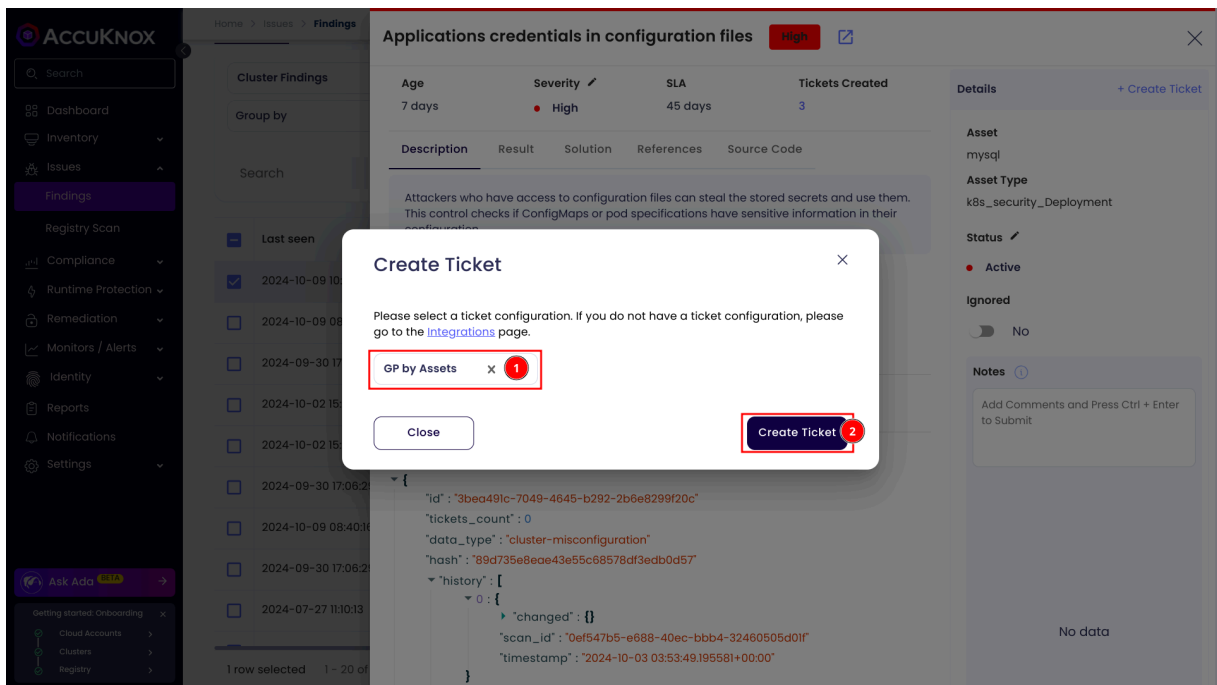
You can streamline vulnerability remediation and lifecycle management by creating Jira tickets directly from the AccuKnox UI.

Follow these steps for creating a ticket.

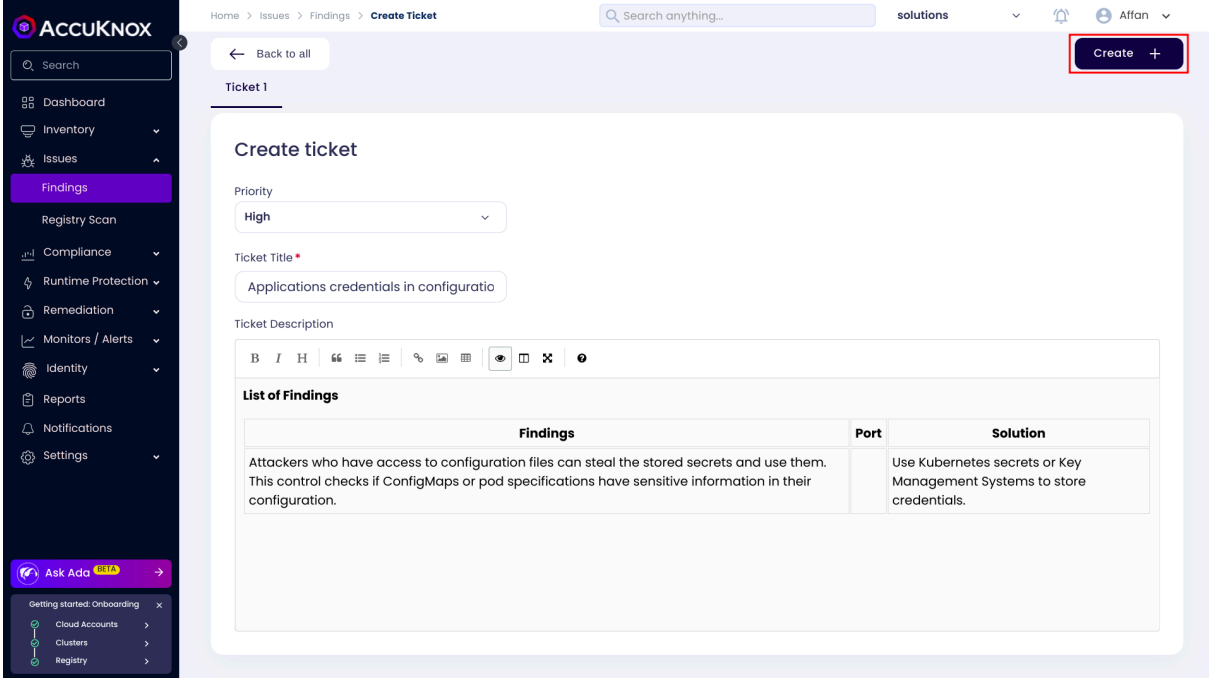
**Step 1.** Select a vulnerability and click on the create ticket button.



**Step 2.** Select your ticket configuration and click on the create ticket button.



**Step 3.** It will open up a new tab where you can review and modify the ticket details. Once you have reviewed the ticket click on the create button.



Home > Issues > Findings > Create Ticket

Search anything...

solutions

Affan

← Back to all

Create +

### Create ticket

Priority: High

Ticket Title: Applications credentials in configuratio

Ticket Description:

**List of Findings**

Findings	Port	Solution
Attackers who have access to configuration files can steal the stored secrets and use them. This control checks if ConfigMaps or pod specifications have sensitive information in their configuration.		Use Kubernetes secrets or Key Management Systems to store credentials.

In conclusion, AccuKnox helps you to detect, remediate and manage the lifecycle of Kubernetes security misconfiguration vulnerabilities.

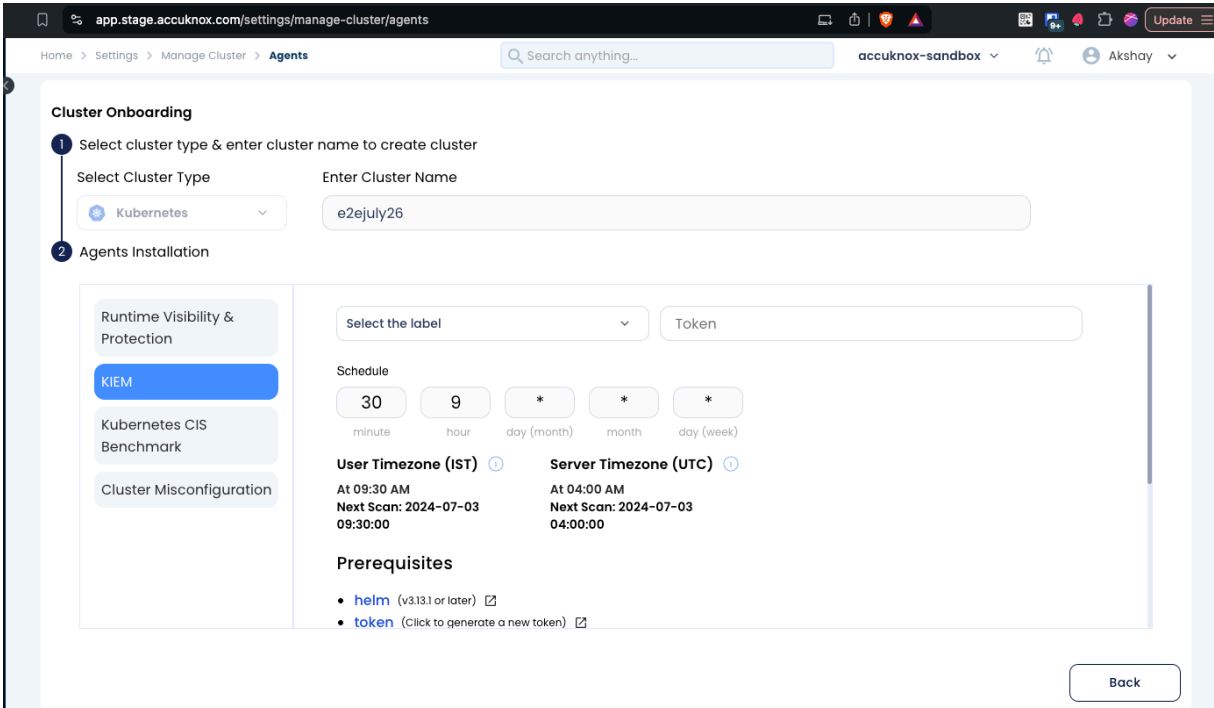
## 12.3 Kubernetes Identity and Entitlement Management (KIEM)

### 12.3.1. Onboarding Process

Follow these steps to set up and start using AccuKnox KIEM:

### 12.3.2 Install KIEM Agents

1. Navigate to the "Manage Cluster" section in your AccuKnox dashboard.
2. Select the target cluster for KIEM installation.
3. Install the KIEM job on the selected cluster.
4. Set up and schedule the cron job for regular scans.



The screenshot shows the "Cluster Onboarding" page in the AccuKnox dashboard. The page is divided into two main sections: "Cluster Onboarding" and "Agents Installation".

**Cluster Onboarding:**

- Step 1: Select cluster type & enter cluster name to create cluster.
  - Select Cluster Type: Kubernetes
  - Enter Cluster Name: e2ejuly26

**Agents Installation:**

- Runtime Visibility & Protection
- KIEM** (Selected)
- Kubernetes CIS Benchmark
- Cluster Misconfiguration

**Schedule Configuration:**

- Select the label: Token
- Schedule: 30 (minute) 9 (hour) \* (day/month) \* (month) \* (day/week)

**Timezone Information:**

- User Timezone (IST):** At 09:30 AM, Next Scan: 2024-07-03 09:30:00
- Server Timezone (UTC):** At 04:00 AM, Next Scan: 2024-07-03 04:00:00

**Prerequisites:**

- helm (v3.13.1 or later)
- token (click to generate a new token)

A "Back" button is located at the bottom right of the configuration area.

### 12.3.3 Post-Onboarding Steps

After completing the onboarding process:

1. Wait for the initial KIEM cron job to complete its first scan.
2. Once the scan is finished, navigate to the "Identity > KIEM" section in your dashboard.
3. Review the initial findings and adjust configurations as necessary.



## 12.3.4 Permissions Overview

- Summarizes all permissions in a unified view.
- Rolebinding and workloads are connected to permissions.
- Filter on constraints such as Role, Resource, ApiGroup, Verbs, Rolebinding, Service Accounts, Workload.
- View distilled permission summary for filtered entities.

Cluster: **kiem-test** | Key Query: **Select from Key Queries** | Entity Type: **Select Entity Type** | Search for any element:

**Overview** | Filter | List | Graph

Subject	RoleBinding	Role	Rule	
			Verb	Resource
<b>service-controller</b> Scope: kube-system	<b>system:controller:service-controller</b> Scope: cluster_wide	<b>system:controller:service-controller</b> Scope: cluster_wide	patch, update	services/status
			create, patch, update	events
			get, list, watch	services
			create, patch, update	events
			list, watch	nodes
<b>ephemeral-volume-controller</b> Scope: kube-system	<b>system:controller:ephemeral-volume-controller</b> Scope: cluster_wide	<b>system:controller:ephemeral-volume-controller</b> Scope: cluster_wide	create, patch, update	events
			get, list, watch	pods
			create, patch, update	events
			create, get, list, watch	persistentvolumes
			update	pods/finalizers
<b>legacy-service-account-token-cleaner</b>	<b>system:controller:legacy-service-account-token-cleaner</b>	<b>system:controller:legacy-service-account-token-cleaner</b>	delete, patch	secrets

Current Page: 1 | [Prev](#) | [Next](#)

Cluster: kiem-test | Overview | Key Query: Select from Key Queries | Entity Type: Select Entity Type | Search for any element: Search...

Filter

Verb: delete | Resource: secrets | Search Namespace: service-

+ Add Query | Apply

Rule	Verb	Resource
system:kube-controller-manager	delete, patch	secrets
system:kube-controller-manager	delete	secrets
token-cleaner	delete, get, list, watch	secrets

Current Page: 1 | Prev | Next

### 12.3.5 Key Queries

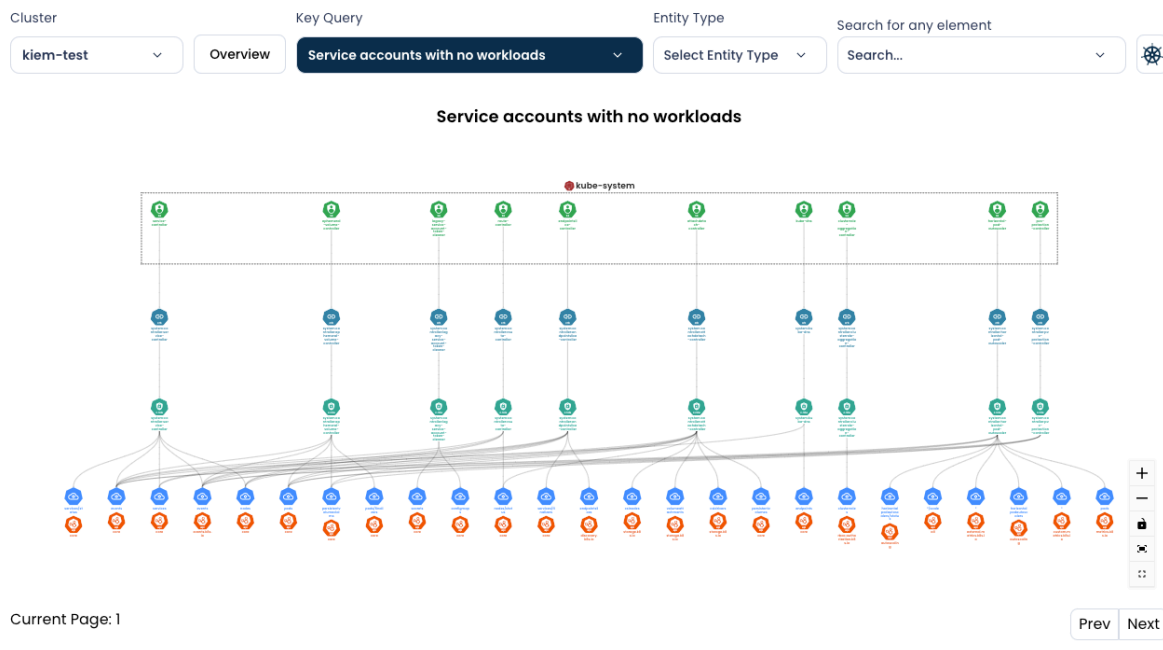
Our KIEM solution includes predefined queries that can detect security risks, misconfigurations, or compliance issues within Kubernetes RBAC configurations. These prebuilt queries aid in maintaining Kubernetes RBAC configurations with security as a primary factor.

#### Examples:

- **Identify Service Accounts not connected to any workloads** (indicator of dormant excessive permissions).
- **Identify principals with excessive privileges.** Excessive privileges in Kubernetes can increase the risk of security breaches, as overprivileged users or processes can misuse their access, leading to data breaches, service disruptions, or unauthorized changes in the cluster.
- **Find roles that have permissions to modify workload resources.** Excessive access rights to Kubernetes workload resources can lead to security vulnerabilities, allowing unauthorized access or modifications to critical applications and data, undermining the cluster's security posture.
- **List roles that have read access to Kubernetes secrets.** Kubernetes secrets, often containing sensitive information like passwords, tokens, or encryption keys, can pose a significant security risk if read access to these

roles is compromised, potentially leading to data leakage or unauthorized system access.

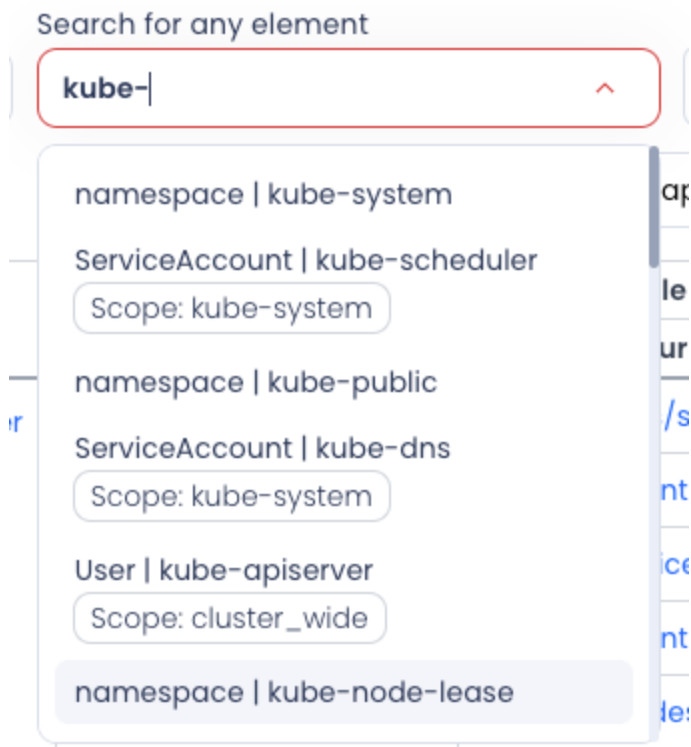
- **Identify roles that are not in use.** Unused roles can pose security risks if not regularly audited and cleaned up, potentially accumulating unnecessary permissions or becoming a target for exploitation by attackers.



## 12.3.6 Full-text Search

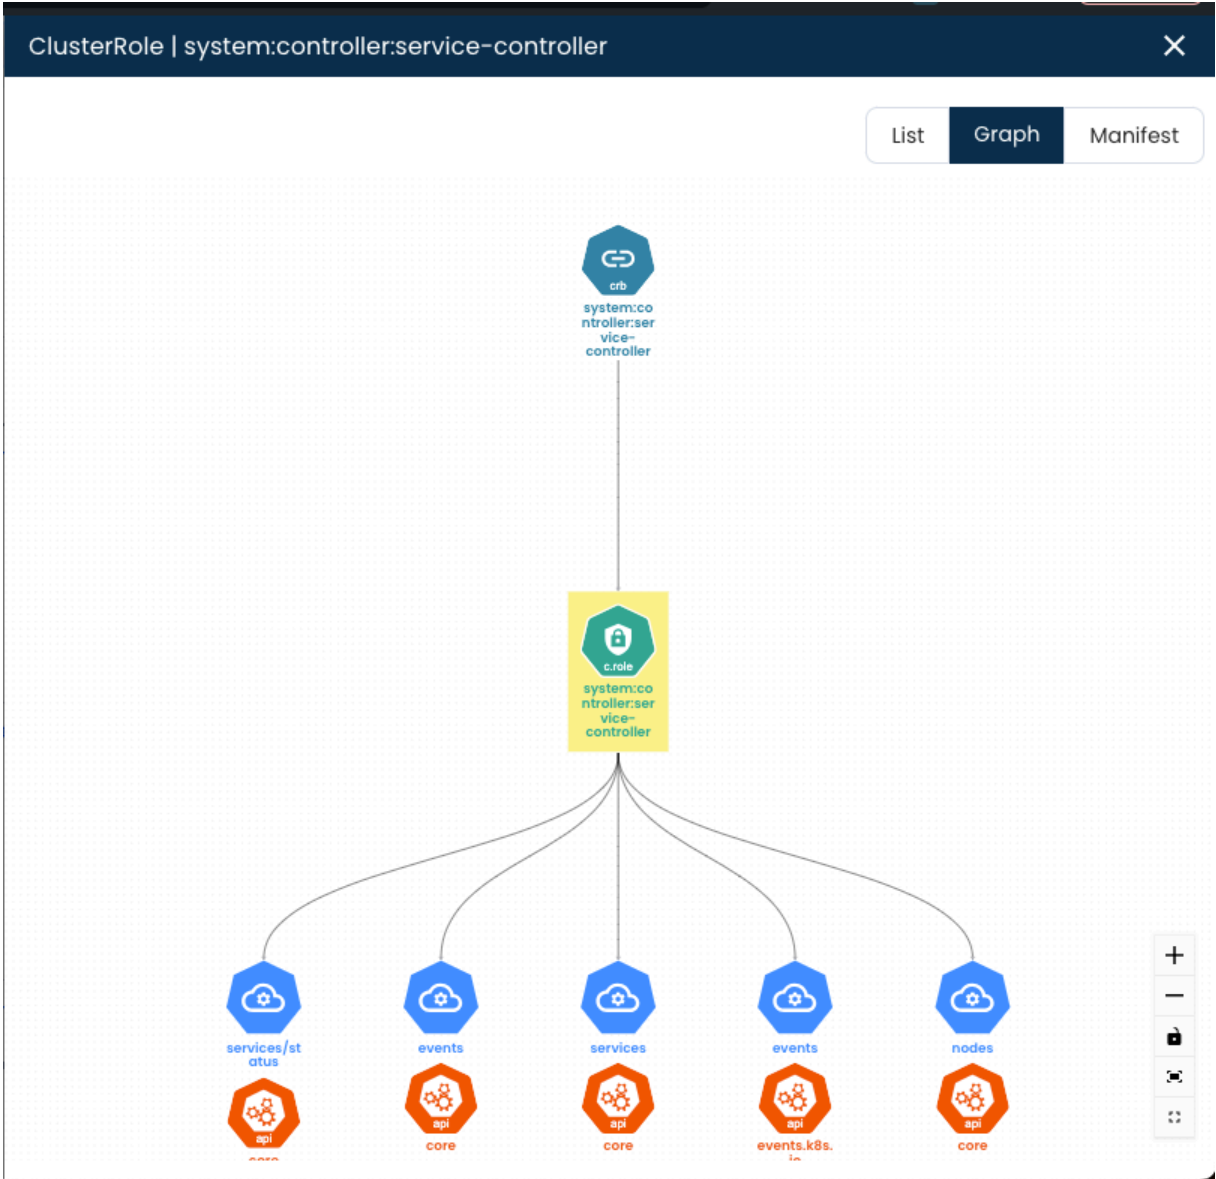
Search across all RBAC entities:

- ServiceAccounts
- RoleBindings
- Roles And more



### 12.3.7 Entity Exploration

- View connections and manifest for select entities.
- Discover excessive permissions.



ClusterRole | system:controller:service-controller ✕

List Graph Manifest

```

1  apiVersion: rbac.authorization.k8s.io/v1
2  kind: ClusterRole
3  metadata:
4    name: system:controller:service-controller
5    labels:
6      kubernetes.io/bootstrapping: rbac-defaults
7  rules:
8    - verbs:
9      - get
10     - list
11     - watch
12     apiGroups:
13       - ''
14     resources:
15       - services
16     - verbs:
17       - patch
18       - update
19     apiGroups:
20       - ''
21     resources:
22       - services/status
23     - verbs:
24       - list
25       - watch
26     apiGroups:
27       - ''
28     resources:
29       - nodes
30     - verbs:
31       - create
32       - patch
33       - update

```

- Explore all RBAC entities:
- Service Accounts
- Users
- Groups
- Roles
- RoleBindings

Cluster: kiem-test | Key Query: Select from Key Queries | Entity Type: ServiceAccount | Search for any element: Search...

Buttons: Overview, List, Graph

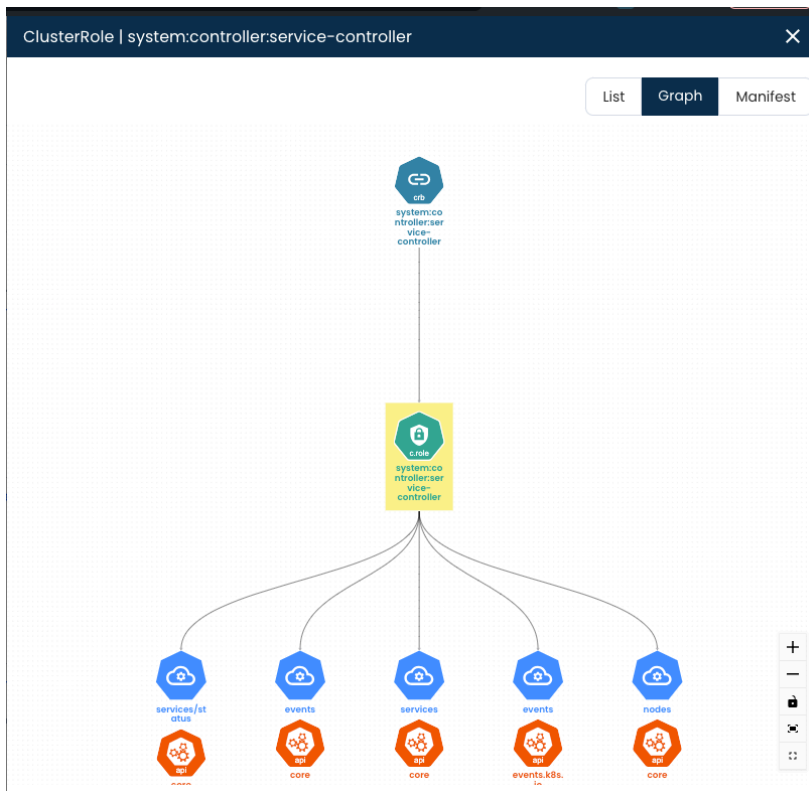
### ServiceAccount

Name	Has Role Binding	Mounted By
metrics-server Scope: kube-system	system:metrics-server Scope: cluster_wide +1	metrics-server-54fd9b65b-prr6n Scope: kube-system
local-path-provisioner-service-account Scope: kube-system	local-path-provisioner-bind Scope: cluster_wide	local-path-provisioner-6c86858495-4cr4k Scope: kube-system
horizontal-pod-autoscaler Scope: kube-system	system:controller:horizontal-pod-autoscaler Scope: cluster_wide	-
pvc-protection-controller Scope: kube-system	system:controller:pvc-protection-controller Scope: cluster_wide	-
replication-controller Scope: kube-system	system:controller:replication-controller Scope: cluster_wide	-
bootstrap-signer Scope: kube-system	system:controller:bootstrap-signer Scope: cluster_wide +1	-

Current Page: 1 | Prev | Next

## 12.3.8 Interactive Visualization

Open any entity and view all its connections by clicking on the link.

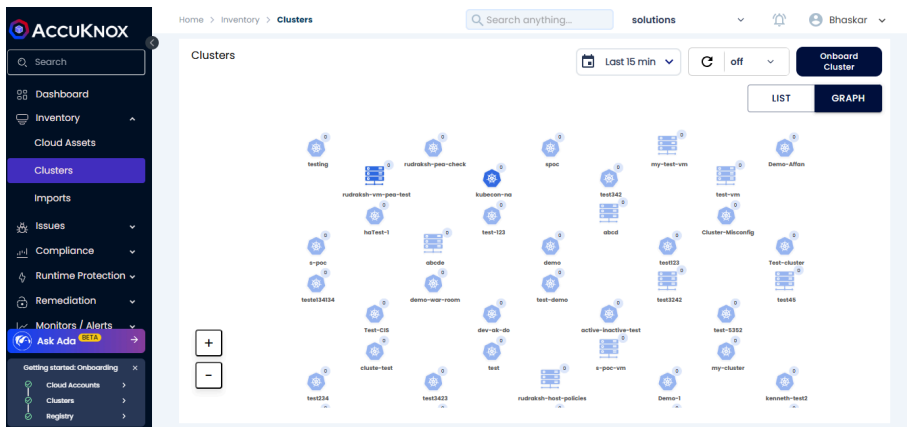


# 13. CWPP (Cloud Workload Protection Platform)

## 13.1 Cloud Workloads

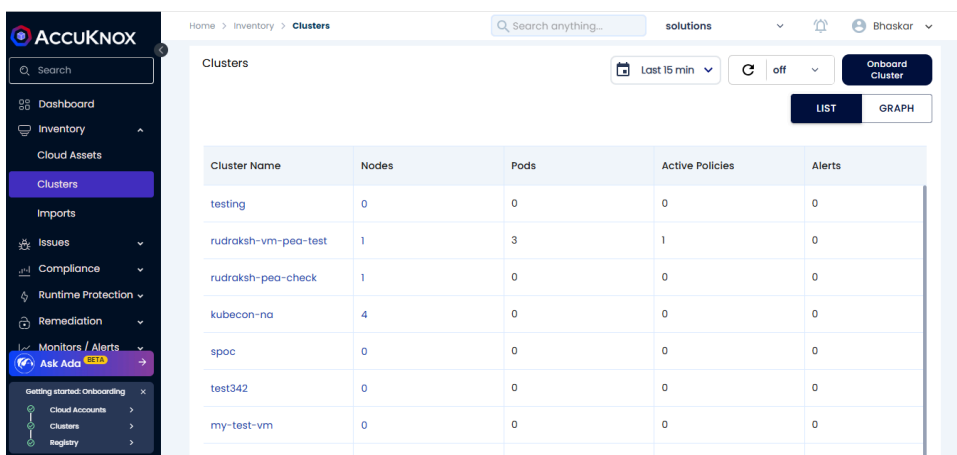
### 13.1.1 How to find graph view of clusters

Navigate to Clusters screen under Inventory to view the clusters that have been onboarded:



### 13.1.2 How to find list view of clusters

Click on the LIST option in the top right of the Cloud Workloads screen to get a list view of all the clusters

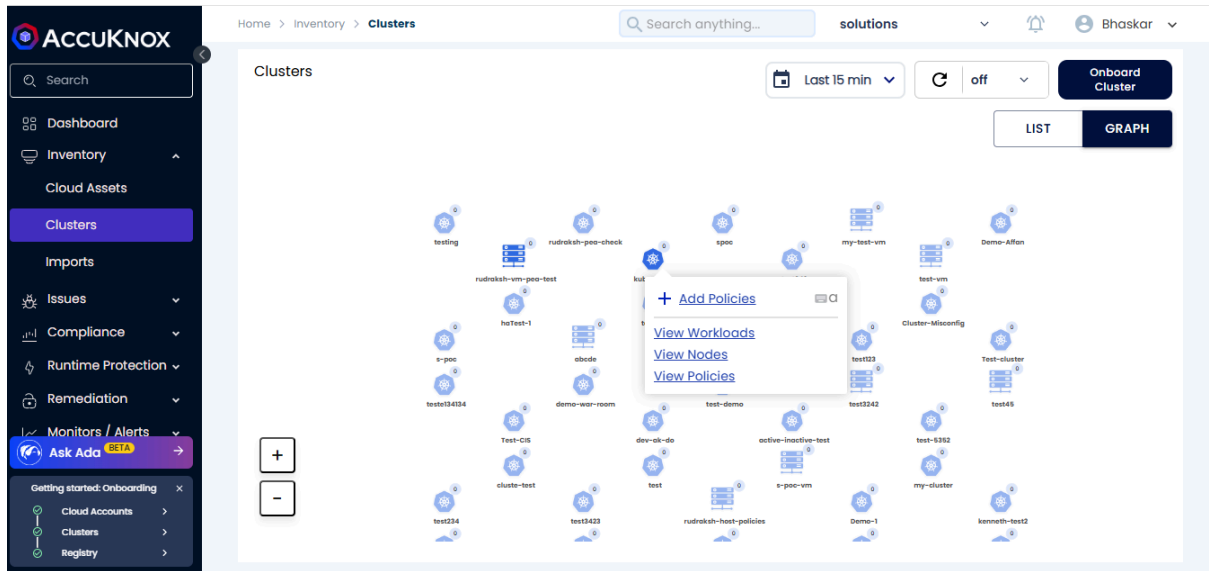


- The view can be freely switched between LIST and GRAPH as required.

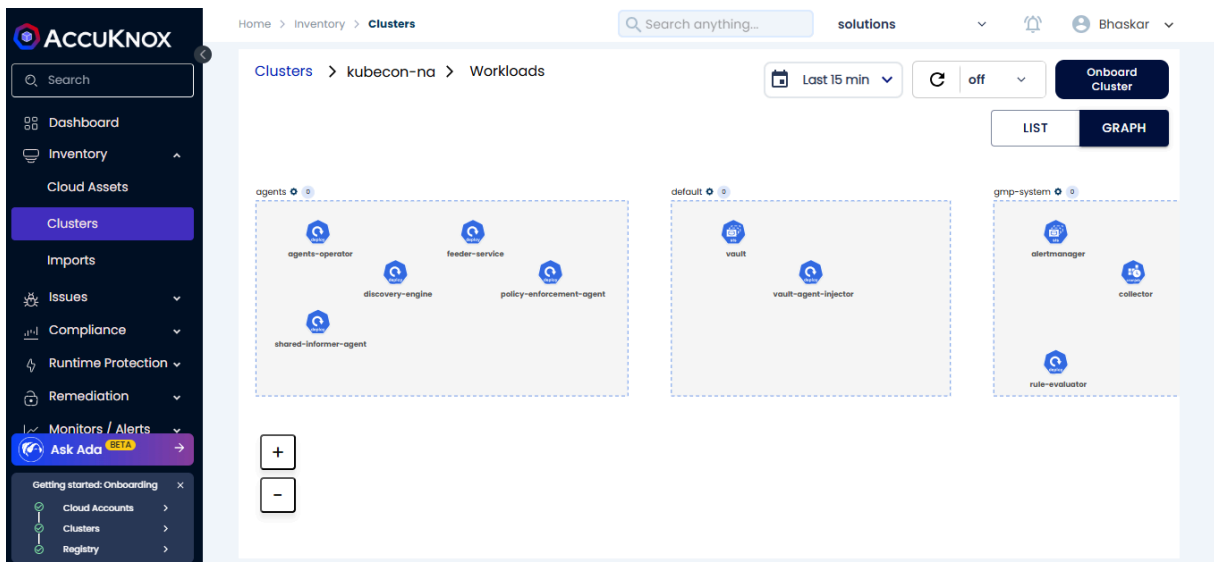


### 13.1.3 How to find details on cluster

- Clicking on any of the clusters gives more information about the cluster:



- Click on View Workloads to view the Pods present in the cluster classified according to the namespaces they are present in:

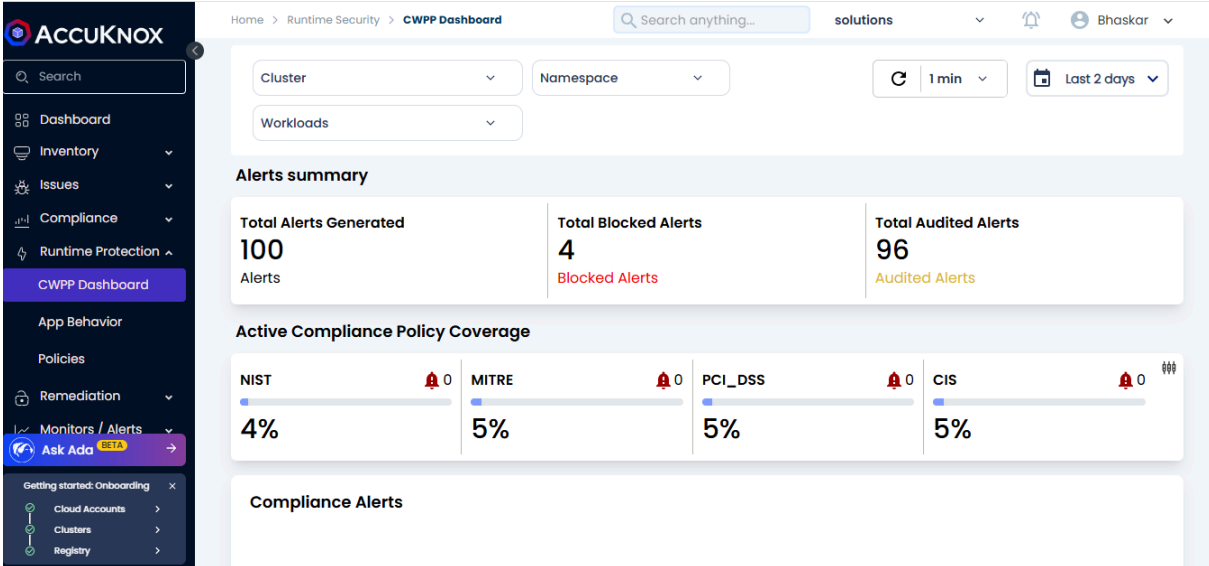


- View Policies can be clicked to jump to the Policies screen to show the policies for the selected cluster or pod. Click on the cluster name and then click on View Policies.

Policy Name	Category	Status	Clusters
<a href="#">autopol-egress-408221164</a> Kubernetes Network	Discovered	Inactive	kubecon-na
<a href="#">autopol-egress-639261075</a> Kubernetes Network	Discovered	Inactive	kubecon-na
<a href="#">autopol-egress-4171005049</a> Kubernetes Network	Discovered	Inactive	kubecon-na
<a href="#">autopol-egress-955124290</a> Kubernetes Network	Discovered	Inactive	kubecon-na
<a href="#">autopol-egress-3741501696</a> Kubernetes Network	Discovered	Inactive	kubecon-na

### 13.1.4 How to get Compliance for Cloud Workload

- AccuKnox leverage KubeArmor to harden your workload by enforcing hardening policies
- These hardening policies are based on different compliance frameworks like NIST, CIS, MITRE etc.



The screenshot displays the AccuKnox CWPP Dashboard interface. The left sidebar contains navigation options such as Dashboard, Inventory, Issues, Compliance, Runtime Protection, and CWPP Dashboard. The main content area shows the following data:

Alerts summary		
<b>Total Alerts Generated</b> 100 Alerts	<b>Total Blocked Alerts</b> 4 Blocked Alerts	<b>Total Audited Alerts</b> 96 Audited Alerts

Active Compliance Policy Coverage			
<b>NIST</b> 4% 0	<b>MITRE</b> 5% 0	<b>PCI_DSS</b> 5% 0	<b>CIS</b> 5% 0

Below these sections is a 'Compliance Alerts' section, which is currently empty.

## 13.2 App Behavior

Application Behavior of the cluster workloads that are onboarded to the AccuKnox SaaS are collected with help of KubeArmor and the AccuKnox Agents that are installed as Daemon sets in the cluster. The informations are collected at the pod level granularity. So that the users can get the information about each pods that are running in each namespace. Application behavior of the cluster workloads are given in two ways, one is the list view and other is the Graphical view.

### 13.2.1 How to interpret network graph

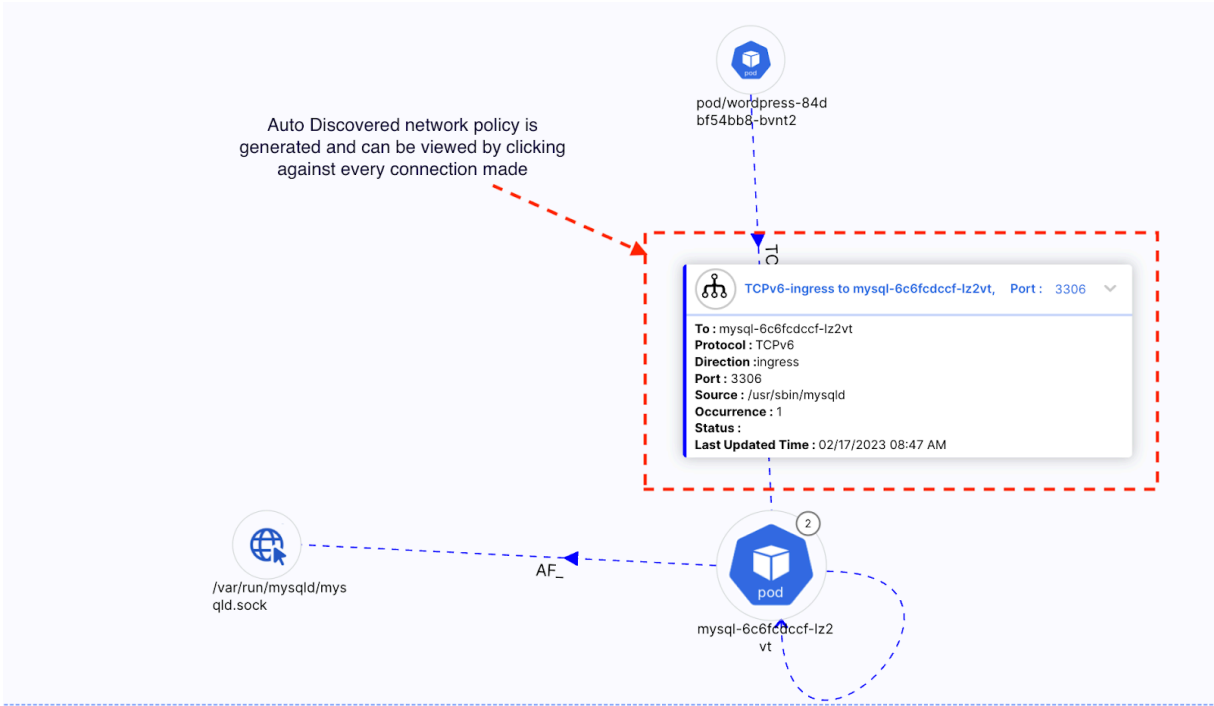
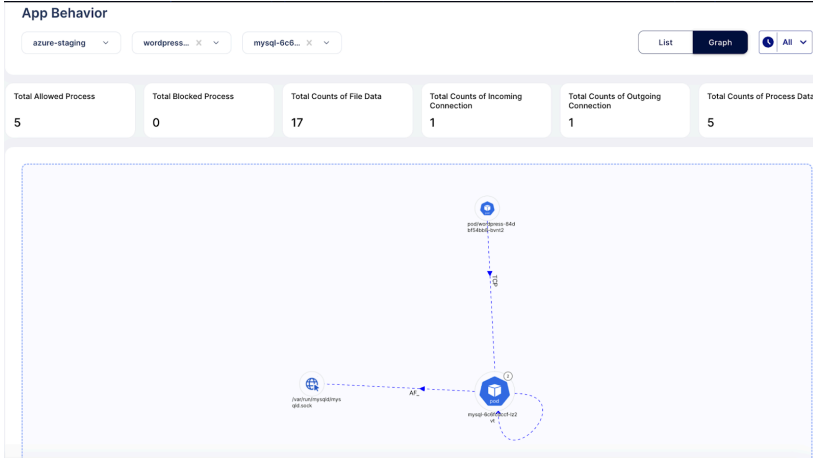
Lets understand this by following use-case example - **Auditing Application Behavior of MySQL application**

1.Install workload:

```
sh kubectl apply -f
https://raw.githubusercontent.com/kubearmor/KubeArmor/main/examples/wordpress-mysql/wordpress-mysql-deployment.yaml
```

2.Showing App behavior screen in the context of the wordpress-mysql application. To see the Application Behavior user must Navigate to the *Runtime Protection->App Behavior* section. Then click on the Cluster and Namespace and pod from the filters to see the Application Behavior.

- Network Graph: This view gives the graphical representation of Ingress and Egress traffic that are occurring in the Pod. When we click on the connections we can get a clear view of the traffic type and port details.



- File Observability: This view gives details about the files that are getting accessed in the pod.

**App Behavior**

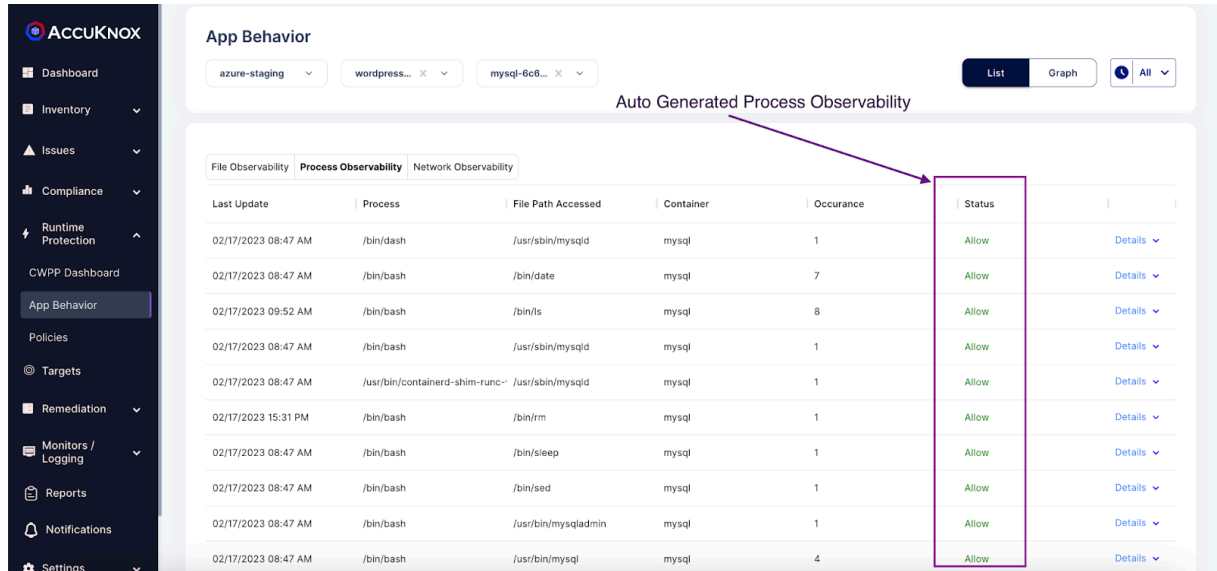
azure-staging | wordpress... | mysql-6c6... | List | Graph | All

Auto Generated Whitelisted Application Behaviour

File Observability | Process Observability | Network Observability | Show Aggregated View

Last Update	Process	File Path Accessed	Container	Occurance	Status	
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/posix/Amer	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/right/Ameri	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/posix/Amer	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/posix/Amer	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/Canada/Ea	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/sbin/mysqld	/var/lib/mysql/mysql/proc.MYI	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/America/Re	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/right/GMT0	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/posix/Amer	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql_tzinfo_to_sq	/usr/share/zoneinfo/right/Europ	mysql	1	Allow	Details

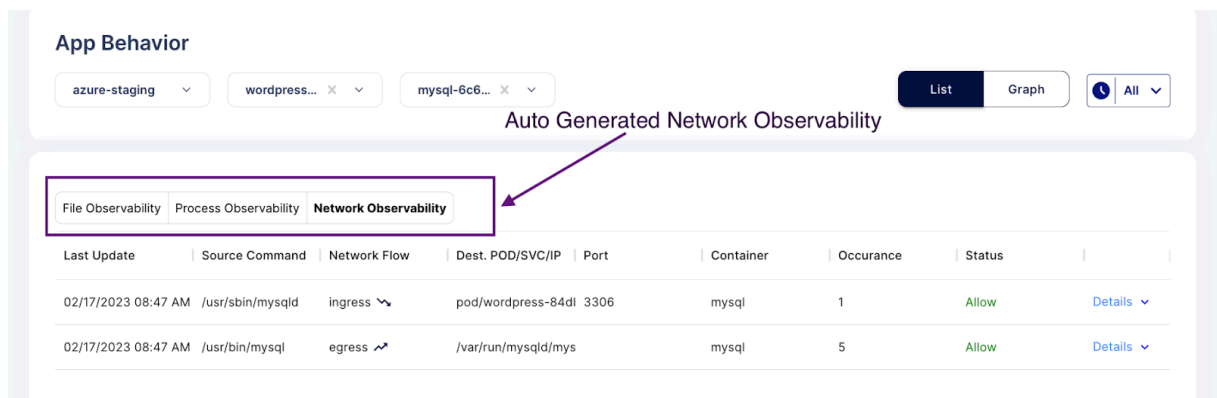
- **Process Observability:** This view gives the details of Processes that are currently running in the Pod.



The screenshot shows the 'App Behavior' section with filters for 'azure-staging', 'wordpress...', and 'mysql-6c6...'. The 'Auto Generated Process Observability' view is active, showing a table with columns: Last Update, Process, File Path Accessed, Container, Occurance, Status, and Details. A purple box highlights the 'Status' column, and a purple arrow points to it from the text 'Auto Generated Process Observability' above the table.

Last Update	Process	File Path Accessed	Container	Occurance	Status	Details
02/17/2023 08:47 AM	/bin/dash	/usr/sbin/mysqld	mysql	1	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/bin/date	mysql	7	Allow	Details
02/17/2023 09:52 AM	/bin/bash	/bin/ls	mysql	8	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/usr/sbin/mysqld	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/containerd-shim-runc-	/usr/sbin/mysqld	mysql	1	Allow	Details
02/17/2023 15:31 PM	/bin/bash	/bin/rm	mysql	1	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/bin/sleep	mysql	1	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/bin/sed	mysql	1	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/usr/bin/mysqldadmin	mysql	1	Allow	Details
02/17/2023 08:47 AM	/bin/bash	/usr/bin/mysql	mysql	4	Allow	Details

- **Network Observability:** The network observability can also be seen in the list here you can see the details of ingress and egress traffic in the list view.

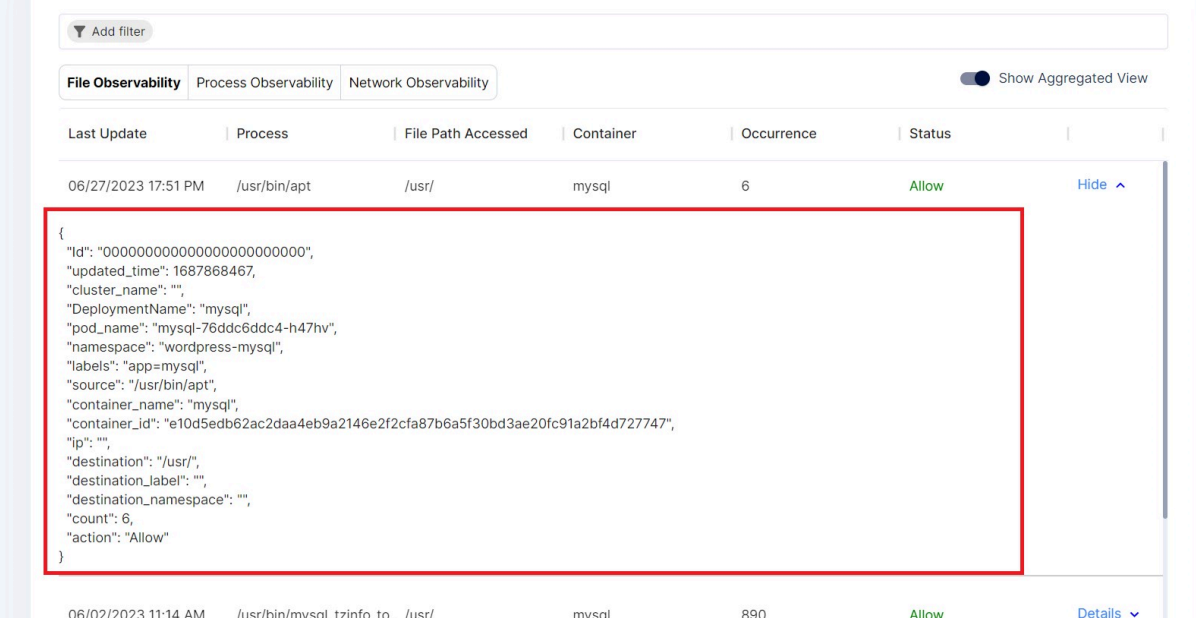


The screenshot shows the 'App Behavior' section with filters for 'azure-staging', 'wordpress...', and 'mysql-6c6...'. The 'Auto Generated Network Observability' view is active, showing a table with columns: Last Update, Source Command, Network Flow, Dest. POD/SVC/IP, Port, Container, Occurance, Status, and Details. A purple box highlights the 'Network Observability' tab, and a purple arrow points to it from the text 'Auto Generated Network Observability' above the table.

Last Update	Source Command	Network Flow	Dest. POD/SVC/IP	Port	Container	Occurance	Status	Details
02/17/2023 08:47 AM	/usr/sbin/mysqld	ingress ↘	pod/wordpress-84dl	3306	mysql	1	Allow	Details
02/17/2023 08:47 AM	/usr/bin/mysql	egress ↗	/var/run/mysqld/mys		mysql	5	Allow	Details

## 13.2.2 How to see App Behavior Telemetry

- To see the contextual information about the File and Network and Process observability user needs to navigate to the *Runtime Protection->App Behavior* Section.
- **File Observability Telemetry:** To see the file observability related telemetry user needs to click the list view and select file observability part and click on any of the file events to see the Telemetry



The screenshot displays the 'File Observability' section of the AccuKnox interface. It features a table with the following columns: Last Update, Process, File Path Accessed, Container, Occurrence, and Status. The first row shows an event from 06/27/2023 at 17:51 PM, with process '/usr/bin/apt', file path '/usr/', container 'mysql', and 6 occurrences, with a status of 'Allow'. A red box highlights the JSON object in the details view for this event:

```
{
  "id": "00000000000000000000000000000000",
  "updated_time": 1687868467,
  "cluster_name": "",
  "DeploymentName": "mysql",
  "pod_name": "mysql-76ddc6ddc4-h47hv",
  "namespace": "wordpress-mysql",
  "labels": "app=mysql",
  "source": "/usr/bin/apt",
  "container_name": "mysql",
  "container_id": "e10d5edb62ac2daa4eb9a2146e2f2cfa87b6a5f30bd3ae20fc91a2bf4d727747",
  "ip": "",
  "destination": "/usr/",
  "destination_label": "",
  "destination_namespace": "",
  "count": 6,
  "action": "Allow"
}
```

Below the highlighted event, another row is partially visible, showing an event from 06/02/2023 at 11:14 AM, with process '/usr/bin/mysql\_tzinfo\_to\_', file path '/usr/', container 'mysql', and 890 occurrences, with a status of 'Allow'.

- **Process Observability Telemetry:** To see the process observability related telemetry user needs to click the list view and select process observability part and click on any of the process events to see the Telemetry



▼ Add filter

File Observability **Process Observability** Network Observability

Last Update	Process	File Path Accessed	Container	Occurrence	Status	
06/02/2023 11:14 AM	/bin/bash	/bin/sed	mysql	1	Allow	Hide ^

```

{
  "id": "6479820c3d63aeb2d1f0a112",
  "updated_time": 1685684643,
  "cluster_name": "aks-demo-prod",
  "DeploymentName": "mysql",
  "pod_name": "mysql-76ddc6ddc4-h47hv",
  "namespace": "wordpress-mysql",
  "labels": "app=mysql",
  "source": "/bin/bash",
  "container_name": "mysql",
  "container_id": "e10d5edb62ac2daa4eb9a2146e2f2cfa87b6a5f30bd3ae20fc91a2bf4d727747",
  "ip": "",
  "destination": "/bin/sed",
  "destination_label": "",
  "destination_namespace": "",
  "count": 1,
  "action": "Allow"
}

```

- **Network observability:** To see the Network observability related telemetry user needs to click the list view and select Network observability part and click on any of the Network events to see the Telemetry

▼ Add filter

File Observability Process Observability **Network Observability**

Last Update	Source Command	Network Flow	Dest. POD/SVC/IP	Port	Container	Occurrence	Status	
06/02/2023 11:14 AM	/usr/bin/mysql	egress ↗	/var/run/mysqlq/mysqlc		mysql	2	Allow	Details v
07/06/2023 15:47 PM	/usr/bin/mysql	egress ↗	svc/wordpress	3306	mysql	2	Allow	Hide ^

```

{
  "id": "64a69475f181075cb4a563a3",
  "updated_time": 1688638674,
  "cluster_name": "",
  "DeploymentName": "mysql",
  "pod_name": "mysql-76ddc6ddc4-h47hv",
  "namespace": "wordpress-mysql",
  "labels": "app=mysql",
  "source": "/usr/bin/mysql",
  "nw_type": "egress",
  "container_name": "mysql",
  "container_id": "e10d5edb62ac2daa4eb9a2146e2f2cfa87b6a5f30bd3ae20fc91a2bf4d727747",
  "ip": "svc/wordpress",
  "port": 3306,
  "protocol": "TCP",
  "destination": "",
  "destination_label": "",
  "destination_namespace": "wordpress-mysql",
  "count": 2,
  "action": "Allow"
}

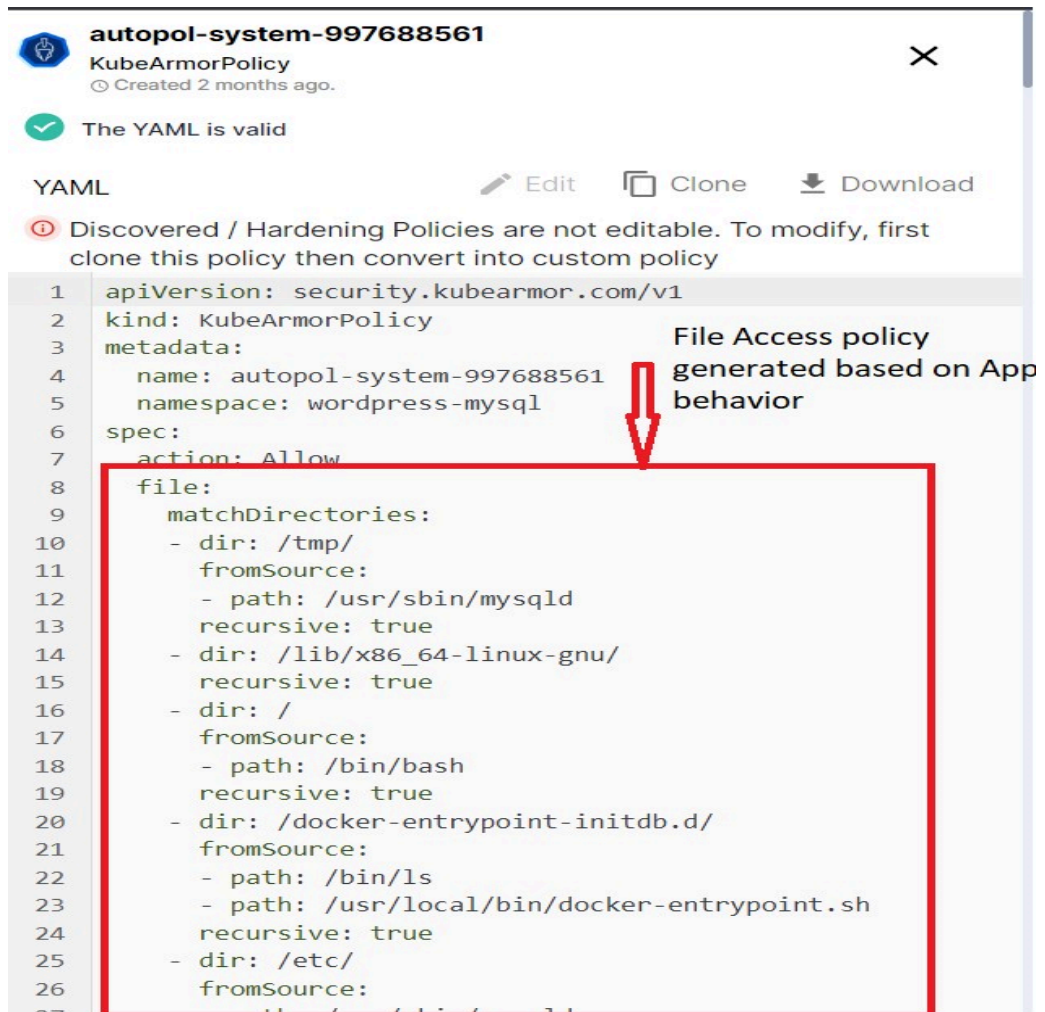
```

## 13.3 Runtime Protection w/ Policy Management

### 13.3.1 How to understand discover policies

Auto Discovered Policies are generated based on the Application Behavior. AccuKnox Runtime Security Engine KubeArmor when deployed as agent will model the default application behavior of the workload and comes up with the Auto discovered policies.

- **File access behavior based policies:** Based on the files that are accessed in pod, the Auto discovered system policies are generated. To view that policy user must navigate to *Runtime Protection->policies* section. Then click on the cluster and pod for which we want to see the auto-discovered policies.



The screenshot shows a KubeArmorPolicy interface for a policy named 'autopol-system-997688561'. It includes a status bar indicating 'The YAML is valid' and a warning that discovered policies are not editable. The main content is a YAML configuration for a KubeArmorPolicy. A red box highlights the 'file' section, which lists matchDirectories with paths like /tmp/, /usr/sbin/mysqld, /lib/x86\_64-linux-gnu/, /bin/bash, /docker-entrypoint-initdb.d/, /bin/ls, /usr/local/bin/docker-entrypoint.sh, and /etc/. A red arrow points to this section with the text 'File Access policy generated based on App behavior'.

```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: autopol-system-997688561
5    namespace: wordpress-mysql
6  spec:
7    action: Allow
8    file:
9      matchDirectories:
10     - dir: /tmp/
11       fromSource:
12         - path: /usr/sbin/mysqld
13           recursive: true
14     - dir: /lib/x86_64-linux-gnu/
15       recursive: true
16     - dir: /
17       fromSource:
18         - path: /bin/bash
19           recursive: true
20     - dir: /docker-entrypoint-initdb.d/
21       fromSource:
22         - path: /bin/ls
23           - path: /usr/local/bin/docker-entrypoint.sh
24             recursive: true
25     - dir: /etc/
26       fromSource:
27

```

- **Process access behavior based policies:** Based on the process that are running in pod, the Auto discovered system policies are generated.

To view that policy user must navigate to *Runtime Protection->policies* section. Then click on the cluster and pod for which we want to see the auto-discovered policies.

```

process:
  matchDirectories:
    - dir: /bin/
      fromSource:
        - path: /bin/bash
          recursive: true
    - dir: /usr/bin/
      fromSource:
        - path: /bin/bash
          recursive: true
  matchPaths:
    - fromSource:
        - path: /usr/bin/mysql_install_db
      path: /bin/sh
    - fromSource:
        - path: /bin/sh
      path: /usr/bin/my_print_defaults
    - path: /usr/local/bin/docker-entrypoint.sh
    - path: /usr/local/bin/gosu
    - fromSource:
        - path: /bin/bash
        - path: /bin/dash
      path: /usr/sbin/mysqld
    - path: /usr/bin/mysql
    - path: /usr/bin/mysqladmin
    - path: /bin/mktemp
    - path: /bin/cat
    - path: /bin/date
  
```

Process access policy generated based on App Behavior

- Network access behavior based Policies:** Based on the Network connections that are Ingress and egress connections that are present in pod, the auto discovered system policies are generated. To view that policy user must navigate to the Runtime *Protection->policies* section. Then click on the cluster and pod for which we want to see the auto-discovered policies.

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: autopol-egress-3275896150
  namespace: wordpress-mysql
spec:
  egress:
  - ports:
    - protocol: UDP
    - ports:
      - port: 443
        protocol: TCP
    - ports:
      - port: 3306
        protocol: TCP
    to:
      - podSelector:
          matchLabels:
            app: mysql
      - ports:
          - port: 8081
            protocol: TCP
      - ports:
          - port: 22
            protocol: TCP
  podSelector:
    matchLabels:
      app: wordpress
  policyTypes:
  - Egress
  
```

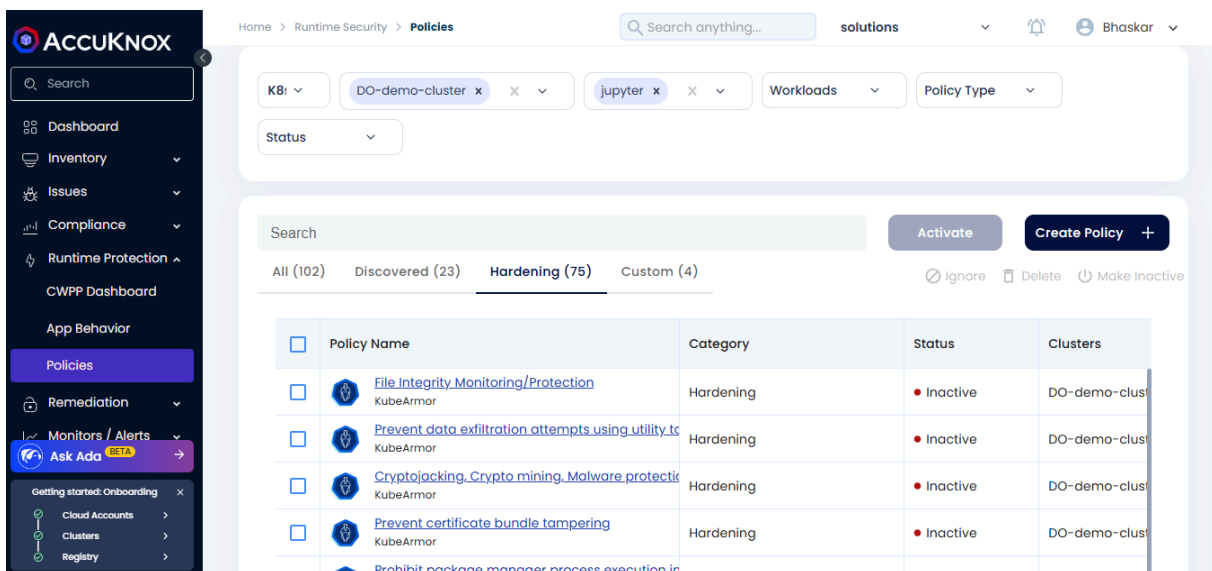
← Egress policy generated based on the application Behavior

### 13.3.2 How to understand Hardening policies

One of the methods to achieve a zero-trust environment is Application Hardening. KubeArmor is a security solution for the Kubernetes and cloud native platforms that helps protect your workloads from attacks and threats. It does this by providing a set of hardening policies which is a block based policies. It is based on industry-leading technical conformance to standard compliance and attack frameworks such as CIS, MITRE, NIST-800-53, and STIGs. These policies are designed to help you secure your workloads in a way that is compliant with these frameworks and recommended best practices.

- Lets understand by taking an use-case example - **Disallowing any binaries execution to prevent from RCE Vulnerability**

1. Select your cluster and namespace from this Policies screen. We will be getting list of hardening policies for the selected Namespace.



The screenshot shows the AccuKnox interface for managing policies. The left sidebar contains navigation options like Dashboard, Inventory, Issues, Compliance, Runtime Protection, CWPP Dashboard, App Behavior, Policies (highlighted), Remediation, and Monitors / Alerts. The main content area shows the 'Policies' screen for a specific cluster and namespace. The breadcrumb trail is 'Home > Runtime Security > Policies'. There are filters for 'K8s' (DO-demo-cluster), 'jupyter', 'Workloads', and 'Policy Type'. A search bar is present. Below the filters, there are buttons for 'Activate' and 'Create Policy +'. The table below shows a list of hardening policies:

Policy Name	Category	Status	Clusters
<a href="#">File Integrity Monitoring/Protection</a> KubeArmor	Hardening	Inactive	DO-demo-clus
<a href="#">Prevent data exfiltration attempts using utility te</a> KubeArmor	Hardening	Inactive	DO-demo-clus
<a href="#">Cryptojacking, Crypto mining, Malware protectio</a> KubeArmor	Hardening	Inactive	DO-demo-clus
<a href="#">Prevent certificate bundle tampering</a> KubeArmor	Hardening	Inactive	DO-demo-clus
<a href="#">Prohibit package manager process execution in</a>			

2. Selecting the below hardening policy to apply. This policy disallows execution of any of the Package management tools inside the pod. This policy is generated based on the Compliance Frameworks like NIST, NIST 800



**harden-wordpress-pkg-mngr-exec**  
KubeArmorPolicy Updated 17days ago

✕

YAML
✎ Edit
📄 Clone
⬇️ Download

ⓘ Discovered / Hardening Policies are not editable. To modify, first clone this policy then convert into custom policy

```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: harden-wordpress-pkg-mngr-exec
5    namespace: wordpress-mysql
6  spec:
7    action: Block
8    message: Alert! Execution of package management process inside
9    process:
10   matchPaths:
11     - path: /usr/bin/apt
12     - path: /usr/bin/apt-get
13     - path: /bin/apt-get
14     - path: /sbin/apk
15     - path: /bin/apt
16     - path: /usr/bin/dpkg
17     - path: /bin/dpkg
18     - path: /usr/bin/gdebi
19     - path: /bin/gdebi
20     - path: /usr/bin/make
21     - path: /bin/make
22     - path: /usr/bin/yum
23     - path: /bin/yum
24     - path: /usr/bin/rpm
25     - path: /bin/rpm
26     - path: /usr/bin/dnf
27     - path: /bin/dnf
28     - path: /usr/bin/pacman
29     - path: /usr/sbin/pacman
30     - path: /bin/pacman
31     - path: /sbin/pacman
32     - path: /usr/bin/makepkg
33     - path: /usr/sbin/makepkg
34     - path: /bin/makepkg
35     - path: /sbin/makepkg
36     - path: /usr/bin/yaourt
37     - path: /usr/sbin/yaourt
38     - path: /bin/yaourt
39     - path: /sbin/yaourt
40     - path: /usr/bin/zypper
41     - path: /bin/zypper
42   selector:
43     matchLabels:
44       app: wordpress
45   severity: 5
46   tags:
47     - NIST
48     - NIST_800-53_CM-7(4)
49     - SI-4
50     - process
51     - NIST_800-53_SI-4
52
```

3. Select this policy and click on the Activate option.

The screenshot shows the ACCUKNOX interface with the 'Policies' section active. The left sidebar contains navigation options like Dashboard, Inventory, Issues, Compliance, Runtime Protection, CWPP Dashboard, App Behavior, Policies (highlighted), Remediation, and Monitors / Alerts. The main content area shows a breadcrumb trail 'Home > Runtime Security > Policies' and a search bar. Below the search bar are filters for 'KB:', 'DO-demo-cluster', 'jupyter', 'Workloads', and 'Policy Type'. A 'Status' dropdown is also present. The main table lists policies with columns for 'Policy Name', 'Category', 'Status', and 'Clusters'. The first policy, 'File Integrity Monitoring/Protection' by KubeArmor, is selected with a checkmark. The 'Activate' button is highlighted in dark blue.

Policy Name	Category	Status	Clusters
<input checked="" type="checkbox"/> File Integrity Monitoring/Protection KubeArmor	Hardening	Inactive	DO-demo-clus
<input type="checkbox"/> Prevent data exfiltration attempts using utility tr KubeArmor	Hardening	Inactive	DO-demo-clus
<input type="checkbox"/> Cryptojacking, Crypto mining, Malware protectio KubeArmor	Hardening	Inactive	DO-demo-clus
<input type="checkbox"/> Prevent certificate bundle tampering KubeArmor	Hardening	Inactive	DO-demo-clus

4. After applying, the policy goes into Active state.

The screenshot shows the same ACCUKNOX interface after the policy activation. A green notification banner at the top reads 'Policies activated successfully'. The 'Activate' button is now greyed out, indicating it is disabled. The table below shows the same list of policies, but the first policy's status is now 'Active' (indicated by a green dot).

Policy Name	Category	Status	Clusters
<input type="checkbox"/> File Integrity Monitoring/Protection KubeArmor	Hardening	Active	DO-demo-clus
<input type="checkbox"/> Prevent data exfiltration attempts using utility tr KubeArmor	Hardening	Inactive	DO-demo-clus
<input type="checkbox"/> Cryptojacking, Crypto mining, Malware protectio KubeArmor	Hardening	Inactive	DO-demo-clus
<input type="checkbox"/> Prevent certificate bundle tampering KubeArmor	Hardening	Inactive	DO-demo-clus

### 13.3.3 How to Audit application and get alerts for that

- AccuKnox Runtime Security Engine kubeArmor can be used for auditing the application with help of audit based security policies. Let us consider the following policy



#### ksp-mysql-audit-dir (v3)




KubeArmorPolicy

🕒 Created a month ago.

✕

 The YAML is valid

YAML

 Edit  Clone  Download

```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: ksp-mysql-audit-dir
5    namespace: wordpress-mysql
6  spec:
7    severity: 5
8    selector:
9      matchLabels:
10     app: mysql
11   file:
12     matchDirectories:
13     - dir: /var/lib/mysql/
14       recursive: true
15   action: Audit
16   message: mysql-audit-policy

```

- This policy helps to audit the access to /var/lib/mysql/ folder. If any modification or any contents of this folder is read user will be intimated with alerts.
- Applying the Audit base policy from SaaS




Home > Runtime Security > Policies partnerdemo

## Policies

K8s aks-demo-prod wordpress-mysql Policy Type Active

Search Audit based policy is applied from AccuKnox SaaS

All (1) Discovered (0) Hardening (0) **Custom (1)** Ignore Delete

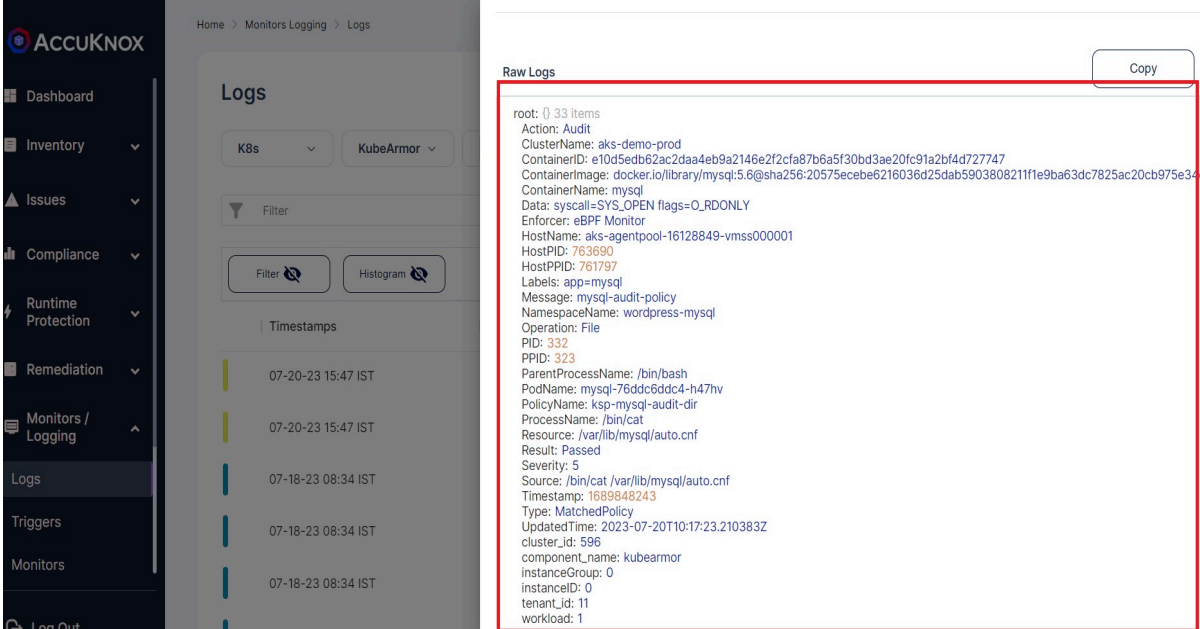
Policy Name	Category	Status	Clusters	Namespace	Selector Labels
<input type="checkbox"/>  <a href="#">ksp-mysql-audit-dir (v3)</a> KubeArmor	Custom Applied a few seco	Active	aks-demo-prod	wordpress-mysql	None

- o Now if we try to read the contents of this `/var/lib/mysql` folder running in a mysql pod by exec into the pod.

```

~$ kubectl exec -it -n wordpress-mysql mysql-76ddc6ddc4-h47hv -- bash
root@mysql-76ddc6ddc4-h47hv:/# cd /var/lib/mysql
root@mysql-76ddc6ddc4-h47hv:/var/lib/mysql# ls
auto.cnf  ib_logfile0  ib_logfile1  ibdata1  mysql  performance_schema
test  wordpress
root@mysql-76ddc6ddc4-h47hv:/var/lib/mysql# cat auto.cnf
[auto]
server-uuid=7ad615d7-0108-11ee-8442-a6440d433e17
  
```

- We can see the Audit based alert in the Monitoring/Logging Section from AccuKnox SaaS as below



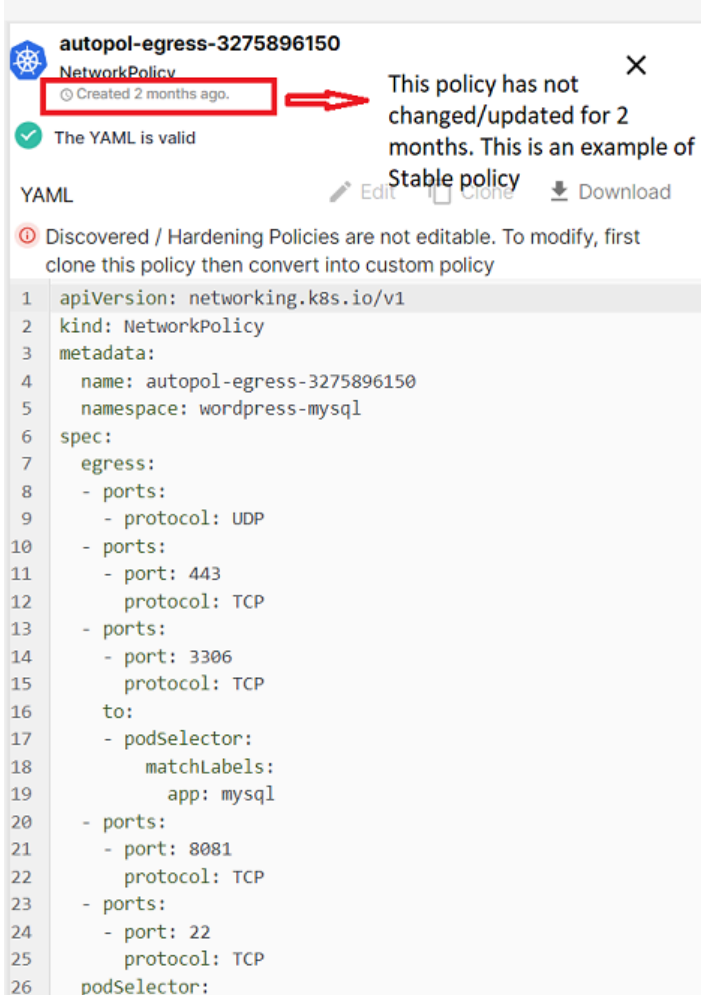
The screenshot shows the AccuKnox interface with the 'Monitors / Logging' section selected. A log entry is visible with a timestamp of 07-20-23 15:47 IST. A 'Raw Logs' window is open, displaying the following audit log entry:

```

root: {} 33 items
Action: Audit
ClusterName: aks-demo-prod
ContainerID: e10d5edb62ac2daa4eb9a2146e2f2fa87b6a5f30bd3ae20fc91a2bf4d727747
ContainerImage: docker.io/library/mysql:5.8@sha256:20575e3e3e6216036d25dab5903808211f1e9ba63dc7825ac20cb975e34cfcc
ContainerName: mysql
Data: syscall=SYS_OPEN flags=O_RDONLY
Enforcer: eBPF Monitor
HostName: aks-agentpool-16128849-vmss000001
HostPID: 763690
HostPPID: 761797
Labels: app=mysql
Message: mysql-audit-policy
NamespaceName: wordpress-mysql
Operation: File
PID: 332
PPID: 323
ParentProcessName: /bin/bash
PodName: mysql-76ddc6ddc4-h47hv
PolicyName: ksp-mysql-audit-dir
ProcessName: /bin/cat
Resource: /var/lib/mysql/auto.cnf
Result: Passed
Severity: 5
Source: /bin/cat /var/lib/mysql/auto.cnf
Timestamp: 1689848243
Type: MatchedPolicy
UpdatedTime: 2023-07-20T10:17:23.210383Z
cluster_id: 596
component_name: kubearmor
instanceGroup: 0
instanceID: 0
tenant_id: 11
workload: 1
  
```

### 13.3.4 When do we say policies are stable?

- AccuKnox Runtime Security Engine KubeArmor will discover the policies based on the Application Behavior. If the Application behavior changes the Policies generated will also be updated.
- When the policy created date or updated date doesn't change for some days then we can say that the policy which was discovered is stable. For example consider the following policy



**autopol-egress-3275896150**  
 NetworkPolicy  
 Created 2 months ago. → This policy has not changed/updated for 2 months. This is an example of Stable policy

The YAML is valid

YAML Edit Clone Download

Discovered / Hardening Policies are not editable. To modify, first clone this policy then convert into custom policy

```

1  apiVersion: networking.k8s.io/v1
2  kind: NetworkPolicy
3  metadata:
4    name: autopol-egress-3275896150
5    namespace: wordpress-mysql
6  spec:
7    egress:
8      - ports:
9        - protocol: UDP
10     - ports:
11       - port: 443
12         protocol: TCP
13     - ports:
14       - port: 3306
15         protocol: TCP
16     to:
17     - podSelector:
18       matchLabels:
19         app: mysql
20     - ports:
21       - port: 8081
22         protocol: TCP
23     - ports:
24       - port: 22
25         protocol: TCP
26     podSelector:
  
```

- The above auto discovered policy has not changed for more than a month. This policy can be called a stable policy as it didn't get any updates or changes.

### 13.3.5 What if something changes in Application?

- AccuKnox Runtime Security Engine KubeArmor will discover the policies based on the Application Behavior. If the Application behavior changes the Policies generated will also be updated.
- For example consider the following auto discovered policy


**autopol-system-1804736057 (v1)**  
 Discovered (Changes Available 2months ago)  
Created 2 months ago.

Update
✕

Updated YAML

```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: autopol-system-1804736057
5    namespace: dvwa
6  spec:
7    action: Allow
8    file:
9      matchDirectories:
10     - dir: /tmp/
11       fromSource:
12         - path: /usr/sbin/apache2
13           recursive: true
14     - dir: /var/www/html/
15       fromSource:
16         - path: /usr/sbin/apache2
17           recursive: true
18     - dir: /lib/x86_64-linux-gnu/
19       recursive: true
20     - dir: /etc/
21       fromSource:
22         - path: /bin/bash
23         - path: /bin/ping
24       recursive: true
25     - dir: /etc/
26       fromSource:
27         - path: /bin/bash

```


- In the above policy there are some changes that are detected after the initial policy discovery due to changes in application behavior. Those changes are highlighted.

```

58     path: /usr/lib/x86_64-linux-gnu/libaprutil-1.so.0
59     - fromSource:
60     - path: /usr/sbin/apache2
61     path: /usr/lib/x86_64-linux-gnu/libuuid.so.1
62 +   - fromSource:
63 +   - path: /bin/bash
64 +   path: /root/.bash_history
65 +   - fromSource:
66 +   - path: /bin/bash
67 +   path: /dev/pts/0
68 +   - fromSource:
69 +   - path: /bin/ls
70 +   path: /etc/ld.so.cache
71 +   - fromSource:
72 +   - path: /bin/ls
73 +   path: /usr/lib/x86_64-linux-gnu/libpcre2-8.so.0
74   process:
75     matchPaths:
76     - path: /usr/sbin/apache2
77     - path: /bin/bash
78     - fromSource:
79     - path: /bin/bash
80     path: /bin/ping
81     - fromSource:
82     - path: /bin/bash
83     path: /usr/sbin/apache2
23     - path: /bin/ping
24     recursive: true

```

- If the user is satisfied with the changes they can accept the change by clicking on the update button



**autopol-system-1804736057 (v1)**

Discovered (Changes Available 2months ago)

Created 2 months ago.

Update

✕

Updated YAML

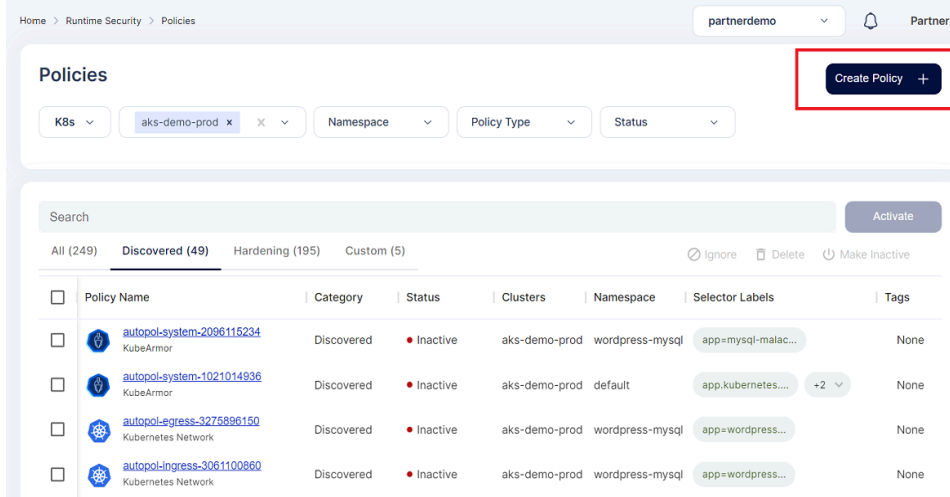
```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: autopol-system-1804736057
5    namespace: dvwa
6  spec:
7    action: Allow
8    file:
9      matchDirectories:
10     - dir: /tmp/
11       fromSource:
12         - path: /usr/sbin/apache2
13           recursive: true
14     - dir: /var/www/html/
15       fromSource:
16         - path: /usr/sbin/apache2
17           recursive: true
18     - dir: /lib/x86_64-linux-gnu/
19       recursive: true
20     - dir: /etc/
21       fromSource:
22         - path: /bin/bash
23         - path: /bin/ping
24       recursive: true
25     - dir: /etc/
26       fromSource:
27         - path: /bin/bash
                
```

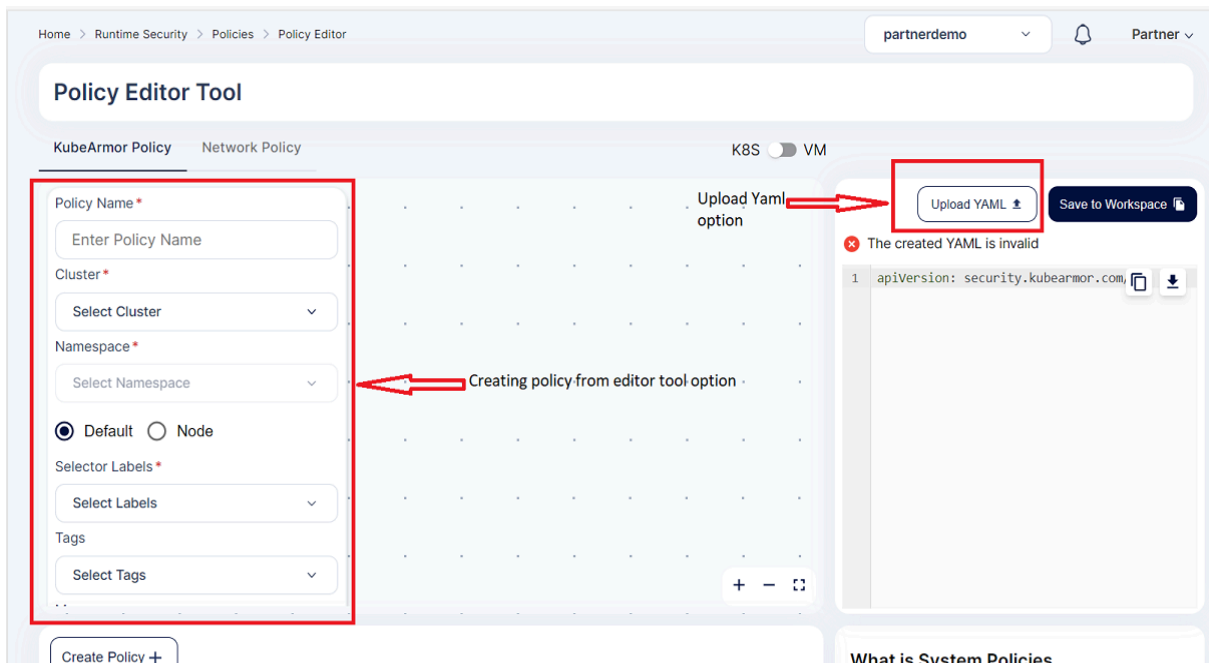
- After the user clicks the update the policy will be updated.

### 13.3.6 How to create a custom Policy

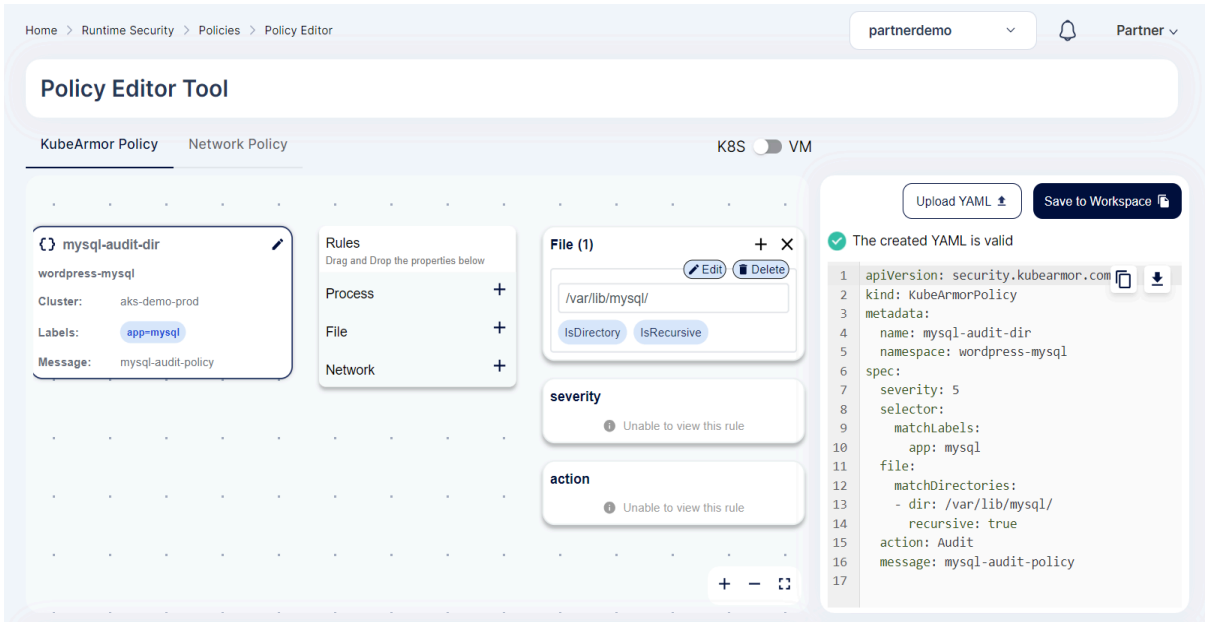
- File restriction Policy
  - To create a file restriction based custom policy user must navigate to *Runtime Protection->Policies* section.
  - To create the policy user needs to click on the create policy option



- Now user has two options either to upload the yaml file or to create the policy from policy editor tool



- Now upload the file access policy yaml from your system. After it is upload some the columns in the left side will be prefilled and user needs to select the cluster and namespace where the policy needs to applied and click save.



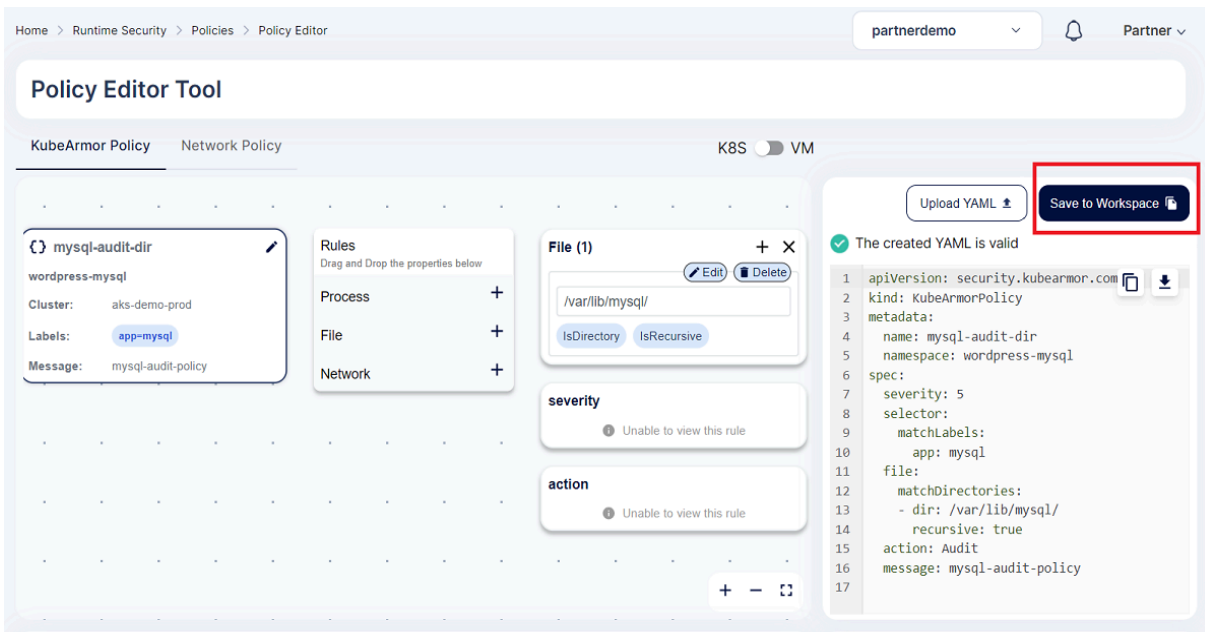
The screenshot shows the 'Policy Editor Tool' interface. On the left, a policy named 'mysql-audit-dir' is being edited for the 'wordpress-mysql' namespace. The 'Cluster' is set to 'aks-demo-prod' and the 'Labels' are 'app=mysql'. The 'Message' is 'mysql-audit-policy'. In the center, the 'File (1)' rule is configured with the path '/var/lib/mysql/' and the 'IsRecursive' checkbox is checked. The 'severity' and 'action' sections are currently disabled. On the right, the 'Save to Workspace' button is highlighted in a red box. Below the main editor, a preview of the generated YAML is shown, which is valid.

```

1  apiVersion: security.kubearmor.com
2  kind: KubeArmorPolicy
3  metadata:
4    name: mysql-audit-dir
5    namespace: wordpress-mysql
6  spec:
7    severity: 5
8    selector:
9      matchLabels:
10       app: mysql
11    file:
12      matchDirectories:
13        - dir: /var/lib/mysql/
14          recursive: true
15    action: Audit
16    message: mysql-audit-policy
17

```

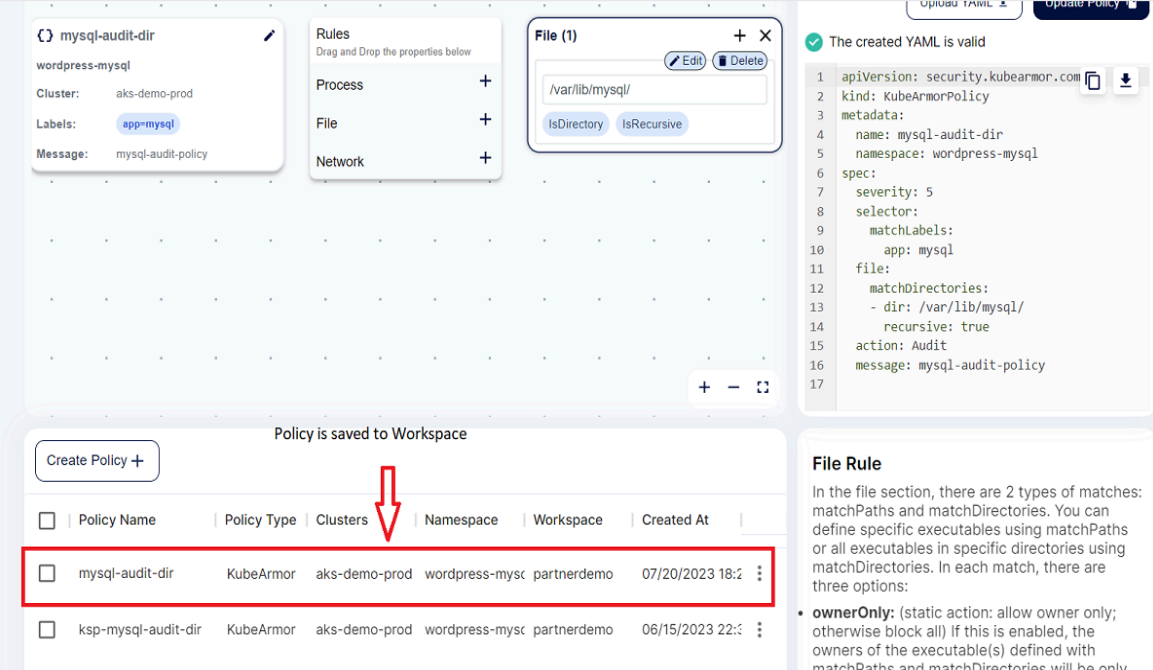
- Now to save the policy user needs to click the *save to workspace* option



This screenshot is identical to the previous one, but the 'Save to Workspace' button is highlighted with a red box to indicate the next step in the process.

- After that policy will be saved to the workspace.





mysql-audit-dir

wordpress-mysql

Cluster: aks-demo-prod

Labels: app=mysql

Message: mysql-audit-policy

Rules

Drag and Drop the properties below

Process +

File +

Network +

File (1)

/var/lib/mysql/

IsDirectory IsRecursive

The created YAML is valid

```

1 apiVersion: security.kubearmor.com
2 kind: KubeArmorPolicy
3 metadata:
4   name: mysql-audit-dir
5   namespace: wordpress-mysql
6 spec:
7   severity: 5
8   selector:
9     matchLabels:
10      app: mysql
11   file:
12     matchDirectories:
13       - dir: /var/lib/mysql/
14         recursive: true
15   action: Audit
16   message: mysql-audit-policy
17

```

Policy is saved to Workspace

Create Policy +

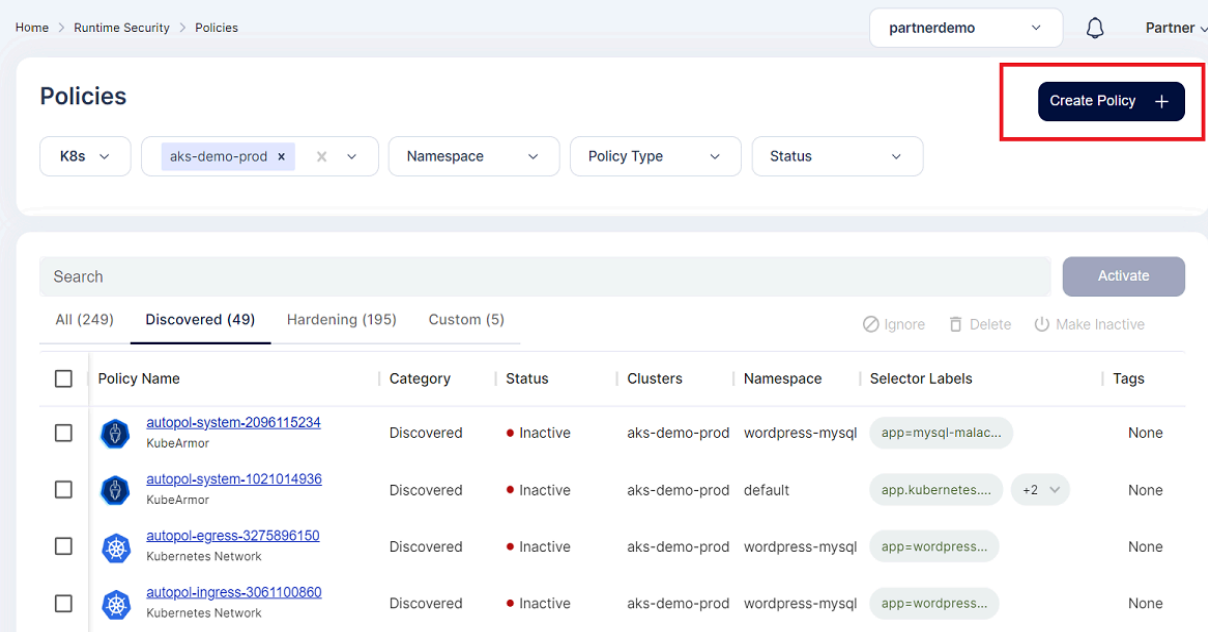
<input type="checkbox"/>	Policy Name	Policy Type	Clusters	Namespace	Workspace	Created At
<input type="checkbox"/>	mysql-audit-dir	KubeArmor	aks-demo-prod	wordpress-mysc	partnerdemo	07/20/2023 18:2
<input type="checkbox"/>	ksp-mysql-audit-dir	KubeArmor	aks-demo-prod	wordpress-mysc	partnerdemo	06/15/2023 22:3

**File Rule**

In the file section, there are 2 types of matches: matchPaths and matchDirectories. You can define specific executables using matchPaths or all executables in specific directories using matchDirectories. In each match, there are three options:

- ownerOnly:** (static action: allow owner only; otherwise block all) If this is enabled, the owners of the executable(s) defined with matchPaths and matchDirectories will be only

- Network access Policy
- To create a Network access policy restriction based custom policy user must navigate to *Runtime Protection->Policies* section.
- To create the policy user needs to click on the create policy option



Home > Runtime Security > Policies

partnerdemo Partner

**Policies**

Create Policy +

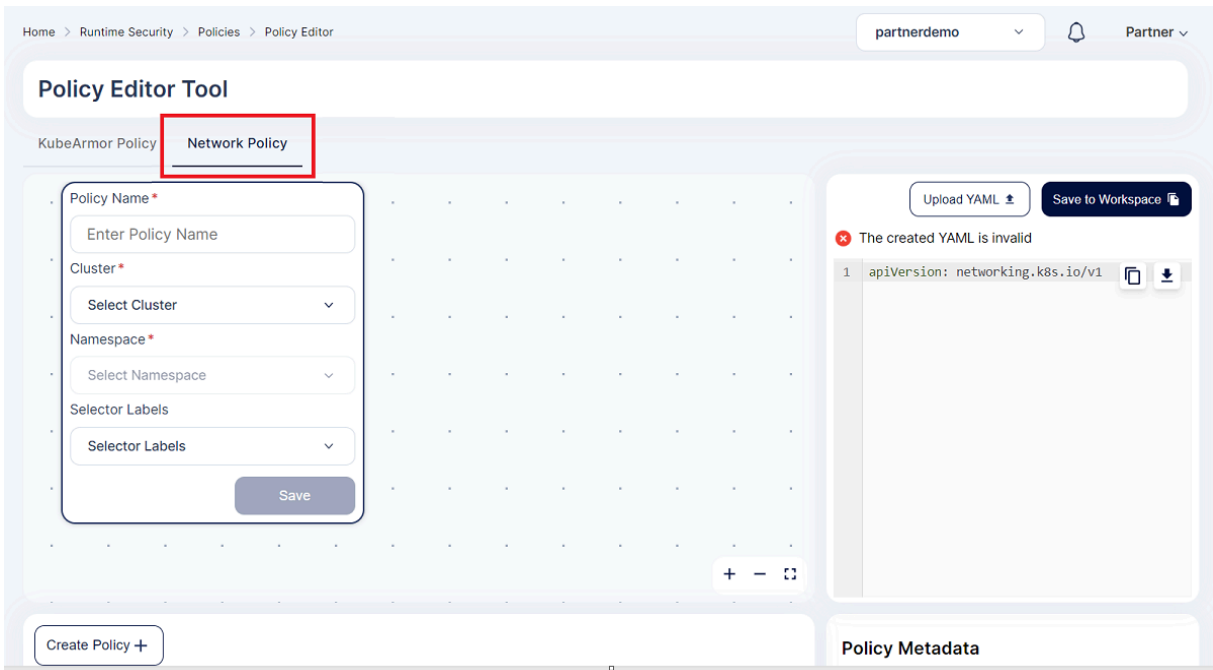
K8s aks-demo-prod Namespace Policy Type Status

Search Activate

All (249) **Discovered (49)** Hardening (195) Custom (5) Ignore Delete Make Inactive

<input type="checkbox"/>	Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input type="checkbox"/>	<a href="#">autopol-system-2096115234</a> KubeArmor	Discovered	Inactive	aks-demo-prod	wordpress-mysql	app=mysql-malac...	None
<input type="checkbox"/>	<a href="#">autopol-system-1021014936</a> KubeArmor	Discovered	Inactive	aks-demo-prod	default	app.kubernetes.... +2	None
<input type="checkbox"/>	<a href="#">autopol-egress-3275896150</a> Kubernetes Network	Discovered	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	None
<input type="checkbox"/>	<a href="#">autopol-ingress-3061100860</a> Kubernetes Network	Discovered	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	None

- In this screen for Network Policy creation user needs to select the Network policy editor tool



Home > Runtime Security > Policies > Policy Editor

partnerdemo Partner

### Policy Editor Tool

KubeArmor Policy **Network Policy**

Policy Name \*  
Enter Policy Name

Cluster \*  
Select Cluster

Namespace \*  
Select Namespace

Selector Labels  
Selector Labels

Save

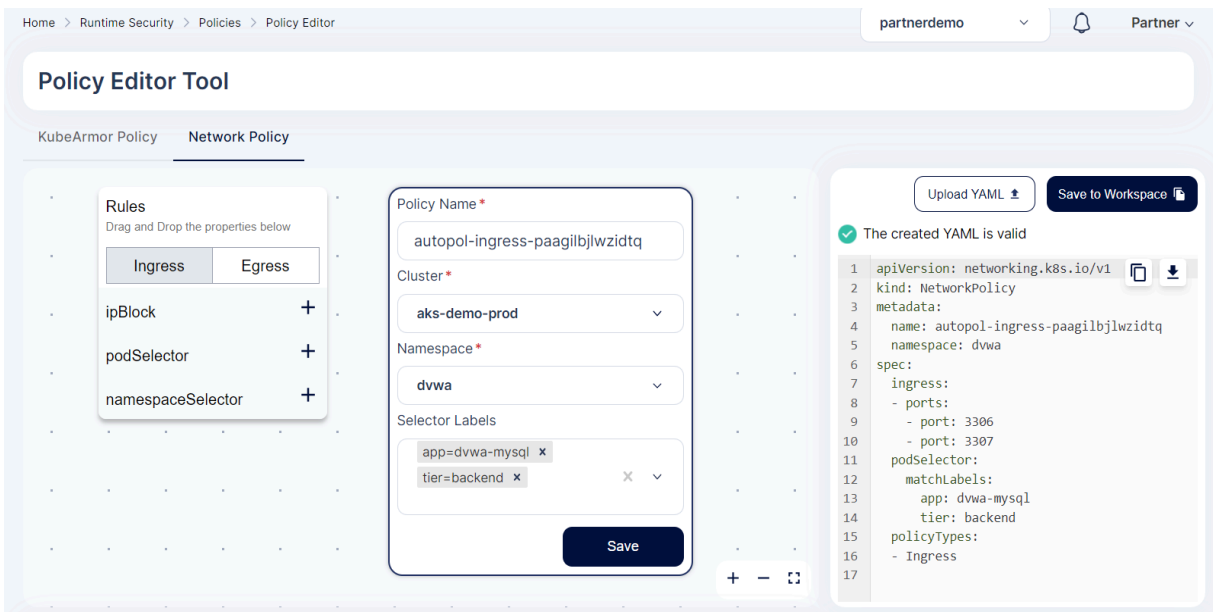
Upload YAML Save to Workspace

**The created YAML is invalid**

```
1 apiVersion: networking.k8s.io/v1
```

Create Policy + Policy Metadata

- Now upload the Network policy yaml from your system by clicking the *upload yaml* option. After it is upload some the columns in the left side will be prefilled and user needs to select the cluster and namespace where the policy needs to applied and click save.



Home > Runtime Security > Policies > Policy Editor

partnerdemo Partner

### Policy Editor Tool

KubeArmor Policy **Network Policy**

Rules  
Drag and Drop the properties below

Ingress Egress

ipBlock +

podSelector +

namespaceSelector +

Policy Name \*  
autopol-ingress-paagilbjlwzidtq

Cluster \*  
aks-demo-prod

Namespace \*  
dvwa

Selector Labels  
app=dvwa-mysql x  
tier=backend x

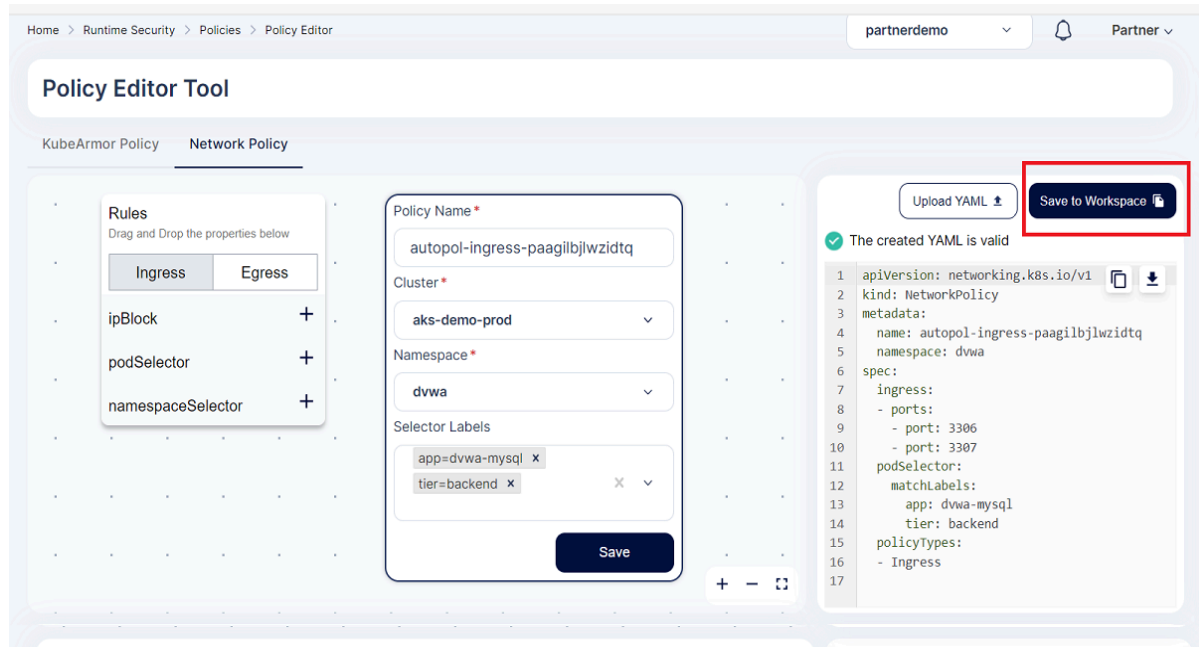
Save

Upload YAML Save to Workspace

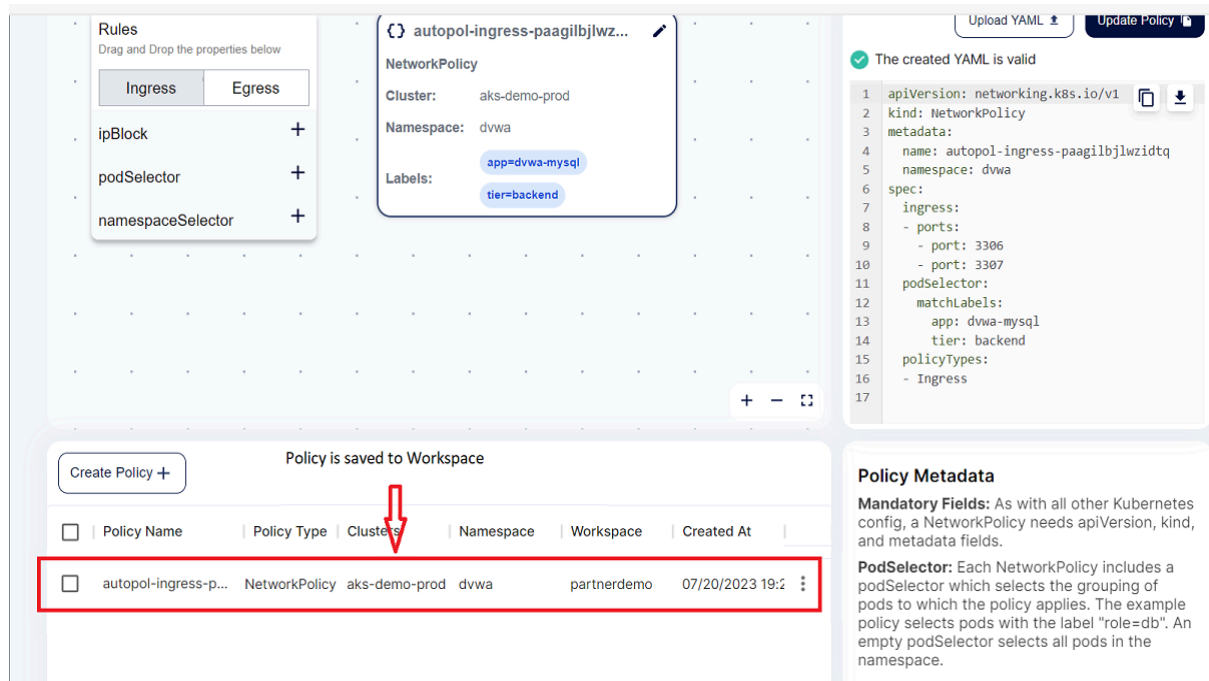
**The created YAML is valid**

```
1 apiVersion: networking.k8s.io/v1
2 kind: NetworkPolicy
3 metadata:
4   name: autopol-ingress-paagilbjlwzidtq
5   namespace: dvwa
6 spec:
7   ingress:
8     - ports:
9       - port: 3306
10      - port: 3307
11   podSelector:
12     matchLabels:
13       app: dvwa-mysql
14       tier: backend
15   policyTypes:
16     - Ingress
17
```

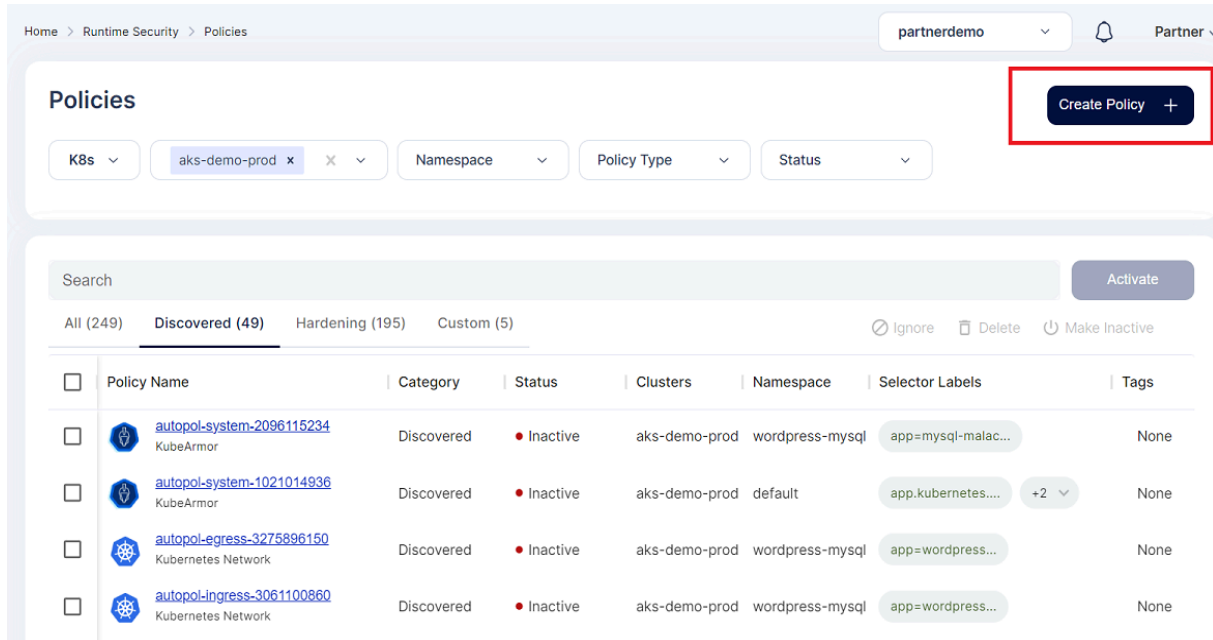
- Now to save the policy user needs to click the *save to workspace* option



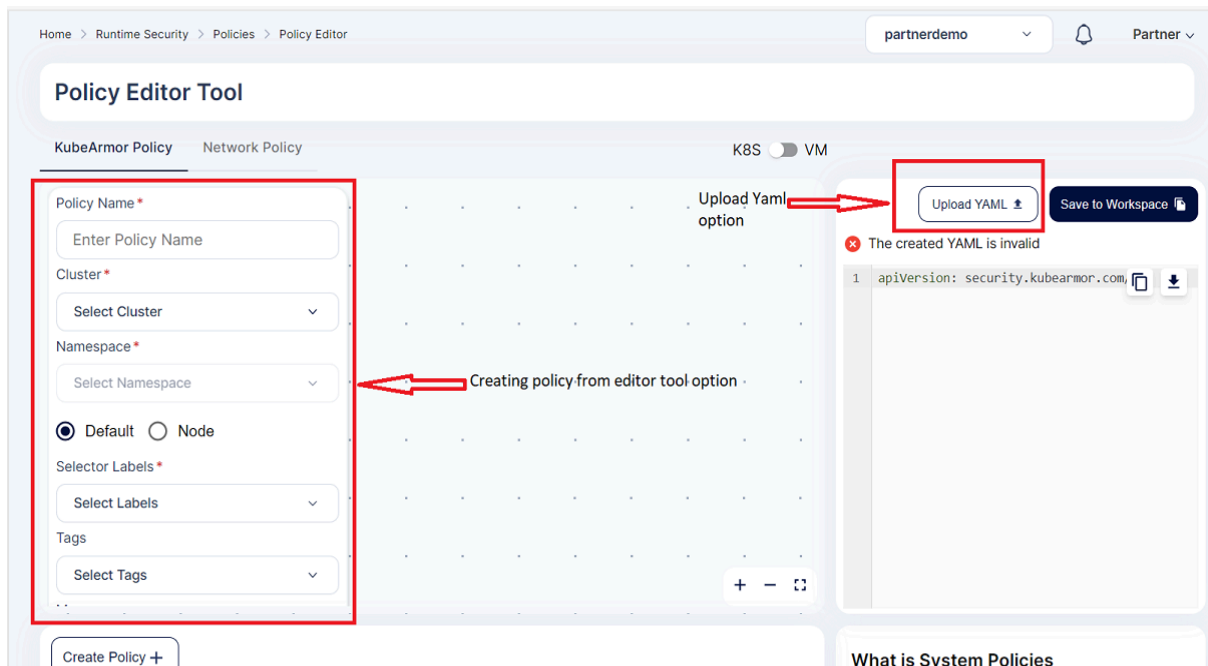
- After that policy will be saved to the workspace.



- Process block restriction Policy
  - To create a Process access restriction based custom policy user must navigate to *Runtime Protection->Policies* section.
  - To create the policy user needs to click on the create policy option



- Now user has two options either to upload the yaml file or to create the policy from policy editor tool



- Now upload the process block policy yaml from your system. After it is upload some the columns in the left side will be prefilled and user needs to select the cluster and namespace where the policy needs to applied and click save.

Home > Runtime Security > Policies > Policy Editor

partnerdemo Partner

### Policy Editor Tool

KubeArmor Policy Network Policy K8S VM

**wordpress-block-process**

wordpress-mysql

Cluster: aks-demo-prod

Labels: app=wordpress

Message: apt process block

**Rules**  
Drag and Drop the properties below

Process +

File +

Network +

**Process (2)**

/usr/bin/apt

/usr/bin/apt-get

**severity**  
Unable to view this rule

**action**  
Unable to view this rule

Upload YAML

**Save to Workspace**

✓ The created YAML is valid

```

1 apiVersion: security.kubearmor.com
2 kind: KubeArmorPolicy
3 metadata:
4   name: wordpress-block-process
5   namespace: wordpress-mysql
6 spec:
7   severity: 3
8   selector:
9     matchLabels:
10      app: wordpress
11   process:
12     matchPaths:
13     - path: /usr/bin/apt
14     - path: /usr/bin/apt-get
15   action: Block
16   message: 'apt process block '
17

```

Now to save the policy user needs to click the *save to workspace* option

Home > Runtime Security > Policies > Policy Editor

partnerdemo Partner

### Policy Editor Tool

KubeArmor Policy Network Policy K8S VM

**wordpress-block-process**

wordpress-mysql

Cluster: aks-demo-prod

Labels: app=wordpress

Message: apt process block

**Rules**  
Drag and Drop the properties below

Process +

File +

Network +

**Process (2)**

/usr/bin/apt

/usr/bin/apt-get

**severity**  
Unable to view this rule

**action**  
Unable to view this rule

Upload YAML

**Save to Workspace**

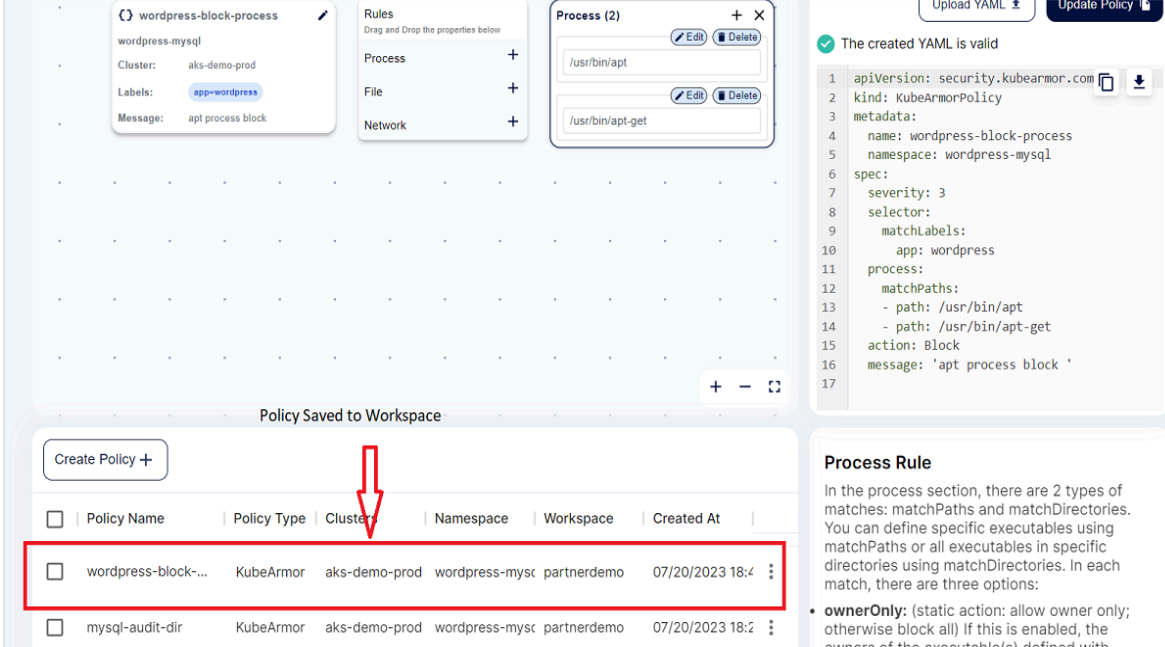
✓ The created YAML is valid

```

1 apiVersion: security.kubearmor.com
2 kind: KubeArmorPolicy
3 metadata:
4   name: wordpress-block-process
5   namespace: wordpress-mysql
6 spec:
7   severity: 3
8   selector:
9     matchLabels:
10      app: wordpress
11   process:
12     matchPaths:
13     - path: /usr/bin/apt
14     - path: /usr/bin/apt-get
15   action: Block
16   message: 'apt process block '
17

```

After that policy will be saved to the workspace.



The screenshot shows the AccuKnox interface for configuring a policy. On the left, a policy named 'wordpress-block-process' is being configured for the 'aks-demo-prod' cluster, with labels 'app=wordpress' and a message 'apt process block'. The 'Rules' section shows 'Process' selected. A 'Process (2)' window shows two paths: '/usr/bin/apt' and '/usr/bin/apt-get'. On the right, the generated YAML is displayed, showing a KubeArmorPolicy for 'wordpress-block-process' in the 'wordpress-mysql' namespace, with a severity of 3 and a block action. Below the configuration, a table lists existing policies, with a red box highlighting the 'wordpress-block-process' entry.

<input type="checkbox"/>	Policy Name	Policy Type	Cluster	Namespace	Workspace	Created At
<input type="checkbox"/>	wordpress-block-...	KubeArmor	aks-demo-prod	wordpress-mysc	partnerdemo	07/20/2023 18:4
<input type="checkbox"/>	mysql-audit-dir	KubeArmor	aks-demo-prod	wordpress-mysc	partnerdemo	07/20/2023 18:2

### 13.3.7 How to enforce Policies and see anomalies

- We can apply any of the Auto Discovered, Hardening or custom policies and see the anomalies getting detected using the Monitoring and Logging section.
- Let us consider the WordPress- MySQL application. In the MySQL application, certain folders will be having certain critical data which can be allowed to access but not modified. So using our AccuKnox hardening policy we are going to prevent the modification of contents inside these critical folders.
- **Before applying the policy:** Currently, any attacker who gets access to the bash or shell of the MySQL pod can modify the contents of the/sbin folder by creating a new file and editing the old files.

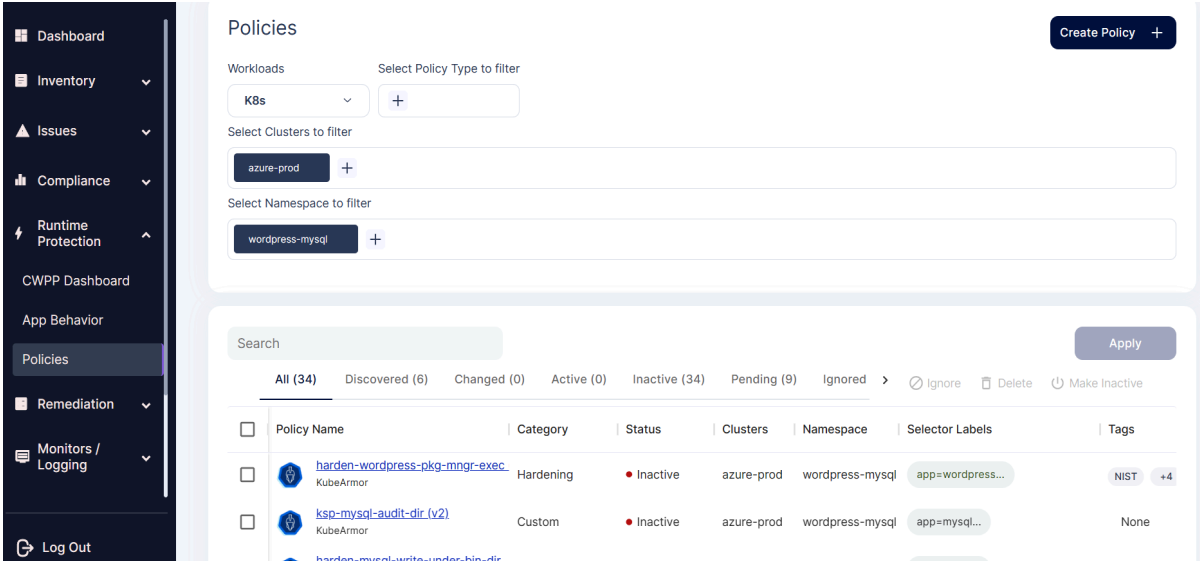
```

root@mysql-6c6fdccf-sk5x2:/# cd sbin
root@mysql-6c6fdccf-sk5x2:/sbin# ls
agetty      dumpe2fs    fsck.ext2   installkernel  mkfs.cramfs  pivot_root  swapoff
badblocks   e2fsck      fsck.ext3   isosize        mkfs.ext2    raw         swapon
blkdiscard  e2image     fsck.ext4   killall5       mkfs.ext3    resize2fs   switch_root
blkid       e2label     fsck.minix  ldconfig       mkfs.ext4    runuser     tune2fs
blockdev    e2undo      fsfreeze    logsave        mkfs.minix   sfdisk      unix_chkpwd
cfdisk      fdisk       fstab-decode  losetup        mkhomedir_helper  shadowconfig  unix_update
chcpu       findfs      fstrim      mke2fs         mkswap       start-stop-daemon  wipefs
ctrlaltdel  fsck        getty       mkfs           pam_tally    sulogin     zramctl
debugfs     fsck.cramfs hwclock     mkfs.bfs       pam_tally2   swaplabel

root@mysql-6c6fdccf-sk5x2:/sbin# touch mks2
root@mysql-6c6fdccf-sk5x2:/sbin# ls
agetty      dumpe2fs    fsck.ext2   installkernel  mkfs.cramfs  pam_tally2  swaplabel
badblocks   e2fsck      fsck.ext3   isosize        mkfs.ext2    pivot_root  swapoff
blkdiscard  e2image     fsck.ext4   killall5       mkfs.ext3    raw         swapon
blkid       e2label     fsck.minix  ldconfig       mkfs.ext4    resize2fs   switch_root
blockdev    e2undo      fsfreeze    logsave        mkfs.minix   runuser     tune2fs
cfdisk      fdisk       fstab-decode  losetup        mkhomedir_helper  sfdisk      unix_chkpwd
chcpu       findfs      fstrim      mke2fs         mks2         shadowconfig  unix_update
ctrlaltdel  fsck        getty       mkfs           mkswap       start-stop-daemon  wipefs
debugfs     fsck.cramfs hwclock     mkfs.bfs       pam_tally    sulogin     zramctl
root@mysql-6c6fdccf-sk5x2:/sbin#

```

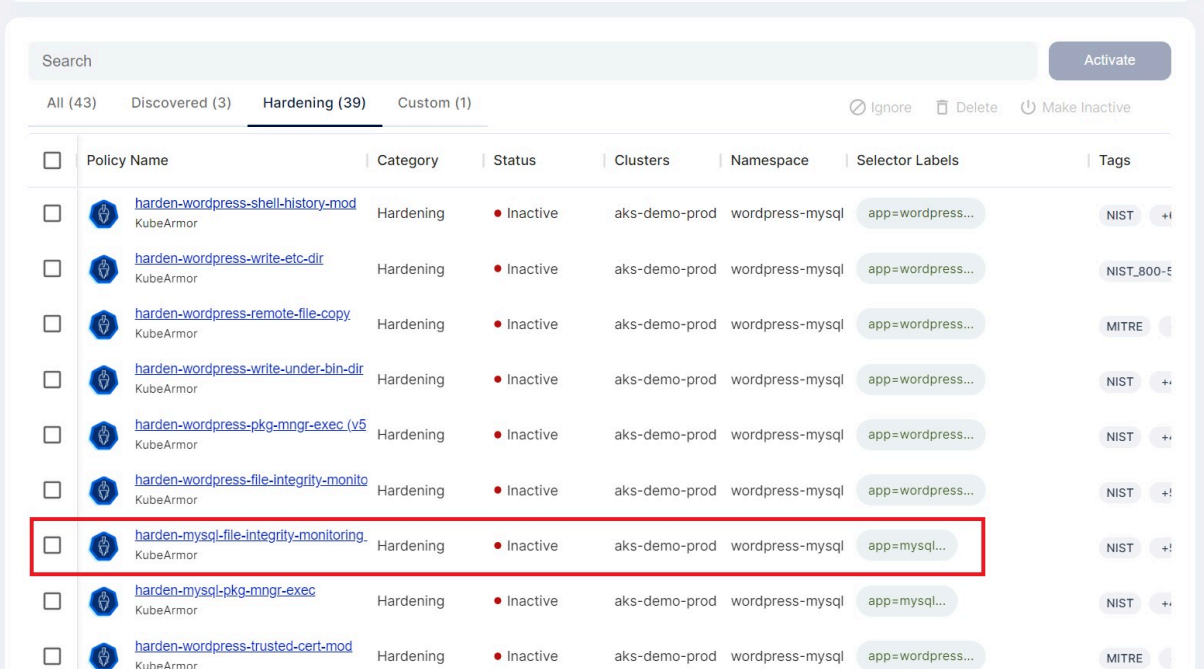
- Now we are going to prevent this using AccuKnox CWPP Solution.
- **Step 1:** Navigate to the Runtime Protection-> Policies and select the cluster and namespace where the WordPress-MySQL application is deployed.



The screenshot shows the AccuKnox CWPP interface. On the left is a navigation sidebar with options like Dashboard, Inventory, Issues, Compliance, Runtime Protection, CWPP Dashboard, App Behavior, Policies (selected), Remediation, and Monitors / Logging. The main area is titled 'Policies' and includes a 'Create Policy +' button. Below this are filters for Workloads (K8s), Clusters (azure-prod), and Namespace (wordpress-mysql). A table lists the policies:

Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
hardened-wordpress-pkg-mngt-exec KubeArmor	Hardening	Inactive	azure-prod	wordpress-mysql	app=wordpress...	NIST +4
ksp-mysql-audit-dir (v2) KubeArmor	Custom	Inactive	azure-prod	wordpress-mysql	app=mysql...	None
hardened-mysql-write-under-bin-dir						


- **Step 2:** In the screen select the hardening policies in the policy filter section to view the hardening policies related to the WordPress-MySQL application.



<input type="checkbox"/>	Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input type="checkbox"/>	<a href="#">harden-wordpress-shell-history-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-write-etc-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST_800-5
<input type="checkbox"/>	<a href="#">harden-wordpress-remote-file-copy</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE
<input type="checkbox"/>	<a href="#">harden-wordpress-write-under-bin-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-pkg-mngr-exec (v5)</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-file-integrity-monitoring</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-mysql-file-integrity-monitoring</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +
<input type="checkbox"/>	<a href="#">harden-mysql-pkg-mngr-exec</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-trusted-cert-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE

- **Step 3:** Click on the MySQL file integrity hardening policy from the list of policies to see the policy





**harden-mysql-file-integrity-monitoring**  
KubeArmorPolicy  
Created 5 days ago.

✕

YAML ✎ Edit 📄 Clone ⬇️ Download

ⓘ Discovered / Hardening Policies are not editable. To modify, first clone this policy then convert into custom policy

```

1  apiVersion: security.kubearmor.com/v1
2  kind: KubeArmorPolicy
3  metadata:
4    name: harden-mysql-file-integrity-monitoring
5    namespace: wordpress-mysql
6  spec:
7    action: Block
8    file:
9      matchDirectories:
10     - dir: /sbin/
11       readOnly: true
12       recursive: true
13     - dir: /usr/bin/
14       readOnly: true
15       recursive: true
16     - dir: /usr/lib/
17       readOnly: true
18       recursive: true
19     - dir: /usr/sbin/
20       readOnly: true
21       recursive: true
22     - dir: /bin/
23       readOnly: true
24       recursive: true
25     - dir: /boot/
26       readOnly: true
27       recursive: true
28  message: Detected and prevented compromise to File integrity
29  selector:
30    matchLabels:

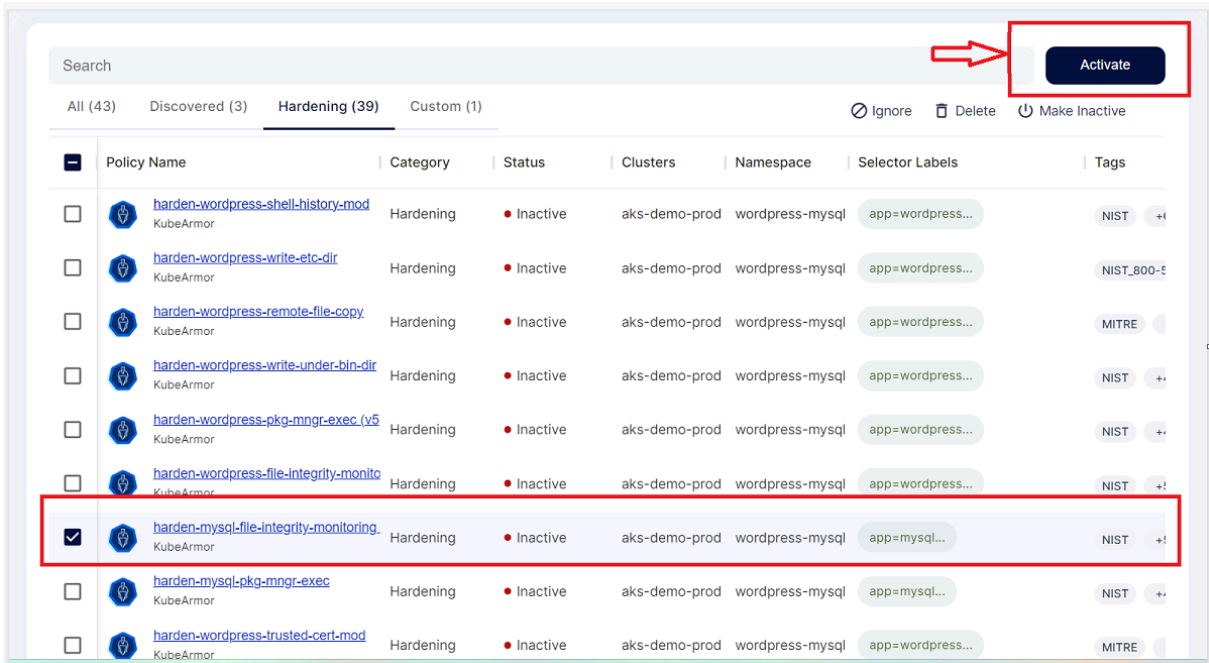
```

- The policy is allowing users to access the critical folders but it is blocking the write or modify access by whitelisting only read access.

```
apiVersion: security.kubearmor.com/v1
kind: KubeArmorPolicy
metadata:
  name: harden-mysql-file-integrity-monitoring
  namespace: wordpress-mysql
spec:
  action: Block
  file:
    matchDirectories:
      - dir: /sbin/
        readOnly: true
        recursive: true
      - dir: /usr/bin/
        readOnly: true
        recursive: true
      - dir: /usr/lib/
        readOnly: true
        recursive: true
      - dir: /usr/sbin/
        readOnly: true
        recursive: true
      - dir: /bin/
        readOnly: true
        recursive: true
      - dir: /boot/
        readOnly: true
        recursive: true
    message: Detected and prevented compromise
    to File integrity
  selector:
    matchLabels:
      app: mysql
    severity: 1
  tags:
    - NIST
    - NIST_800-53_AU-2
    - NIST_800-53_SI-4
    - MITRE
```

- MITRE\_T1036\_masquerading
- MITRE\_T1565\_data\_manipulation

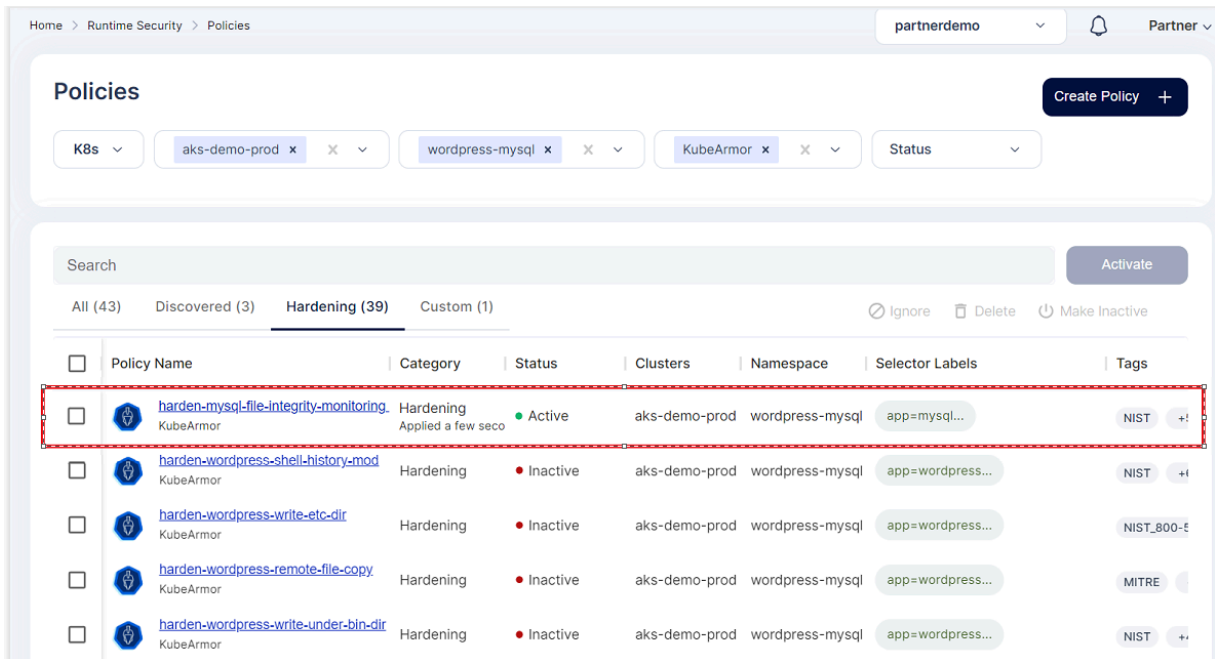
- **Step 4:** To apply this policy, select the policy checkbox and click Activate option



The screenshot shows the AccuKnox console interface. At the top, there is a search bar and a navigation menu with tabs for 'All (43)', 'Discovered (3)', 'Hardening (39)', and 'Custom (1)'. Below the navigation, there are action buttons: 'Ignore', 'Delete', and 'Make Inactive'. The main content is a table of policies. The table has columns for Policy Name, Category, Status, Clusters, Namespace, Selector Labels, and Tags. The policy 'activate-mysql-file-integrity-monitoring' is highlighted in blue and has its checkbox checked. A red box highlights the 'Activate' button in the top right corner, with a red arrow pointing to it.

Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input type="checkbox"/> <a href="#">hardened-wordpress-shell-history-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +!
<input type="checkbox"/> <a href="#">hardened-wordpress-write-etc-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST_800-5
<input type="checkbox"/> <a href="#">hardened-wordpress-remote-file-copy</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE
<input type="checkbox"/> <a href="#">hardened-wordpress-write-under-bin-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST ++
<input type="checkbox"/> <a href="#">hardened-wordpress-pkg-mngr-exec (v5)</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST ++
<input type="checkbox"/> <a href="#">hardened-wordpress-file-integrity-monitoring</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +!
<input checked="" type="checkbox"/> <a href="#">activate-mysql-file-integrity-monitoring</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +!
<input type="checkbox"/> <a href="#">hardened-mysql-pkg-mngr-exec</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=mysql...	NIST ++
<input type="checkbox"/> <a href="#">hardened-wordpress-trusted-cert-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE

- **Step 5:** Now the policy is active and applied on the cluster



Home > Runtime Security > Policies

partnerdemo Partner

Policies Create Policy +

K8s aks-demo-prod x wordpress-mysql x KubeArmor x Status

Search Activate

All (43) Discovered (3) **Hardening (39)** Custom (1) Ignore Delete Make Inactive

<input type="checkbox"/>	Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input type="checkbox"/>	<a href="#">harden-mysql-file-integrity-monitoring</a> KubeArmor	Hardening Applied a few seco	Active	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +!
<input type="checkbox"/>	<a href="#">harden-wordpress-shell-history-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +!
<input type="checkbox"/>	<a href="#">harden-wordpress-write-etc-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST_800-5
<input type="checkbox"/>	<a href="#">harden-wordpress-remote-file-copy</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE
<input type="checkbox"/>	<a href="#">harden-wordpress-write-under-bin-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +!

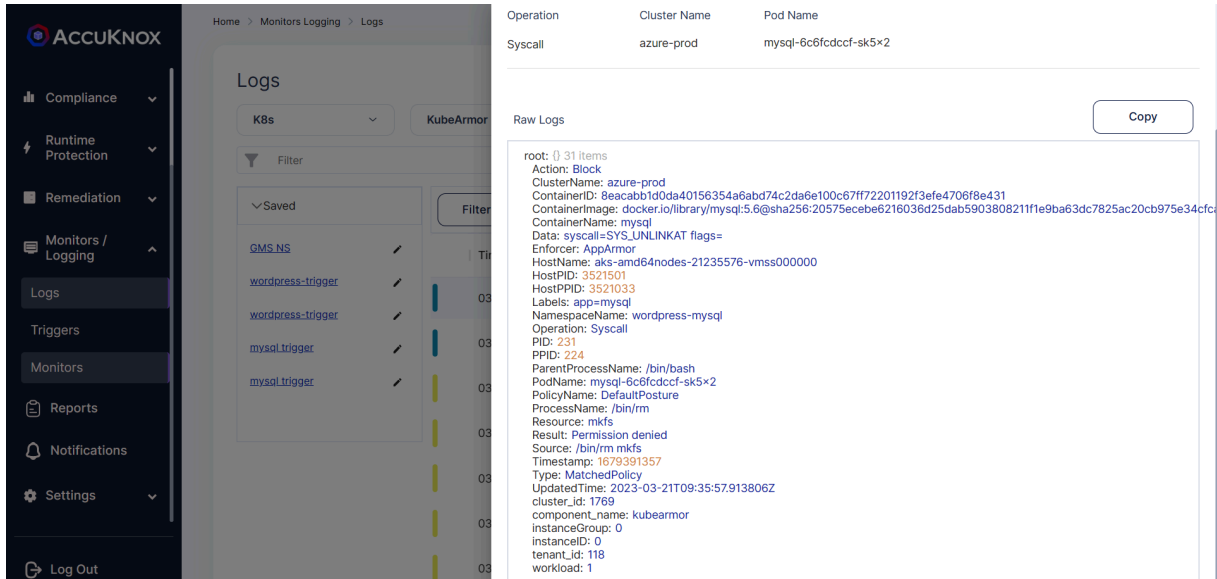
- **Step 6:** If any attacker now tries to modify the content of the critical folders it will be blocked.

```

root@mysql-6c6fcdccf-sk5x2:/# cd sbin
root@mysql-6c6fcdccf-sk5x2:/sbin# ls
agetty      dumpe2fs    fsck.ext2   installkernel  mkfs.cramfs    pam_tally2    swaplabel
badblocks   e2fsck      fsck.ext3   isosize        mkfs.ext2      pivot_root    swapoff
blkdiscard  e2image     fsck.ext4   killall5       mkfs.ext3      raw           swapon
blkid       e2label     fsck.minix  ldconfig       mkfs.ext4      resize2fs    switch_root
blockdev    e2undo      fsfreeze    logsave        mkfs.minix     runuser      tune2fs
cfdisk      fdisk       fstab-decode  losetup        mkhomedir_helper  sfdisk       unix_chkpwd
chcpu       findfs      fstrim      mke2fs         mks2           shadowconfig  unix_update
ctrlaltdel  fsck        getty        mkfs           mkswap         start-stop-daemon  wipefs
debugfs     fsck.cramfs hwclock     mkfs.bfs       pam_tally      sulogin      zramctl
root@mysql-6c6fcdccf-sk5x2:/sbin# rm mkfs
rm: cannot remove 'mkfs': Permission denied
root@mysql-6c6fcdccf-sk5x2:/sbin#

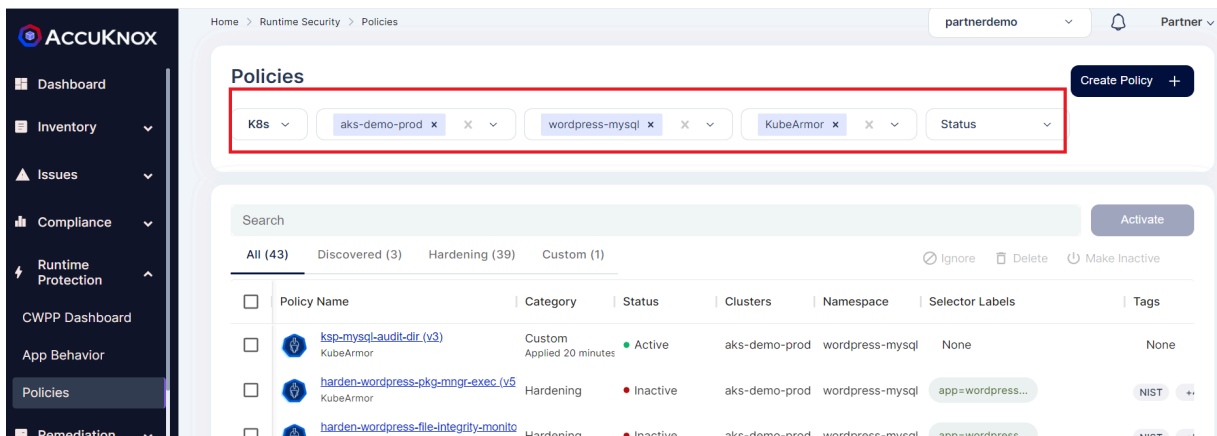
```

- **Step 7:** To see the logs Navigate to the Monitoring/logging->logs



### 13.3.8 How to perform bulk operation on applying policies

- AccuKnox SaaS supports applying multiple policies at one time. To perform this user must navigate to the *Runtime Protection->Policies* Section.
- From the Filters shown in the Screen user must select the Cluster and Namespace for which they are going to apply multiple policies



- To apply multiple policies in single go select the all policies from the screen and click Activate button

Home > Runtime Security > Policies

partnerdemo Partner

### Policies

Create Policy +

K8s aks-demo-prod wordpress-mysql KubeArmor Status

Applying Multiple policies

Search Activate

All (43) Discovered (3) Hardening (39) Custom (1)

<input checked="" type="checkbox"/>	Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input checked="" type="checkbox"/>	<a href="#">ksp-mysql-audit-dir (v3)</a> KubeArmor	Custom	Inactive	aks-demo-prod	wordpress-mysql	None	None
<input checked="" type="checkbox"/>	<a href="#">autopol-system-2096115234</a> KubeArmor	Discovered	Inactive	aks-demo-prod	wordpress-mysql	app=mysql-malac...	None
<input checked="" type="checkbox"/>	<a href="#">harden-wordpress-shell-history-mod</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input checked="" type="checkbox"/>	<a href="#">harden-wordpress-write-etc-dir</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST_800-8
<input checked="" type="checkbox"/>	<a href="#">harden-wordpress-remote-file-copy</a> KubeArmor	Hardening	Inactive	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE

- Now after activating all the policies they will be made active and applied in the cluster.

ACCUKNOX

Dashboard Inventory Issues Compliance Runtime Protection CWPP Dashboard App Behavior Policies Remediation Monitors / Logging Log Out

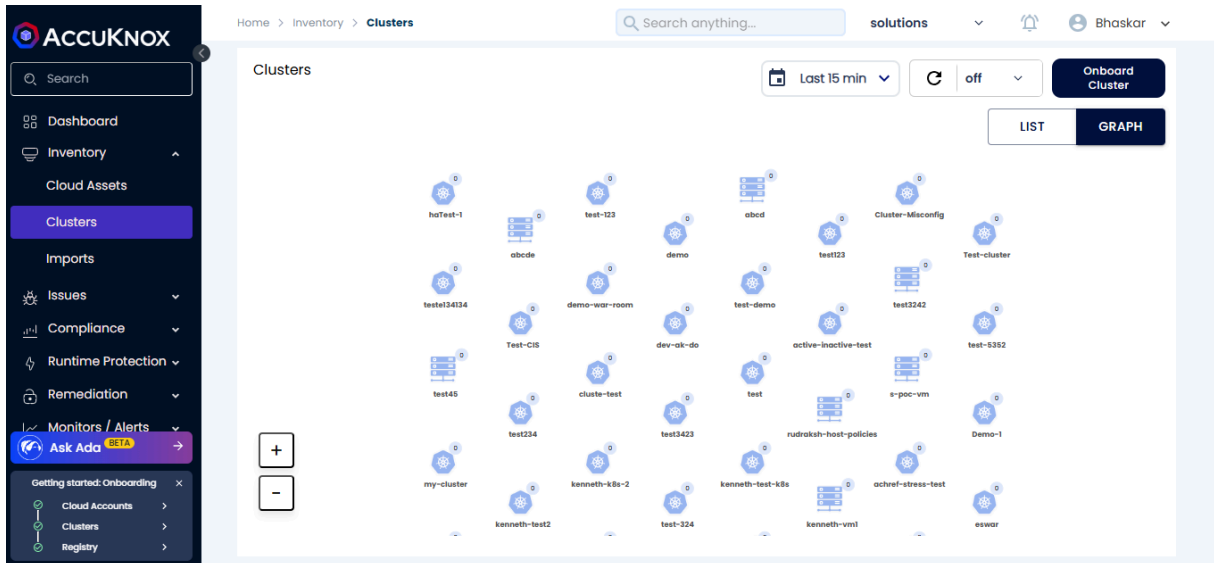
Search Activate

All (43) Discovered (3) Hardening (39) Custom (1)

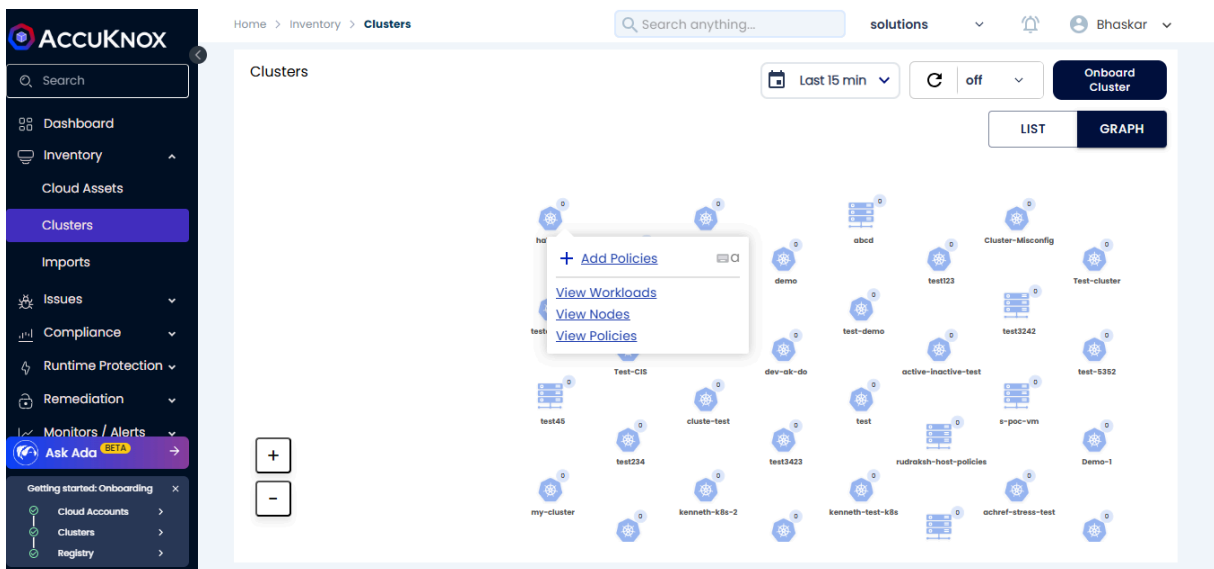
<input type="checkbox"/>	Policy Name	Category	Status	Clusters	Namespace	Selector Labels	Tags
<input type="checkbox"/>	<a href="#">harden-mysql-file-integrity-monitoring</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-remote-file-copy</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	MITRE
<input type="checkbox"/>	<a href="#">harden-wordpress-write-etc-dir</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST_800-8
<input type="checkbox"/>	<a href="#">harden-wordpress-shell-history-mod</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">autopol-system-2096115234</a> KubeArmor	Discovered	Active	aks-demo-prod	wordpress-mysql	app=mysql-malac...	None
<input type="checkbox"/>	<a href="#">harden-mysql-pkg-mngr-exec</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=mysql...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-pkg-mngr-exec (v5)</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-write-under-bin-dir</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +
<input type="checkbox"/>	<a href="#">harden-wordpress-file-integrity-monito</a> KubeArmor	Hardening	Active	aks-demo-prod	wordpress-mysql	app=wordpress...	NIST +

### 13.3.9 How to Find Nodes of a VM cluster

Under Inventory -> Clusters -> you can see the graphical view of all the clusters.



To see nodes of your cluster, click on the cluster name, click on “View Nodes”



You can see the Nodes of that cluster.

Home > Inventory > Clusters

Search anything...

solutions

Bhaskar

Clusters > Nodes

Last 15 min

off

Onboard Cluster

LIST GRAPH

+

-

Getting started: Onboarding

- Cloud Accounts
- Clusters
- Registry



## 14. Host Security

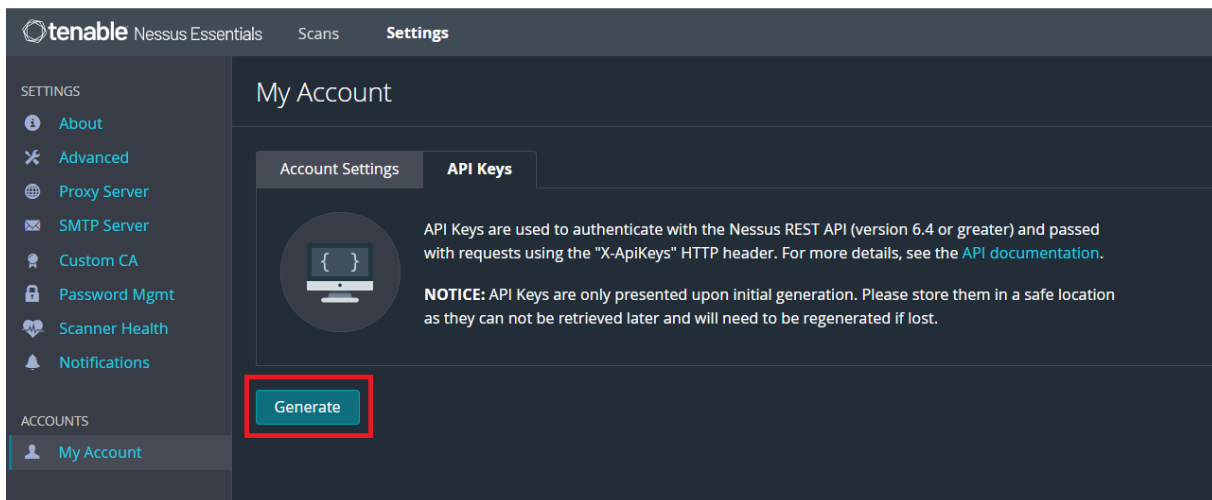
### 14.1 Host Scan

AccuKnox provides host scanning capabilities through an integration with tools like Nessus, Masscan and Zeek.

### 14.2 Prerequisites for Nessus Integration

To integrate Nessus with AccuKnox, the Nessus scanner is required to be deployed securely in the target environment with the nessus port accessible by AccuKnox SaaS. To fetch the results from the Nessus deployment, AccuKnox requires the Nessus Deployment URL and the API keys.

To generate the API keys, switch to the **Settings** tab, navigate to **My Account** → **API keys** and click on **Generate**

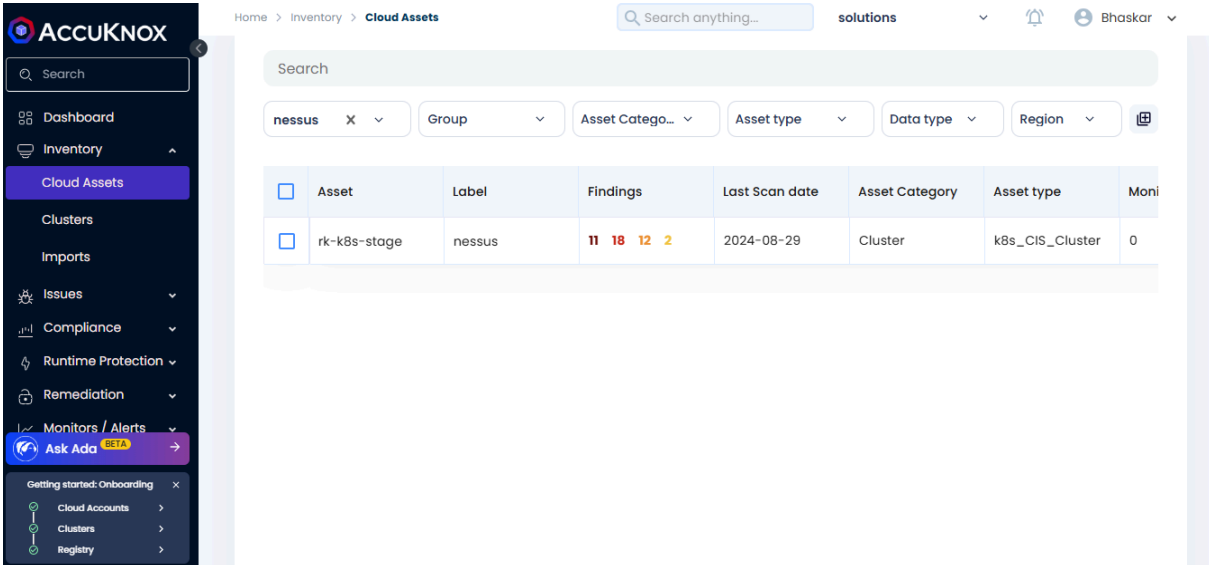


Copy the Access and Secret keys that are generated for the integration.

## 14.3 Asset Inventory

As Nessus scans all the hosts in the environment, the data is processed to improve the Asset Inventory of AccuKnox with additional data about the Hosts scanned through Nessus.

**Step 1:** Navigate to Inventory → Cloud Assets



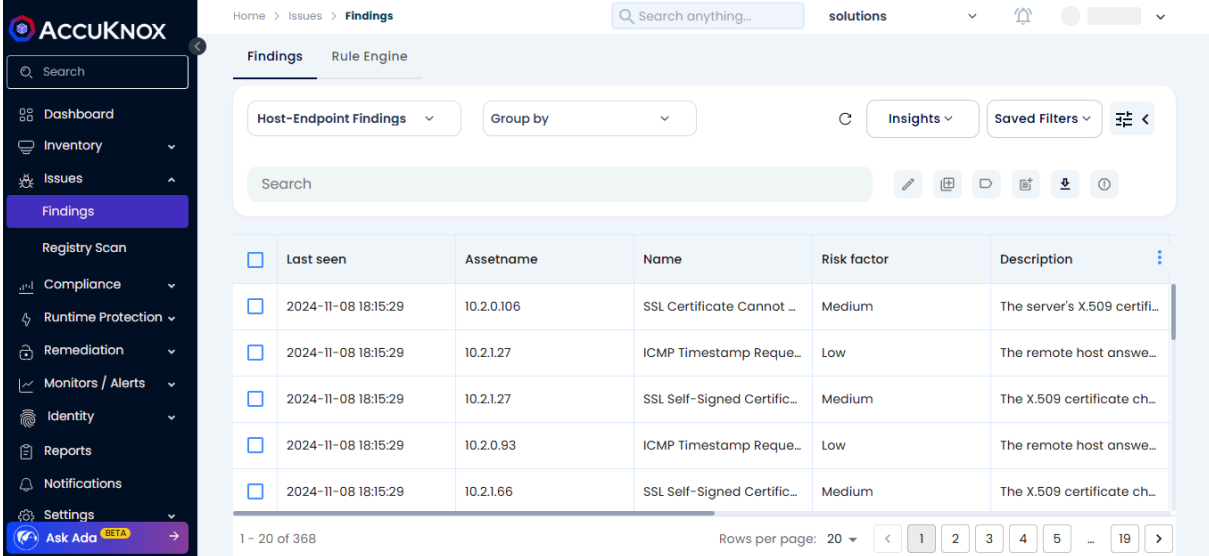
Asset	Label	Findings	Last Scan date	Asset Category	Asset type	Monitors
<input type="checkbox"/>	rk-k8s-stage	nessus	2024-08-29	Cluster	k8s_CIS_Cluster	0

**Step 2:** Filter using the **Label** name used for integrating Nessus to view only the assets that were identified using Nessus scanner and their associated findings.

## 14.4 Vulnerability Management

The vulnerabilities found using Nessus are populated in the AccuKnox SaaS in addition to the findings from other tools for easy management on a single platform.

**Step 1:** Navigate to Issues → Findings -> select Host-Endpoint Findings from the Findings filter.



Home > Issues > Findings

Search anything...

solutions

Findings Rule Engine

Host-Endpoint Findings Group by Insights Saved Filters

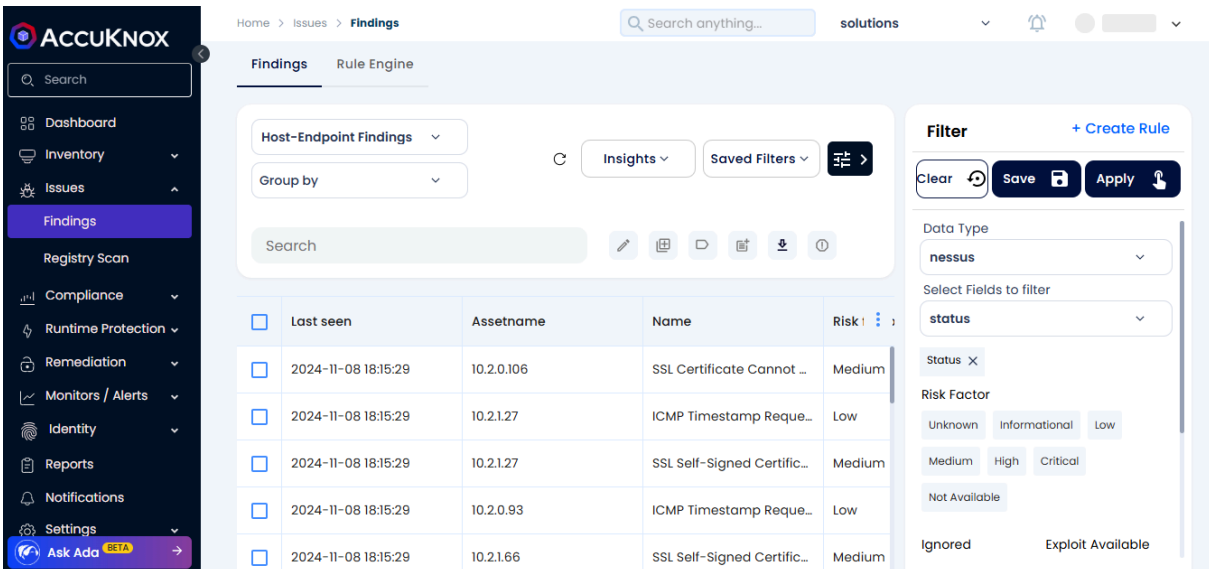
Search

<input type="checkbox"/>	Last seen	Assetname	Name	Risk factor	Description
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.0.106	SSL Certificate Cannot ...	Medium	The server's X.509 certifi...
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.27	ICMP Timestamp Reque...	Low	The remote host answe...
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.27	SSL Self-Signed Certific...	Medium	The X.509 certificate ch...
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.0.93	ICMP Timestamp Reque...	Low	The remote host answe...
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.66	SSL Self-Signed Certific...	Medium	The X.509 certificate ch...

1 - 20 of 368 Rows per page: 20

**Step 2:** To view only the findings from Nessus, choose **Nessus** from the **Data Type** filter.

This will list the issues identified by Nessus such as SSL certificate issues, vulnerable software versions in use, etc...



Home > Issues > Findings

Search anything...

solutions

Findings Rule Engine

Host-Endpoint Findings Group by Insights Saved Filters

Search

<input type="checkbox"/>	Last seen	Assetname	Name	Risk
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.0.106	SSL Certificate Cannot ...	Medium
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.27	ICMP Timestamp Reque...	Low
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.27	SSL Self-Signed Certific...	Medium
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.0.93	ICMP Timestamp Reque...	Low
<input type="checkbox"/>	2024-11-08 18:15:29	10.2.1.66	SSL Self-Signed Certific...	Medium

Filter + Create Rule

Clear Save Apply

Data Type  
nessus

Select Fields to filter  
status

Status X

Risk Factor  
Unknown Informational Low  
Medium High Critical  
Not Available  
Ignored Exploit Available

To learn more about how AccuKnox's advanced Vulnerability Management can be leveraged, refer here

AccuKnox can also leverage other tools such as Masscan and Zeek to provide host scanning. By integrating these tools, AccuKnox will be able to scan the on-premise environment and help in identifying issues on the hosts.

**Masscan** helps red teamers doing offensive research (penetration testing) as well as blue teamers and IT managers doing defensive research (to find attack vectors within the network).

**Zeek** is a Network Security Monitor (NSM) to support detection and investigations of suspicious or malicious activity. Zeek also supports a wide range of traffic analysis tasks beyond the security domain, including performance measurement and troubleshooting.

## 15. Admission Controller Support Using Knoxguard

As Kubernetes adoption continues to surge, securing your clusters becomes critical. Knoxguard, the latest security feature, aims to bolster Kubernetes environment security and compliance through robust policy enforcement. Knoxguard operates independently of any policy engine, offering the flexibility to integrate with your preferred enforcement add on. Currently, Knoxguard supports **Kyverno** as the policy enforcement engine.

### 15.1 Introduction

#### Key Features of Knoxguard/Admission Controller

##### Registry Restrictions

Registry Restrictions allow you to define rules that either restrict or whitelist specific container registries or patterns at the **cluster** and **namespace** levels. This feature ensures that only trusted images are deployed within your Kubernetes clusters, reducing the risk of deploying vulnerable or malicious containers.

##### Vulnerability Scan Thresholds (Pipeline Feature)

Knoxguard enables you to set thresholds for the maximum number of **critical** or **high-level vulnerabilities** that an image can have. This feature will block the deployment of images with known vulnerabilities, maintaining a high security posture for your applications.

##### Security Posture Rules

Enforcing security policies like **privileged container restrictions** and **capabilities constraints** helps maintain a secure Kubernetes environment. Knoxguard currently supports **denying privileged mode containers**, with more security rules expected to be added soon.

## 15.2 Prerequisite for Knoxguard Admission Controller

Before deploying Knoxguard in your Kubernetes environment, ensure the following prerequisite is met:

- **Accuknox Agent Installation:** Install Accuknox Agents on your Kubernetes cluster. These agents facilitate SaaS integration, alerting, and enforcement.

### Info

Refer to Cluster On-boarding guide for Accuknox Agents Installation.

Verify the agents' status using the following command:

```
userx@fedora:~$ kubectl get pods -n accuknox-agents
```

NAME	READY	STATUS
agents-operator-d8585d594-55s29 72d	1/1	Running 0
discovery-engine-59c69ff787-scrnj 72d	4/4	Running 0
feeder-service-765d8f7d65-d4vq2 (2d21h ago) 4d	1/1	Running 13
policy-enforcement-agent-f5c5f87b-9fw79 ago) 40d	1/1	Running 84 (2d21h ago)
shared-informer-agent-77569db588-c944p ago) 40d	1/1	Running 1090 (2m36s ago)

## 15.3 Deployment of Knoxguard

### Deploy Kyverno:

First, you need to deploy Kyverno, a policy engine for Kubernetes, which Knoxguard utilizes for policy enforcement.

```
helm repo add kyverno https://kyverno.github.io/kyverno/  
helm repo update  
helm install kyverno kyverno/kyverno -n kyverno --create-namespace
```

### Step 2: Deploy Knoxguard:

Next, deploy Knoxguard in your Kubernetes cluster. Knoxguard will work in tandem with Kyverno to enforce the defined policies.

```
helm upgrade --install kxnoxguard oci://public.ecr.aws/k9v9d5v2/knoxguard-chart  
--version=v0 -n kxnoxguard --create-namespace
```

Verify the deployments:

```
userx@fedora:~$ kubectl get deployments -n kxnoxguard  
NAME                                READY   UP-TO-DATE   AVAILABLE   AGE  
accuknox-knoxguard-controller-manager 1/1     1             1           16s  
userx@fedora:~$ kubectl get pods -n kyverno  
NAME                                READY   STATUS      RESTARTS   AGE  
kyverno-admission-controller-78d5464dbc-p2248 1/1     Running     1           (49m ago) 52m  
kyverno-background-controller-5f96748b4c-mrcxm 1/1     Running     0           52m  
kyverno-cleanup-admission-reports-28796130-mzg8t 0/1     Completed   0           4m2s  
kyverno-cleanup-cluster-admission-reports-28796130-9nkb7 0/1     Completed   0           4m2s
```

kyverno-cleanup-cluster-ephemeral-reports-28796130-drsmn 4m2s	0/1	Completed	0
kyverno-cleanup-controller-7b5fb595d6-x57g7 52m	1/1	Running	0
kyverno-cleanup-ephemeral-reports-28796130-mxnkx 4m2s	0/1	Completed	0
kyverno-reports-controller-76cd67fb8d-v66wm (49m ago) 52m	1/1	Running	1

## 15.4 Policy Enforcement

Once Knoxguard is deployed, you can start enforcing policies within your cluster. This involves Creating, uploading and activating your custom admission policies.

To enforce the admission policy, follow these steps with example:

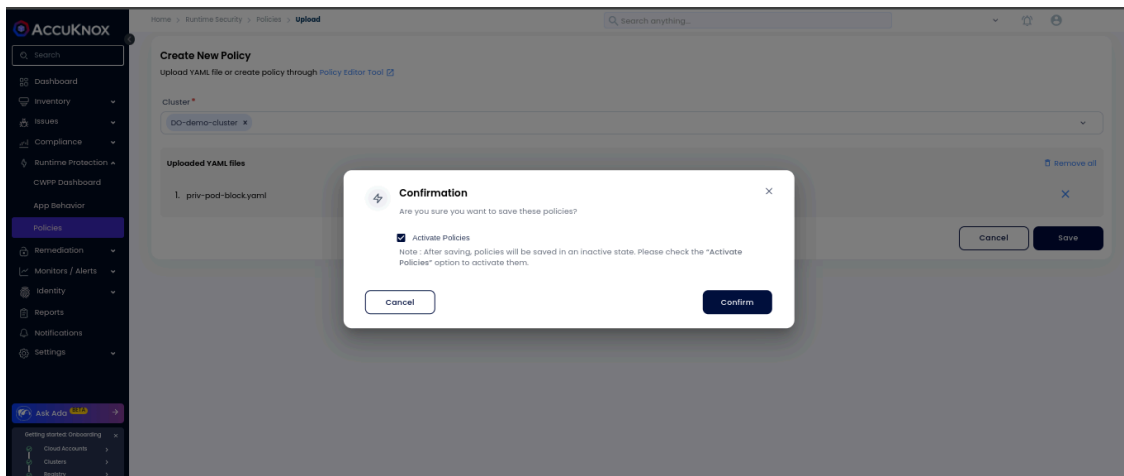
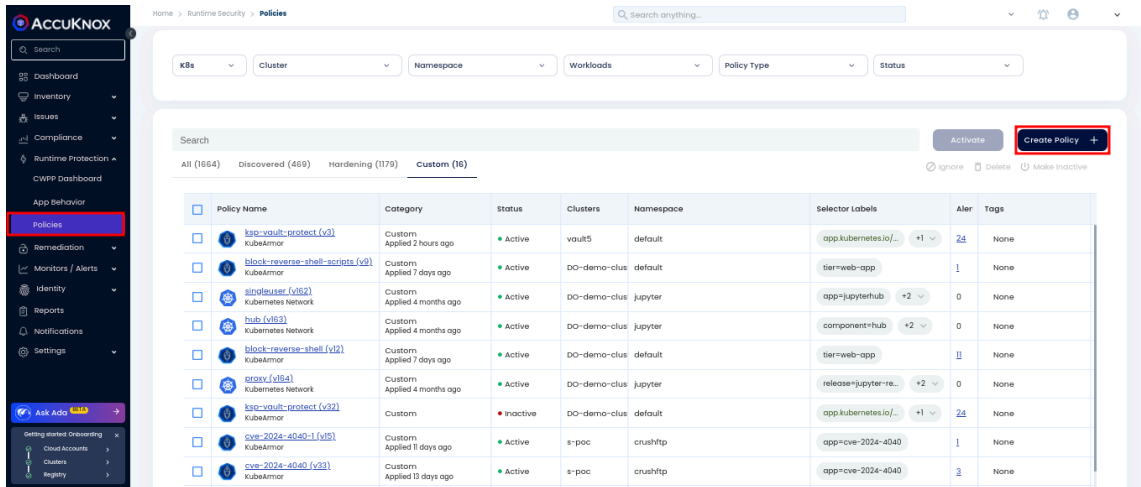
1. **Define the Admission Policy:** Create an **AdmissionPolicy** resource based on the requirement. Below is the configuration to block privileged pod admission in the default namespace:

```
apiVersion: admission.accuknox.com/v1
kind: AdmissionPolicy
metadata:
  labels:
    app.kubernetes.io/name: admission-controller
    app.kubernetes.io/managed-by: kustomize
  name: test-priv-pod-policy
spec:
  denyPrivilegedPod:
    action: Block
    targetNamespaces:
      - default
```



## 1. Upload and Activate Admission Policies:

Use the upload YAML feature to upload your custom admission policies by clicking on Create Policy. This allows you to define and enforce policies specific to your security requirements.



After uploading and activating the policy, you can verify its status with the following command:

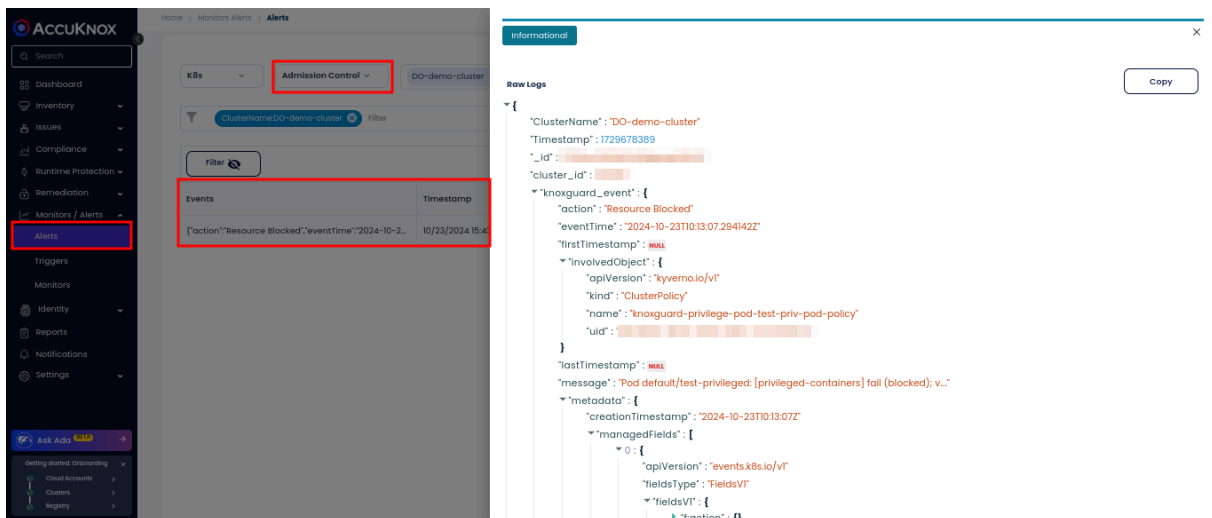
```
userx@fedora:~$ kubectl get admissionpolicy
NAME                                READY    MESSAGE
OWNED_POLICIES
test-priv-pod-policy               True     clusterpolicy has been updated successfully
["knoxguard-privilege-pod-test-priv-pod-policy"]
```



Navigate to **Monitors > Alerts** in the AccuKnox dashboard. Change the alert type to **Admission Controller** to view alerts related to admission policy violations. The system provides comprehensive logging to help you quickly identify and address any security concerns.

## Info

These logs can be forwarded to SIEM tools or notification tools by setting up triggers for improved security monitoring. Refer to the guide here for more details.



## 15.6 Pod Security Admission Controller

Pod Security Admission (PSA) enforces security standards on a Pod's Security Context and related fields based on three levels defined by the Pod Security Standards:

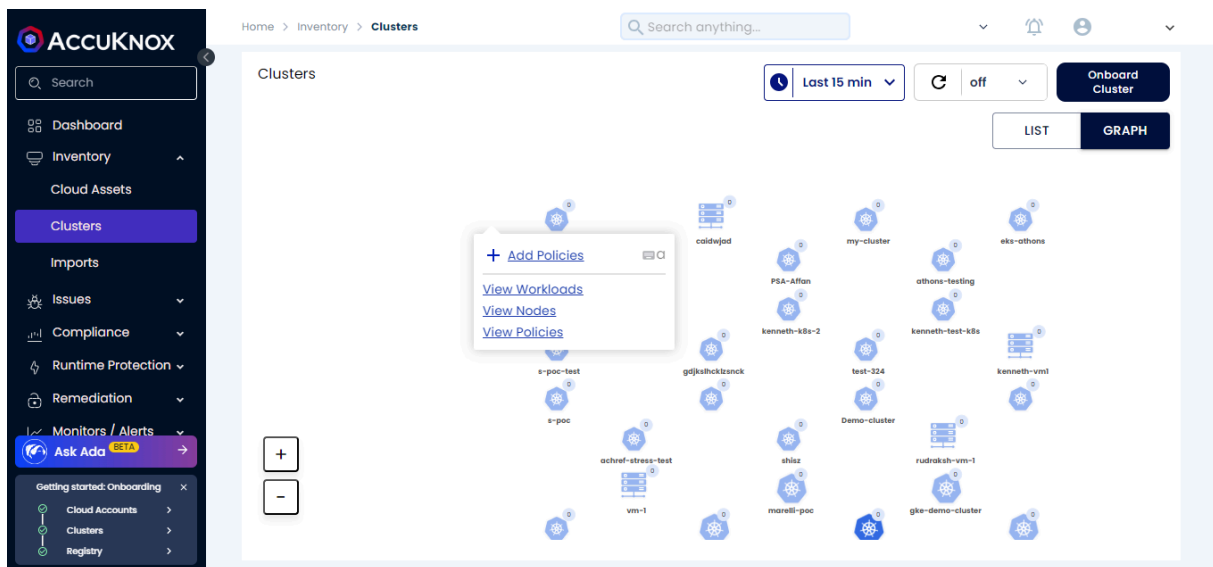
1. **Privileged:** Unrestricted policy, allows known privilege escalations.
2. **Baseline:** Minimally restrictive policy, allows default (minimally specified) Pod configuration.
3. **Restricted:** Heavily restricted policy, adhering to current Pod hardening best practices.

PSA can be enabled in two modes:

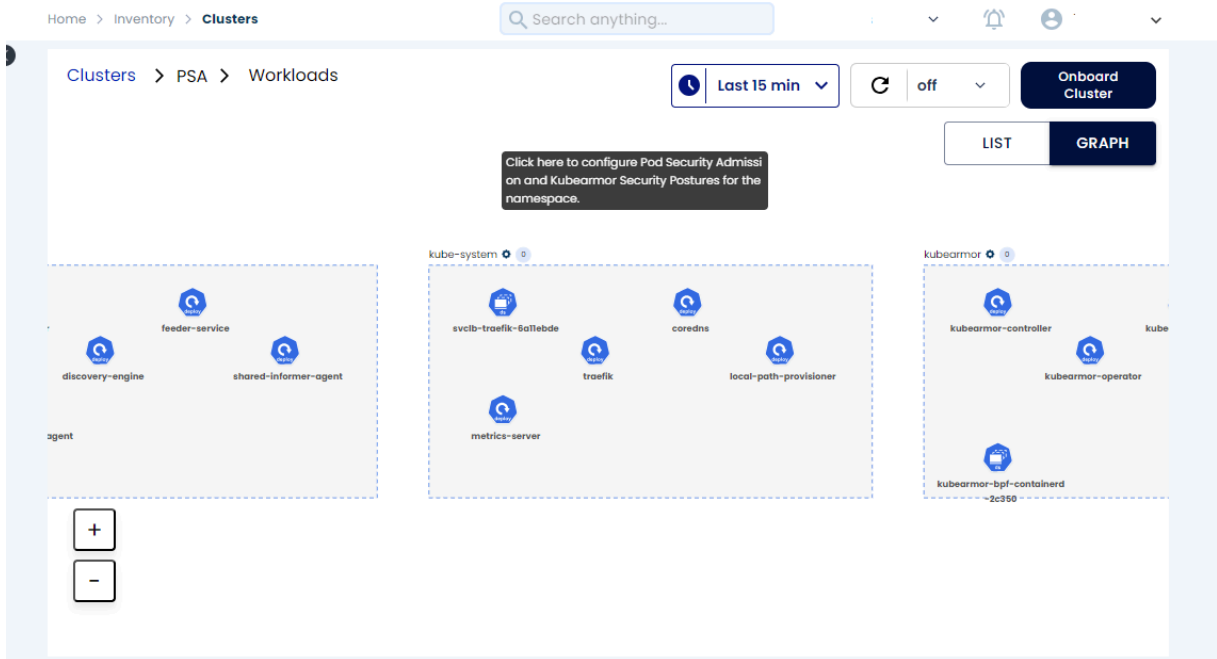
1. **Enforce:** Policy violations will cause the pod to be rejected.
2. **Audit:** Policy violations will trigger an alert but still allow the pod.

## 15.7 Enabling Pod Security Admission (PSA)

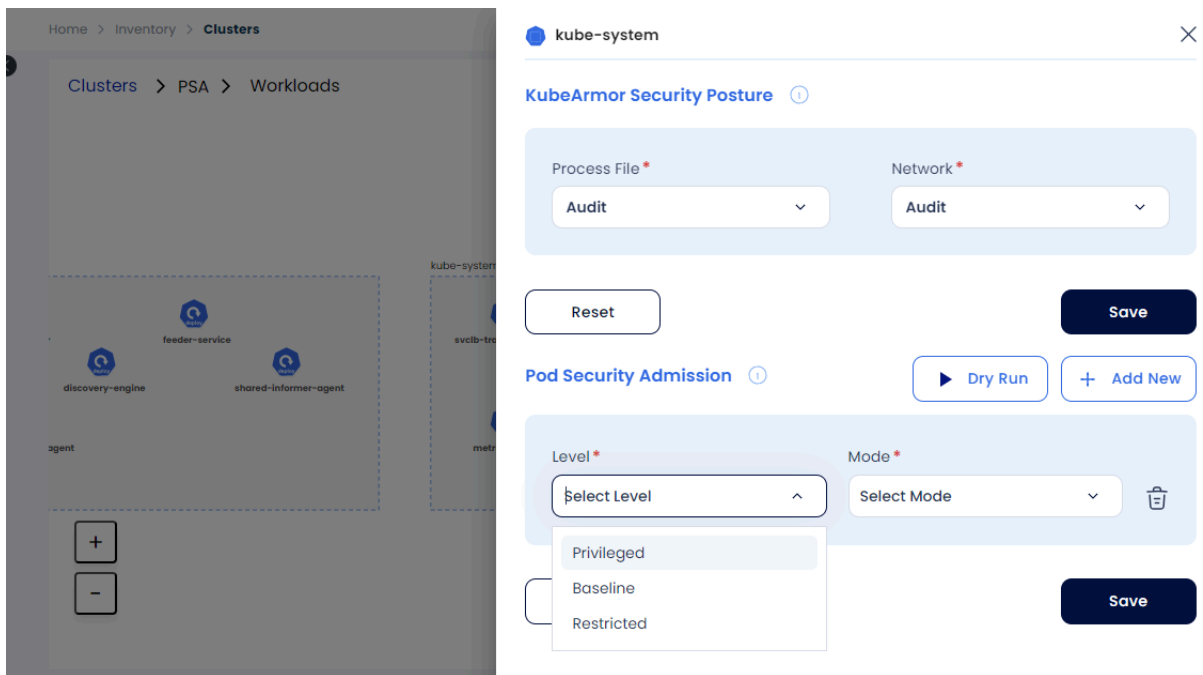
Navigate to **Inventory** → **Clusters** and click on the cluster, then select **View Workloads**.




1. Click on the **cog icon** next to the namespace.



2. Select the desired **PSA Level** and click **Save**.



3. Select the desired **Mode** for PSA.

 kube-system
✕

---

### KubeArmor Security Posture ⓘ

Process File \*

Audit

Network \*

Audit

Reset

Save

### Pod Security Admission ⓘ

▶ Dry Run

+ Add New


Level \*

Select Level

Mode \*

Select Mode

Audit  
Enforce



Reset

Save

4. If using **Enforce** mode, click on **Dry Run** to preview potential effects before applying.

## Dry Run ✕

As you can imagine, enforcing security standards on a namespace with existing workloads could be disruptive. You can use dry run on namespace to evaluate existing workloads against the policy and determine which workloads will need to be modified so they won't violate the policy.

Level

Restricted ▼

Mode

Enforce ▼

Dry Run

Warning: existing pods in namespace "default" violate the new PodSecurity enforce level "restricted:latest" Warning: kiem-job-28806330-rdxwt (and 2 other pods): allowPrivilegeEscalation != false, unrestricted capabilities, runAsNonRoot != true, seccompProfile Warning: vault-0: unrestricted capabilities, seccompProfile Warning: vault-agent-injector-8667c5945-fv4nd: seccompProfile

The **Dry Run** mode allows users to confirm potential effects of the PSA. Once reviewed and acceptable, click **Save** to apply the PSA.

## 15.8 PSA Protection Example

After setting PSA to enforce the **restricted** level, attempt to run a privileged Pod in the cluster:

```
root@demo:~# kubectl run nginx --image=nginx
```

An error will be returned as shown below:

```
Error from server (Forbidden): pods "nginx" is forbidden: violates PodSecurity "restricted:latest": allowPrivilegeEscalation != false (container "nginx" must set securityContext.allowPrivilegeEscalation=false), unrestricted capabilities (container "nginx" must set securityContext.capabilities.drop=["ALL"]), runAsNonRoot != true (pod or container "nginx" must set securityContext.runAsNonRoot=true), seccompProfile (pod or container "nginx" must set securityContext.seccompProfile.type to "RuntimeDefault" or "Localhost")
```

Since the **restricted** PSA label was applied to the namespace, attempting to create a pod with excessive privileges results in this error, successfully preventing the privileged pod from running.



## 16. CWPP Report Generation

### Understand the Regex to Select the Cluster Name and Namespace

The CWPP report generation utilizes regular expressions (regex) to specify and filter cluster names and namespaces. The syntax for regex follows a particular pattern to ensure accurate selection.

### 16.1 Regex

**Regex Syntax Format:** Cluster Name Selection / Namespace Selection

#### 16.1.1 Rules for Regular Expression

##### Excluding

- To exclude a specific cluster or namespace, prefix it with a hyphen (-).

##### NOTE

To exclude any cluster or namespace, it must be included in the selection first.

##### Select all

- Use an asterisk (\*) to select all clusters or namespaces.

##### Delimiter

- A forward slash (/) is used to delimit the cluster name selection from the namespace selection.

##### Examples

- cluster1/ns1: Include only namespace ns1 from cluster cluster1.

- cluster1/\*: Include all namespaces from cluster cluster1.
- cluster1/ns\*: Include namespaces starting with ns from cluster cluster1.
- -cluster1/ns3: Exclude namespace ns3 from cluster cluster1.
- \*/ns1: Include namespace ns1 from all clusters.
- \*/\*: Include all namespaces from all clusters.

## 16.2 Reports Configuration

Reports can be configured in two ways: On Demand and Scheduled.

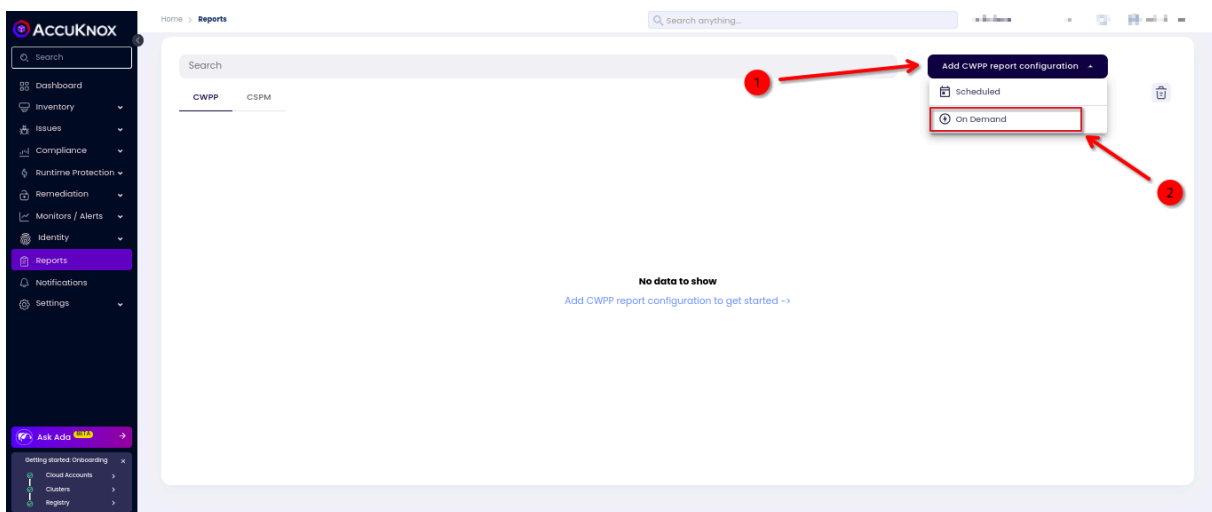
### 16.2.1 On Demand Report Configuration

In On Demand Report, you can generate the report for the clusters shortly after the configuration is completed.

To generate On Demand reports:

#### Step 1: Add CWPP Report Configuration

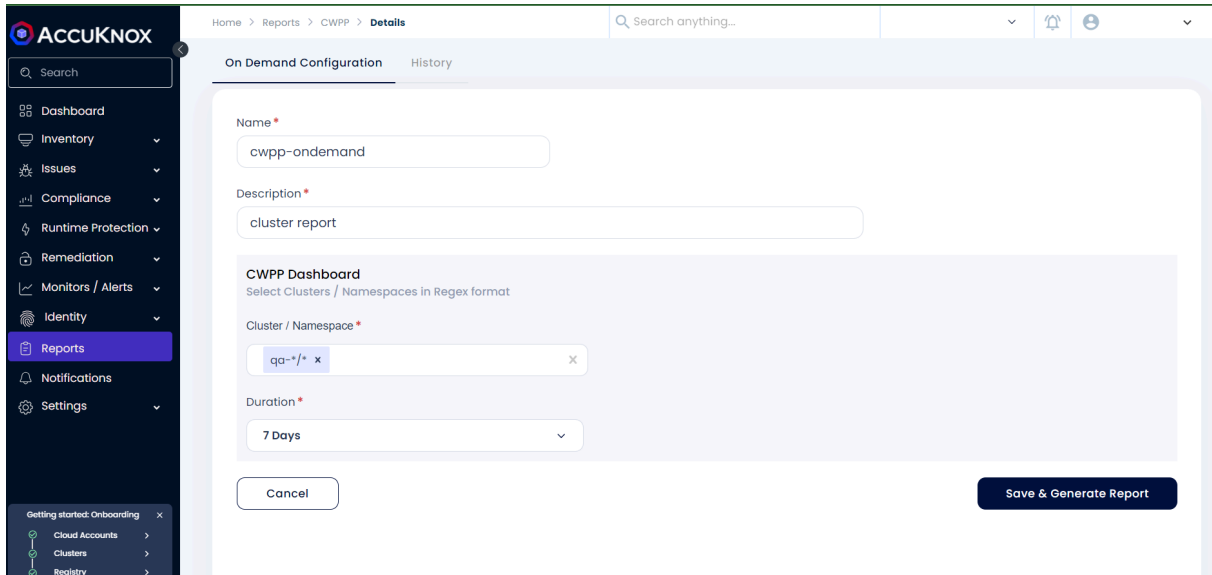
- Go to the Reports section in AccuKnox SaaS.
- Under CWPP Tab, click on “Add CWPP report configuration” and choose "On Demand" from the drop-down menu.



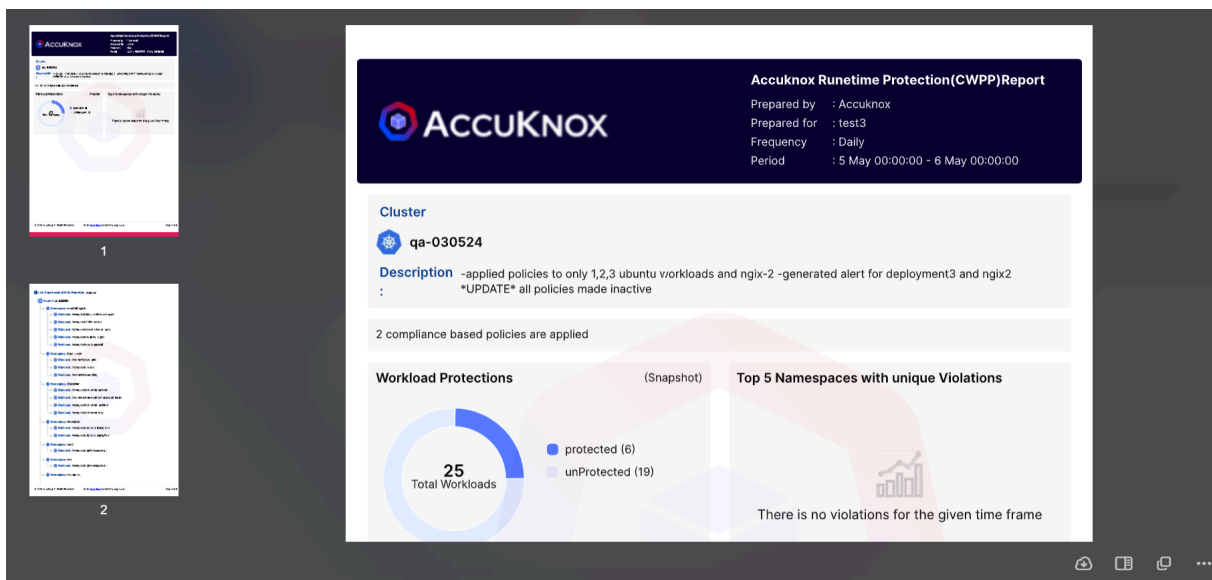
**Step 2:** In the Configuration user needs to provide the details about Name, Description and Cluster and NameSpace.

#### NOTE

The cluster field drop-down will show all the clusters that are active during the report generation.



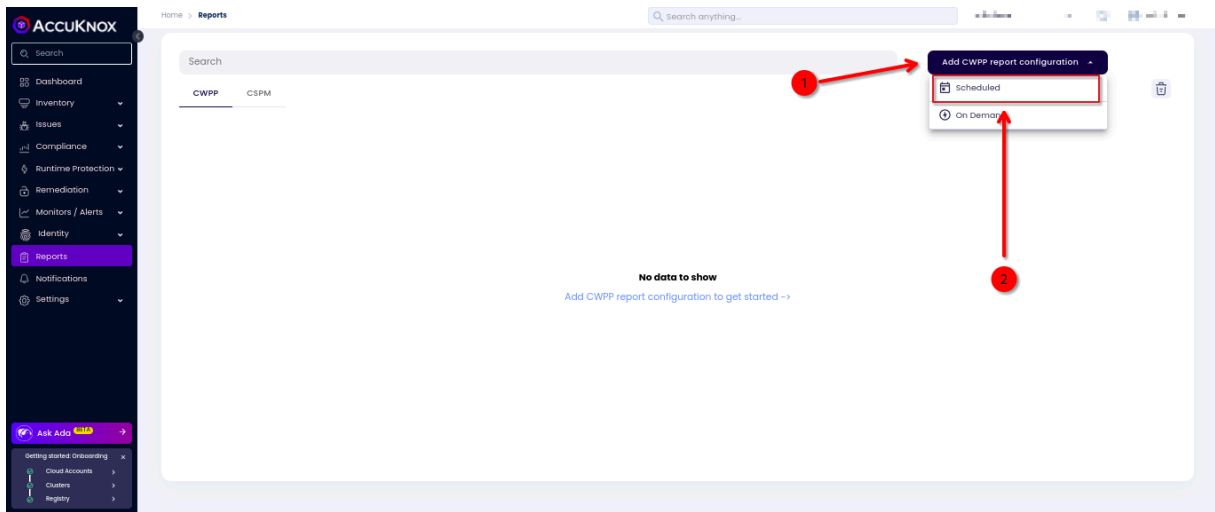
By clicking Save and Generate Report it will generate the report in the PDF format as per the selected duration.



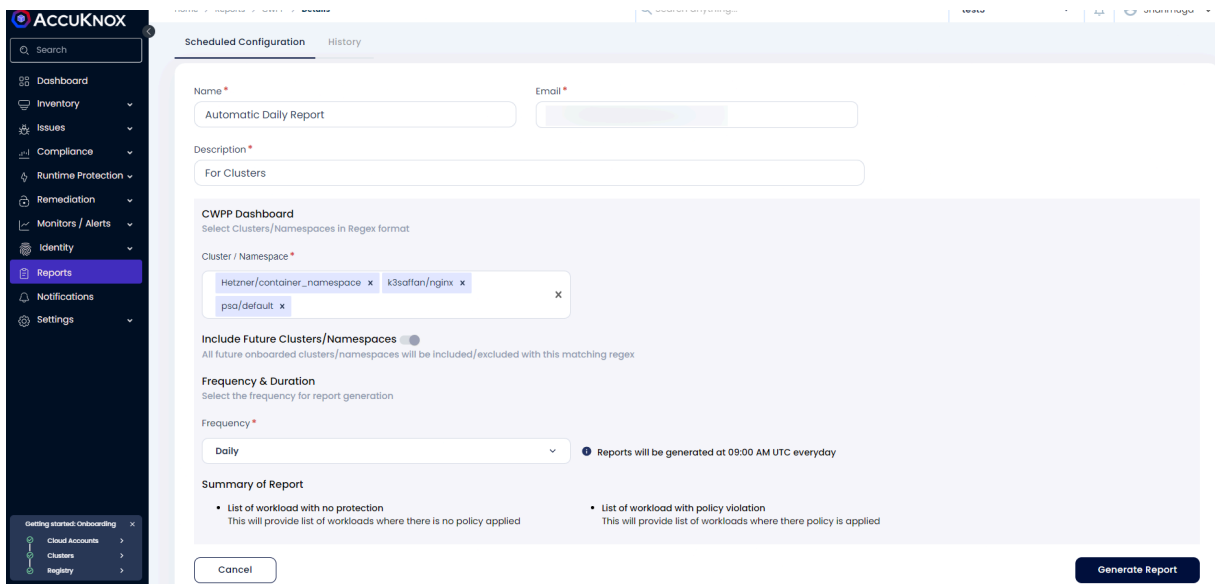
## 16.2.2 Scheduled Report Configuration

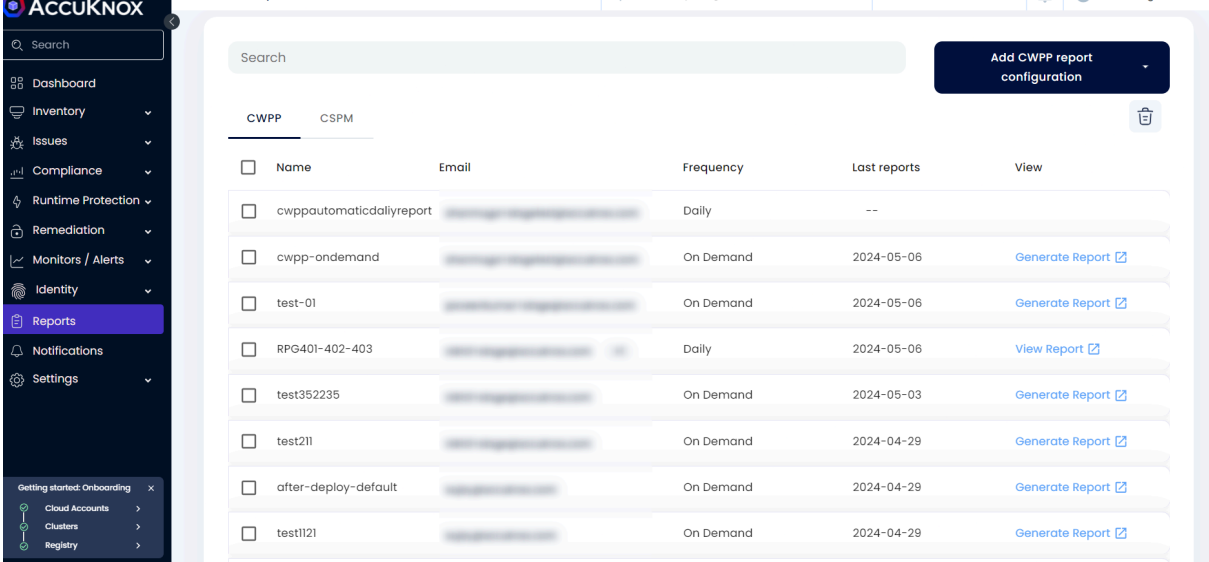
To get the report of the clusters automatically as per the frequency that is chosen .i.e by weekly or by monthly or daily this is the go to way:

**Step 1:** Add CWPP report configuration as Scheduled and choose the Scheduled option from the drop down.



**Step 2:** In the Configuration user needs to provide the details about their Name, Email, Selecting the Cluster, Namespace in the regex format and Frequency of the report then click the Generate Report.

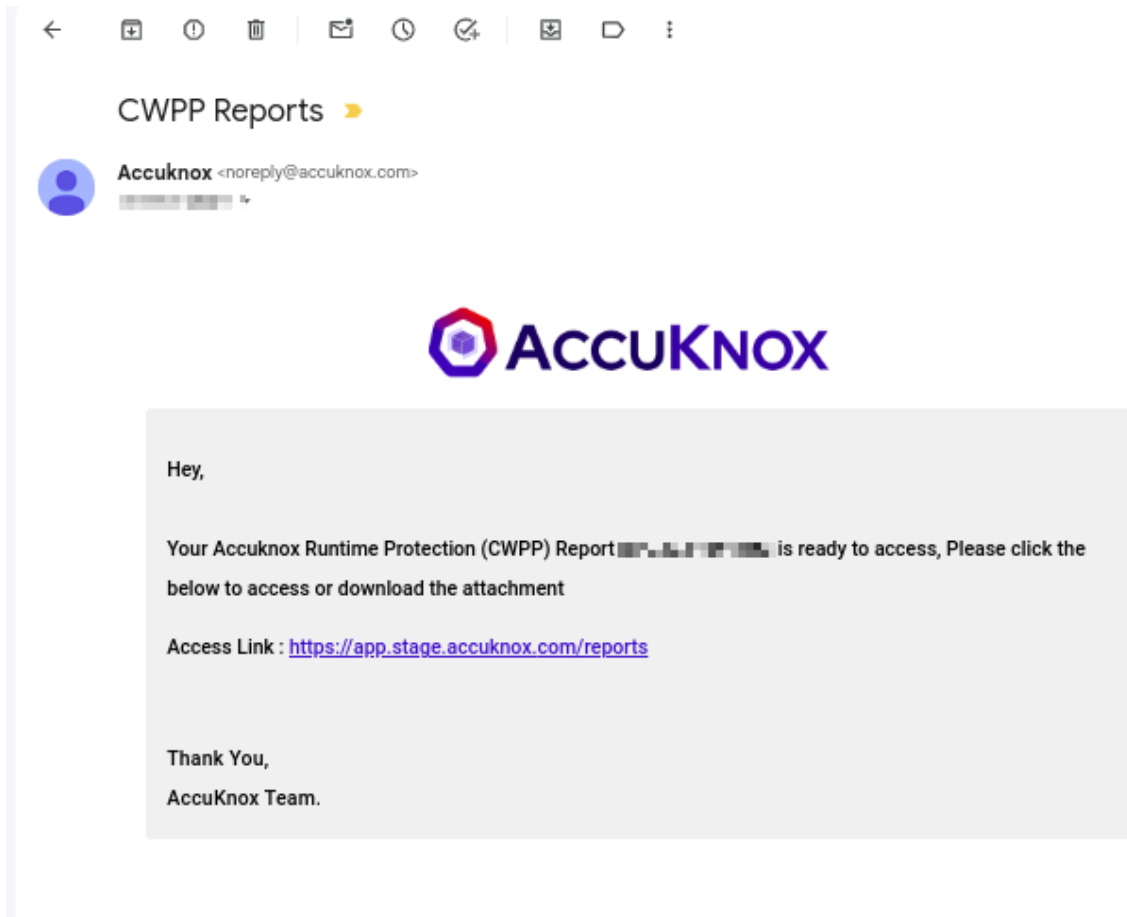




The screenshot displays the AccuKnox interface for CWPP configurations. The main content area features a table with the following data:

<input type="checkbox"/>	Name	Email	Frequency	Last reports	View
<input type="checkbox"/>	cwppautomaticdailyreport	[Redacted]	Daily	--	
<input type="checkbox"/>	cwpp-ondemand	[Redacted]	On Demand	2024-05-06	<a href="#">Generate Report</a>
<input type="checkbox"/>	test-01	[Redacted]	On Demand	2024-05-06	<a href="#">Generate Report</a>
<input type="checkbox"/>	RPG401-402-403	[Redacted]	Daily	2024-05-06	<a href="#">View Report</a>
<input type="checkbox"/>	test352235	[Redacted]	On Demand	2024-05-03	<a href="#">Generate Report</a>
<input type="checkbox"/>	test211	[Redacted]	On Demand	2024-04-29	<a href="#">Generate Report</a>
<input type="checkbox"/>	after-deploy-default	[Redacted]	On Demand	2024-04-29	<a href="#">Generate Report</a>
<input type="checkbox"/>	test1121	[Redacted]	On Demand	2024-04-29	<a href="#">Generate Report</a>

**Step 3:** After finishing the configuration the report would be scheduled to be sent to you in the email. Users can reconfigure the past configurations by clicking on them to edit the configuration.



## 17. Integrations

### 17.1 Integrate SIEM tools

- SPLUNK
- AWS Cloud Watch
- Rsyslog

#### 17.1.1 Splunk

Splunk is a software platform to search, analyze, and visualize machine-generated data gathered from websites, applications, sensors, and devices.

AccuKnox integrates with Splunk and monitors your assets and sends alerts for resource misconfigurations, compliance violations, network security risks, and

anomalous user activities to Splunk. To forward the events from your workspace you must have Splunk Deployed and HEC URL generated first for Splunk Integration.

### a. Prerequisites:

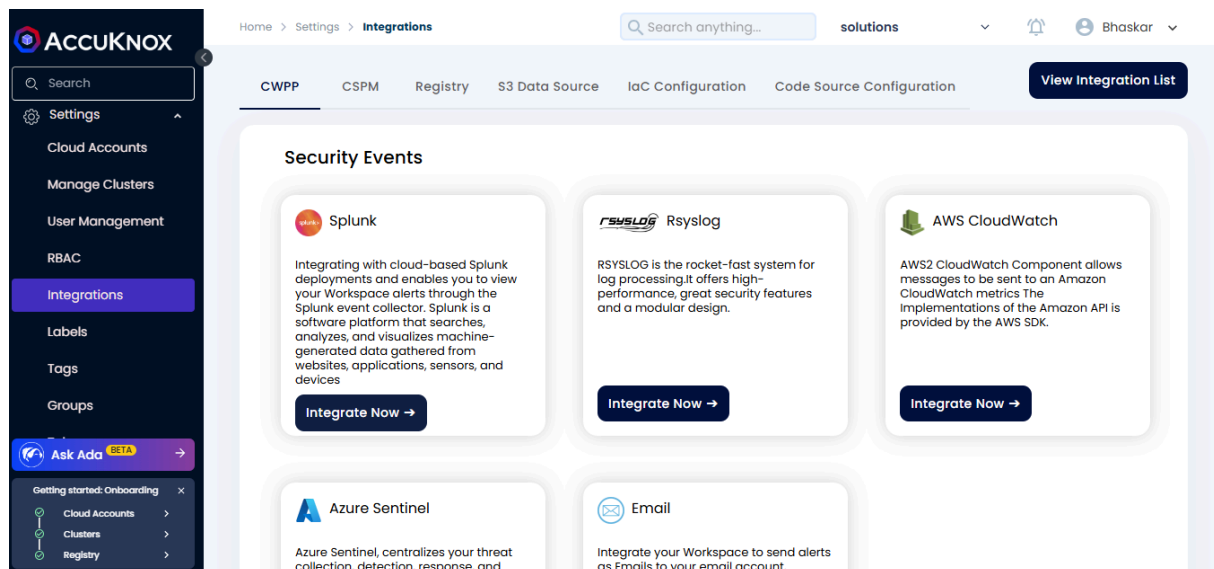
Set up Splunk HTTP Event Collector (HEC) to view alert notifications from AccuKnox in Splunk. Splunk HEC lets you send data and application events to a Splunk deployment over the HTTP and Secure HTTP (HTTPS) protocols.

To set up HEC, use instructions in [Splunk documentation](#). For source type, `_json` is the default; if you specify a custom string on AccuKnox, that value will overwrite anything you set here.

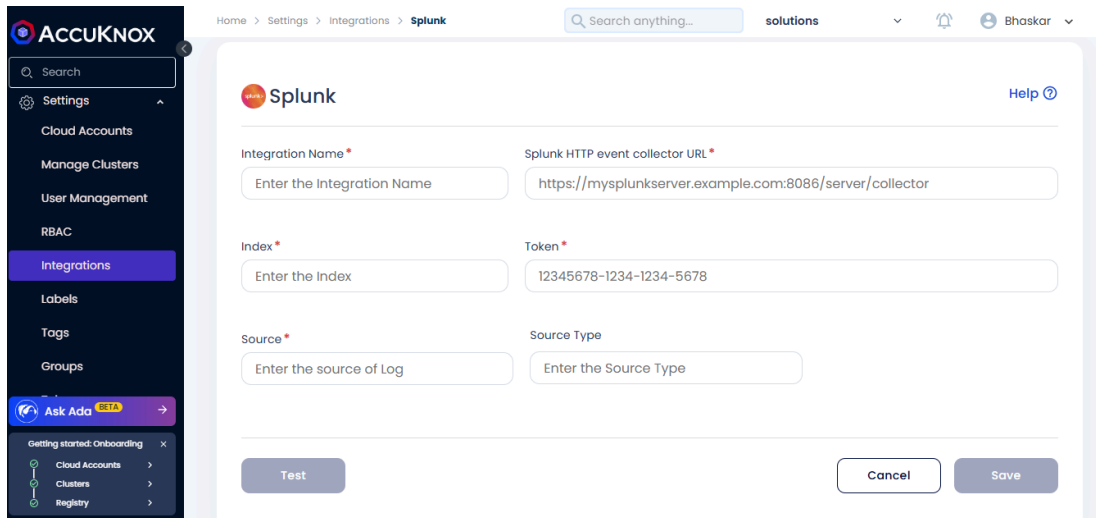
Select Settings > Data inputs > HTTP Event Collector and make sure you see HEC added in the list and that the status shows that it is Enabled.

### b. Steps to Integrate:

- Go to Settings->Integration.
- Click Integrate now on Splunk.







- Enter the following details to configure Splunk.
- Select the Splunk App: From the dropdown, Select Splunk Enterprise.
  - Integration Name: Enter the name for the integration. You can set any name. e.g., sh Test Splunk
  - Splunk HTTP event collector URL: Enter your Splunk HEC URL generated earlier.e.g., sh   
https://splunk-xxxxxxxxx.com/services/collector
  - Index: Enter your Splunk Index, once created while creating HEC. e.g., sh main
  - Token: Enter your Splunk Token, generated while creating HEC URL. e.g., sh x000x0x0x-0xxx-0xxx-xxxx-xxxxx00000
  - Source: Enter the source as http: sh Kafka
  - Source Type: Enter your Source Type here, this can be anything and the same will be attached to the event type forwarded to Splunk. e.g., sh \_json
  - Click Test to check the new functionality, You will receive the test message on the configured slack channel. e.g.,sh Test Message host = xxxxxx-deployment-xxxxxx-xxx00 source = http:kafka sourcetype = trials

- Click Save to save the Integration. You can now configure Alert Triggers for Slack Notifications.

## 17.1.2 AWS Cloudwatch

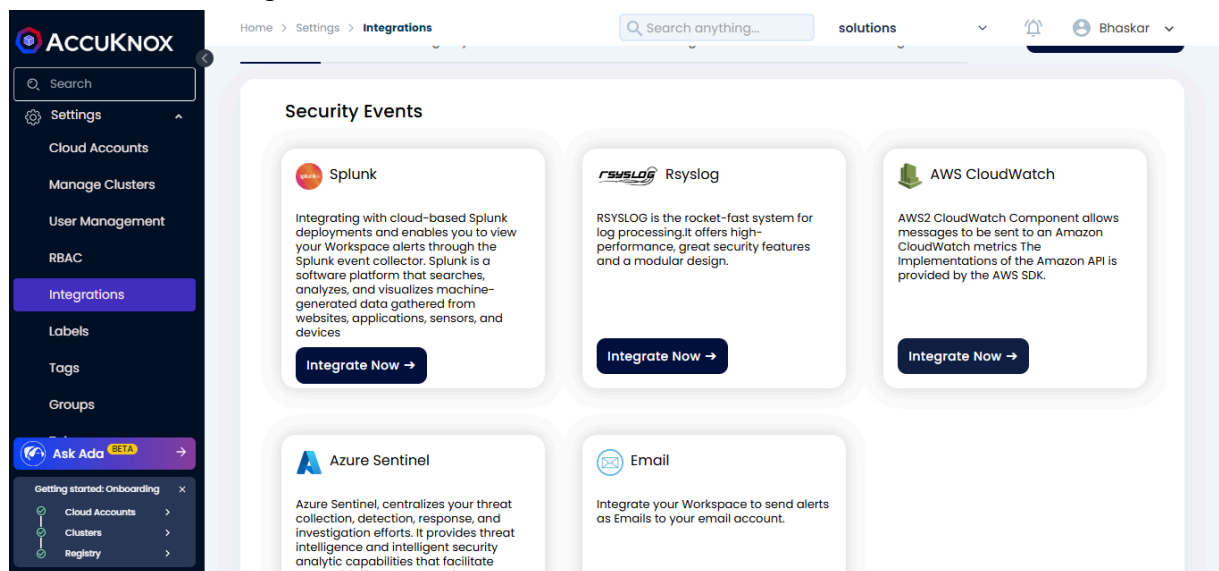
Navigate to Settings->Integrations. Choose "AWS CloudWatch" services and click the Integrate Now button.

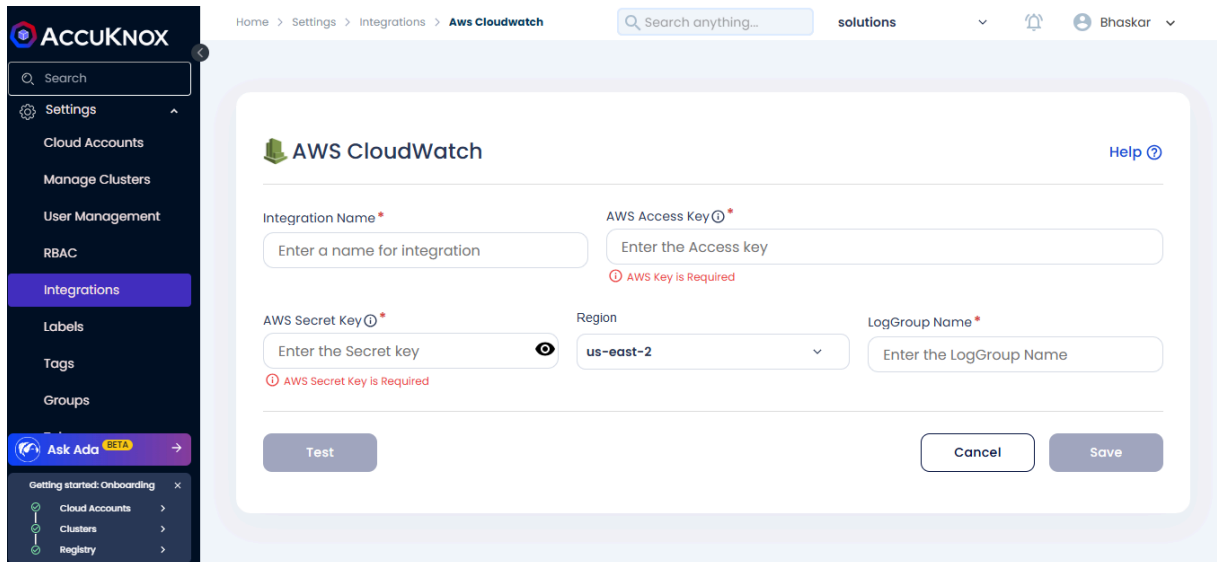
### a. Prerequisites

- AWS Access Key / AWS Secret Key is required for this Integration.
- [Note]: Please refer to this link to create an access keys [link](#)

### b. Steps to Integrate:

- Go to Settings -> Integration
- Click the Integrate Now button under AWS CloudWatch.





- Here you'll be able to see these entries:
  - Integration Name: Enter the name for the integration. You can set any name.
  - AWS Access Key: Enter your AWS Access Key here.
  - AWS Secret Key: Enter your AWS Secret Key here.
  - Region Name: Enter your AWS Region Name here.
- Once you fill in every field and then click the button this will test whether your integration is working or not.
- Click the Save button.

#### c. Configuration of Alert Triggers:

- On the Logs page, after choosing a specific log filter click on the 'Create Trigger' button.
- The below fields need to be entered with appropriate data:
- Name: Enter the name of the trigger. You can set any name without special characters.
- When to Initiate: The frequency of the trigger as Real Time /.
- Status: Enter the severity of the trigger.
- Search Filter Data: The filter log chosen is automatically populated here. This is optional.
- Predefined queries: The list of predefined queries for this workspace is shown as default.
- Notification Channel: Select the integration channel that needs to receive logs. This should be AWS CloudWatch. (Note: Channel Integration is done on the previous step)

- Save: Click on Save for the trigger to get stored in the database.

#### d. Logs Forwarding:

- For each Enabled Trigger, please check the AWS platform to view the logs.
- Based on Frequency (Real Time / Once in a Day / Week)
- The Rule Engine matches the real-time logs against the triggers created.

### 17.1.3 Azure Sentinel Integration

To forward the events to Azure Sentinel you must first set up the Azure Sentinel Integration.

#### a. Prerequisites:

- Azure Logic App - Webhook.
- Azure Sentinel Subscription.

#### b. Steps to Integrate:

- Go to Settings → Integrations → CWPP(Tab).
- Click integrate now on Azure Sentinel.
- Fill up the following fields:
- **Integration Name:** Enter the name for the integration. You can set any name of your choice. e.g., Container Security Alerts
- **Webhook URL:** Enter your Azure Logic App's Webhook URL here. e.g., <https://xyz.xxxxxx.logic.azure.com:443/workflows/xxxxxxx>
- **Group Name:** You can specify any group name based on your preference, this can be used to filter the events. This works as a key value pair, where key is Group Name and Group Value is the value for the Key Group Name. e.g., K8s Cluster
- **Group Value:** You can add any value to this group value. e.g., Dev Team Cluster
- Click **Test** to check the new functionality, You will receive the test message on configured Azure Sentinel. -Test message Please ignore !!
- Click **Save** to save the Integration. You can now configure Alert Triggers for Azure Sentinel Events

## 17.1.4 Creating webhook using the Azure Logic App

### a. About the logic app:

Azure Logic Apps is a cloud platform where you can create and run automated workflows with little to no code. Using the visual designer and selecting from prebuilt operations, you can quickly build a workflow that integrates and manages your apps, data, services, and systems. To create a webhook using the logic app.

- **Step 1:** Search for the logic app in the Azure portal.
- **Step 2:** Add the new logic app and fill in the relevant details.
- **Step 3:** After creating the logic it will appear in the logic app dashboard.
- **Step 4:** Open the app and click on the go-to resource button.
- **Step 5:** Select the http request to receive the logs.
- **Step 6:** Click on the new step and click HTTP after that click on the Azure log analytics to receive the alert data.
- **Step 7:** Add the connection name, workspaceID, and workspace key.

You can get the workspace id and key in the log analytics workspace tab.

- **Step 8:** Click on the Integration and click on the Agents tab.
- **Step 9:** Click on the Azure log analytics data collector and click the JSON request body as the body and log name.

After the setup is done you will receive a webhook URL.

### b. To see Logs in the Sentinel:

- **Step 1:** Open Microsoft Sentinel in the portal.
- **Step 2:** Click on the integrations.
- **Step 3:** Click on the logs tab and go to custom logs and select the time range and click on run the query to get the logs.

## 17.1.5 Rsyslog

To forward the events to Rsyslog you must first set up the Rsyslog Integration.

### a. Prerequisites:

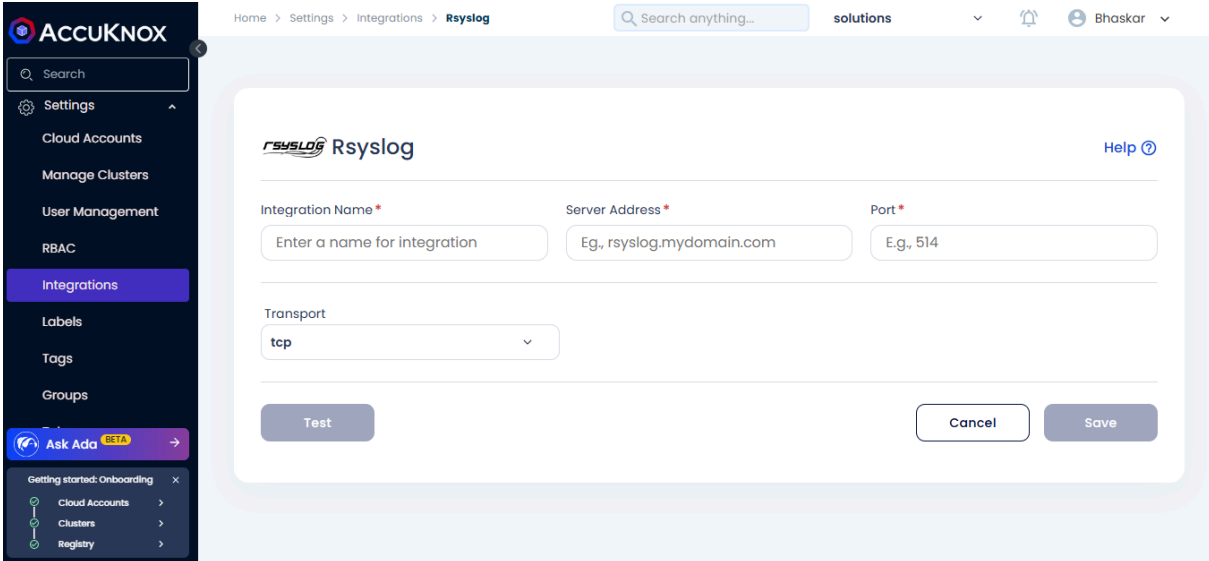
- A running Rsyslog server.

- Host name/IP, Port number, Transport type(TCP or UDP)

Note: To deploy the RSyslog server, follow [RSyslog Documentation](#).

b. Steps to Integrate:

- Go to Settings → Integrations.
- Click integrate now on RSyslog.



The screenshot shows the AccuKnox web interface. On the left is a dark sidebar with a search bar and a menu including Settings, Cloud Accounts, Manage Clusters, User Management, RBAC, Integrations (highlighted), Labels, Tags, and Groups. Below the menu is a 'Get started: Onboarding' section with links for Cloud Accounts, Clusters, and Registry. The main content area is titled 'RSyslog Rsyslog' and contains a form with the following fields: 'Integration Name\*' (placeholder: 'Enter a name for integration'), 'Server Address\*' (placeholder: 'Eg., rsyslog.mydomain.com'), 'Port\*' (placeholder: 'E.g., 514'), and 'Transport' (dropdown menu with 'tcp' selected). At the bottom of the form are three buttons: 'Test', 'Cancel', and 'Save'. The top navigation bar includes 'Home > Settings > Integrations > Rsyslog', a search bar, and user information 'Bhaskar'.

- Fill up the following fields:
  - Integration Name: Enter the name for the integration. You can set any name of your choice. e.g., Container Security Alerts
  - Server Address: Enter your RSyslog Server address here, IP address or fully qualified domain name (FQDN) of the RSyslog server e.g.,rsyslog.mydomain.com or 35.xx.xx.xx
  - Port: The port number to use when sending RSyslog messages (default is UDP on port 514); you must use the same port number. e.g., 514
  - Transport: Select UDP, or TCP as the method of communication with the RSyslog server
- Click Test to check the new functionality, You will receive the test message on configured RSyslog Server. -Test message Please ignore !!
- Click Save to save the Integration. You can now configure Alert Triggers for RSyslog Events

## 17.2 Integrate Notifications Tools

### 17.2.1 Slack

To send an alert notification via Slack you must first set up the Slack notification Channel.

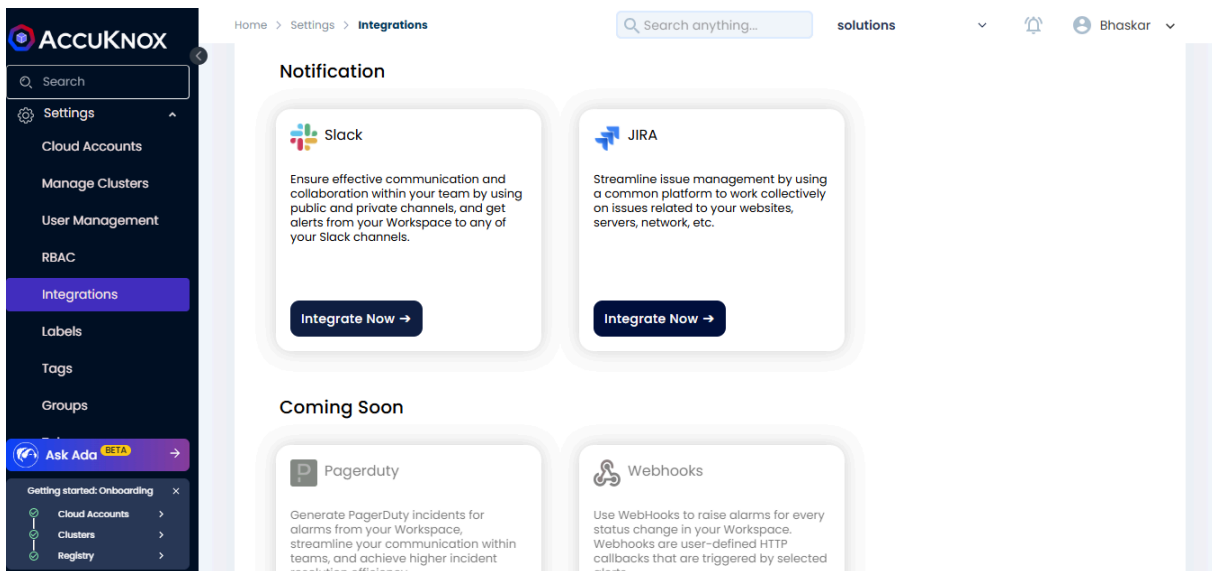
a. Prerequisites:

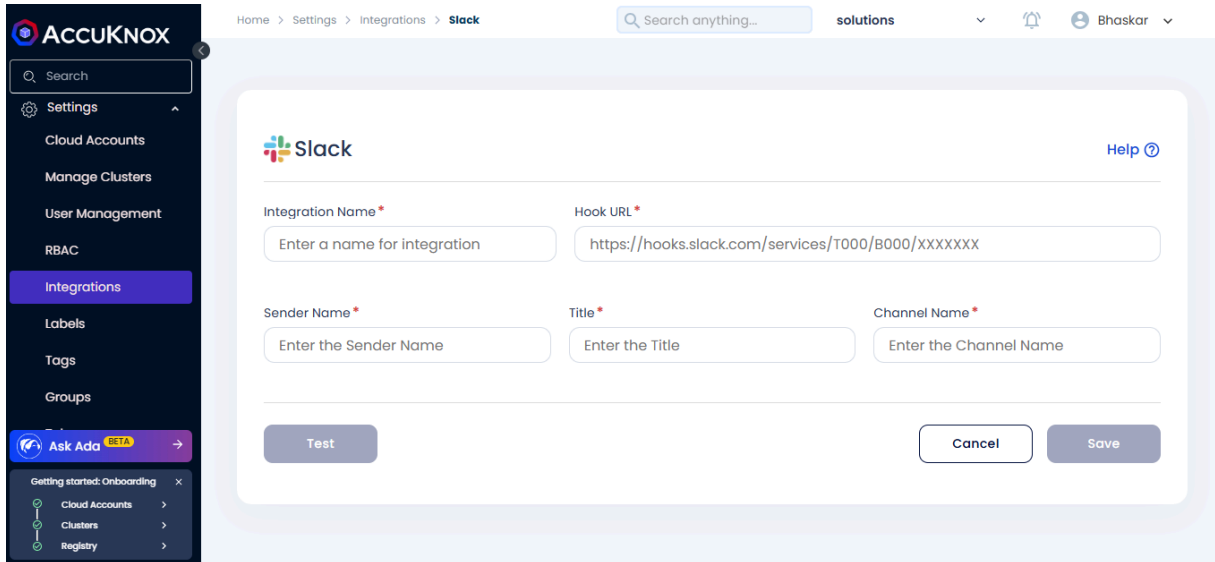
You need a valid and active account in Slack. After logging into your Slack channel, you must generate a Hook URL.

Note: To generate a Hook URL follow the steps, [Webhooks-for-Slack](#).

b. Steps to Integrate:

- Go to Settings -> Integration.
- Click “Integrate Now” under Slack.





- Fill up the following fields:
- Integration Name: Enter the name for the integration. You can set any name. e.g., Container Security Alerts
- Hook URL: Enter your generated slack hook URL here. e.g., https://hooks.slack.com/services/T000/B000/XXXXXXX
- Sender Name: Enter the sender name here. e.g., AccuKnox User
- Channel Name: Enter your slack channel name here. e.g., livealertsforcontainer
- Click Test to check the new functionality, You will receive the test message on configured slack channel. Test message Please ignore !!
- Click Save to save the Integration. You can now configure Alert Triggers for Slack Notifications.

## 17.3 Integrate Ticketing Tools

### 17.3.1 Jira Integration

Integrate AccuKnox with Jira and receive AccuKnox alert notifications in your Jira accounts. With this integration, you can automate the process of generating Jira tickets with your existing security workflow.



To set up this integration, you need to coordinate with your Jira administrator and gather the inputs needed to enable communication between AccuKnox and Jira.

## Integration of JIRA:

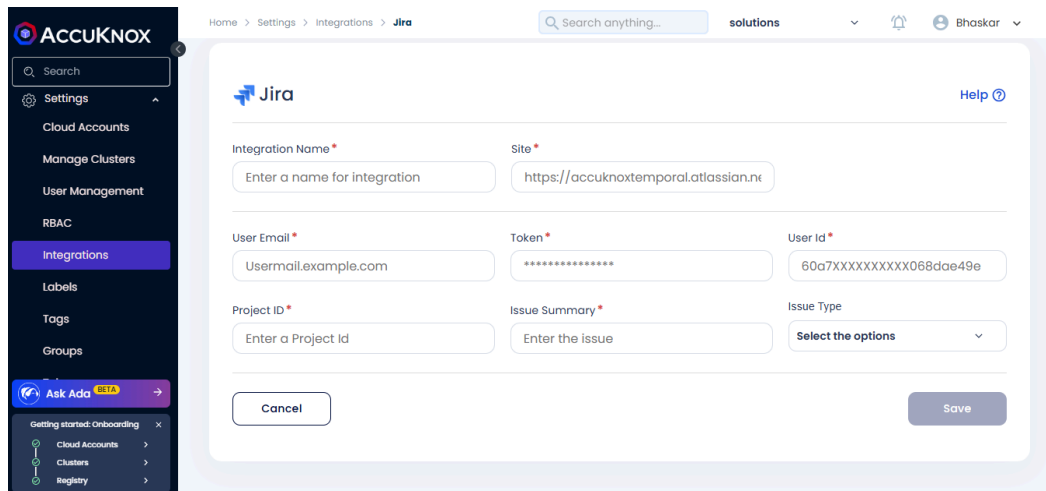
### a. Prerequisites

- You need a Jira Site URL, Email, UserID & API token, and Project key for this integration.
- To create a JIRA token go to <https://id.atlassian.com/manage-profile/security/api-tokens>, and click on create an API token.

### b. JIRA integration for CWPP:

Steps to Integrate:

- Go to Settings -> Integration.
- Click “Integrate Now” under JIRA



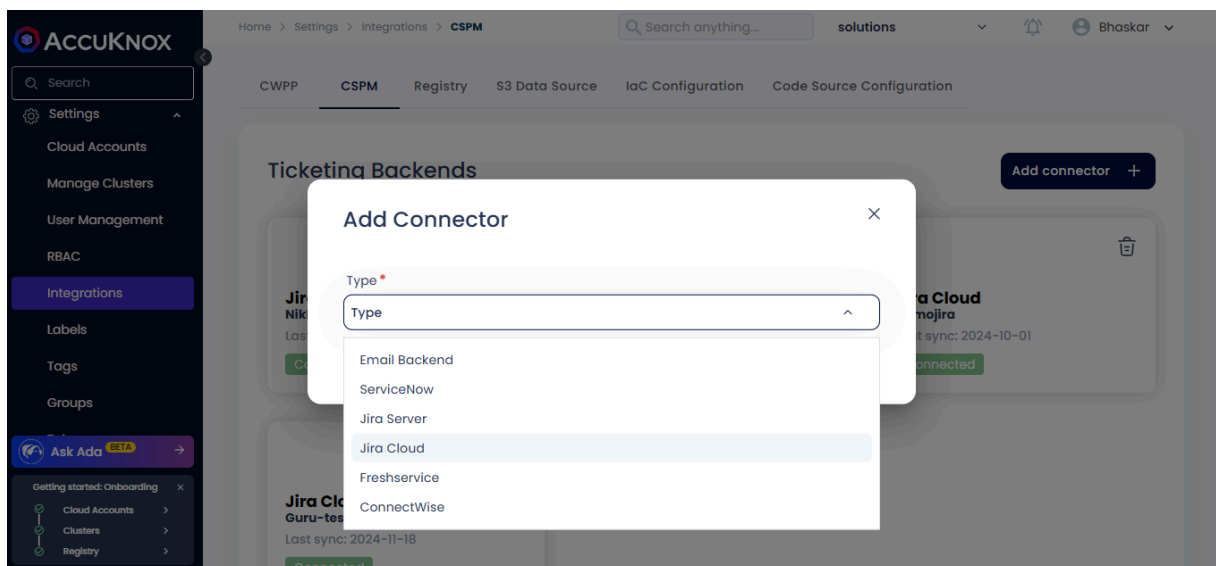
- Enter the following details to configure JIRA.
- Integration Name: Enter the name for the integration. You can set any name. e.g., Test JIRA
- Site: Enter the site name of your organization. e.g., <https://jiratest.atlassian.net/>
- User Email: Enter your Jira account email address here.e.g., [jira@organisation.com](mailto:jira@organisation.com)

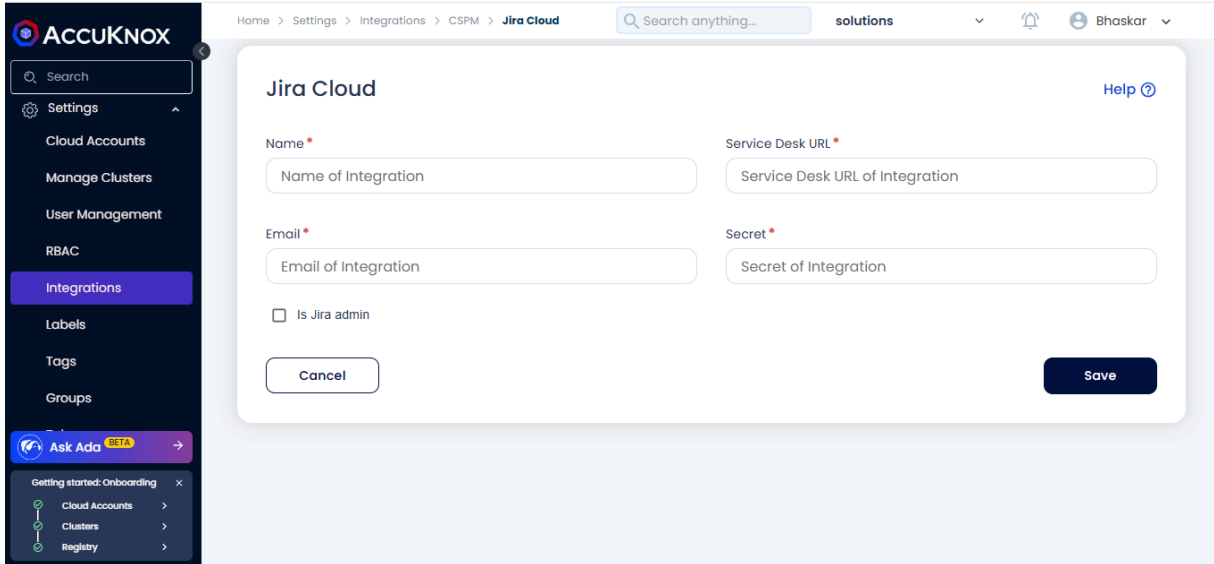
- Token: Enter the generated Token here from <https://id.atlassian.com/manage-profile/security/api-tokens>. .e.g., kRVxxxxxxxxxxxxx39
- User ID: Enter your Jira user ID here. You can visit the people section and search your name to see the User ID. For more details check here. e.g., 5bbxxxxxxxxxx0103780
- Project ID: Enter your Project key here, each project in an organization starts with some key value and is case-sensitive. Breakdown of a Jira ticket to identify Project ID: [https://\[JIRA-SITE\]/browse/\[PROJECT ID\]-1414](https://[JIRA-SITE]/browse/[PROJECT ID]-1414), e.g., DEVSECOPS
- Issue Summary: Enter the summary for the JIRA tickets to be viewed in each JIRA ticket created. e.g., Issues generated from High Severity Incidents on the onboarded cluster.
- Issue Type: You can choose from the dropdown. i.e., Story and Bug
- Click Test to check if the entered details are being validated, If you receive Test Successful, you have entered valid JIRA credentials.
- Click Save to save the Integration.

### 17.3.2 JIRA integration for CSPM:

#### Steps to Integrate:

- Go to Channel Integration -> CSPM.
- Click on “Add connector” and select JIRA Cloud





Home > Settings > Integrations > CSPM > Jira Cloud

Search anything...

solutions

Bhaskar

### Jira Cloud

Help ?

Name \*

Service Desk URL \*

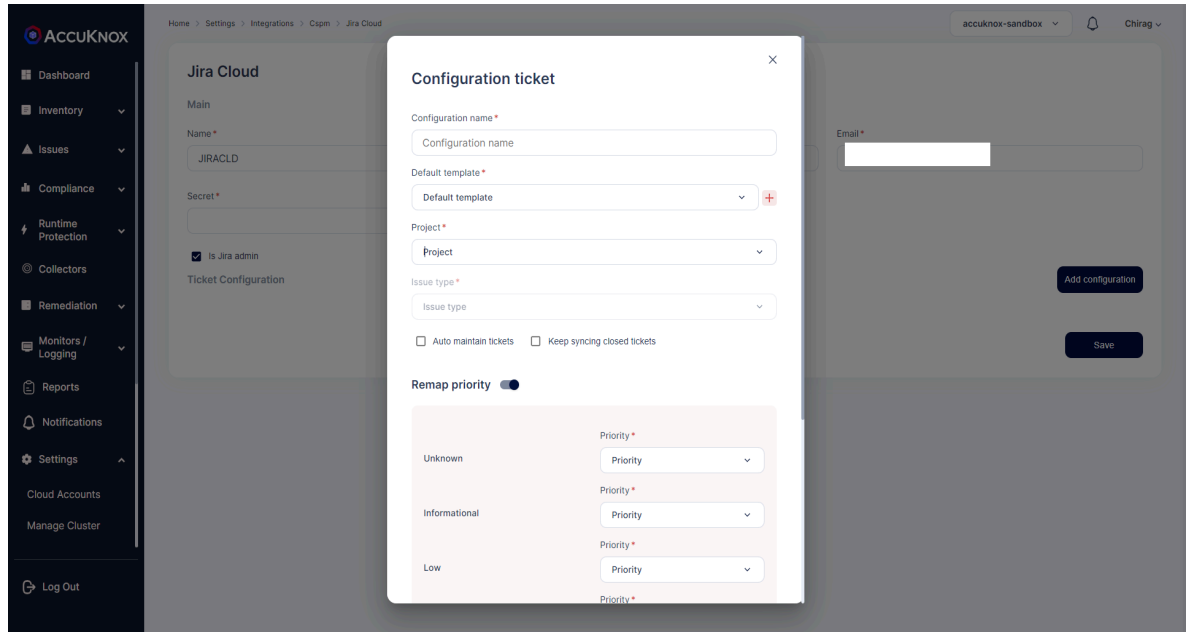
Email \*

Secret \*

Is Jira admin

Enter the following details to configure JIRA.

- Integration Name: Enter the name for the integration. You can set any name. e.g., Test JIRA
- Site: Enter the site name of your organization. e.g., <https://jiratest.atlassian.net/>
- User Email: Enter your Jira account email address here.e.g., [jira@organisation.com](mailto:jira@organisation.com)
- Token: Enter the generated Token here from <https://id.atlassian.com/manage-profile/security/api-tokens>. .e.g., kRVxxxxxxxxxxxx39



Click on the Jira ticketing backend to add config. Here Enter the following details:

- Configuration name: this name will be displayed under ticket configuration while creating tickets.
- Default template: to specify the data that this configuration will be used for making tickets.
- Project name: From the list of projects select the project where you want your tickets to be created.
- Issue Type: You can choose from the dropdown.
- Fill in the priority mapping according to your choice and press save.

You can now configure Alert Triggers for JIRA.

### 17.3.3 ServiceNow Integration

Integrate AccuKnox with ServiceNow and receive AccuKnox alert notifications in your ServiceNow account. With this integration, you can automate the process of generating ServiceNow tickets with your existing security workflow.

To set up this integration, you need to coordinate with your ServiceNow administrator and gather the inputs needed to enable communication between AccuKnox and ServiceNow

#### a. Prerequisites

- The ServiceNow Integration requires the following: Instance URL, Instance Username and Instance Password.
  - Please refer to the ServiceNow Documentation for how to create an instance and obtain the required credentials.

#### b. Steps for integration

- Navigate to Settings → Integrations → CSPM tab
- Click on Add Connector and select ServiceNow, click on Next.

### 17.3.4 Freshservice Integration

Integrate AccuKnox with Freshservice and receive AccuKnox alert notifications in your Freshservice accounts. With this integration, you can automate the process of generating Freshservice “Problem alerts” with your existing security workflow.

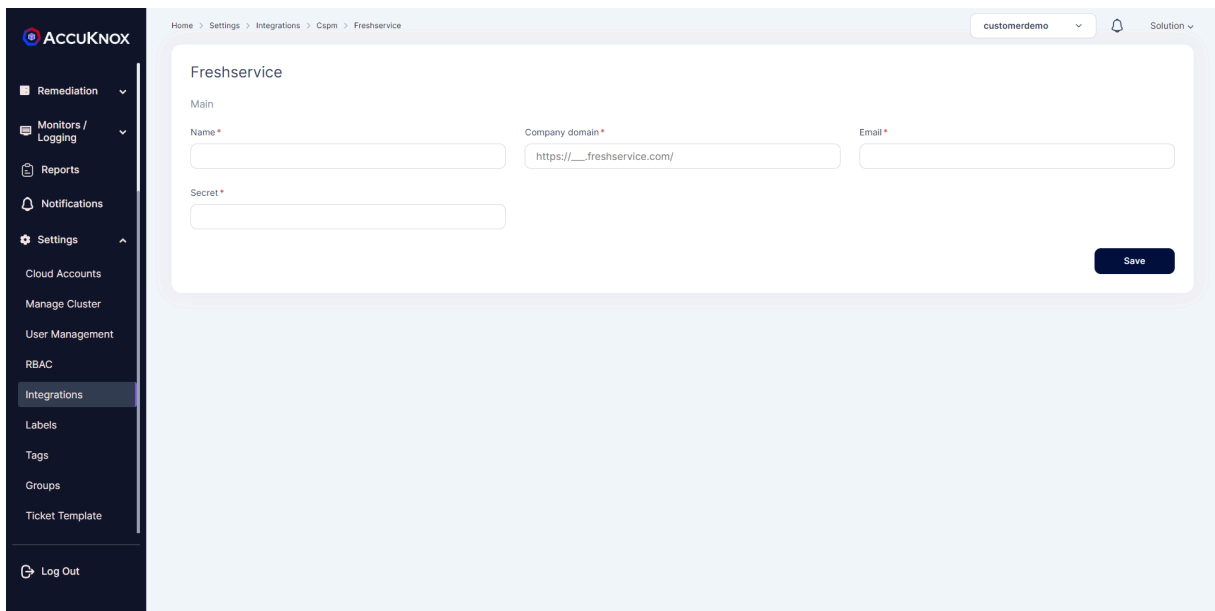
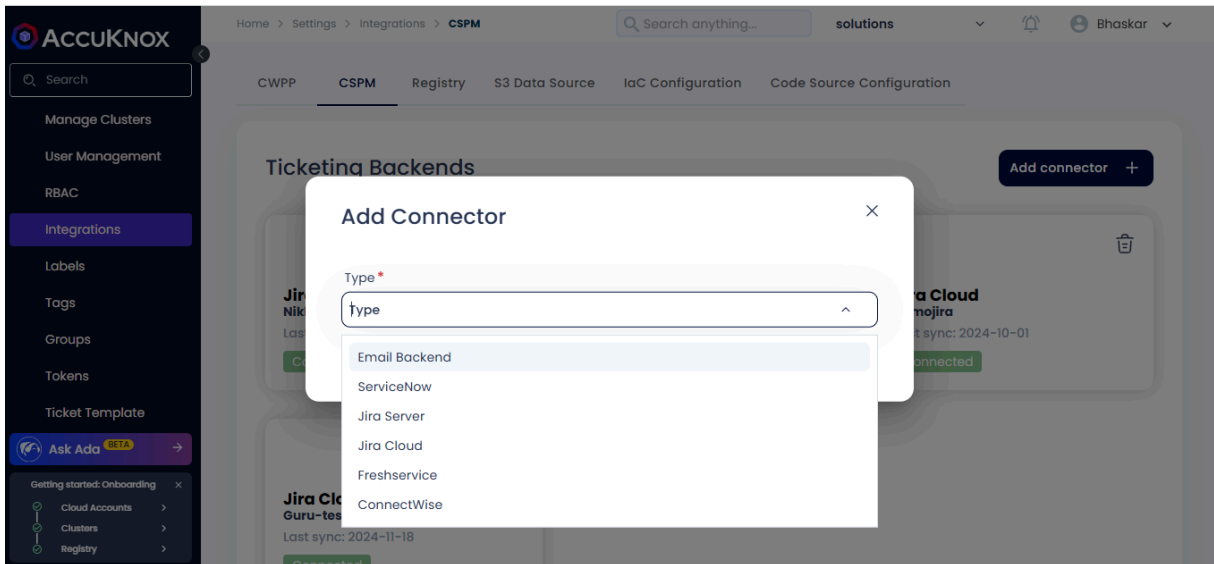
To set up this integration, you need to coordinate with your Freshservice administrator and gather the inputs needed to enable communication between AccuKnox and Freshservice.

#### a. Prerequisites

- You need a Company domain, Email & API key (secret) for this integration.
- You can find your API key in profile settings in the right side column.

b. Steps to Integrate:

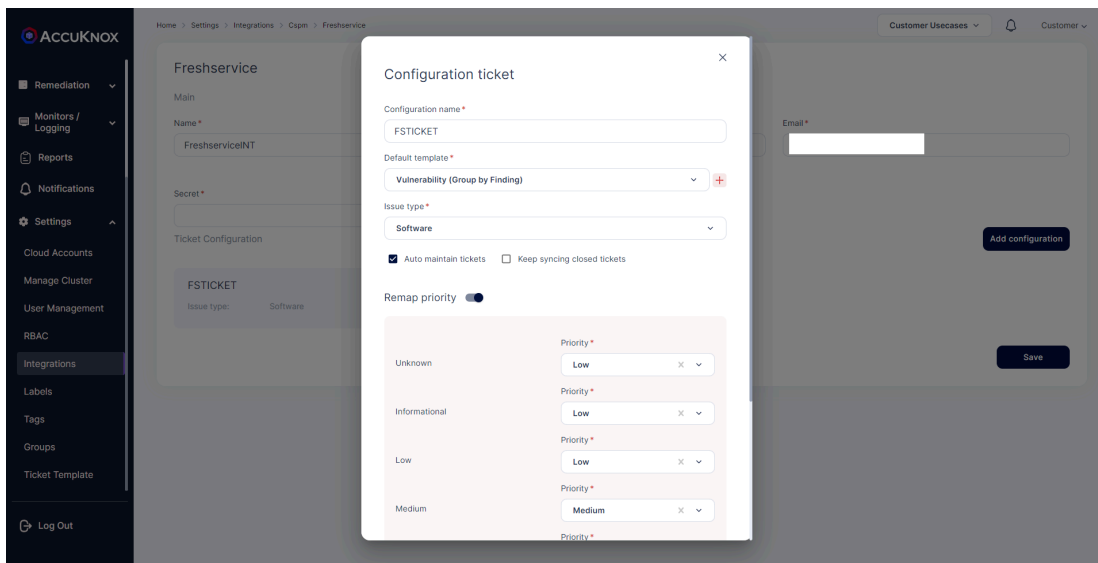
- Go to Channel Integration -> CSPM.
- Click on Add the connector and select Freshservice



Enter the following details to configure Fresh Service.

- Integration Name: Enter the name for the integration. You can set any name. e.g., TestFreshservice
- Domain Name: Enter the site name of your organization as shown in your URL. e.g., for <https://accuknoxexample.freshservice.com/> enter the domain name as accuknoxexample.

- User Email: Enter your Freshservice account email address here. e.g., freshservice@organisation.com
- Secret: Enter the API key Here. This can be found in profile settings.
- Click Save to save the Integration.

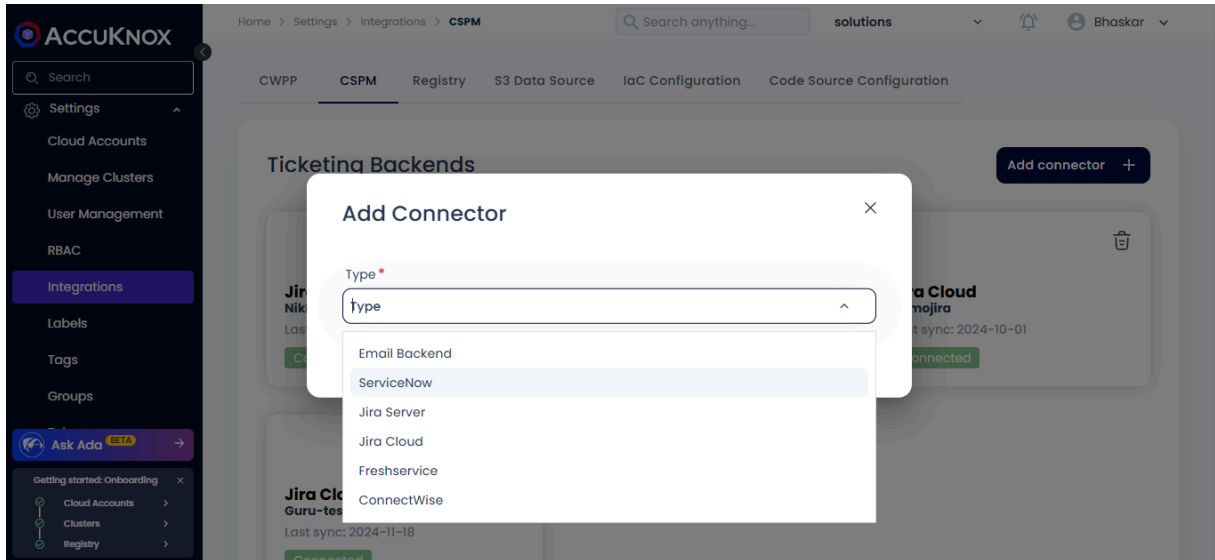


Click on the Freshservice ticketing backend to add configuration.

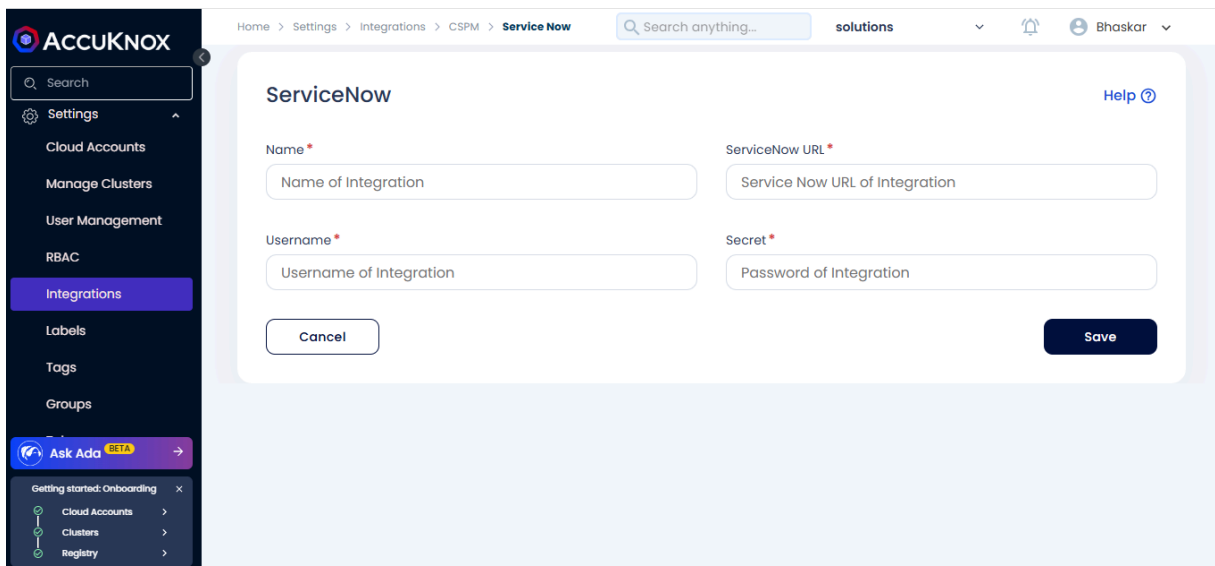
Here Enter the following details:

- Configuration name: this name will be displayed under ticket configuration while creating tickets.
- Default template: to specify the data that this configuration will be used for making tickets.
- Issue Type: You can choose from the dropdown.
- Fill in the priority mapping according to your choice and press save.

You can now configure Alert Triggers for Freshservice.



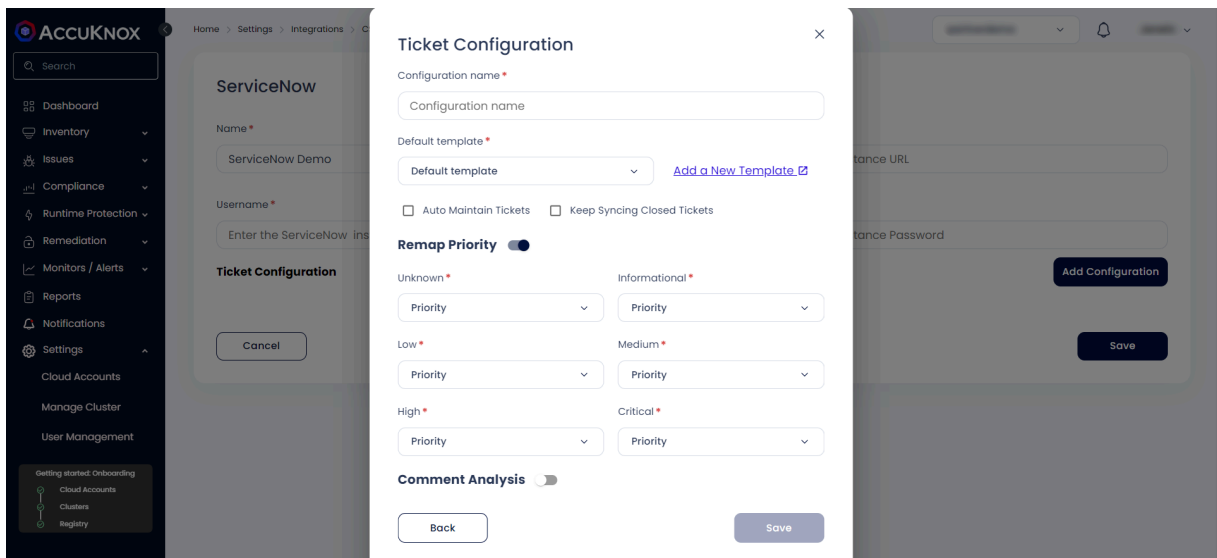
- Enter the following details to configure the ServiceNow Integration:
  - **Integration Name:** Enter the name for the integration. You can set any name. e.g., MyServiceNow
  - **ServiceNow Instance URL:** The URL of the ServiceNow instance. e.g., https://my-instance.service-now.com
  - **Instance Username:** The Username associated with the instance. e.g., admin
  - **Secret:** The current password of the instance.



- Click on the ServiceNow ticketing backend



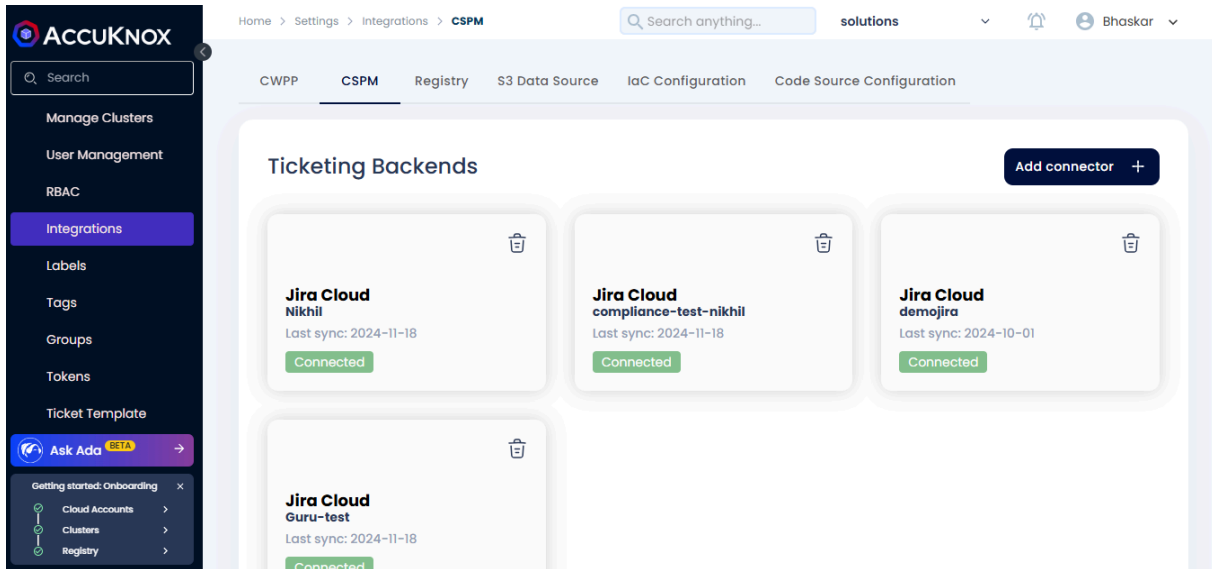
- Click on Add Configuration and enter the following details:
  - **Configuration name:** this name will be displayed under ticket configuration while creating tickets.
  - **Default template:** to specify the of data that this configuration will be used for making tickets.
  - **Issue Type:** You can choose from the dropdown.
  - Fill the priority mapping according to your choice and press **Save**.



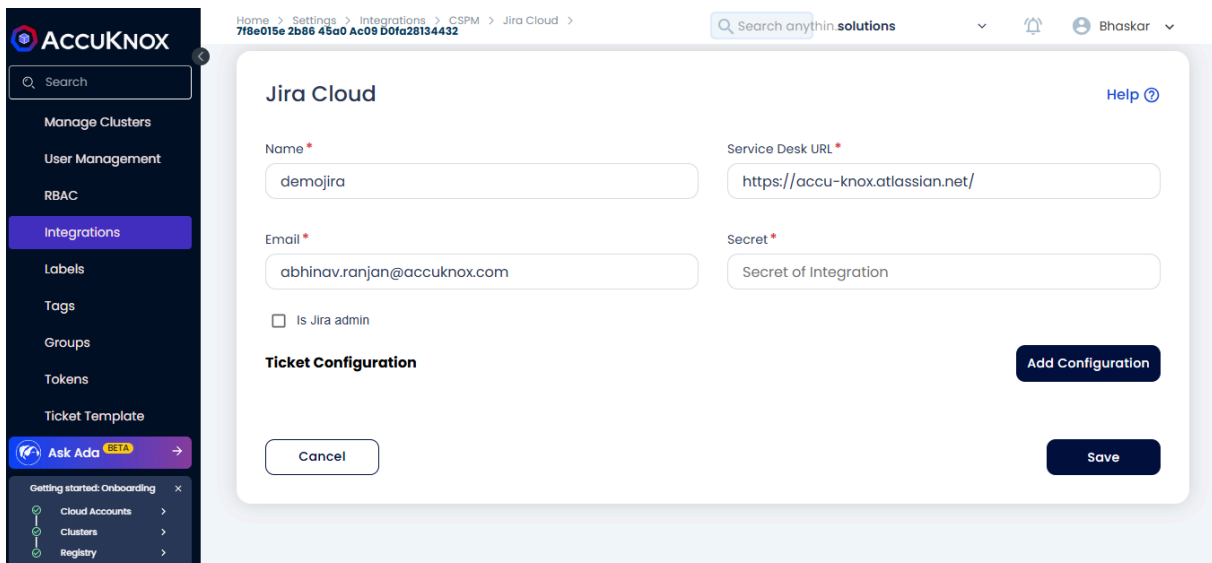
You can now create tickets on ServiceNow through the ticketing configuration.

## 17.4 Creating Ticket Configuration

- To create a ticket configuration, navigate to Integrations under Settings and click on the CSPM tab. This will show all the ticketing backends that have been integrated:

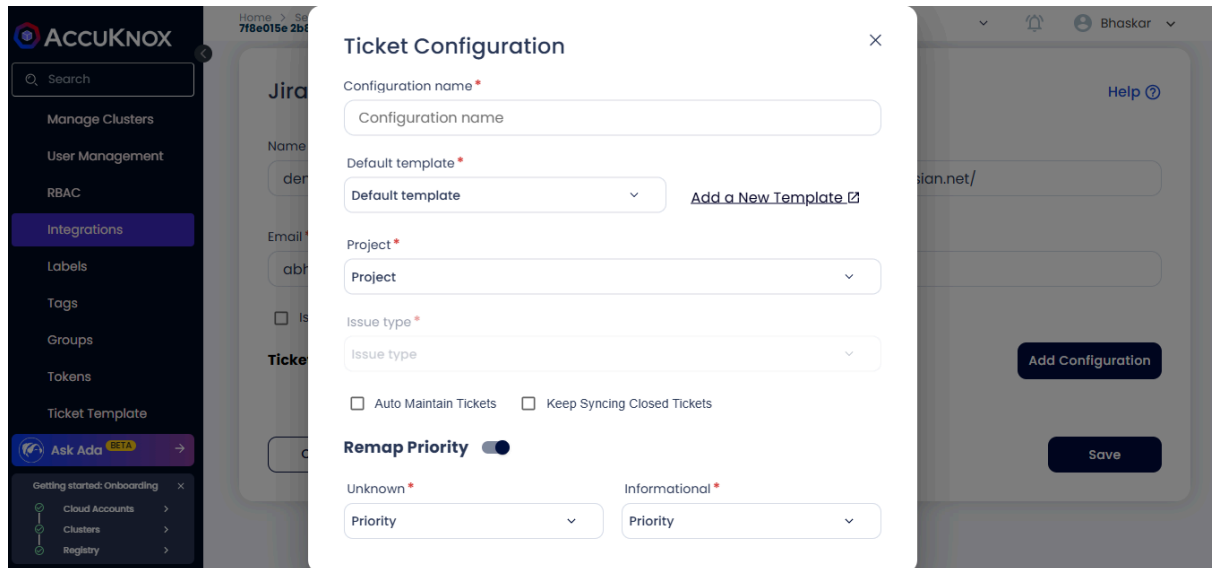


- Click on one of the integrated Ticketing backends and click on Add Configuration button in the subsequent screen:



- Enter a name for the configuration and select template for the ticket. The selected template will make it available in the respective screen as a ticket

configuration. Eg. Selecting Vulnerability will make it available as a ticket configuration to select under Issues -> Vulnerabilities for creating tickets.

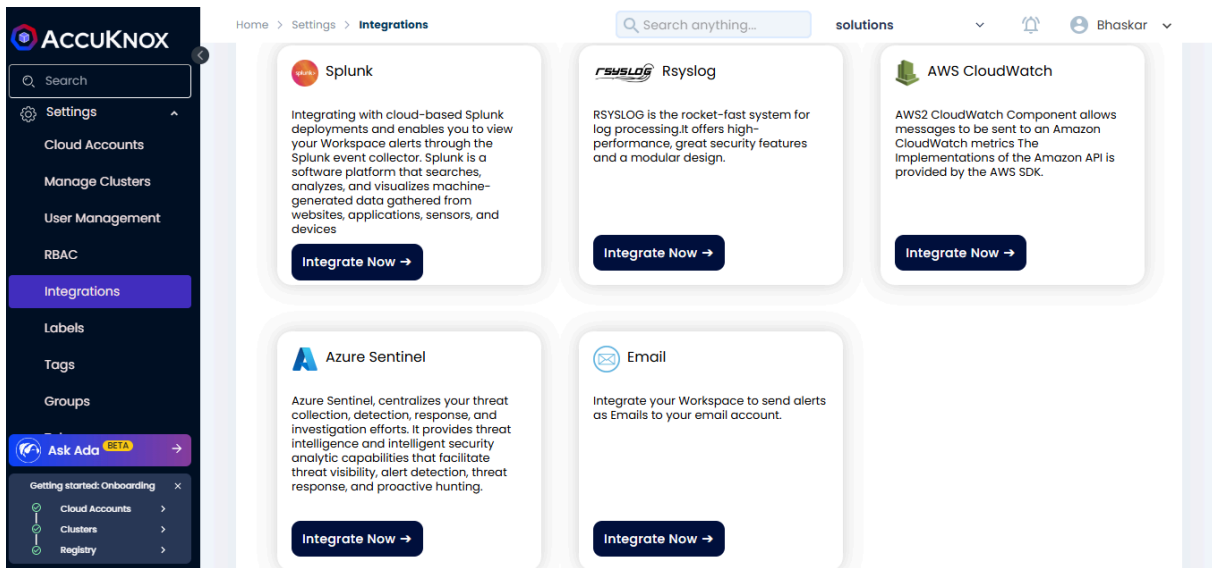


- Enter the relevant data in the remaining fields and click on Save. The ticket configuration is created successfully

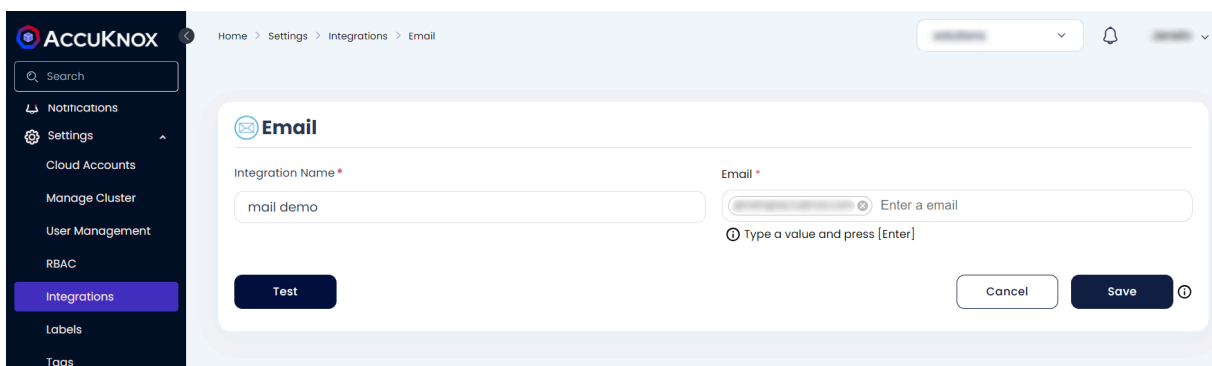
## 17.5 Email Integration

To send an alert notification via mail you must first set up the Email notification Channel.

- Navigate to Settings → Integrations → CWPP tab
- Click on the **Integrate Now** button under Email.



- Fill the following fields:
  - **Integration Name:** Enter the name for the integration. You can set any name. e.g., Container Security Alerts
  - **Email:** Enter the Email that will receive the notification and press ENTER. You can specify multiple email addresses in this field by pressing ENTER after each email address. e.g.,demo@organization.com



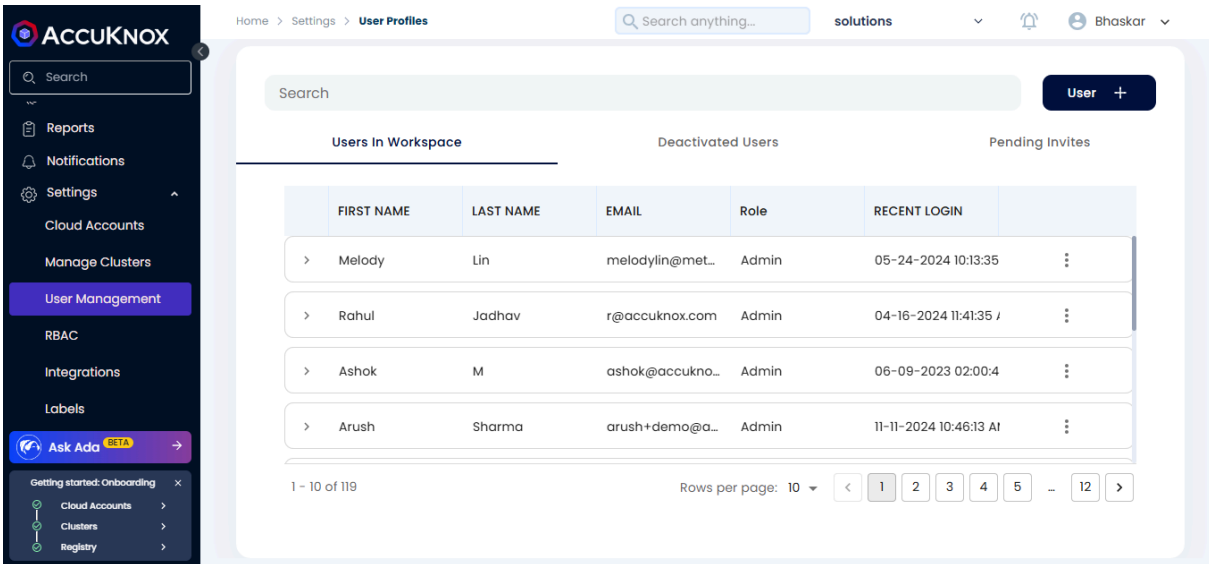
- Click **Test** to check the new functionality, You will receive a test mail on the specified mail addresses with subject "Test email"
- Click **Save** to save the Integration. You can now configure Alert Triggers for Email Notifications.

## 18. User Management

AccuKnox SaaS provides the ability to authenticate and authorize multiple users to access and utilize the SaaS platform. Inside the user management section user can create profiles for other users and these profiles are displayed in the form of a list. From the list, users can View Permissions, Edit, Deactivate, and delete user profiles. Permission is given to users by assigning roles while creating a user profile. These roles are created in the RBAC section. Deactivated users can be viewed under the Deactivated Users subsection. Creating a user sends an invite to their email id, invites that are not yet accepted are present inside the Pending Invites subsection.

### 18.1 Inviting a New User

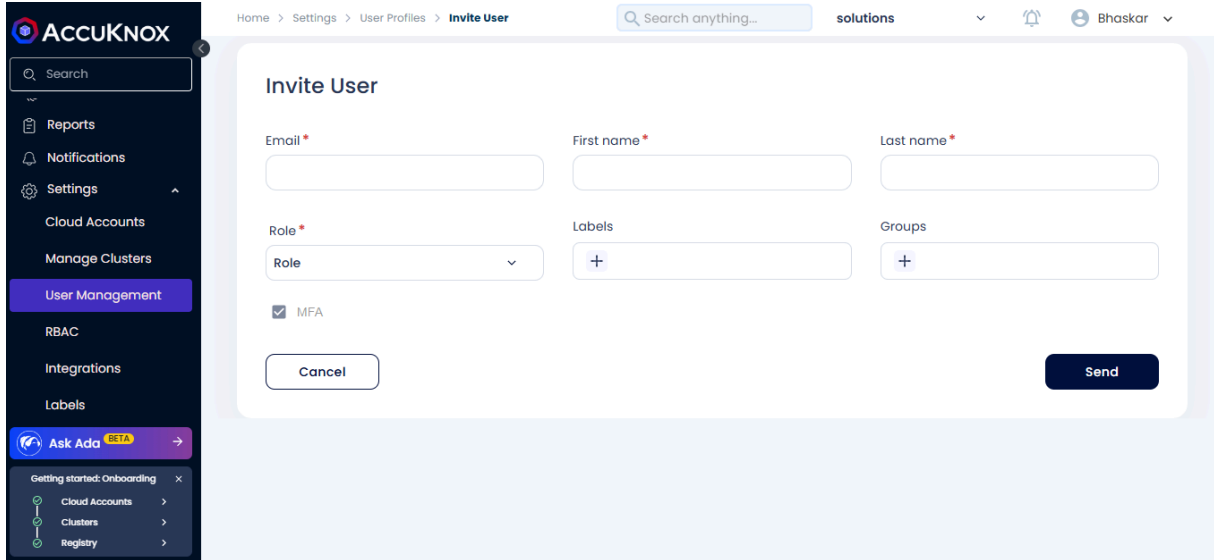
Log in to your AccuKnox dashboard. Navigate to "User Management" in the left sidebar menu. Click the "User +" button in the top right corner of the Users page.



The screenshot shows the AccuKnox dashboard with the "User Management" section active. The interface includes a sidebar menu on the left with options like Reports, Notifications, Settings, Cloud Accounts, Manage Clusters, User Management (highlighted), RBAC, Integrations, and Labels. The main content area displays a table of users in the workspace. At the top right, there is a search bar and a "User +" button. Below the search bar, there are tabs for "Users In Workspace", "Deactivated Users", and "Pending Invites". The "Users In Workspace" tab is selected, showing a table with columns for First Name, Last Name, Email, Role, and Recent Login. The table contains four rows of user data. At the bottom, there is a pagination control showing "1 - 10 of 119" and "Rows per page: 10".

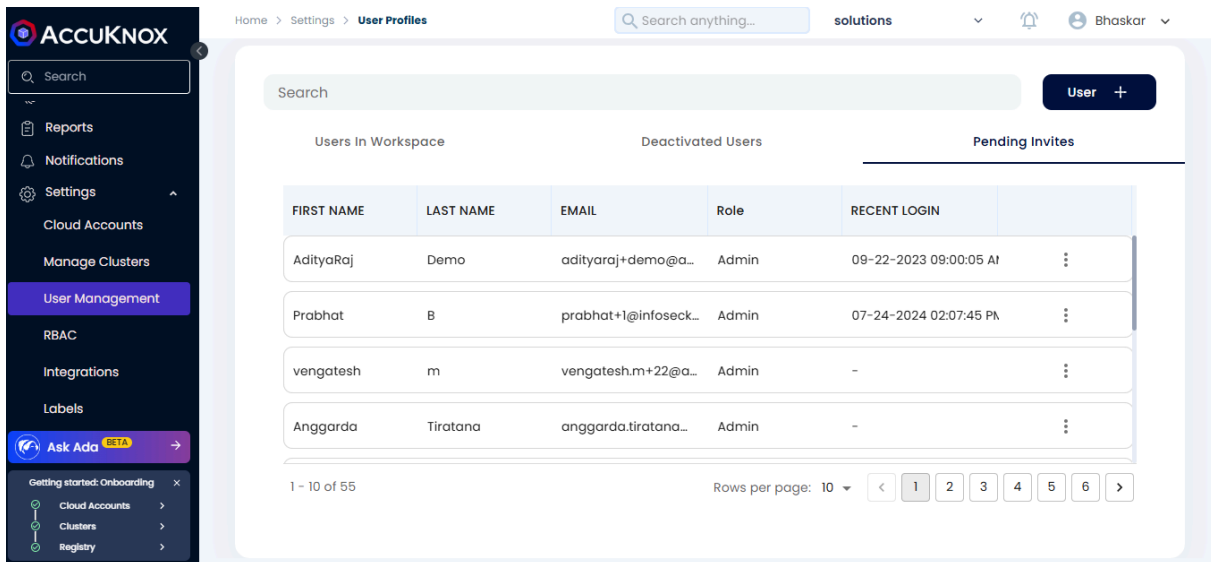
FIRST NAME	LAST NAME	EMAIL	Role	RECENT LOGIN
Melody	Lin	melodylin@met...	Admin	05-24-2024 10:13:35
Rahul	Jadhav	r@accuknox.com	Admin	04-16-2024 11:41:35 /
Ashok	M	ashok@accukno...	Admin	06-09-2023 02:00:4
Arush	Sharma	arush+demo@a...	Admin	11-11-2024 10:46:13 At

In the "Invite User" form, fill out the following details and hit Send.



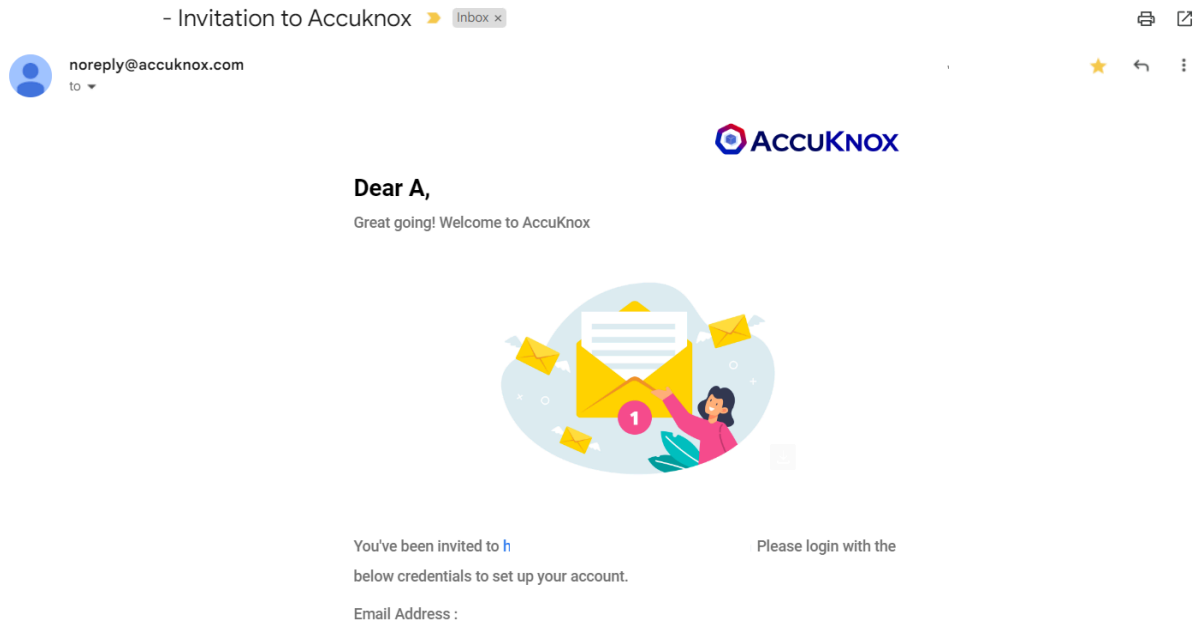
### Note

You can view pending invitations in the "Pending Invites" tab on the Users page. You can resend or revoke invitations from this tab. Viewing all permissions of a user is possible via the main tab.



## 18.2 User Receives Invitation

The invited user will receive a link on their mail to accept the invitation and set up their account if they haven't already done so.



## 18.3 User Login Options

Users can log in to AccuKnox using two methods:

### Option A: Traditional Login

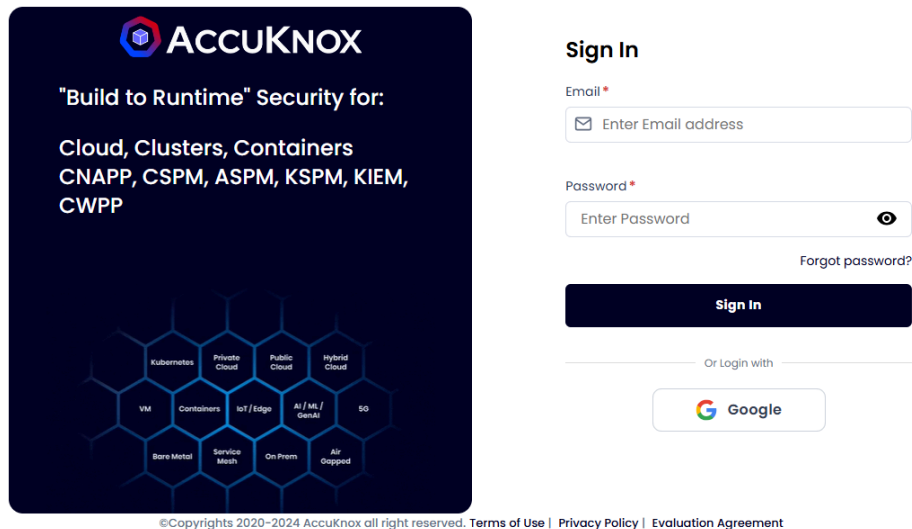
1. Go to the AccuKnox login page.
2. Enter the email address and password.
3. Click "Sign In".

### Note

This requires you to use the MFA (multi-factor authentication) code if it was enabled during the invitation process. MFA is required for every sign-in attempt.



## Option B: Single Sign-On (SSO) with Google



1. Go to the AccuKnox login page.
2. Look for "Or login with" at the bottom of the form.
3. Click on the "Google" button.
4. If not already signed in to Google, enter Google account credentials.
5. Grant any necessary permissions for AccuKnox.

### Note

If you are already signed in to Google, you will be automatically logged in to AccuKnox. No need for MFA in this case.

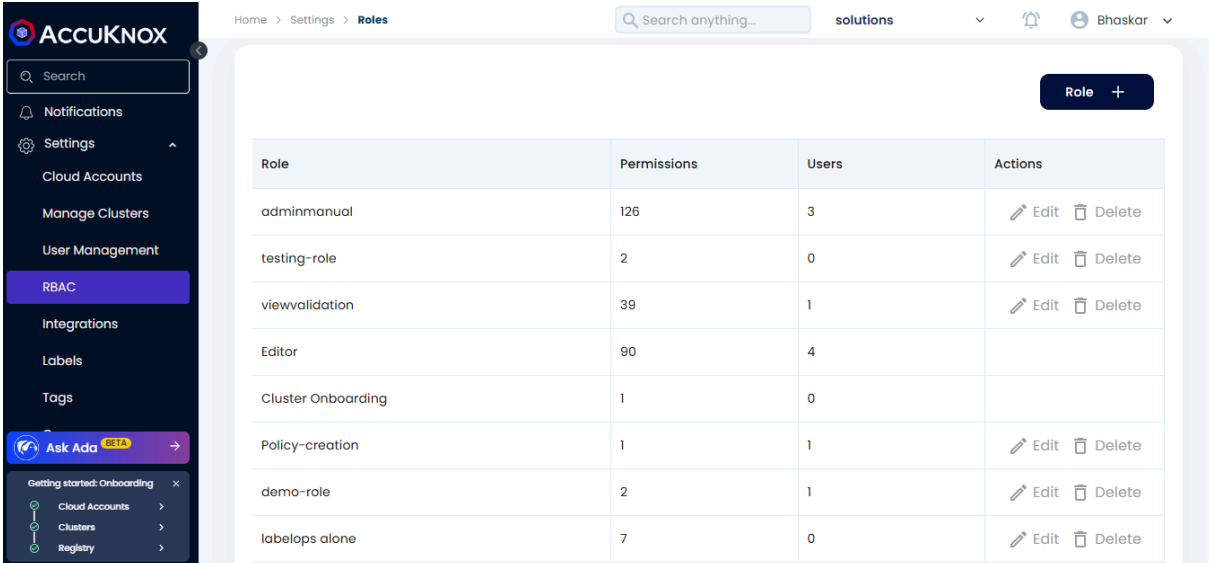
### 3. Notes

- SSO is currently only supported for Google accounts.
- Users must be invited with their Gmail address to use Google SSO.
- For the best experience, use the same email address for invitation and login.
- If you encounter any issues, contact your AccuKnox administrator or support team.

- Emails with + modifiers (e.g., test+stable@gmail.com or example+solutions@gmail.com) are not supported for SSO. Please use a base email address.

## 18.4 Assign RBAC

The role-Based Access Control option gives the option of creating users with different roles. we can create and manage roles that will be assigned to user profiles for their authorization. Users can select a set of permissions for each role like access to the Dashboard, Inventory, Issues, Runtime Protection, Compliance, Remediation, Monitors, and Settings. Roles can be created by clicking add roles or by cloning the existing roles. Roles are of two types, default roles come prebuilt and cannot be edited or deleted, and all other roles are custom roles.

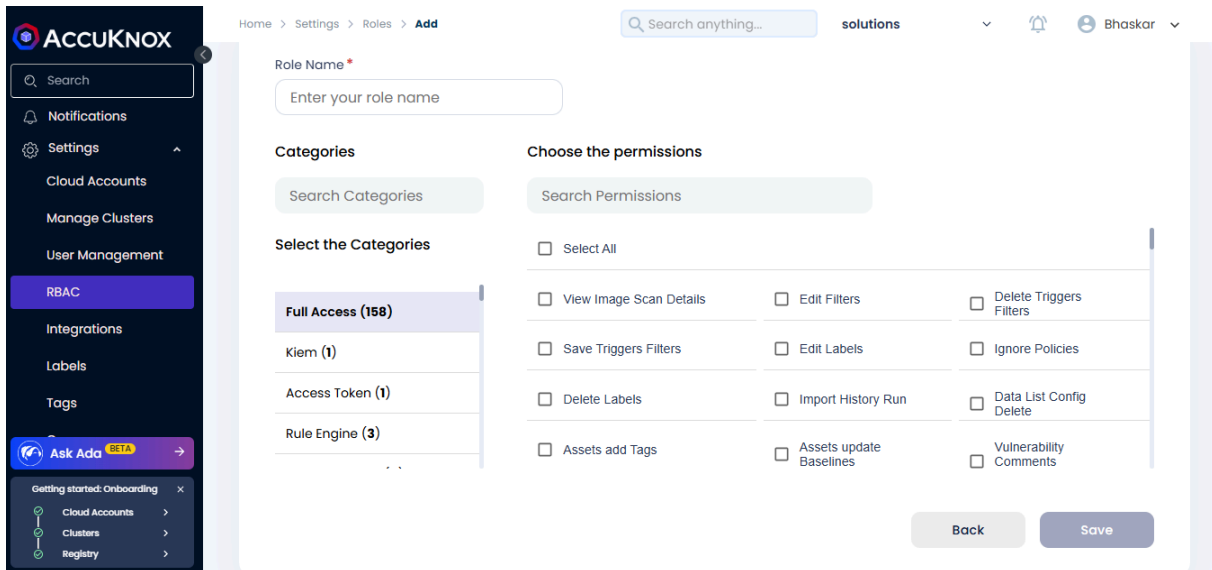


Role	Permissions	Users	Actions
adminmanual	126	3	Edit  Delete
testing-role	2	0	Edit  Delete
viewvalidation	39	1	Edit  Delete
Editor	90	4	
Cluster Onboarding	1	0	
Policy-creation	1	1	Edit  Delete
demo-role	2	1	Edit  Delete
labelops alone	7	0	Edit  Delete

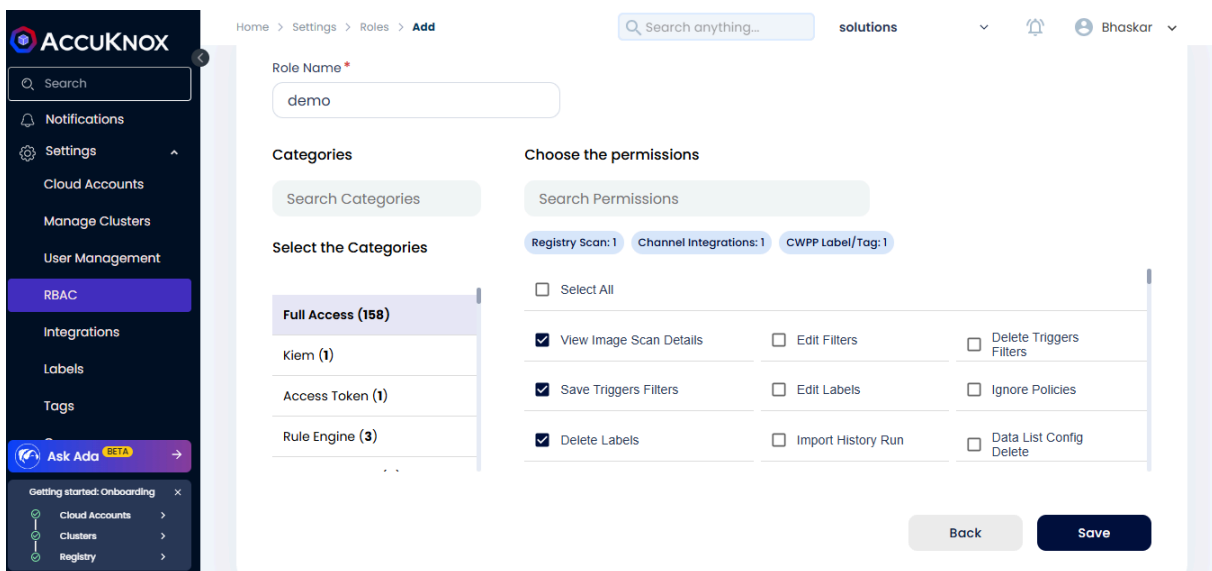
## 18.5 Create Roles and Assign Users

Steps:

- Click on Add Role



- Enter the name for Role along with it specify the role permission



- Click on Save
- Navigate to User Management > Add User > Choose the role created
- Send to the new user with custom role and permission

**AccuKnox**

Search

- Notifications
- Settings
  - Cloud Accounts
  - Manage Clusters
  - User Management**
  - RBAC
  - Integrations
  - Labels
  - Tags
- Ask Ada BETA
- Getting started: Onboarding
  - Cloud Accounts
  - Clusters
  - Registry

### Invite User

Email\*

First name\*

Last name\*

Role\* 

- adminmanual
- testing-role
- viewvalidation
- Editor

Labels

Groups

**Send**

## 19. Ticketing Procedures

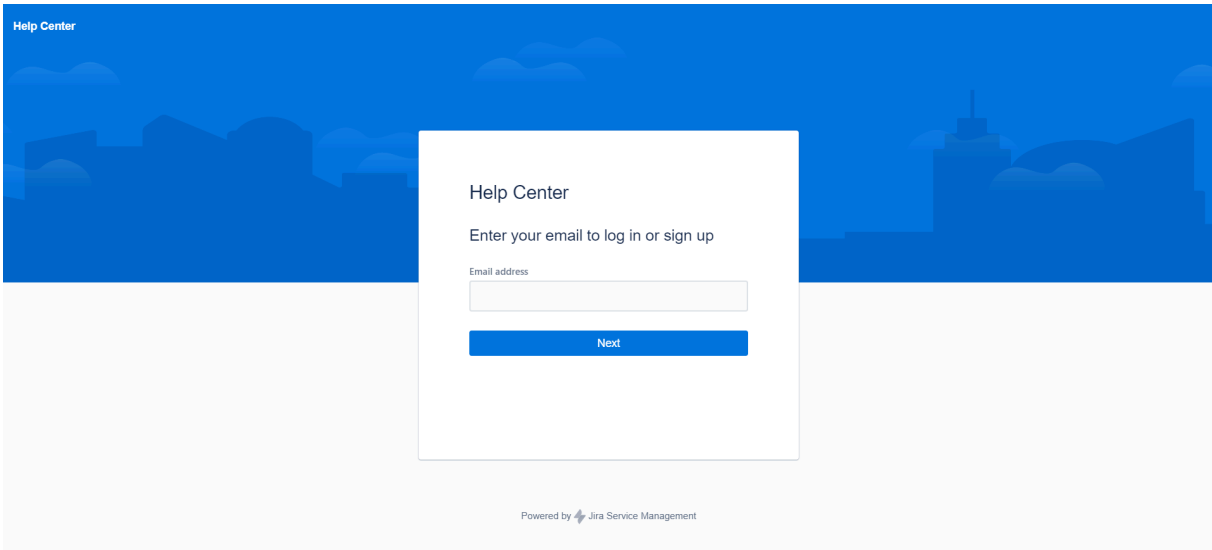
By following these steps, you can quickly and effectively raise a Jira ITSM support ticket for major platform issues, ensuring that your problem is addressed promptly and efficiently.

### 19.1 How to raise an AccuKnox support ticket?

**Step 1:** Please click the following URL for raising the ticket:

<https://accu-knox.atlassian.net/servicedesk/customer/portal/1>

**Step 2:** The page will ask for you to input the mail ID for signup



Help Center

Help Center

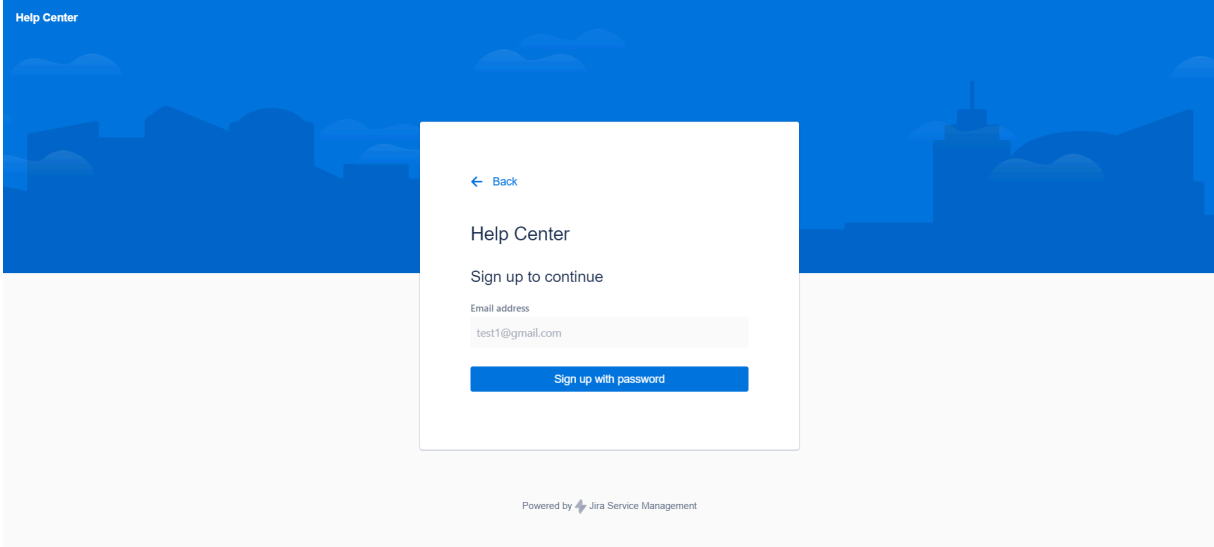
Enter your email to log in or sign up

Email address

Next

Powered by Jira Service Management

**Step 3:** After giving the email ID and selecting next will ask the user to sign in with a password



Help Center

← Back

Help Center

Sign up to continue

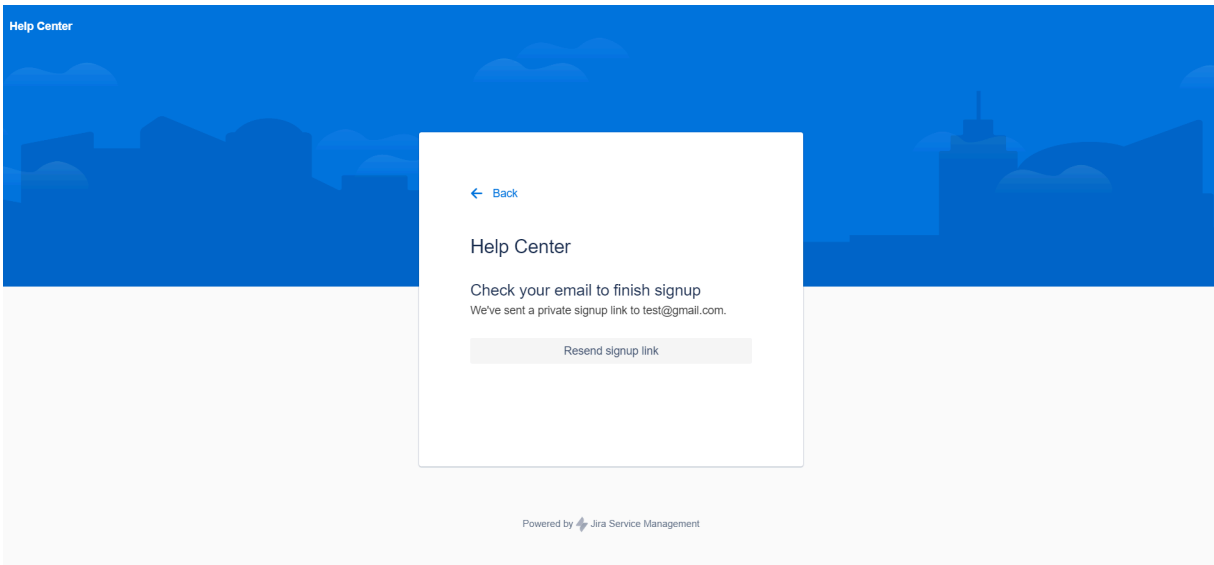
Email address

test1@gmail.com

Sign up with password

Powered by Jira Service Management

**Step 4:** Once users click the sign up with password, they will get an email for setting the password to the registered email id.



Help Center

← Back

Help Center

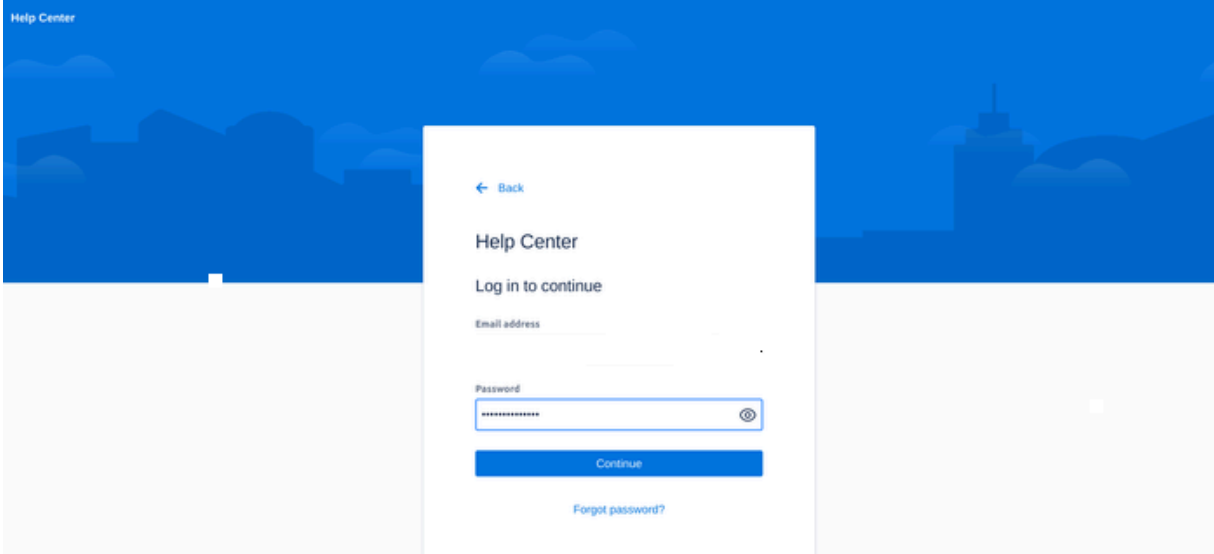
Check your email to finish signup

We've sent a private signup link to test@gmail.com.

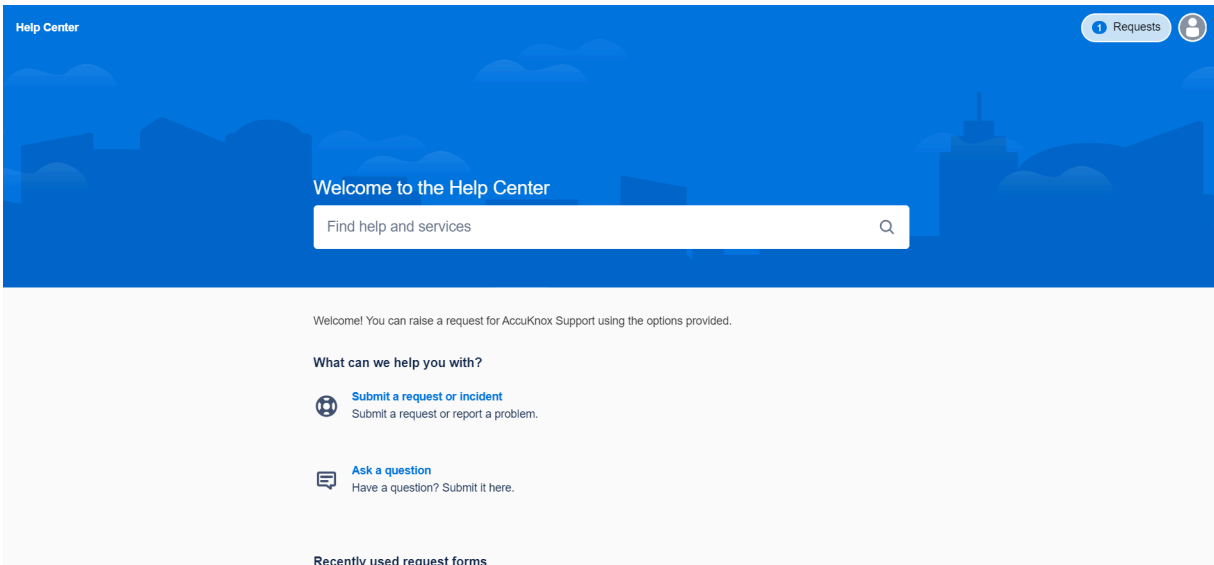
Resend signup link

Powered by Jira Service Management

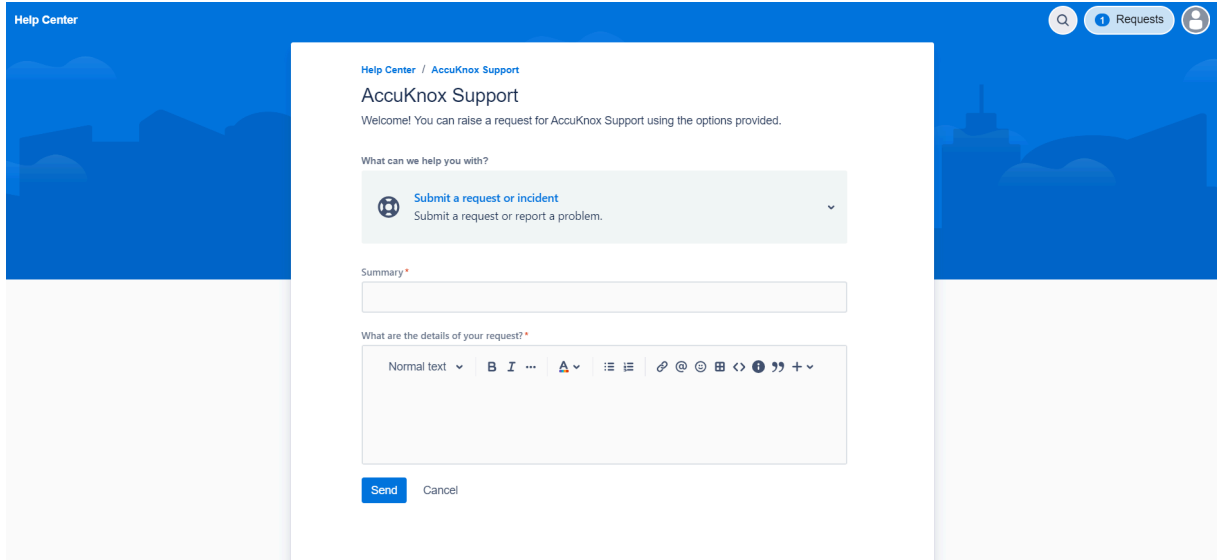
**Step 5:** After clicking the link and setting up the password and username login into the customer portal again  
<https://accu-knox.atlassian.net/servicedesk/customer/portal/1>



**Step 6:** Click on the Submit a request or incident option to create the issue



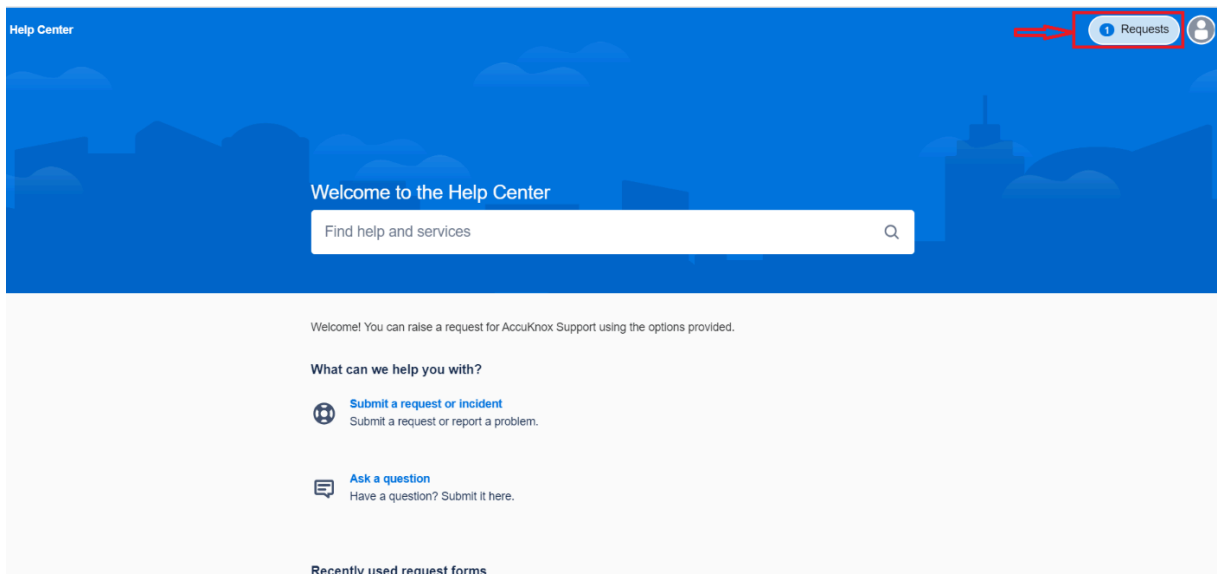
**Step 7:** To create an issue fill out this form and click send. Once it is clicked, the issue is created, and you will get a confirmation email to your registered email ID.



## 19.2 How to track the issue resolution status?

To track the issue raised by the user they can log into the customer service portal using the link <https://accu-knox.atlassian.net/servicedesk/customer/portal/1>

**Step 1:** Click on the requests section in the top left corner of the screen



**Step 2:** Here you will find the list of issues created by the user and their status





Help Center

Help Center

### Requests

Request contains...   **Status: Open requests**  **Created by me**  Request type

Type	Reference :	Summary	Status	Service project	Requester :
	SUPPORT-2	Issue in Registry scan	<b>TO DO</b>	AccuKnox Support	

Powered by  Jira Service Management

## 20. FAQs

### 20.1 AccuKnox FAQs

#### **1. Does AccuKnox CNAPP support only agent-based scanning or does it support agentless scanning ?**

For CSPM, AccuKnox supports agentless scanning for Public Cloud Infrastructure. For Infrastructure behind a firewall or Private Cloud, AccuKnoxCSPM leverages open source based agent to manage remote nodes for Automated reporting, Error log Delivery, Microservice Monitoring, User Shell Activity, Resource Monitoring.

For CWPP, AccuKnox leverage open source CNCF sandbox project KubeArmor for scanning and in-line mitigation from known attacks. Together we provide a complete static and runtime security for a variety of workloads whether they are on Public/Private Cloud, VM, Baremetal or pure-containerized workload. Thus we require agents to be installed to support scanning the workloads.

#### **2. What is the differentiation of AccuKnox in Static Security?**

In the Static Security solution, unlike other CSPM tools, AccuKnox provides flexibility to integrate a variety of open source and commercial security scanning tools through built-in parsers to provide you a composite security posture of your infrastructure. We also correlate and normalize results from a variety of security scanning tools and provide detailed results of vulnerabilities across infrastructure.

#### **3. How does AccuKnox help to achieve static security?**

AccuKnox Cloud Security Posture Management (CSPM) tool scans the Cloud Account to assess Vulnerabilities, Misconfigurations that are present in the cloud infrastructure based on security best practices & benchmarks. AccuKnox also enables you to handle Vulnerabilities with the ability to mark false positives, Waiting for 3<sup>rd</sup> party or Accepted risk and many more, so that you get to act on findings that are remediable and containing the SLA. We also give comprehensive compliance reports based on various security governance for third party assessment operators (3PAO) auditing.

#### **4. How does AccuKnox help to achieve Runtime security?**

AccuKnox's Cloud Workload Protection Platform (CWPP) achieves runtime security by leveraging CNCF sandbox project, KubeArmor, which is a cloud-native runtime security enforcement system by AccuKnox that restricts and have more granular control over the application behavior such as process execution, file access, and networking operation of containers and nodes at the system level.

#### **5. What is the differentiation of AccuKnox in Runtime Security?**

AccuKnox leverages KubeArmor, which is a cloud-native runtime security enforcement system that leverages Linux Security Modules to secure the workloads. LSMs are really powerful but they weren't built with modern workloads including Containers and Orchestrators in mind. Hence, eBPF has provided us with the ability to extend capabilities and BPF LSM provide us with the ability to load our custom programs with

decision-making into the kernel seamlessly helping us protect modern workloads. Therefore, KubeArmor helps to enforce security posture wherein any malicious attacks will be stopped before execution, known as in-line mitigation (mentioned by Forrester report)

## **6. What does KubeArmor leverage for enforcement and what are its advantages?**

KubeArmor leverages best of breed Linux Security Modules (LSMs) such as AppArmor, BPF-LSM, and SELinux for inline mitigation to reduce the attack surface of the pod/container/VM. LSMs have several advantages over any other techniques. By using LSMs, KubeArmor does not have to disturb pods/containers and also doesn't require change at host or CRI level to apply security policies.

KubeArmor deploys as a non-privileged daemonset with certain capabilities that allows it to monitor other pods/containers and host. A given cluster can have multiple nodes utilizing different LSMs so KubeArmor abstracts away the complexities of the LSMs and provides an easy way for policy enforcement.

## **7. What are the integration tools and registries that are supported by AccuKnox?**

AccuKnox can integrate multiple Cloud Account, Registries, SIEM platform, Ticketing or Notifications Tools and the list is ever growing. AccuKnox is pretty flexible to support the progression of the list with the customer's request as our roadmap item. Some of the supported today are as follows:

- Security Events/SIEM : Splunk, Rsyslog, AWS CloudWatch, Elastic Search, Webhooks, Azure Sentinel
- Notification Tools: Slack, Jira, PagerDuty, Emails
- Ticketing Tools: Jira, FreshService, Connectwise, Zendesk
- Registries: Nexus, ECR, GCR, DockerHub, ACR, Harbor

## **8. How AccuKnox helps in Policy Version Control for Runtime Security?**

AccuKnox enables DevSecOps teams to embed security policies as code into their GitOps workflow. This provides a unified, collaborative view of the policies and enables them to be shipped and deployed along with the applications they are protecting. Hence, utilizing Gitops based policy version control, it will be easy to enforce changes to policies and keep track of versions in case of audit or rollback requirement alongwith approval mechanisms.

## **9. How AccuKnox helps to achieve Microsegmentation?**

AccuKnox CWPP provides micro-segmentation at the lowest possible granularity level which is also a smallest execution unit in Kubernetes i.e. Pods. We will help you to identify process execution request from the pods, network connections the pods are trying to make internally or externally and files-system the pods are accessing. By observing the behavior of a particular pod and restricting that behavior so that it functions according to the expected flow of process/events/traffic, one can develop a least permissive security posture from creating a whitelisting policies and auditing/denying everything else.

## **10. How AccuKnox helps to recommend Auto-Discovered Policies?**

AccuKnox CWPP solution provide Discovery Engine agent that assesses the security posture of your workloads and auto-discovers the policy-set required to put the workload in least-permissive mode. We also provide Shared Informer Agent which collects information about cluster like pods, nodes, namespaces etc. The Policy Discovery Engine discovers the policies using the workload and cluster information that is relayed by Shared Informer Agent.

## **11. What are Hardening Policies?**

KubeArmor is a security solution for the Kubernetes and cloud native platforms that helps protect your workloads from attacks and threats. It does this by providing a set of hardening policies that are based on industry-leading compliance and attack frameworks such as CIS, MITRE, NIST-800-53, and STIGs. These policies are designed to help you secure your workloads in a way that is compliant with these frameworks and recommended best practices.

## **12. What is Network Segmentation?**

In Kubernetes, the network policy resource is a set of network traffic rules that are applied to a group of pods in a Kubernetes cluster. The network policy specifies how a pod is allowed to communicate with others. Network policy controllers (running as pods in the Kubernetes cluster) convert the requirements and restrictions of the network policies that are retrieved from the Kubernetes API into the network infrastructure.

## **13. How AccuKnox helps to implement Zero Trust?**

By implementing a zero trust posture with KubeArmor, organizations can increase their security posture and reduce the risk of unauthorized access or activity within their Kubernetes clusters. This can help to protect sensitive data, prevent system breaches, and maintain the integrity of the cluster. KubeArmor supports allow-based policies which result in specific actions to be allowed and denying/auditing everything else. For example, a specific pod/container might only invoke a set of binaries at runtime. As part of allow-based rules you can specify the set of processes that are allowed and everything else is either audited or denied based on the default security posture.

## **14. Does KubeArmor only support Kubernetes or it can support on-prem deployments like legacy VM, pure containerized workload as well?**

KubeArmor supports following types of workloads:

- K8s orchestrated workloads: Workloads deployed as k8s orchestrated containers. In this case, KubeArmor is deployed as a k8s daemonset. Note, KubeArmor supports policy enforcement on both k8s-pods (KubeArmorPolicy) as well as k8s-nodes (KubeArmorHostPolicy).
- VM/Bare-Metals workloads: Workloads deployed on Virtual Machines or Bare Metal i.e. workloads directly operating as host processes. In this case, KubeArmor is deployed in systemd mode.

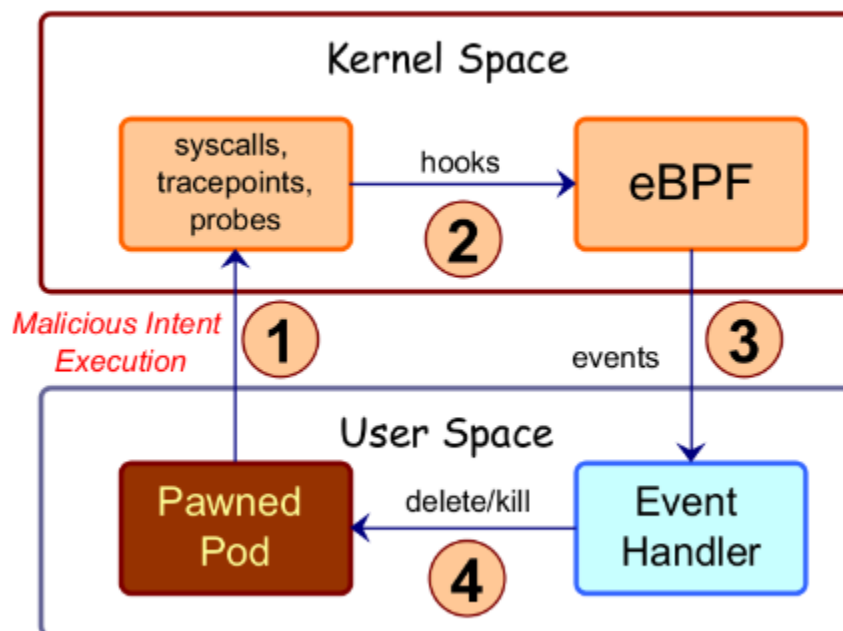
## **15. How AccuKnox helps achieve protection for Edge, 5G workloads?**

With edge computing shifting towards containerized workloads and in few cases to orchestrated kubernetes workloads, it becomes important to have a security solution which can not only provide enforcement into different forms of deployment but can also provide real-time container-rich observability. KubeArmor

supporting un-orchestrated containers, k8s workloads and bare metal VMs makes it an ideal universal engine. Its kernel-level runtime security enforcement and container aware observability brings the best of both worlds..

## 16. What is the difference between Post-attack mitigation and in-line mitigation and which is better?

Post-exploit Mitigation works by killing the suspicious process in response to an alert indicating malicious intent. In this case attacker will be allowed to is able to execute its binary and could possibly disable the security controls, access logs, etc to circumvent the attack detection. By the time the malicious process is killed, it might have already deleted, encrypted, or transmitted the sensitive contents.



Inline Mitigation on the other hand prevents the malicious attack at the time of happening itself. It doesn't allow the attack to happen by protecting the environment with security policy or firewall. AccuKnox's open source tool KubeArmor provides Inline Mitigation. KubeArmor uses inline mitigation to reduce the attack surface of pod/container/VM. KubeArmor leverages best of breed Linux Security Modules (LSMs) such as AppArmor, BPF-LSM, and SELinux (only for host protection) for inline mitigation

## 17. What are the platforms and environments that AccuKnox supports?

AccuKnox supports the following environments: + SaaS + PaaS + IaaS

AccuKnox supports the following cloud platforms: + AWS + GCP + Azure

AccuKnox support for the different platforms are as follows: + Kubernetes - Fully supported + Linux - Supported distributions + Serverless - Fargate and ECS supported, others are on roadmap + Windows - On roadmap

## 18. What role does AccuKnox Agents play in runtime-security?

AccuKnox Enterprise version consists of various agents such as

**KubeArmor:** KubeArmor is a cloud-native runtime security enforcement system that restricts the behavior (such as process execution, file access, and networking operation) of containers and nodes at the system level. KubeArmor dynamically set the restrictions on the pod. KubeArmor leverages Linux Security Modules (LSMs) to enforce policies at runtime.

**Feeder Service:** It collects the feeds from kubeArmor and relays to the app.

**Shared Informer Agent:** It collects information about the cluster like pods, nodes, namespaces etc.,

**Policy Discovery Engine:** It discovers the policies using the workload and cluster information that is relayed by a shared informer Agent.

## 19. Does AccuKnox provide auto discovery of assets and workloads?

Yes, AccuKnox can auto discover assets in the cloud by leveraging the cloud native tools.

For workloads, AccuKnox agents will provide the visibility data.

## 20. Can AccuKnox help in Monitoring?

- With Accuknox, you can create monitors for assets or group of assets to get alerts for changes observed in their Metadata (software version etc)
- Our Drift detection capability is inherently doing monitoring of the compliance checks (pass/fail) that have changed between scans.
- We collect alerts and telemetry generated by Kubearmor and cillium. These alerts are part of our CWPP offering. These alerts are generated for the events that have violated/complied with a policy.
- For these alerts you can have notification enabled as well through channels like Slack, email etc.

## 21. Do I need to enable native security services for AWS to get data into Accuknox?

AccuKnox only requires an IAM role to be created with read only access to be able to get data from AWS. Security Hub and Macie can be optionally enabled for AccuKnox to gather richer telemetry data with more context.

## 22. What are the Hypervisors or Virtualized Environments that are supported by AccuKnox?

AccuKnox technology does not integrate at the VM virtualization layer. AccuKnox tech integrates at the operating system layer and ensures that the right hardening/enforcement for process executions, network access, and file access is in place. Thus AccuKnox can operate on any virtualization tech provided that the underlying VM uses Linux as its operating system.

### **23. What is the differentiation of AccuKnox in ASPM Security?**

In the ASPM Security solution, unlike other tools, AccuKnox provides flexibility to integrate a variety of open source and commercial security scanning tools through built-in parsers to provide you a composite security posture of your infrastructure. This is mainly done for the following two context:

- Remove dependencies and scoped results from one tool
- Bring in contextual understanding of vulnerabilities and prioritization based on that

Further on this, We also correlate and normalize results from a variety of security scanning tools and provide detailed results of vulnerabilities across infrastructure.

## 20.2 Bonus Questions

### 1. What are the modules supported by AccuKnox CNAPP currently?

- CSPM
- ASPM
- DevSecOps security in CI/CD pipeline
- CWPP
- Container Images Scanning
- CDR (Cloud Detection or Response) or CDM (Continuous Diagnostic & Mitigation)

### 2. What are all the compliance frameworks that AccuKnox is covering?

AccuKnox's CNAPP tool checks for compliance and governance from various benchmarks like STIG, CIS, NIST CSF, HIPAA, MITRE, SOC2, ISO 27001.

### 3. Does Inline remediation slowdown the process?

LSMs are already enabled in the environment and use host based LSM security. Since the attacker usually has direct access to the pod, AccuKnox uses Inline remediation to stop the processes before executing. Therefore, inline remediation does not slow down the process

### 4. What does AccuKnox measure, while doing security posture observation and how does it help in securing using policies?

- Compliance Frameworks (MITRE, CIS, NIST) for hardening workloads are used to create hardening policies
- Understanding the Application behaviour using LSMs enables creation of behavioural policies
- Hardening policies are block based policies
- Behavioural policies are allow based policies
- An example of policies is FIM (File Integrity Monitoring) policy

### 5. Do you have any standard hardening rules onboarded and will the hardening policy show what is getting blocked?

Yes, it can show up in terms of Application Behaviour & Logs

### 6. What is the deployment architecture?

- Applications -
  - For Kubernetes - Daemonset
  - For Containers, VM - Systemd mode
- Infrastructure -
  - Public Cloud - Agentless (API Scan) for SaaS based usage
  - On-Prem or Datacenter - On-prem deployment using Helm-charts

### 7. Where is AccuKnox SAAS is located?

Currently it is located in US region

### 8. Is there a support for CIEM?



It is a part of the roadmap, like IOT edge, 5G Security

### **9. What will happen to my application running on a VM?**

You get hardening policies via AccuKnox enforcement engine KubeArmor

### **10. What is AccuKnox's licensing model?**

If it is an end customer - SLA

If it is a MSSP model, it is a revenue share

### **11. How do you work with resellers and partnership models?**

We have a 100% partner aligned go to market approach. to this goal, we provide our partners the following  
+ Free training, certification + Joint marketing + Lead sharing

### **12. Current AccuKnox's marketplace listing?**

AccuKnox is currently listed on

- VMWare
- AWS
- Azure
- RedHat Openshift

We are in the process of listing on

- GCP
- Oracle

### **13. Who are current AccuKnox's partners and resellers?**

- We have a global partnership with TCS
- We have a reseller partnership with Ambisure

## References:

<https://help.accuknox.com/introduction/home/>

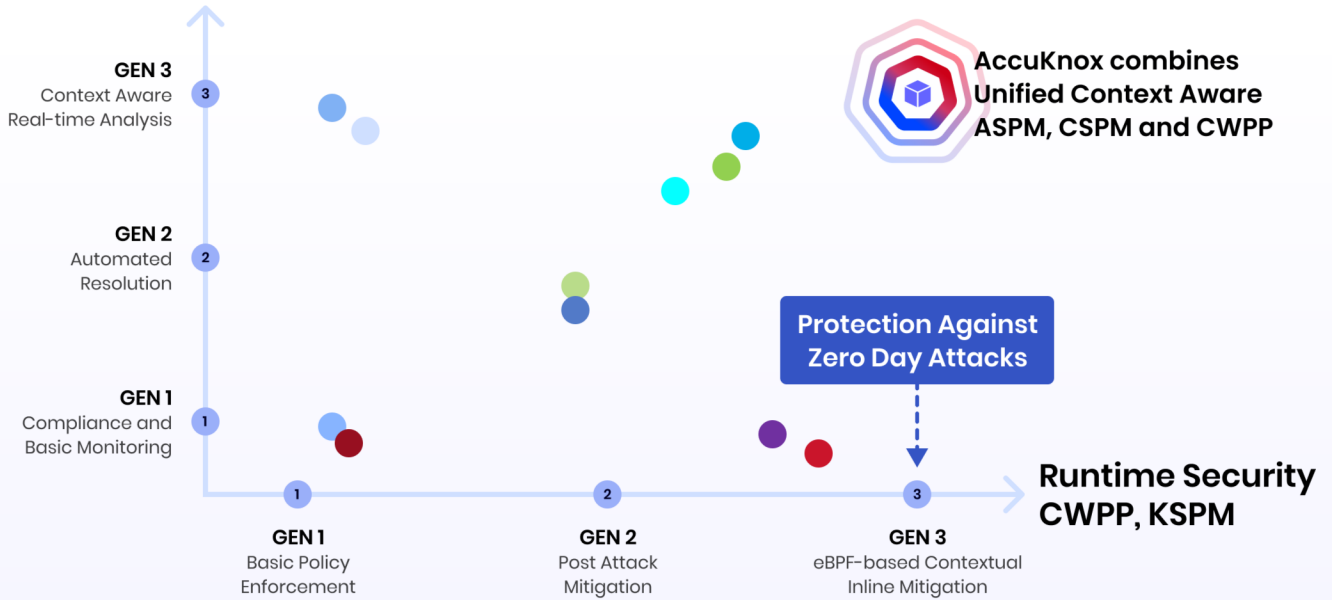
<https://kubearmor.io/>

<https://docs.kubearmor.io/kubearmor>

Featured by



### Static Security ASPM, CSPM



## Extra 30 Days Free Trial



\* No strings attached, Limited period offer!



Scan for Demo

## About AccuKnox

AccuKnox is a Zero Trust CNAPP Cloud Security protects Public clouds, Private clouds, Kubernetes, VMs, Bare metals, IoT Edge, and 5G security.



in [linkedin.com/accuknox](https://linkedin.com/accuknox)

X @AccuKnox

