



ZSCALER AND AWS DEPLOYMENT GUIDE

Contents

Terms and Acronyms	4
Trademark Notice	5
About This Document	6
Zscaler Overview	6
AWS Overview	6
Audience	6
Software Versions	6
Request for Comments	6
Zscaler and AWS Introduction	7
ZIA Overview	7
AWS Overview	8
AWS Resources	8
Overview	9
Cloud App Control Policy	10
Cloud App Control Policies Available via Individual Amazon Web Services	12
File Type Control for AWS	15
Firewall Control Rules for AWS	17
DNS Control	17
ZIA Components that Work on AWS Infrastructure	18
Nanolog Streaming Service	19
Virtual Service Edge	19
Cloud Connector	20
DLP Incident Receiver	20
DLP Index Tool	21
Amazon WorkSpaces Supporting Zscaler Client Connector	21

ZIA Integrations Inside AWS	22
Cloud NSS and S3 Buckets	22
Workflow Automation	22
SaaS Security API for S3 Buckets	22
Contextualizing Risk using AWS and Avalor UVM	26
Creating a Role ARN and an External ID in AWS	26
Output for the RoleARNID and ExternalID	29
Configure the AWS UVM Data Connectors	30
Configure the AWS Accounts Data Source	30
Configure the AWS EC2 Data Source	33
Configure the AWS ECR Data Source	36
Configure the AWS ECR Findings Data Source	39
Configure the AWS EKS Clusters Data Source	42
Configure the AWS Inspector Findings Data Source	45
Configure the AWS RDS Data Source	48
Configure the AWS S3 Buckets Data Source	51
Configure the AWS Security Hub API Data Source	54
Review and Adjust Data Model Mapping	57
Create a Crown Jewel Tag for an EC2 Instance	57
Map the AWS EC2 Data Source	59
Review and Adjust Risk Scoring	62
Appendix A: Requesting Zscaler Support	64

Terms and Acronyms

The following table defines acronyms used in this deployment guide. When applicable, a Request for Change (RFC) is included in the Definition column for your reference.

Acronym	Definition
ARN	Amazon Resource Name (Amazon)
AWS	Amazon Web Services (Amazon)
CA	Central Authority (Zscaler)
CSV	Comma-Separated Values
DLP	Data Loss Prevention
DNS	Domain Name Service
DPD	Dead Peer Detection (RFC 3706)
GRE	Generic Routing Encapsulation (RFC2890)
IAM	Identity and Access Management
ICAP	Internet Content Adaptation Protocol
ICMP	Internet Control Message Protocol
IdP	Identity Provider
IKE	Internet Key Exchange (RFC2409)
IPS	Intrusion Prevention System
IPSec	Internet Protocol Security (RFC2411)
PFS	Perfect Forward Secrecy
PSK	Pre-Shared Key
S3	Simple Storage Service (Amazon)
SaaS	Software as a Service
SNS	Simple Notification Service (Amazon)
SQS	Simple Queue Service (Amazon)
SSL	Secure Socket Layer (RFC6101)
TLS	Transport Layer Security
VDI	Virtual Desktop Infrastructure
XFF	X-Forwarded-For (RFC7239)
ZDX	Zscaler Digital Experience (Zscaler)
ZIA	Zscaler Internet Access (Zscaler)
ZPA	Zscaler Private Access (Zscaler)

Trademark Notice

© 2024 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at zscaler.com/legal/trademarks are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.

About This Document

The following sections describe the organizations and requirements of this deployment guide.

Zscaler Overview

Zscaler (NASDAQ: [ZS](#)) enables the world's leading organizations to securely transform their networks and applications for a mobile and cloud-first world. Its flagship Zscaler Internet Access (ZIA) and Zscaler Private Access (ZPA) services create fast, secure connections between users and applications, regardless of device, location, or network. Zscaler delivers its services 100% in the cloud and offers the simplicity, enhanced security, and improved user experience that traditional appliances or hybrid solutions can't match. Used in more than 185 countries, Zscaler operates a massive, global cloud security platform that protects thousands of enterprises and government agencies from cyberattacks and data loss. To learn more, see [Zscaler's website](#) or follow Zscaler on Twitter @zscaler.

AWS Overview

Amazon Web Services (AWS) (NASDAQ: [AMZN](#)) is the world's most comprehensive and broadly adopted cloud platform, offering over 200 fully featured services from data centers globally. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—are using AWS to lower costs, become more agile, and innovate faster. To learn more, refer to [Amazon's website](#).

Audience

This guide is for network administrators, endpoint and IT administrators, and security analysts responsible for deploying, monitoring, and managing enterprise security systems. For additional product and company resources, see:

- [Zscaler Resources](#)
- [AWS Resources](#)
- [Appendix A: Requesting Zscaler Support](#)

Software Versions

This document was authored using the latest version of Zscaler software.

Request for Comments

- **For prospects and customers:** Zscaler values reader opinions and experiences. Contact partner-doc-support@zscaler.com to offer feedback or corrections for this guide.
- **For Zscaler employees:** Contact z-bd-sa@zscaler.com to reach the team that validated and authored the integrations in this document.

Zscaler and AWS Introduction

Overviews of the Zscaler and AWS applications are described in this section.



If you are using this guide to implement a solution at a government agency, some of the content might be different for your deployment. Efforts are made throughout the guide to note where government agencies might need different parameters or input. If you have questions, contact your Zscaler Account team.

ZIA Overview

ZIA is a secure internet and web gateway delivered as a service from the cloud. Think of ZIA as a secure internet on-ramp—just make Zscaler your next hop to the internet via one of the following methods:

- Setting up a tunnel (GRE or IPSec) to the closest Zscaler data center (for offices).
- Forwarding traffic via our lightweight Zscaler Client Connector or PAC file (for mobile employees).

No matter where users connect—a coffee shop in Milan, a hotel in Hong Kong, or a VDI instance in South Korea—they get identical protection. ZIA sits between your users and the internet and inspects every transaction inline across multiple security techniques (even within SSL).

You get full protection from web and internet threats. The Zscaler cloud platform supports Cloud Firewall, IPS, Sandboxing, DLP, and Isolation, allowing you to start with the services you need now and activate others as your needs grow.

Zscaler Resources

The following table contains links to Zscaler resources based on general topic areas.

Name	Definition
ZIA Help Portal	Help articles for ZIA.
Zscaler Tools	Troubleshooting, security and analytics, and browser extensions that help Zscaler determine your security needs.
Zscaler Training and Certification	Training designed to help you maximize Zscaler products.
Submit a Zscaler Support Ticket	Zscaler Support portal for submitting requests and issues.

The following table contains links to Zscaler resources for government agencies.

Name	Definition
ZIA Help Portal	Help articles for ZIA.
Zscaler Tools	Troubleshooting, security and analytics, and browser extensions that help Zscaler determine your security needs.
Zscaler Training and Certification	Training designed to help you maximize Zscaler products.
Submit a Zscaler Support Ticket	Zscaler Support portal for submitting requests and issues.

AWS Overview

Amazon Web Services (AWS) is the world's most comprehensive and broadly adopted cloud, offering over 200 fully featured services from data centers globally. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—are using AWS to lower costs, become more agile, and innovate faster.

AWS Resources

The following table contains links to AWS support resources.

Name	Definition
AWS Account	Create an AWS account
AWS Support	Support for all Amazon Web Services.
AWS Community	AWS online community forum.

Overview

This guide helps AWS users to enable and deploy ZIA for AWS tenants. After configuring AWS to work with Zscaler, AWS traffic passes through Zscaler's cloud and Zscaler enforces security policies on AWS traffic.

The guide demonstrates and explains how to implement ZIA functionality while leveraging the AWS cloud. Policy is enforced and audited for the following:

1. Integration of ZIA services in AWS:
 - a. Cloud App Control
 - b. File Type Control and Data Loss Prevention (DLP)
 - c. Firewall Control
 - d. DNS Control
2. ZIA components that work on AWS infrastructure:
 - a. Virtual Service Edge
 - b. Cloud Connector
 - c. DLP Incident Receiver
 - d. DLP Exact Data Match (EDM) Index Tool
 - e. Amazon Workspaces supporting Zscaler Client Connector:
 - i. Microsoft Windows
 - ii. Ubuntu for AWS
3. ZIA integrations inside AWS:
 - a. Cloud NSS to S3
 - b. Workflow Automation
 - c. SaaS Security API (S3)

Cloud App Control Policy

The following section describes Zscaler Cloud App Control. To learn more, see [About Cloud App Control](#) (government agencies, see [About Cloud App Control](#)).

The Cloud App Control policy provides granular control over popular websites and applications. Policies are organized by function into categories for easy reference and to define rules for similar apps.

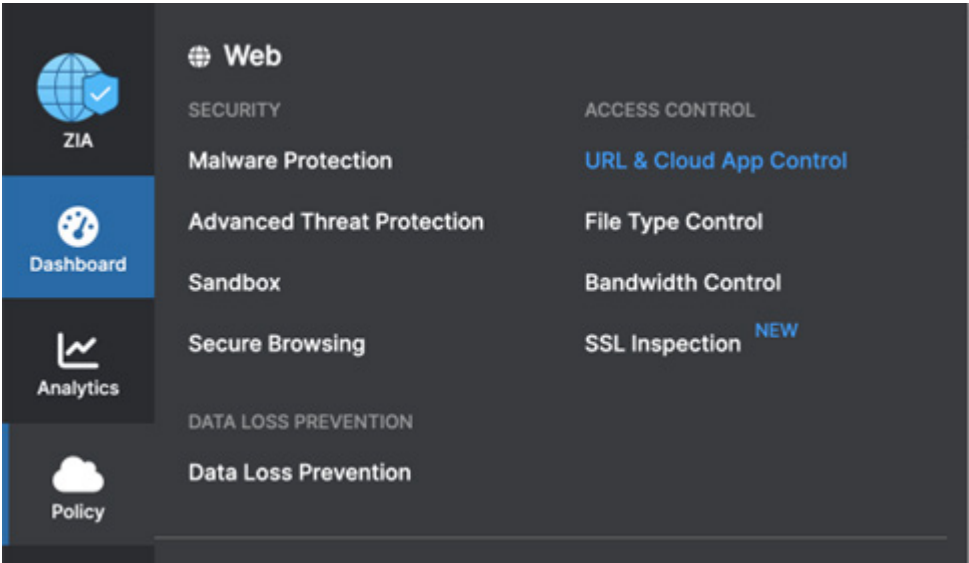


Figure 1. Cloud App Control

All polices can have the following actions:

- Allow: Allows traffic.
- Caution: Allows traffic, but provides the user a caution message before they continue.
- Block: Denies access.
- Isolate: Launches a web browser in a Zscaler cloud that runs the application in isolation (normally, the process runs locally).

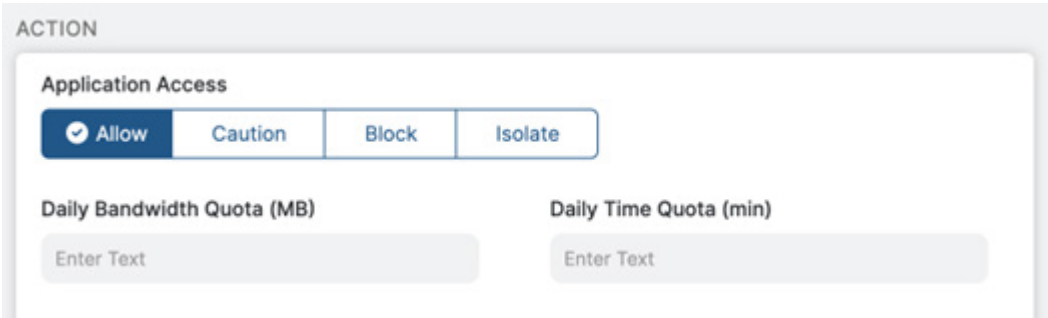


Figure 2. Policy Actions

You can also provide a Bandwidth Quota or Daily Time Quota. These are useful when bandwidth is costly or limited.

You can add a rule for Amazon Chime under the Criteria section.

Figure 3. Add meetings rule

When a user attempts to access Amazon Chime, they are blocked (since the block is enabled). The following shows the blocked access message.

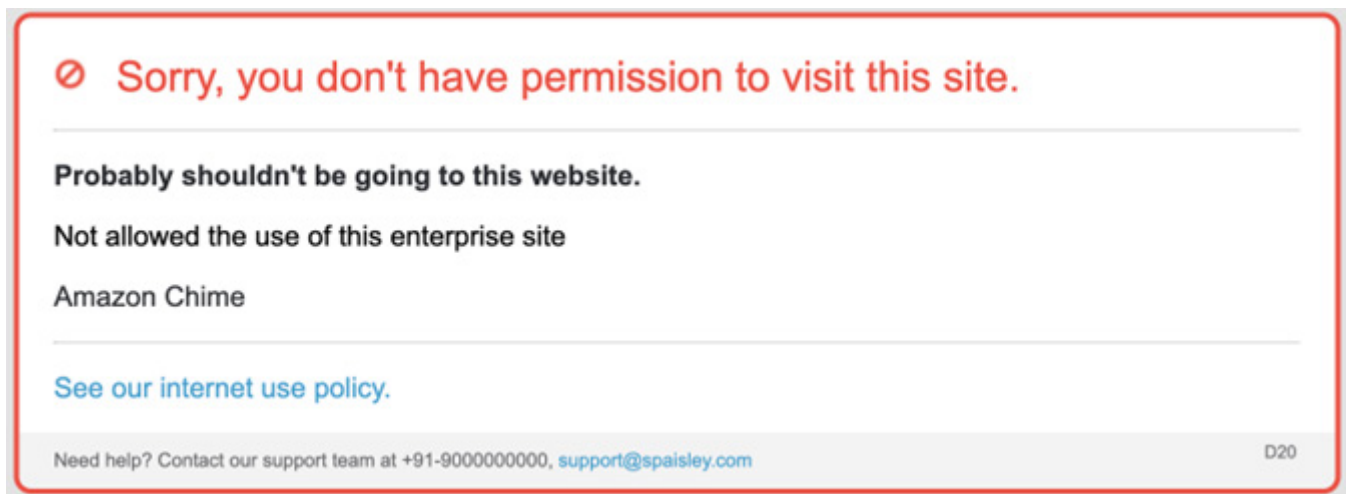


Figure 4. Blocked access message

Another example of Cloud App Control used as policy enforcement is a rule to limit access to AWS Cloud Financial Management.

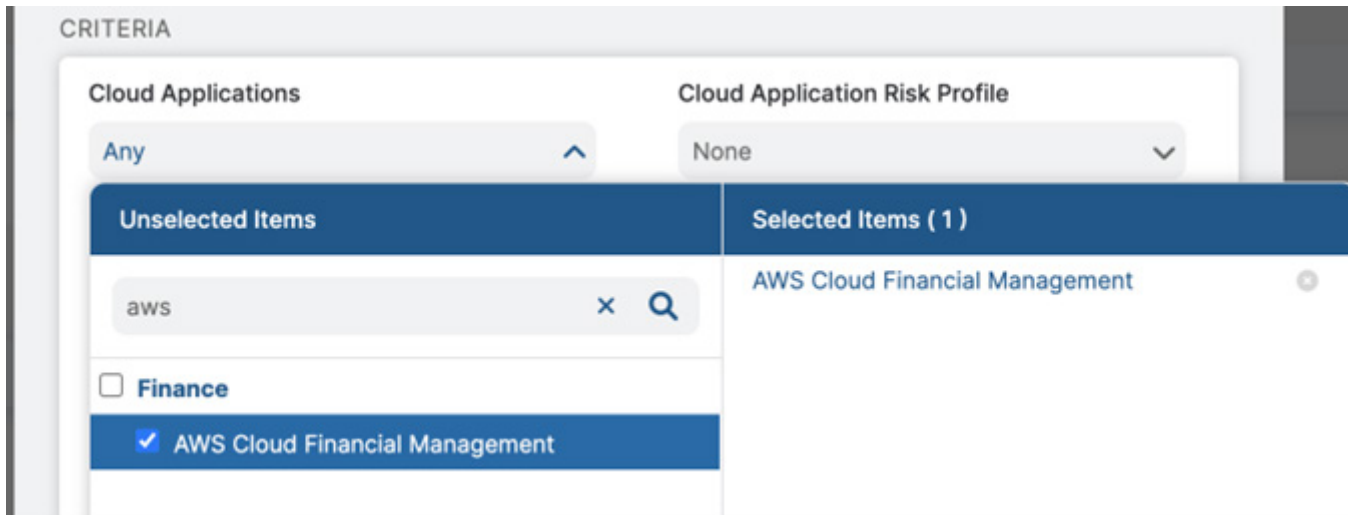


Figure 5. Policy criteria

This removes access from a user that should not have access to AWS Cloud Financial Management. If a user is on a remote network, you could use Isolation to help isolate any threats from the remote network or prevent a user from cutting and pasting sensitive corporate information (such as usage statistics).

You can create a policy for individuals or a group of users. You can Allow, Caution (which provides the user a caution message before they choose to continue), Block (deny access) or Isolate. Isolate launches a web browser in a Zscaler cloud that runs the application in isolation.

Cloud App Control Policies Available via Individual Amazon Web Services

The following is a table of all the individual Amazon Web Services available for the Cloud App Control policies.

AWS Service	Definition
AWS Auto Scaling	AWS Auto Scaling monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost.
Amazon Braket	Amazon Braket is a fully managed quantum computing service designed to help speed up scientific research and software development for quantum computing.
Amazon Chime	Meet, chat, and place business phone calls with a single, secure application.
Amazon Cloud Directory	Amazon Cloud Directory enables you to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions.
Amazon CloudSearch	Amazon CloudSearch is a managed service in the AWS Cloud that makes it simple and cost-effective to set up, manage, and scale a search solution for your website or application.
Amazon DynamoDB	Fast and flexible NoSQL database service for any scale.
Amazon Elastic Block Store	Easy to use, high performance block storage at any scale.
Amazon Elastic Container Service	Amazon ECS is a fully managed container orchestration service that helps you to more efficiently deploy, manage, and scale containerized applications.
Amazon Elastic Kubernetes Service	Amazon EKS is a managed Kubernetes service that makes it easy for you to run Kubernetes on AWS and on-premises.

AWS Service	Definition
Amazon Elastic Load Balancing	Elastic Load Balancing (ELB) automatically distributes incoming application traffic across multiple targets and virtual appliances in one or more Availability Zones (AZs).
Amazon Elasticsearch Service	Elasticsearch is a distributed search and analytics engine built on Apache Lucene.
Amazon EMR	Amazon EMR is the industry-leading cloud big data solution for petabyte-scale data processing, interactive analytics, and machine learning using open-source frameworks such as Apache Spark, Apache Hive, and Presto.
Amazon EventBridge	Build event-driven applications at scale using events generated from your applications, integrated SaaS applications, and AWS services.
Amazon Fraud Detector	Build, deploy, and manage fraud detection models without previous machine learning (ML) experience.
Amazon FSx	Amazon FSx makes it cost effective to launch, run, and scale feature-rich, high-performance file systems in the cloud.
Amazon Kendra	Find information faster with an intelligent enterprise search service powered by ML.
Amazon Lightsail	Get started for free with Amazon Lightsail, a powerful virtual cloud server built for reliability and performance.
Amazon Advertising Console	The advertising console is a self-service tool used to set up and manage sponsored ads campaigns.
Amazon MSK	With Amazon Managed Streaming for Apache Kafka (Amazon MSK), you can ingest and process streaming data in real time with fully managed Apache Kafka.
Amazon Partner Network	The AWS Partner Network (APN) is a global community of partners that leverages programs, expertise, and resources to build, market, and sell customer offerings.
Amazon S3	Amazon Simple Storage Service (Amazon S3) is an object storage service offering industry-leading scalability, data availability, security, and performance.
Amazon SES	Amazon Simple Email Service (Amazon SES) lets you reach customers confidently without an on-premises Simple Mail Transfer Protocol (SMTP) email server using the Amazon SES API or SMTP interface.
Amazon Simple Queue Service	Amazon Simple Queue Service (Amazon SQS) lets you send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available.
Amazon SNS	Amazon Simple Notification Service (Amazon SNS) sends notifications two ways: application-to-application (A2A) and application-to-person (A2P).
Amazontrust	Amazon Trust Services is a certificate authority created and operated by Amazon Web Services.
Amazon WorkDocs	Amazon WorkDocs is a fully managed platform for creating, sharing, and enriching digital content.
AWS Data Exchange	AWS Data Exchange makes the world's third-party data easy to find in one data catalog, simple to subscribe to, and seamless to use with any AWS data and analytics and ML services.
AWS Identity and Access Management	With AWS Identity and Access Management (IAM), you can specify who or what can access services and resources in AWS, centrally manage fine-grained permissions, and analyze access to refine permissions across AWS.
AWS Key Management Service	AWS Key Management Service (AWS KMS) lets you create, manage, and control cryptographic keys across your applications and Amazon services.

AWS Service	Definition
AWS Managed Services	AWS Managed Services (AMS) helps you adopt AWS at scale and operate more efficiently and securely.
AWS Network Firewall	With AWS Network Firewall, you can define firewall rules that provide fine-grained control over network traffic.
AWS Resource Access Manager	AWS RAM helps you securely share your resources across AWS accounts, within your organization or organizational units (OUs), and with IAM roles and users for supported resource types.
AWS Snow Family	Process data at the edge or move petabytes of data to and from AWS.
AWS Storage Gateway	AWS Storage Gateway is a set of hybrid cloud storage services that provide on-premises access to virtually unlimited cloud storage.
AWS VPN	Connect your on-premises networks and remote workers to the cloud.



The Cloud App Control section demonstrates only one topic to show how to create the policy. All policies are very similar. They are included here to provide a searchable list of AWS-supported features and functions that can be enforced and viable with ZIA.

There is also one more Cloud App for all of Amazon Web Services. You can use this category with Tenant restrictions. This enables you to enable specific tenant IDs. For example, you could allow office users, but not personal accounts.



Special Note for the Amazon Web Services category: All the sections for Cloud App Control enable you to provide policy to all the Amazon Cloud Applications. However, you can enable an additional restriction called Tenant Restriction for the Hosting Providers section, which includes Amazon Web Services as a whole. This allows you to provide access only to specific tenant IDs.

Tenant restrictions are useful if you want to restrict a user or device to only be able to access AWS from a corporate account, for example. Thus, if a user has their own AWS account, they cannot access AWS.

[Zscaler's tenancy restriction](#) (government agencies, see [Zscaler's tenancy restriction](#)) feature allows you to restrict access either to personal accounts, business accounts, or both for AWS. It consists of two parts: creating tenant profiles and associating them with the Cloud App Control policy rules.

File Type Control for AWS

Zscaler File Type Control enables organizations to regulate and monitor the types of files that you can upload, download, or transfer for AWS, Chime, and S3 buckets. The feature allows administrators to define policies that restrict or allow specific file types, thereby preventing the transmission of potentially harmful or non-compliant files.

You can create File Type Control policies for the same Amazon services shown in the [Cloud App Control Policies Available via Individual Amazon Web Services](#).

To add the File Type Control policy, go to **Policy > File Type Control**.

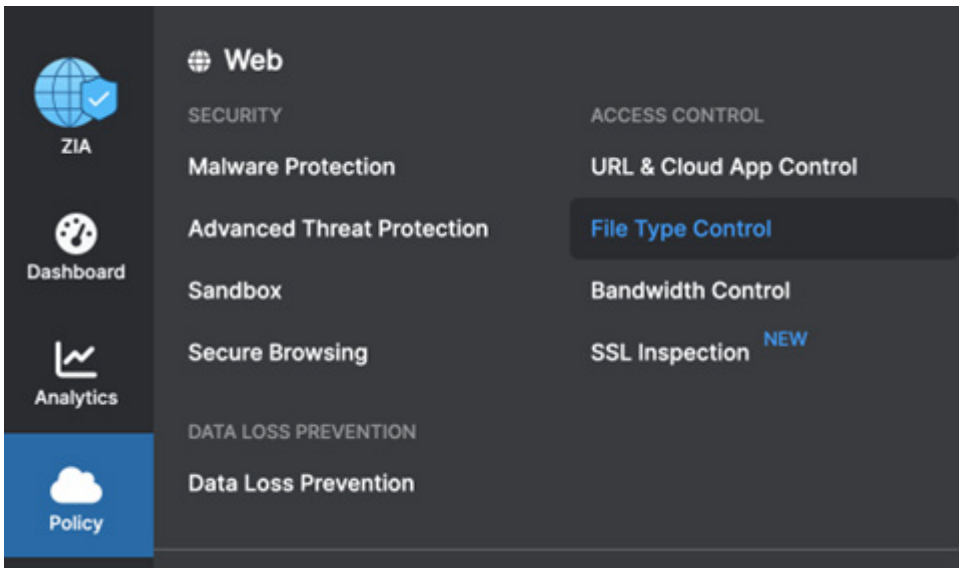


Figure 6. File Type Control

The following is an example of such a policy to prevent and block any ZIP files from being uploaded to any of the services listed earlier:

No Zip files UL or DL to AWS	ACTIVE CONTENT	Block Upload/Download
	Disabled	
	CLOUD APPLICATIONS	
	Amazon Macie; Amazon Partner Central; Amazon Asin; Amazon DynamoDB; ...	
	FILE TYPES	
	ZIP (zip)	
UNSCANNABLE FILE		
Disabled		
PROTOCOLS		
FTP over HTTP; Native FTP; HTTPS; HTTP		

Figure 7. File Type Control policy

As with File Type Control, you can add Data Loss Prevention (DLP) policies specific to all the Amazon Web Services. You can add the sections by adding a DLP rule similar to the following image:

Add DLP Rule [X]

CRITERIA

DLP Engines ClassificationConfidential; Credit Cards [v]	URL Categories Any [v]
Cloud Applications Amazon - Elastic Container Service; Am... [v]	ZPA Application Segment Any [v]
File Type Any [v]	Minimum Data Size (KB) 0
Users Any [v]	Groups Any [v]
Departments Any [v]	User Risk Profile Any [v]
Locations Any [v]	Location Groups Any [v]
Time Always [v]	Protocols HTTP; HTTPS; Native FTP [v]

DLP INCIDENT RECEIVER

Incident Receiver

ICAP [] Zscaler Incident Receiver [x]

Figure 8. Add DLP Rule

You can send the DLP violation to a DLP Zscaler Incident Receiver, which can run on an Amazon EC2 instance. You can send a DLP violation file to the AWS customer cloud instance for later review. To learn more, see [ZIA Components that Work on AWS Infrastructure](#).

Firewall Control Rules for AWS

The [Zscaler firewall](#) (government agencies, see [Zscaler firewall](#)) provides protection policies specific to AWS as well as all traffic. The Zscaler firewall service provides integrated cloud-based next-generation firewall capabilities that allow granular control over your organization's outbound TCP, UDP, and ICMP traffic.

As an example, you can create a firewall rule that covers these specific Amazon Web Services. You can create a specific firewall rule that combines the Who, Where, When, Services, Applications, Source IP, and Destination IP.

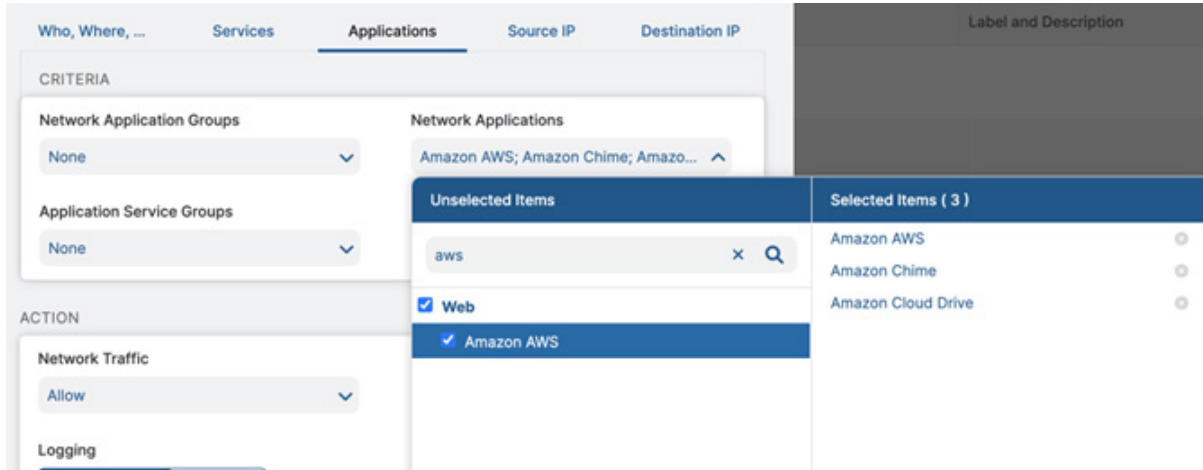


Figure 9. Firewall policy

DNS Control

Zscaler DNS Control monitors and applies policies to all DNS requests. It can also make specific DNS rules to apply specifically to AWS traffic.

DNS Control provides the following benefits:

- Monitor and apply policies to all DNS requests and responses, regardless of the protocol and the encryption used. This includes UDP, TCP, and DNS over HTTPS (DoH).
- Define granular DNS filtering rules using several DNS conditions such as users, groups, departments, client locations, categorization of domains and IP addresses, DNS record types, the location of resolved IPs, etc.
- Enforce condition-based actions on DNS traffic, such as allowing or blocking traffic, redirecting requests to specific DNS servers, redirecting users by overwriting DNS responses, etc.
- Detect and prevent DNS-based attacks and data exfiltration through DNS tunnels.
- Enhance your security posture by using Zscaler Trusted DNS Resolver for domain resolution.

You can apply your Zscaler DNS Control rules specifically to Amazon and Amazon AWS traffic or all traffic.

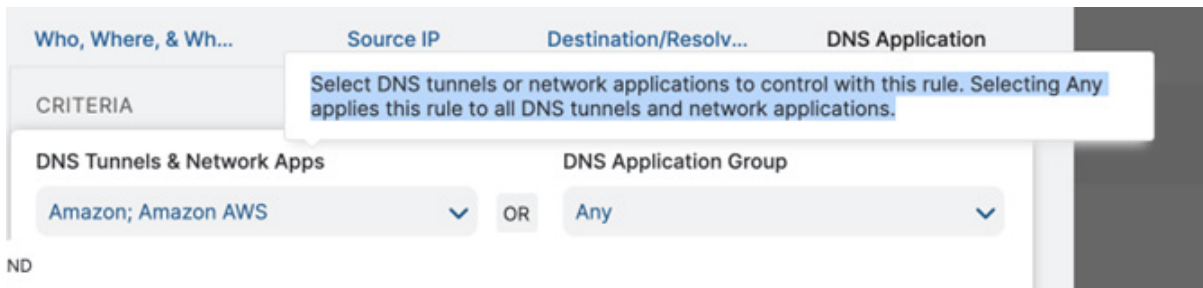


Figure 10. DNS policy

ZIA Components that Work on AWS Infrastructure

The following services can run directly inside the AWS cloud. You can acquire some services from the AWS marketplace or install the services directly on EC2 instances. Each provide a unique solution to provide Zscaler ZIA cloud services inside the AWS cloud.

The following diagram shows the integrations of NSS, Cloud Connector, and the Virtual Service Edge running on AWS. In addition, this guide also covers two DLP instances and the Zscaler Client Connector, which can run on Amazon WorkSpaces.

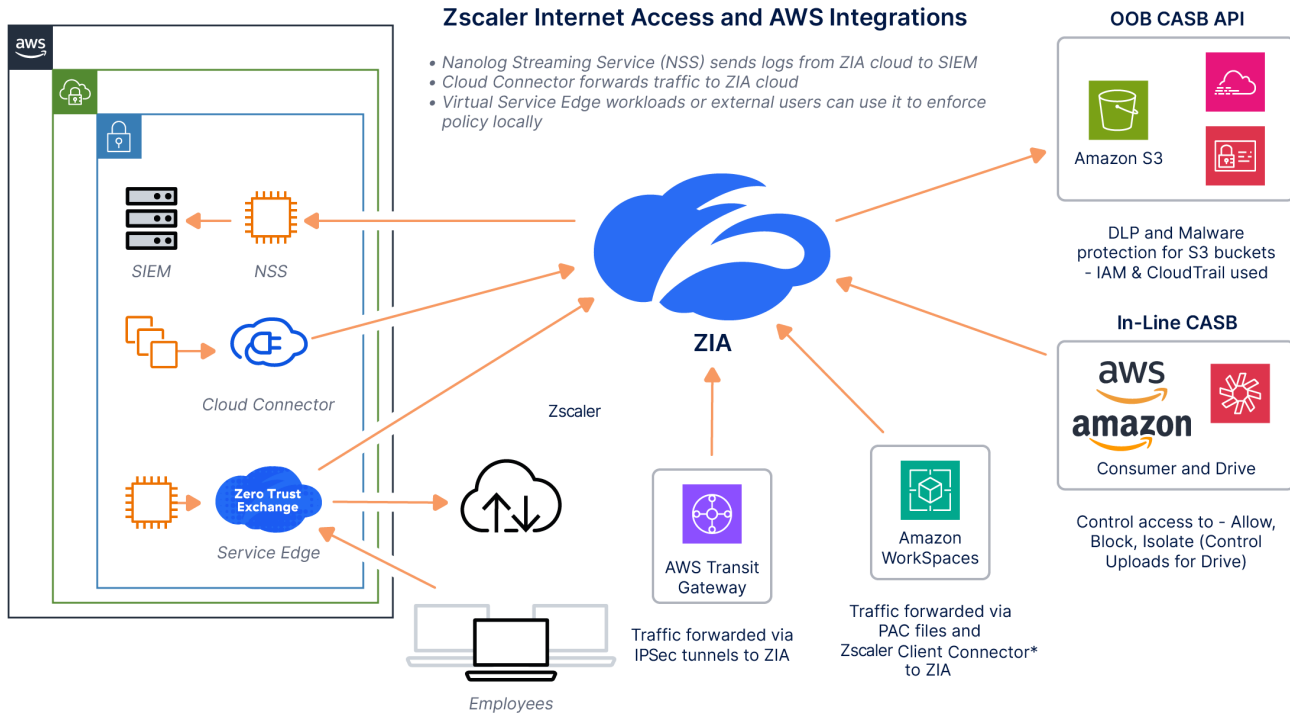


Figure 11. ZIA and AWS integration

Nanolog Streaming Service

Zscaler Nanolog Streaming Service (NSS) provides a method for streaming of all logs from Zscaler Nanolog to your security information and event management (SIEM) system.

You can deploy the NSS instance directly on an EC2 instance on AWS. When an organization deploys one NSS for web and mobile logs and another NSS for firewall logs, each NSS opens a secure tunnel to Nanolog in the Zscaler cloud. Nanolog then streams copies of the logs to each NSS in a highly compressed format to reduce bandwidth footprint. The original logs are retained on Nanolog. To learn more, see [NSS deployment documentation for AWS](#) (government agencies, see [NSS deployment documentation for AWS](#)).

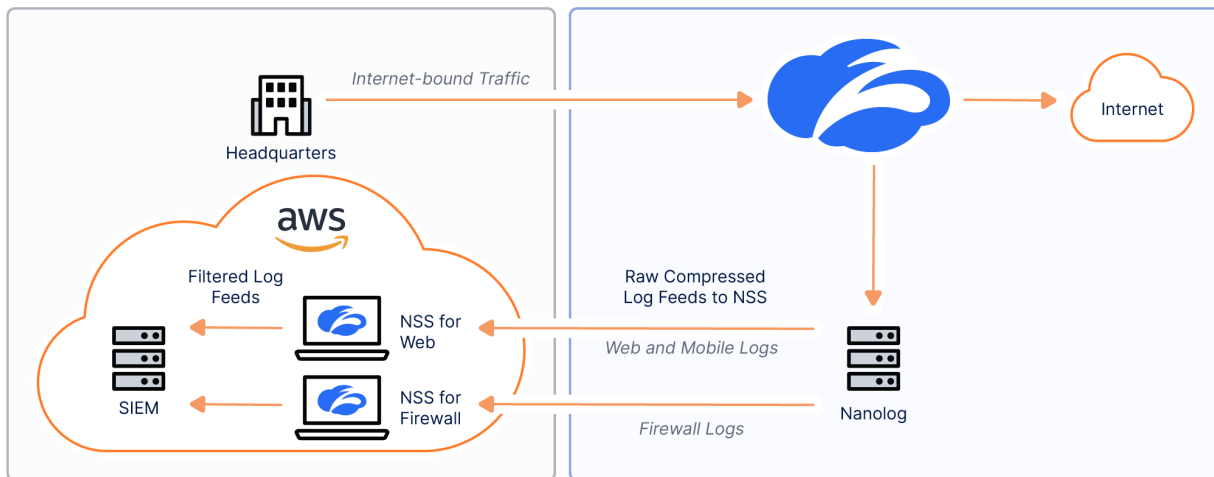


Figure 12. Zscaler NSS and AWS integration

Virtual Service Edge

Zscaler supports standalone ZIA Virtual Service Edge for production deployments on AWS. An organization can deploy the Virtual Service Edge instance on an EC2 Instance. The Virtual Service Edge acts as an extension of the Zscaler data centers into the AWS cloud itself, which keeps traffic local and ensures that IP address ranges remain local. This helps with IP anchoring, where remote sites require specific IP addresses.

To learn more, see [Zscaler Virtual Service Edge for AWS](#) (government agencies, see [Zscaler Virtual Service Edge for AWS](#)).

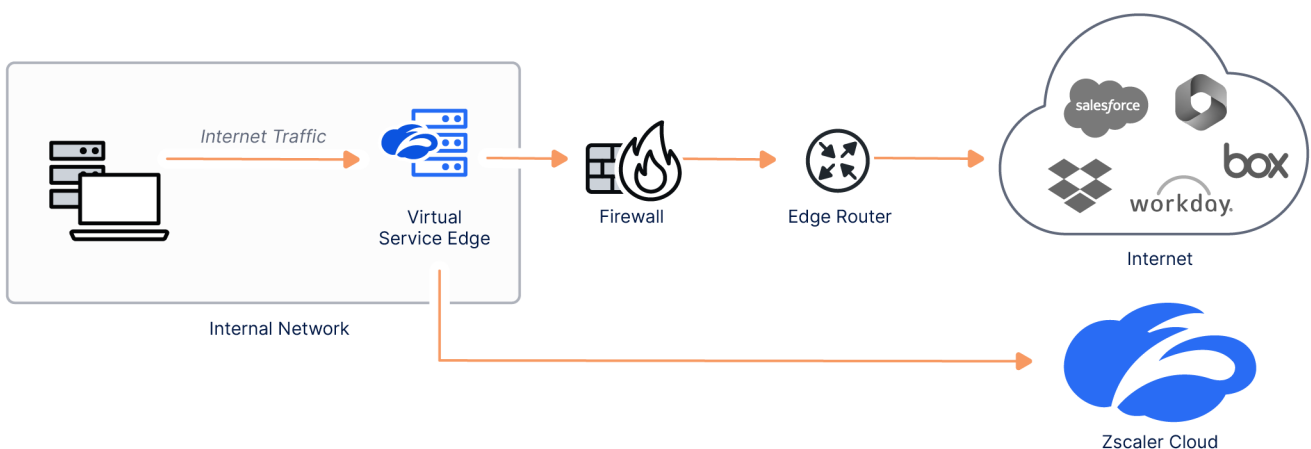


Figure 13. Virtual service edge

Cloud Connector

Cloud Connector ensures that cloud workloads adhere to organizational security policy when accessing both public and private endpoints. Cloud Connector intelligently forwards traffic to the ZIA and ZPA platforms. Cloud Connector also enables multi-cloud connectivity and enforces a security policy for cloud-to-cloud traffic. Cloud Connector identifies egress traffic and sends it to the Zscaler Zero Trust Exchange (ZTE) without the need for the network components behind the Cloud Connector to have their own configuration.

To learn more, see [Cloud Connector Reference Architecture](#) and [Step-by-Step Configuration Guide for Zscaler Cloud Connector](#) (government agencies, see [Cloud Connector Reference Architecture](#) and [Step-by-Step Configuration Guide for Zscaler Cloud Connector](#)).

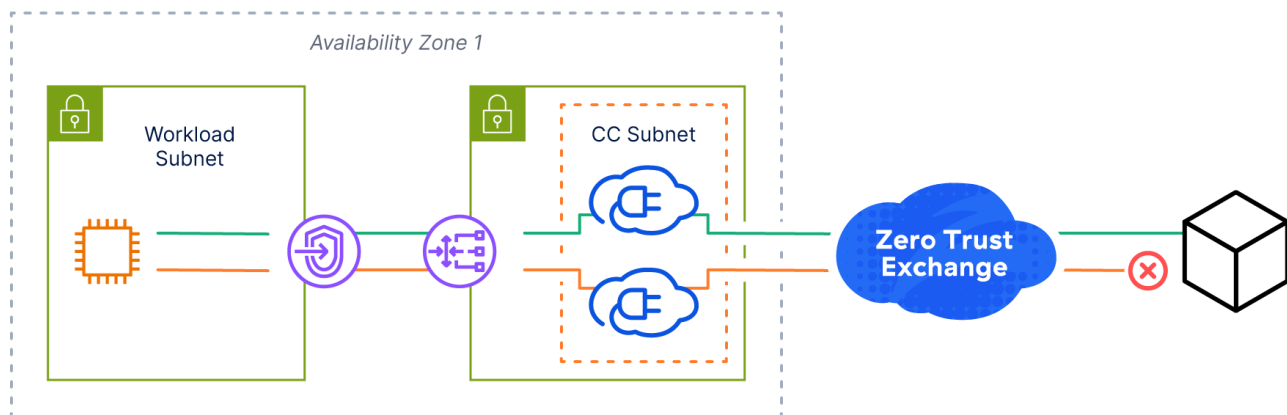


Figure 14. Cloud Connector

When traffic has reached the Cloud Connector, there are four Traffic Forwarding options available to direct traffic out of the AWS cloud:

- Direct: Traffic matching the criteria defined bypasses the Cloud Connector and is routed out of the service interface, where it follows AWS route tables towards the destination.
- Zscaler Internet Access (ZIA): Traffic matching the criteria defined is forwarded to the ZIA cloud for inspection.
- Zscaler Private Access (ZPA): Traffic matching the criteria defined is forwarded to the ZPA cloud for inspection.
- Drop: Traffic matching the criteria is dropped by the Cloud Connector.

Each of the four options permits the administrator to define a range of match criteria. In general, you can define macro forwarding logic within the Cloud & Branch Connector Portal, whereas ZIA or ZPA can perform more granular inspection.

DLP Incident Receiver

The Zscaler Incident Receiver runs as an EC2 instance, and allows you to securely receive information about DLP policy violations. The Zscaler service sends information about policy violations via the secure ICAP protocol to the Incident Receiver. This tool sends the policy-violating content and a JSON file containing the metadata for the inline web and DLP policy scan (e.g., the URL, Collaborators, DLP dictionaries, DLP engines, etc.)

To learn more, see [AWS Incident Receiver Installation](#) documents (government agencies, see [AWS Incident Receiver Installation](#)).

DLP Index Tool

The Zscaler Index Tool allows you to create and modify Exact Data Match (EDM) and Indexed Document Match (IDM) index templates, as well as see a dashboard view of your EDM and IDM index templates.

To learn more, see [Configuring the Index Tool with AWS](#) (government agencies, see [Configuring the Index Tool with AWS](#)).

Amazon WorkSpaces Supporting Zscaler Client Connector

The Zscaler Client Connector is an agent software that runs on an OS such as Windows or Ubuntu. It is part of Zscaler's cloud security platform, designed to provide seamless and secure access to the internet and corporate resources for users, regardless of their location.

This software solution acts as a secure gateway, routing traffic through the Zscaler cloud, which enables advanced threat protection and policy enforcement. The Zscaler Client Connector ensures consistent security and policy enforcement, making it a very useful tool to deploy in Amazon WorkSpaces. Currently, Zscaler supports the Zscaler Client Connector on Microsoft Windows and Ubuntu for AWS.

To learn more, see [Installing the Zscaler Client Connector](#) (government agencies, see [Installing the Zscaler Client Connector](#)).

ZIA Integrations Inside AWS

The following sections detail integrating ZIA inside of AWS.

Cloud NSS and S3 Buckets

This integration enables organizations to stream Zscaler ZIA logs directly to Amazon S3. Zscaler's logs are conveniently stored in S3 buckets, facilitating streamlined monitoring, auditing, and compliance reporting. This comprehensive approach to cloud security bolsters protection and simplifies management and compliance efforts. To learn more, see [AWS S3 Zscaler SaaS Deployment Guide](#) (government agencies, see the [AWS S3 Zscaler SaaS Deployment Guide](#)).

Workflow Automation

Workflow Automation is an application that enables governance analysts to manage and resolve the different Data Protection incidents that occur in their organization. Workflow Automation integrates with ZIA to capture those Data Protection incidents generated from the different DLP policies defined in ZIA.

To learn more, see [Configuring the DLP Application Integration Using Amazon Web Services](#).

AWS integration requires three AWS resources:

- S3 Bucket: The S3 bucket names share a common prefix.
- SNS Topic: The metadata S3 bucket pushes notifications to the SNS topic which is subscribed by the Workflow Automation SQS Queue.
- Cross Account IAM Role: The cross-account IAM role allows read-only access to the Workflow Automation AWS account to the data and metadata buckets.

SaaS Security API for S3 Buckets

Zscaler's SaaS Security API for AWS S3 enhances organizations' cloud security posture on the AWS platform. Leveraging the power of Zscaler's extensive security research and DLP technologies, this API secures data and applications hosted in AWS environments. It provides real-time threat protection, data loss prevention, and secure access controls, ensuring that businesses can maintain the highest level of security while embracing the scalability and flexibility of AWS.

To learn more, see [AWS S3 Zscaler SaaS Deployment Guide](#) (government agencies, see the [AWS S3 Zscaler SaaS Deployment Guide](#)).


There are two major reasons to implement this API for AWS:

1. Scanning for DLP violations: The [SaaS Security API Data Loss Prevention \(DLP\) policy](#) (government agencies, see [SaaS Security API Data Loss Prevention \(DLP\) policy](#)) allows you to create rules to discover and protect sensitive data at rest in an Amazon S3 bucket. You can configure criteria, such as file type or collaboration scope, to specify the type of content for the policy to scan. You can also configure actions for the policy to take if it detects content that matches the criteria. This is also available as a service of the SaaS Security API for S3 buckets.
2. Scanning for Malware Threats: With the Zscaler SaaS Security API you can scan your AWS S3 bucket for threats and malware. This ensures your S3 buckets are free from malware and have not been compromised. Adding a malware policy for a SaaS application provides the benefit to maintain individualized malware policies for each SaaS application tenant in your organization, and detects and remove malware threats to extend comprehensive web security to your SaaS applications.

To set up SaaS Security API for S3 buckets, ensure you have the S3 tenant enabled in your Zscaler ZIA tenant. If not, you can create a support ticket to request the S3 tenant to be enabled in your instance.

Add SaaS Application Tenant

- Choose the SaaS Application Provider


- Name the SaaS Application Tenant

Tenant Name

The tenant name must be unique
- Onboard SaaS Application for

DLP and Malware scanning SaaS API
- Authorize the SaaS Application

To give Zscaler access to Amazon S3, you must configure an IAM role for the Zscaler S3 Connector.

Zscaler Connector Account Number	Zscaler Connector User ARN	External ID
011284747002 Copy	arn:aws:iam::011284747002:user/ZscalerSaaSConnectorZScloud01 Copy	aOd67ORLLEblgNHQ Copy

[Go to AWS](#)

Figure 15. Add SaaS Application Tenant

After you have enabled the Amazon S3 Tenant on both Amazon and ZIA, you can start enabling policies.

To enable a DLP policy, for the out-of-band scan (SaaS Security API):

- Create a policy in the **Policy > SaaS Security API Control**.

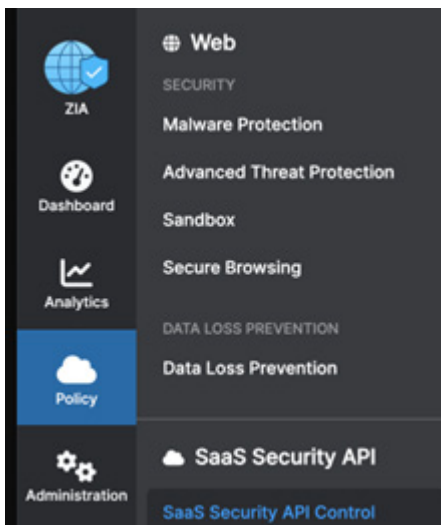


Figure 16. SaaS Security API

The following is an example of an Amazon S3 Tenant policy that scans all configured buckets for a confidential key word, and credit card information that is stored in the S3 bucket. You can alert the DLP policy violation or even remove publicly sharable links if they exist.




Figure 17. SaaS Security API Policy

- 2. Add the DLP rule to scan for DLP violations in Amazon S3 buckets.

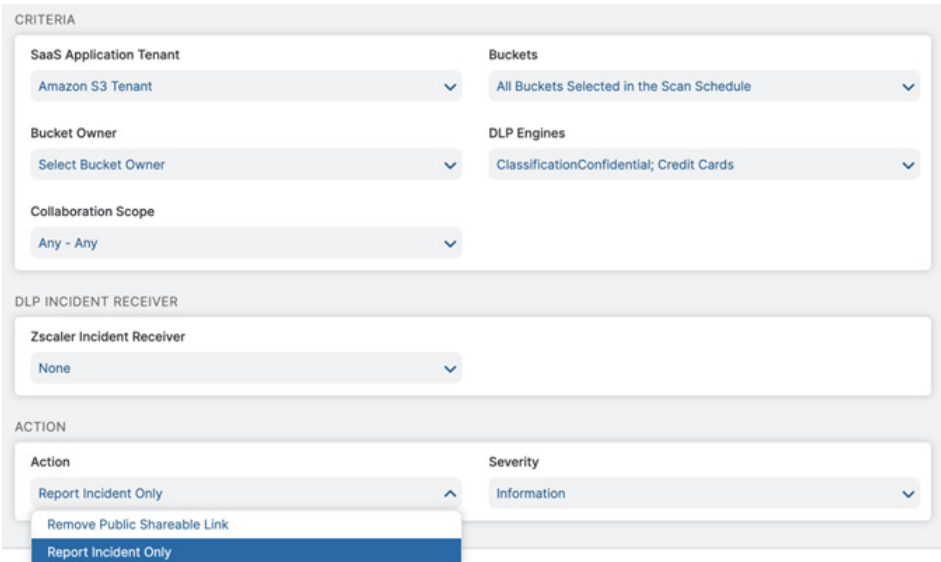


Figure 18. DLP rule

- 3. To search and identify files that contain malware, click the **Malware Detection** tab.

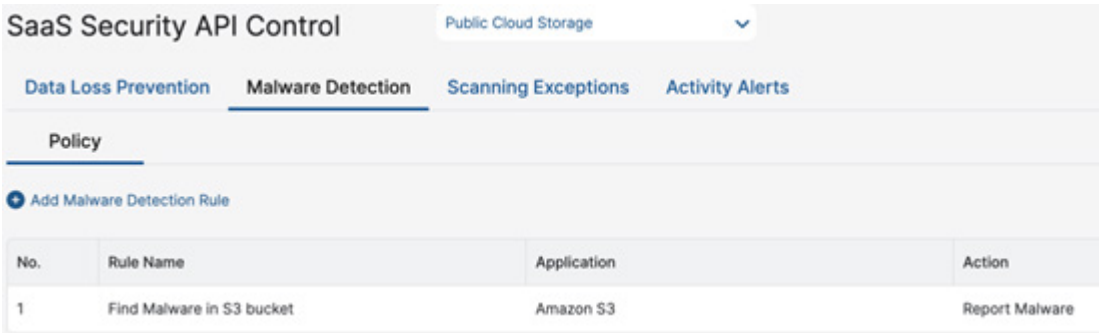


Figure 19. Malware Detection

To learn more, see [About SaaS Security Malware Detection](#) (government agencies, see [About SaaS Security Malware Detection](#)).

Figure 20. Add Malware Detection Rule

4. (Optional) Add Scanning Exceptions and Activity Alerts.

Figure 21. Scanning Exceptions

Zscaler includes four default activity alerts:

- Default Impossible Travel Alert
- Default Multiple Failed Logins Alert
- Default Bulk Upload of Data Alert
- Default Bulk Download of Data Alert

You can alert on many other activities. To learn more, see [About SaaS Security Activity Alerts](#) (government agencies, see [About SaaS Security Activity Alerts](#)).

Contextualizing Risk using AWS and Avalor UVM

Avalor's Data Fabric for Security and Unified Vulnerability Management (UVM) solution integrates, normalizes, and unifies data from various enterprise security and business systems to provide actionable insights, analytics, and operational efficiencies.

Avalor offers preconfigured connectors for the following AWS services, which you can add as Assets:

- EC2
- Relational Database Service (RDS)
- Elastic Container Registry (ECR)
- Elastic Kubernetes Service (EKS) Clusters API
- S3 Buckets
- AWS Accounts

In addition, you can add the following as Findings:

- AWS Inspector Findings
- AWS Security Hub
- AWS Elastic Container Registry (ECR) Findings

The following steps outline how to start ingesting data from these sources, while also (optionally) combining EC2 data with Avalor vulnerability information to provide a more contextualized and personalized risk assessment for your organization.

Creating a Role ARN and an External ID in AWS

This process takes you through creating a Role ARN and External ID for a Single AWS account. To use the alternative options of a Secret Key or Multiple Accounts, refer to the [Avalor documentation](#).

1. Open the [cloudformation.json](#) file and copy its contents into a text editor.
2. Determine which roles ARN permissions you must add to the cloudformation.json file from the following table:

Connector Name	Data Retrieved	Permissions Required
Security Hub API	Findings	securityhub:GetFindings
Inspector Findings	Findings	inspector2:ListFindings
ECR Findings	Findings	ecr:DescribeImageScanFindings
EC2	Resources	ec2:DescribeInstances
Relational Database Service (RDS)	Resources	rds:DescribeDBInstances
Elastic Container Registry (ECR)	Resources	ecr:ListImages ecr:DescribeImages ecr:DescribeRepositories
Elastic Kubernetes Service (EKS) Clusters API	Resources	eks:ListClusters eks:DescribeCluster
S3 Buckets	Resources	s3:ListAllMyBuckets

Connector Name	Data Retrieved	Permissions Required
Accounts	Retrieves your organization's accounts details.	organizations:DescribeAccount organizations:ListAccounts organizations:ListTagsForResource Note: Attach this permission to the root/ organization account.

3. Under the second Action in AvalorPolicy, edit the permissions list to cover those necessary for the data you want to retrieve:

```

"AvalorPolicy": {
  "Properties": {
    "PolicyDocument": {
      "Statement": [
        {
          "Sid": "AllowSQSReceiveMessage",
          "Effect": "Allow",
          "Action": [
            "sqs:ReceiveMessage",
            "sqs>DeleteMessage",
            "sqs:ChangeMessageVisibility"
          ],
          "Resource": "arn:aws:sqs:*:*:*avalor*"
        },
        {
          "Action": [
            "securityhub:GetFindings ",
            "inspector2:ListFindings ",
            "ecr:DescribeImageScanFindings ",
            "ec2:DescribeInstances ",
            "rds:DescribeDBInstances ",
            "ecr:ListImages",
            "ecr:DescribeImages",
            "ecr:DescribeRepositories",

```

```

    "eks:ListClusters",
    "eks:DescribeCluster",
    "s3:ListAllMyBuckets",
    "organizations:DescribeAccount",
    "organizations:ListAccounts",
    "organizations:ListTagsForResource"
  ],
  "Effect": "Allow",
  "Resource": "*"

```

4. Save this CloudFormation file locally as `avalor-aws-connector.json`.
5. Generate a UUID to use in the next step. You can use this [UUID Generator](#).
6. Install the `aws-cli` if it's not installed on your system already. For instructions, refer to the [AWS documentation](#).
7. Run the following CloudFormation Role Stack command:

```

aws cloudformation create-stack \
--region <REGION> \
--stack-name AvalorStackIntegration \
--capabilities CAPABILITY_NAMED_IAM \
--template-body file://avalor-aws-connector.json \
--parameters ParameterKey=ExternalId,ParameterValue=<Generated UUID>

```

Before running the command, ensure:

- a. You replace `<REGION>` with the region of the AWS service from which you're retrieving data.
 - b. The `avalor-aws-connector.json` file is in the present working directory.
 - c. Replace `<Generated UUID>` with the UUID you created in the previous step.
8. Look for the confirmation that the stack was created with a response of a StackID, such as:

```

{
  "StackId": "arn:aws:cloudformation:ap-southeast-2:*****459973:stack/AvalorStackIntegration/*****-****-11ef-bb5b-023b19c7266f"
}

```

Output for the RoleARNID and ExternalID

Run the command `aws cloudformation describe-stacks --stack-name AvalorStackIntegration` to get the RoleARNID and ExternalID. The output includes the following:

```
"Outputs": [  
  {  
    "OutputKey": "RoleARNID",  
    "OutputValue": "arn:aws:iam:: :*****459973:role/  
AvalorAccess-Role",  
    "Description": "Your Role ARN ID"  
  },  
  {  
    "OutputKey": "ExternalID",  
    "OutputValue": "*****-****-****-aac9-b7d80eb9ac6e",  
    "Description": "Your External ID"  
  }  
],
```

Configure the AWS UVM Data Connectors

The following sections describe how to configure AWS UVM data connectors.

Configure the AWS Accounts Data Source

To configure the AWS accounts data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.

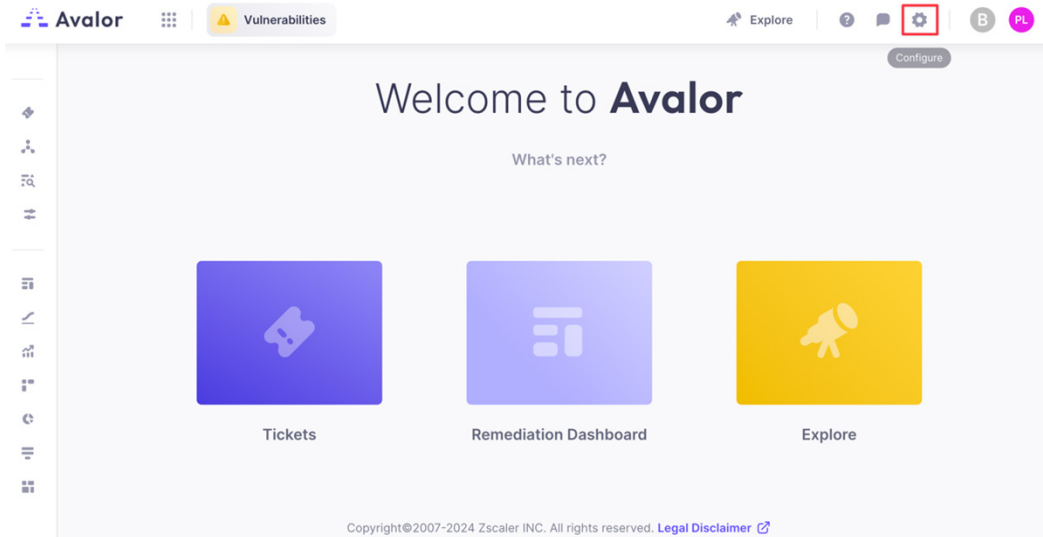


Figure 22. Avalor UVM Platform

3. Click **Create**, then search for AWS Accounts.

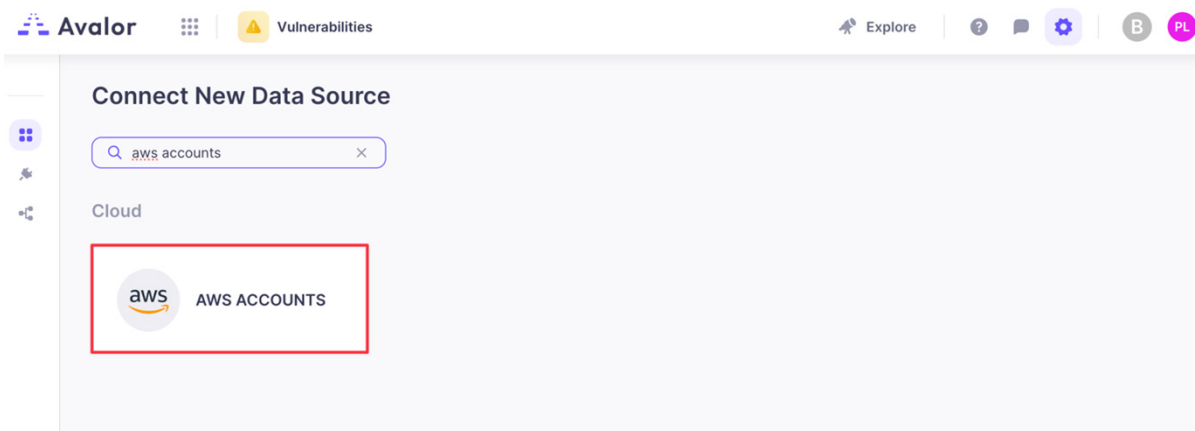


Figure 23. Connect New Data Source

4. Click the **AWS Accounts** application.
5. On the **Create AWS Accounts Source** window, complete the following:
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names to which this data source will apply.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically become undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click Test. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.

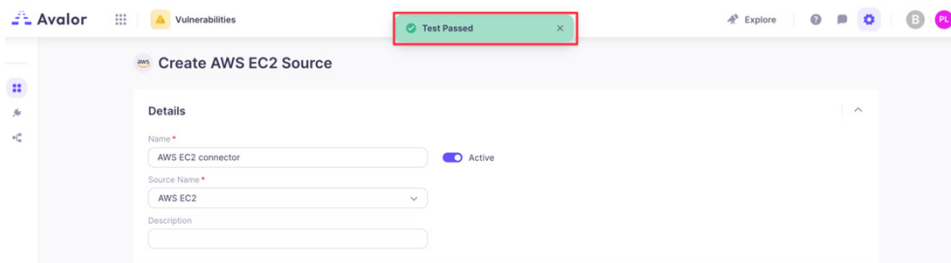


Figure 24. Test Passed

7. Click **Save**.

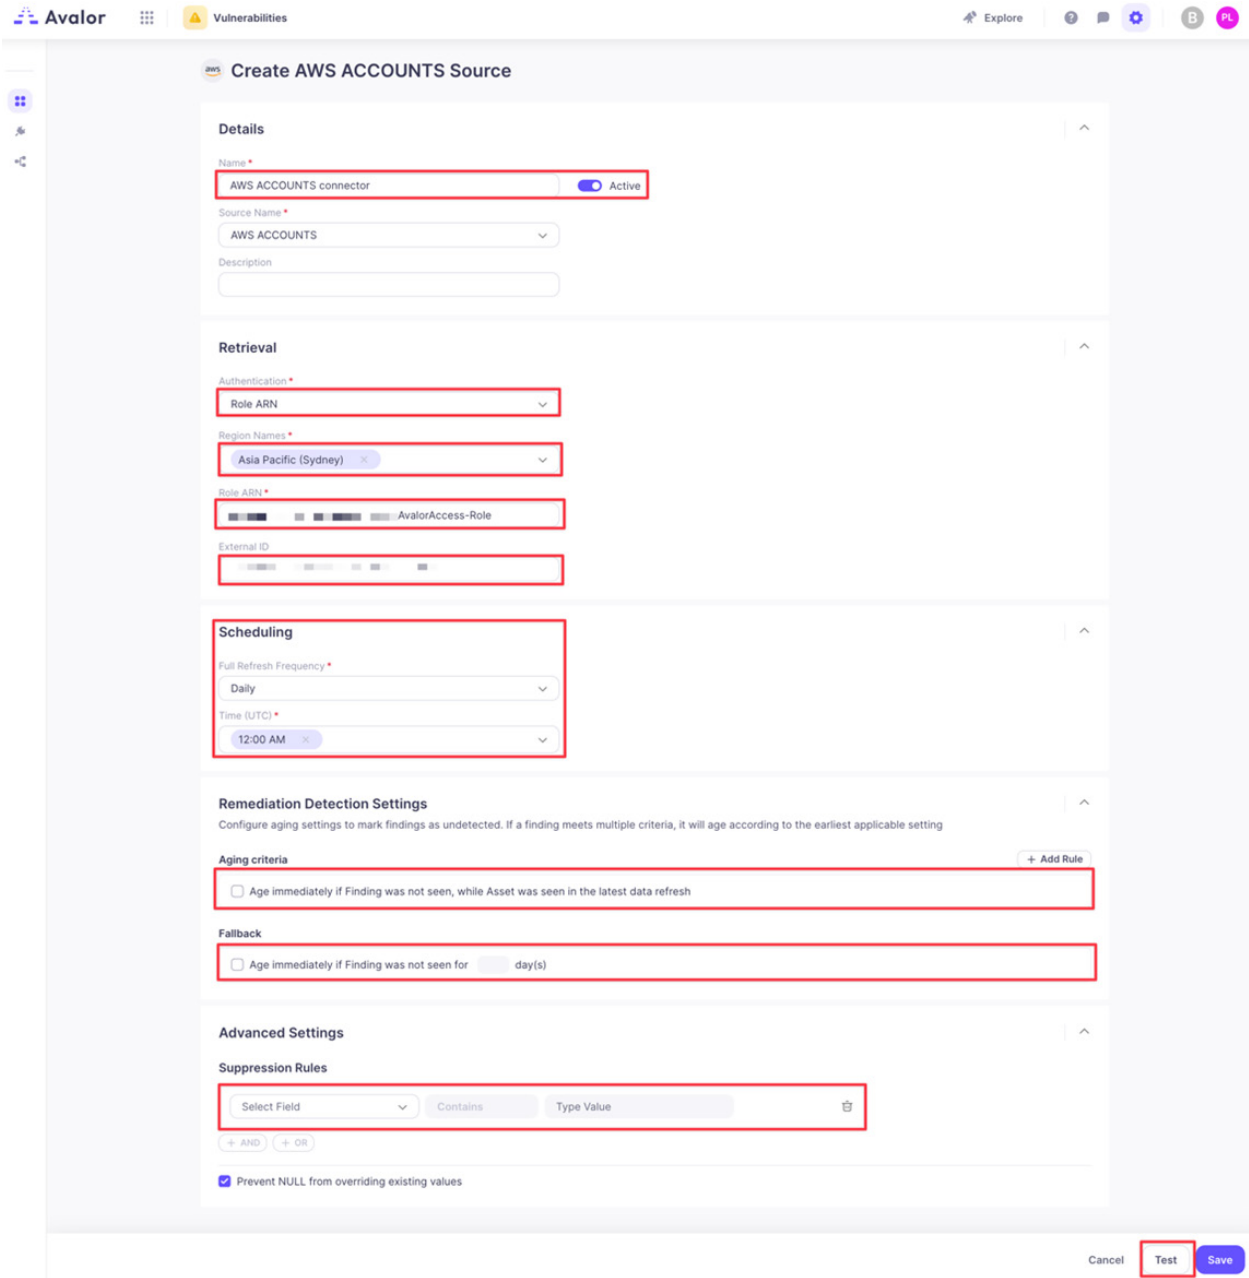


Figure 25. Create AWS Accounts Source

Configure the AWS EC2 Data Source

To configure the AWS EC2 data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.




Figure 26. Avalor UVM Platform

3. Click **Create**, then search for AWS EC2.




Figure 27. Connect New Data Source

4. Click the **AWS EC2** application.
5. On the **Create AWS EC2 Source** page, complete the following:
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names to which this data source applies.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.

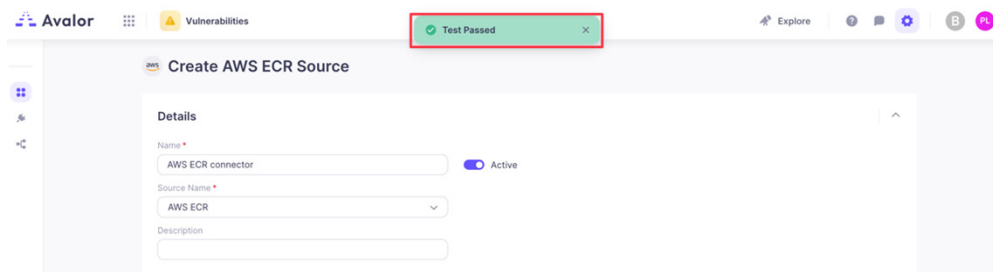


Figure 28. Test Passed

7. Click **Save**.

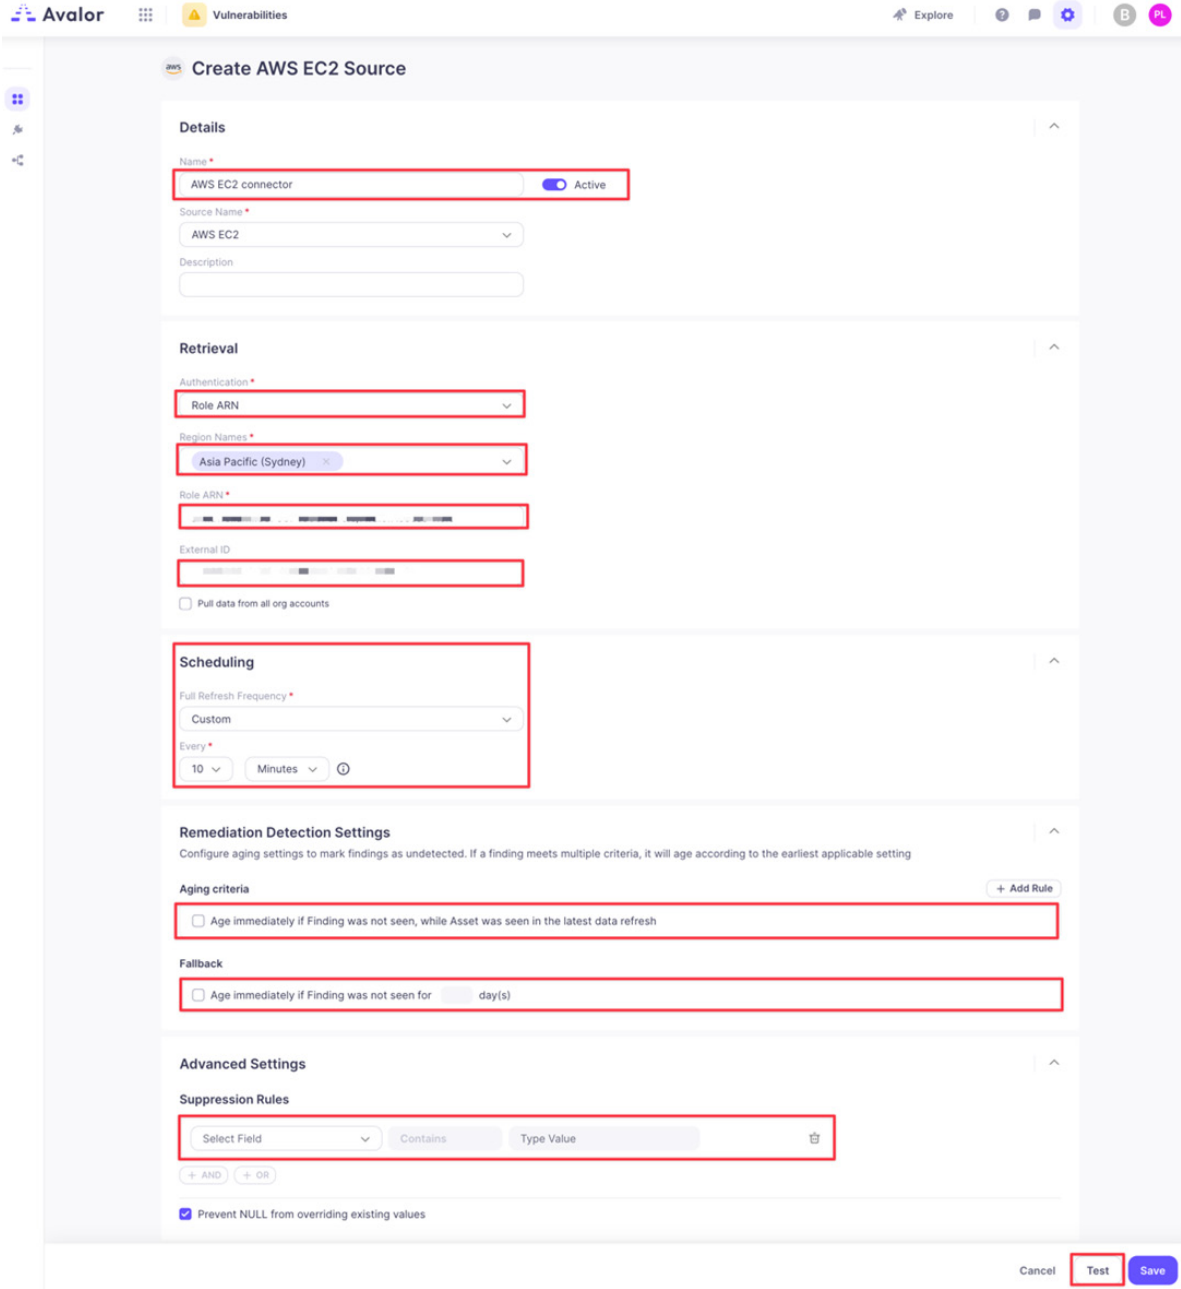


Figure 29. Create AWS EC2 Source

Configure the AWS ECR Data Source

To configure the ZWS ECR data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.




Figure 30. Avalor UVM Platform

3. Click **Create**, then search for AWS ECR.




Figure 31. Connect New Data Source

4. Click the **AWS ECR** application.
5. On the **Create AWS ECR Source** page, complete the following
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names this data source will apply to.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.

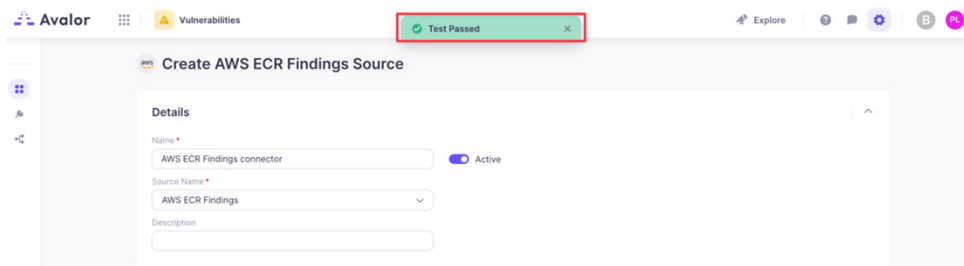


Figure 32. Test Passed

7. Click **Save**.

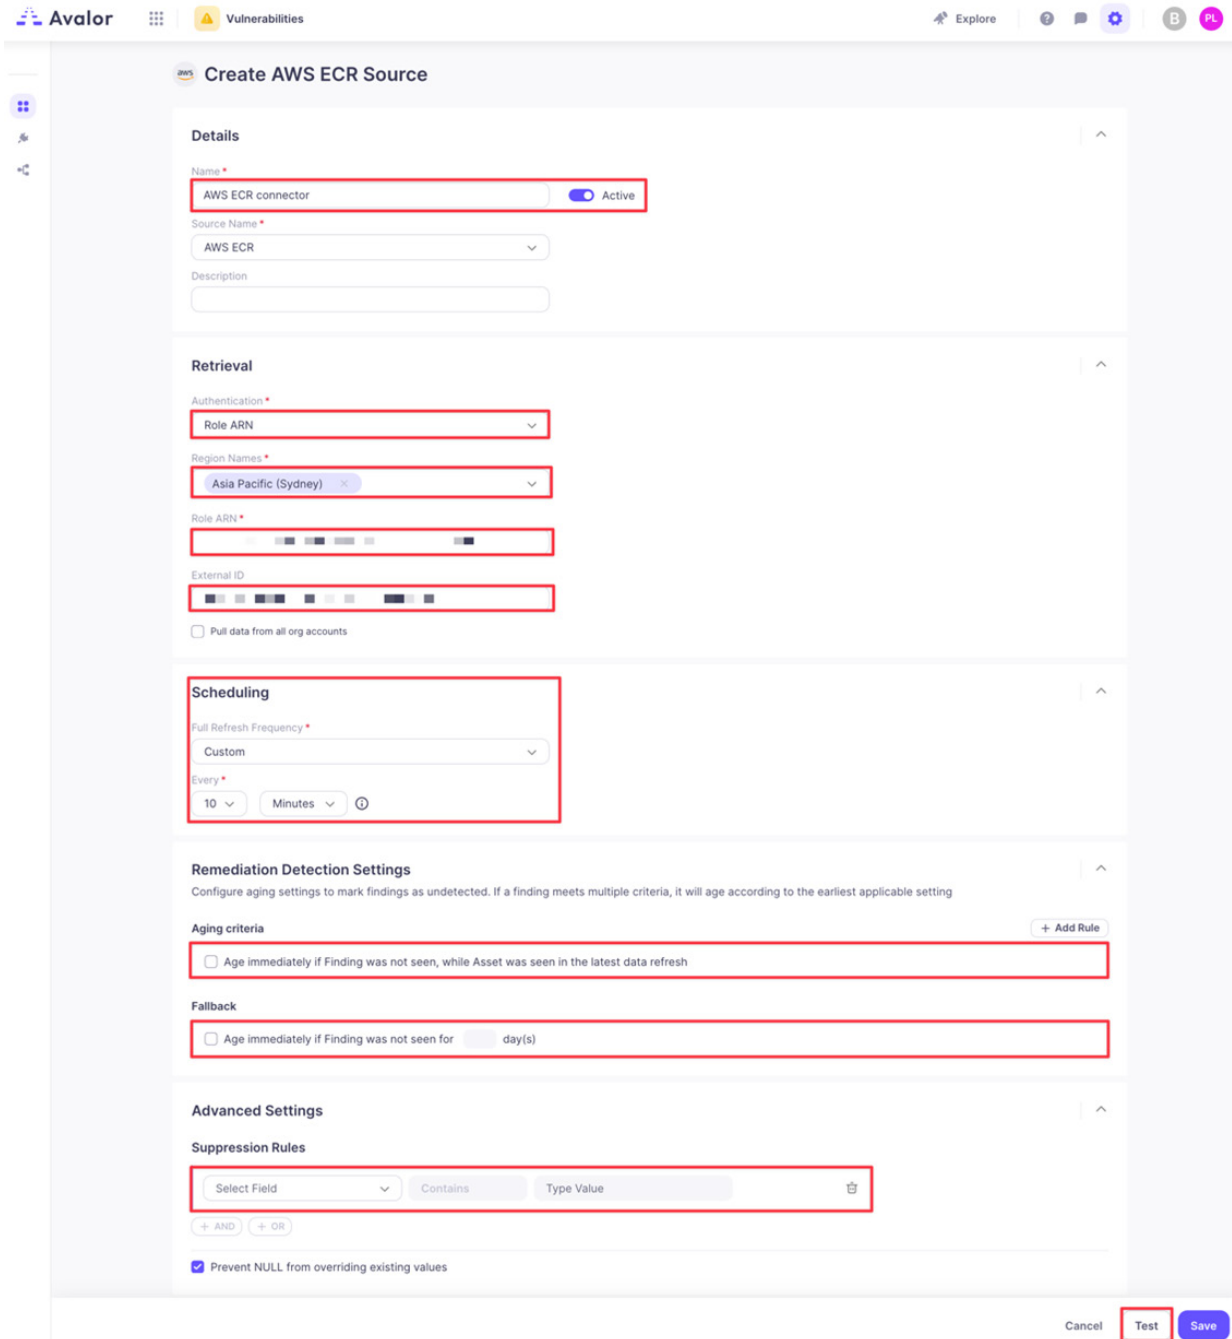


Figure 33. Create AWS ECR Source

Configure the AWS ECR Findings Data Source

To configure the AWS ECR findings data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.




Figure 34. Avalor UVM Platform

3. Click **Create**, then search for AWS ECR Findings.




Figure 35. Connect New Data Source

4. Click the **AWS ECR Findings** application.
5. On the **Create AWS ECR Findings Source** page, complete the following
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names this data source will apply to.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system respond withs Test Passed.




Figure 36. Test Passed

7. Click **Save**.

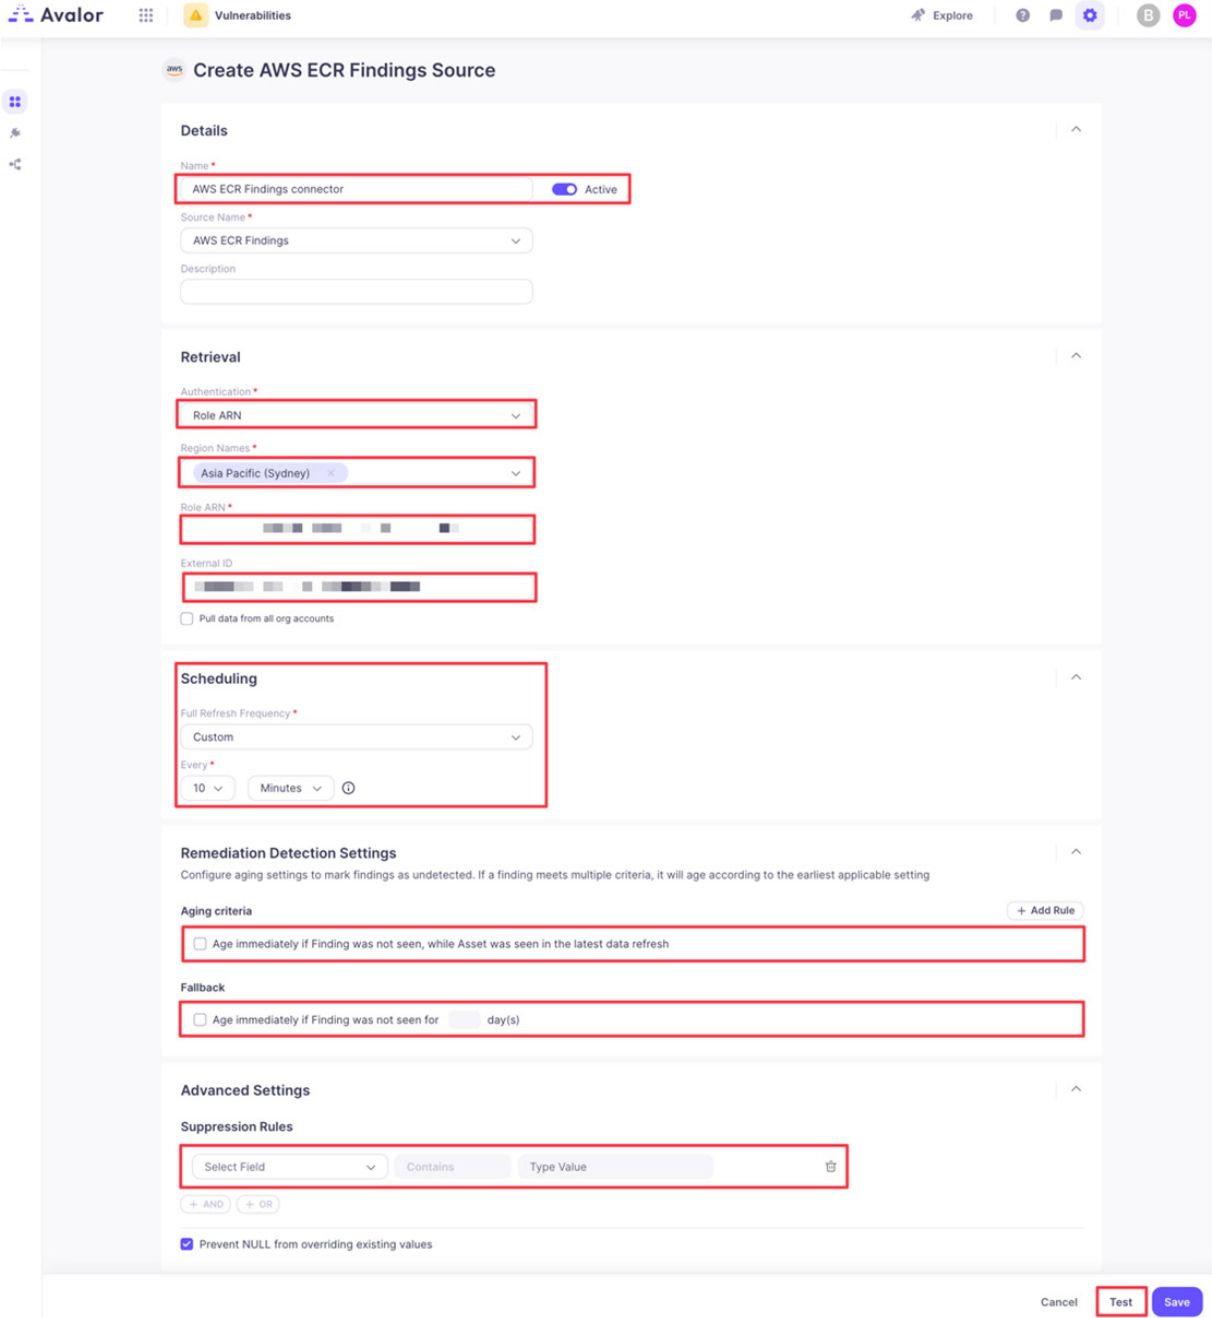


Figure 37. Create AWS ECR Findings Source

Configure the AWS EKS Clusters Data Source

To configure the AWS EKS clusters data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.

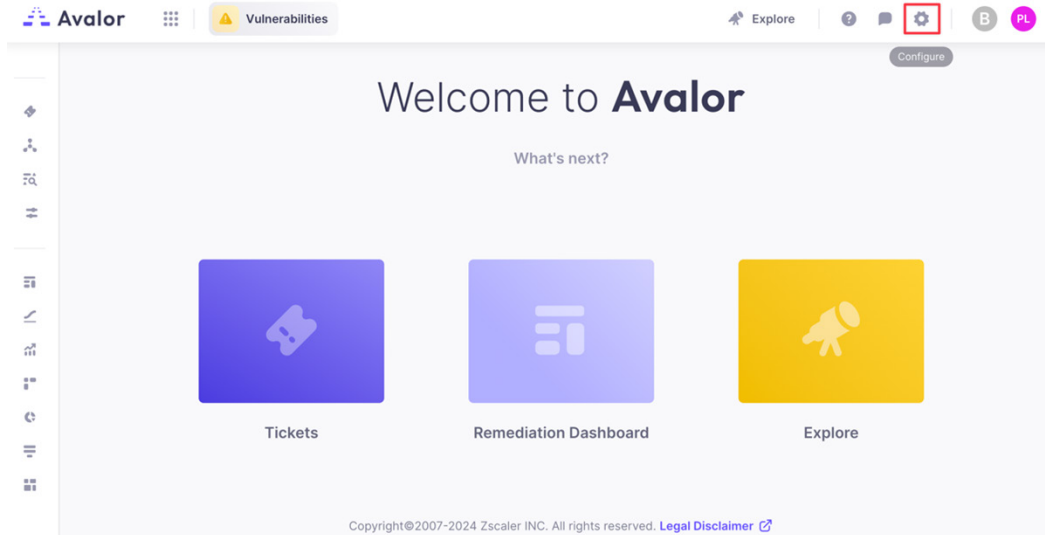


Figure 38. Avalor UVM Platform

3. Click **Create**, then search for AWS EKS Clusters API.




Figure 39. Connect New Data Source

4. Click on the **AWS EKS Clusters API** application.
5. On the **Create AWS EKS Clusters API Source** page, complete the following
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names this data source will apply to.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.




Figure 40. Test Passed

7. Click **Save**.

Avalor Vulnerabilities Explore ? B PL

Create AWS EKS Clusters API Source

Details

Name * AWS EKS Clusters API connector Active

Source Name * AWS EKS Clusters API

Description

Retrieval

Authentication * Role ARN

Region Names * Asia Pacific (Sydney)

Role ARN *

External ID

Pull data from all org accounts

Scheduling

Full Refresh Frequency * Custom

Every * 10 Minutes

Remediation Detection Settings

Configure aging settings to mark findings as undetected. If a finding meets multiple criteria, it will age according to the earliest applicable setting

Aging criteria + Add Rule

Age immediately if Finding was not seen, while Asset was seen in the latest data refresh

Fallback

Age immediately if Finding was not seen for day(s)

Advanced Settings

Suppression Rules

Select Field Contains Type Value

+ AND + OR

Prevent NULL from overriding existing values

Cancel **Test** Save

Figure 41. Create AWS EKS Clusters API Source

Configure the AWS Inspector Findings Data Source

To configure the AWS inspector findings data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.

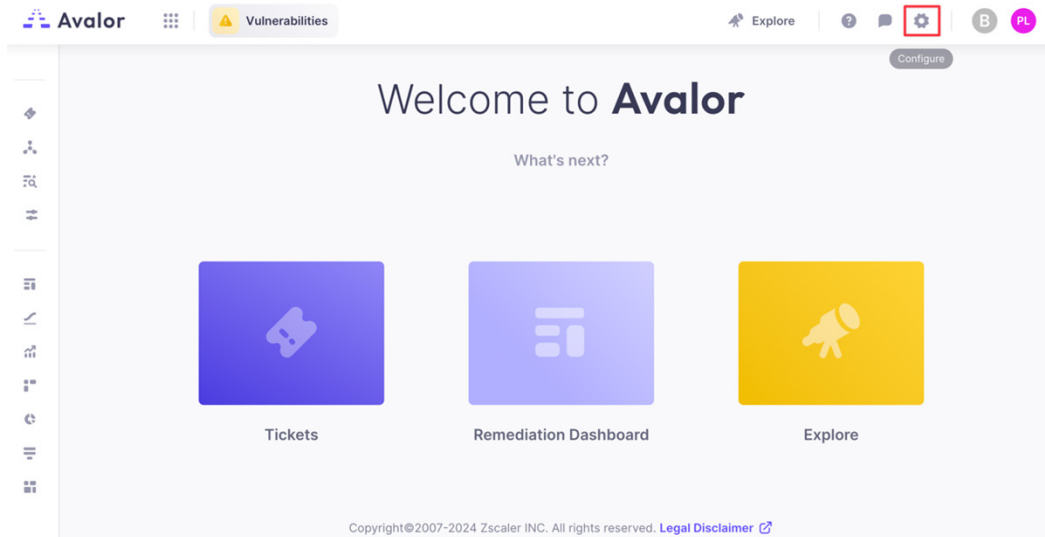


Figure 42. Avalor UVM Platform

3. Click **Create**, then search for AWS Inspector Findings.




Figure 43. Connect New Data Source

4. Click on the **AWS Inspector Findings** application.
5. On the **Create AWS Inspector Findings Source** page, complete the following:
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names this data source will apply to.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, see the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.




Figure 44. Test Passed

7. Click **Save**.

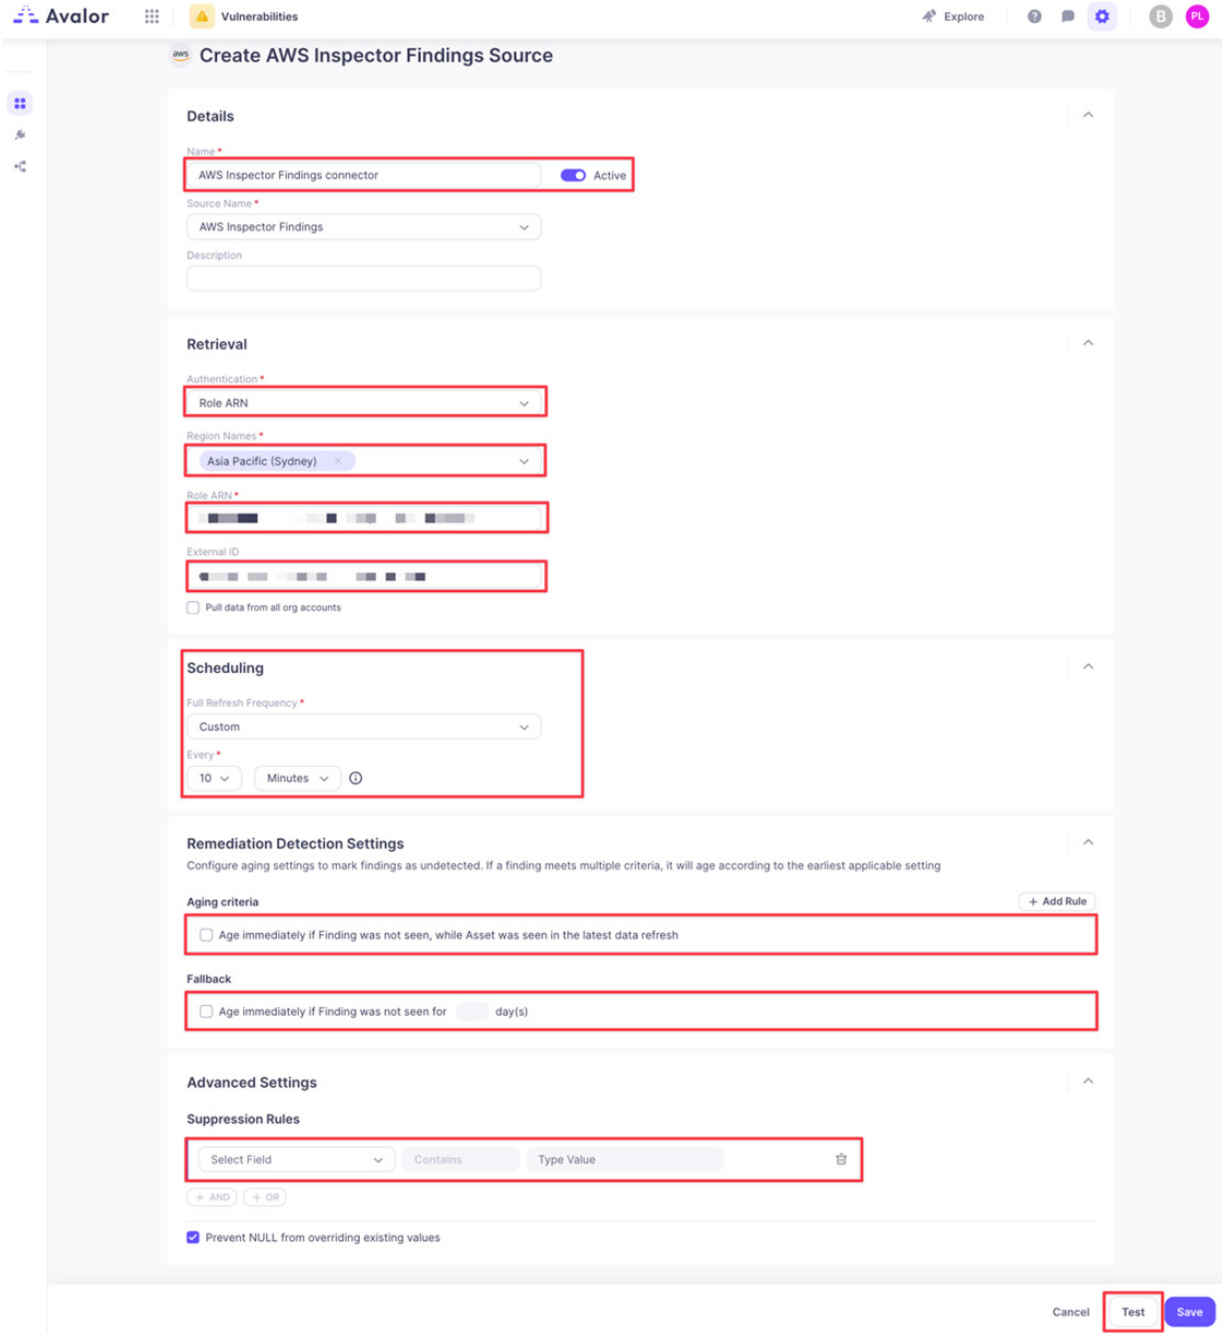


Figure 45. Create AWS Inspector Findings Source

Configure the AWS RDS Data Source

To configure the AWS RDS data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.

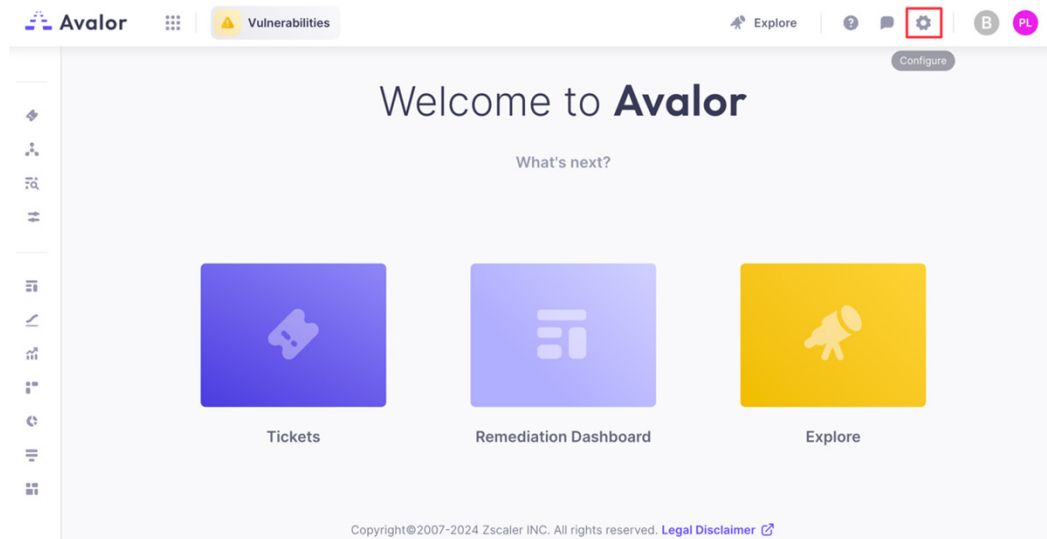


Figure 46. Avalor UVM Platform

3. Click **Create**, then search for AWS RDS.




Figure 47. Connect New Data Source

4. Click the **AWS RDS** application.
5. On the **Create AWS RDS Source** page, complete the following:
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names to which this data source applies.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.




Figure 48. Test Passed

7. Click **Save**.

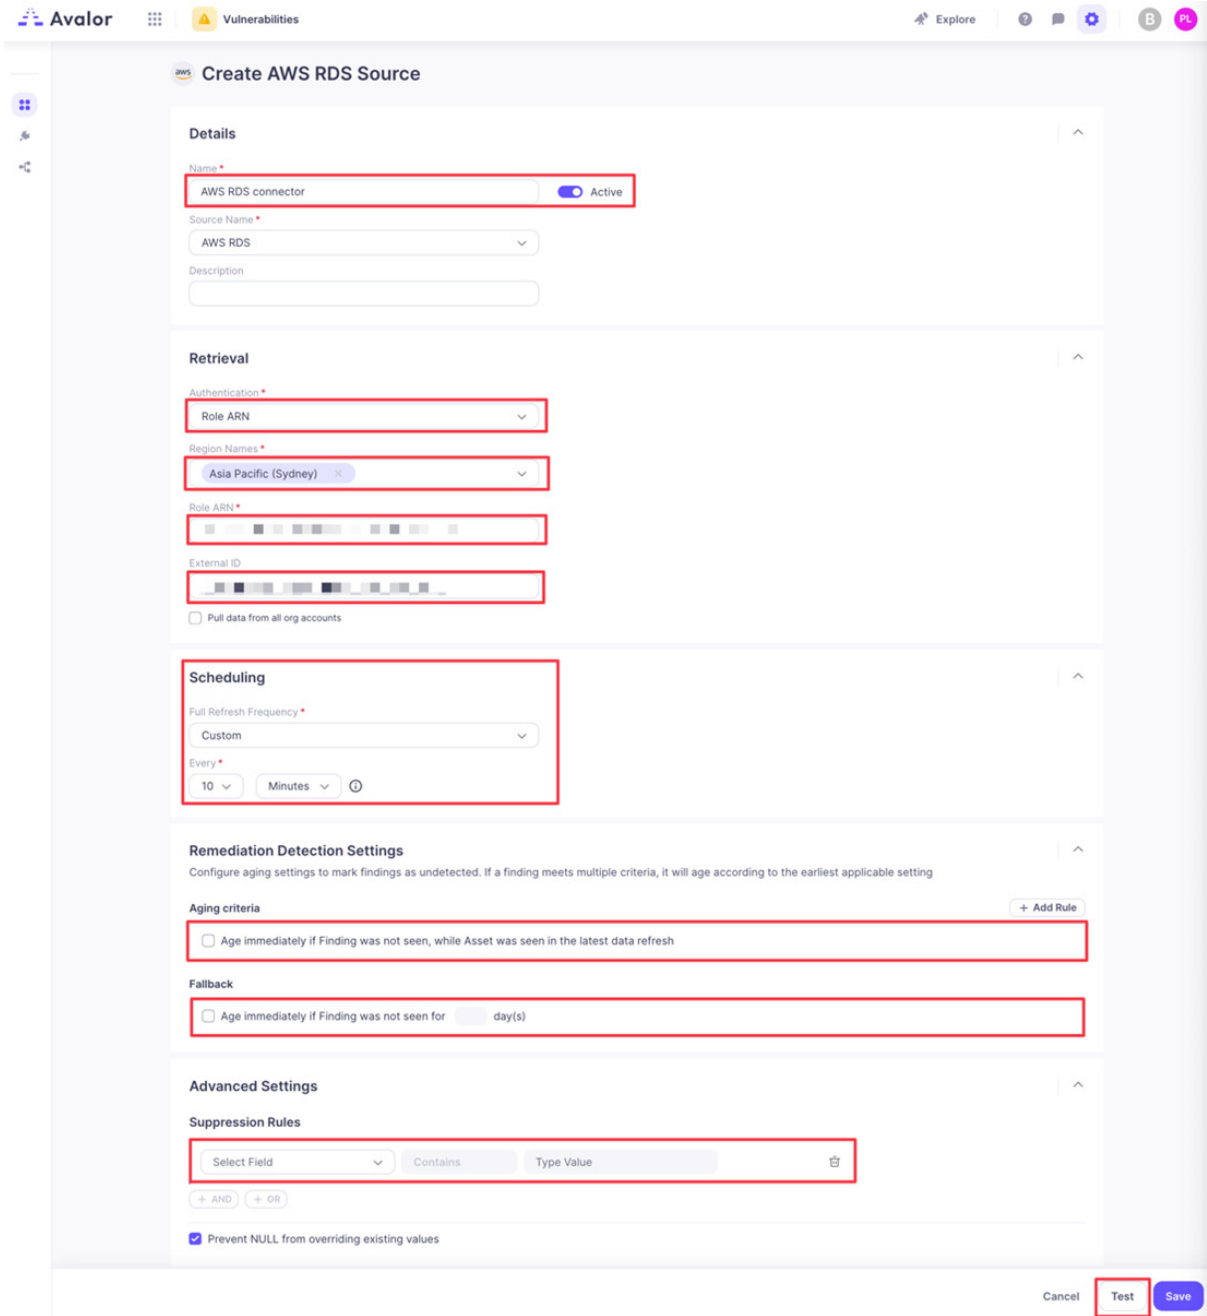


Figure 49. Create AWS RDS Source

Configure the AWS S3 Buckets Data Source

To configure the AWS S3 buckets data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.




Figure 50. Avalor UVM Platform

3. Click **Create**, then search for AWS S3 Buckets.




Figure 51. Connect New Data Source

4. Click on the **AWS S3 Buckets** application.
5. On the **Create AWS S3 Buckets Source** page, complete the following
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names this data source will apply to.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings will automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.




Figure 52. Passed Test

7. Click **Save**.

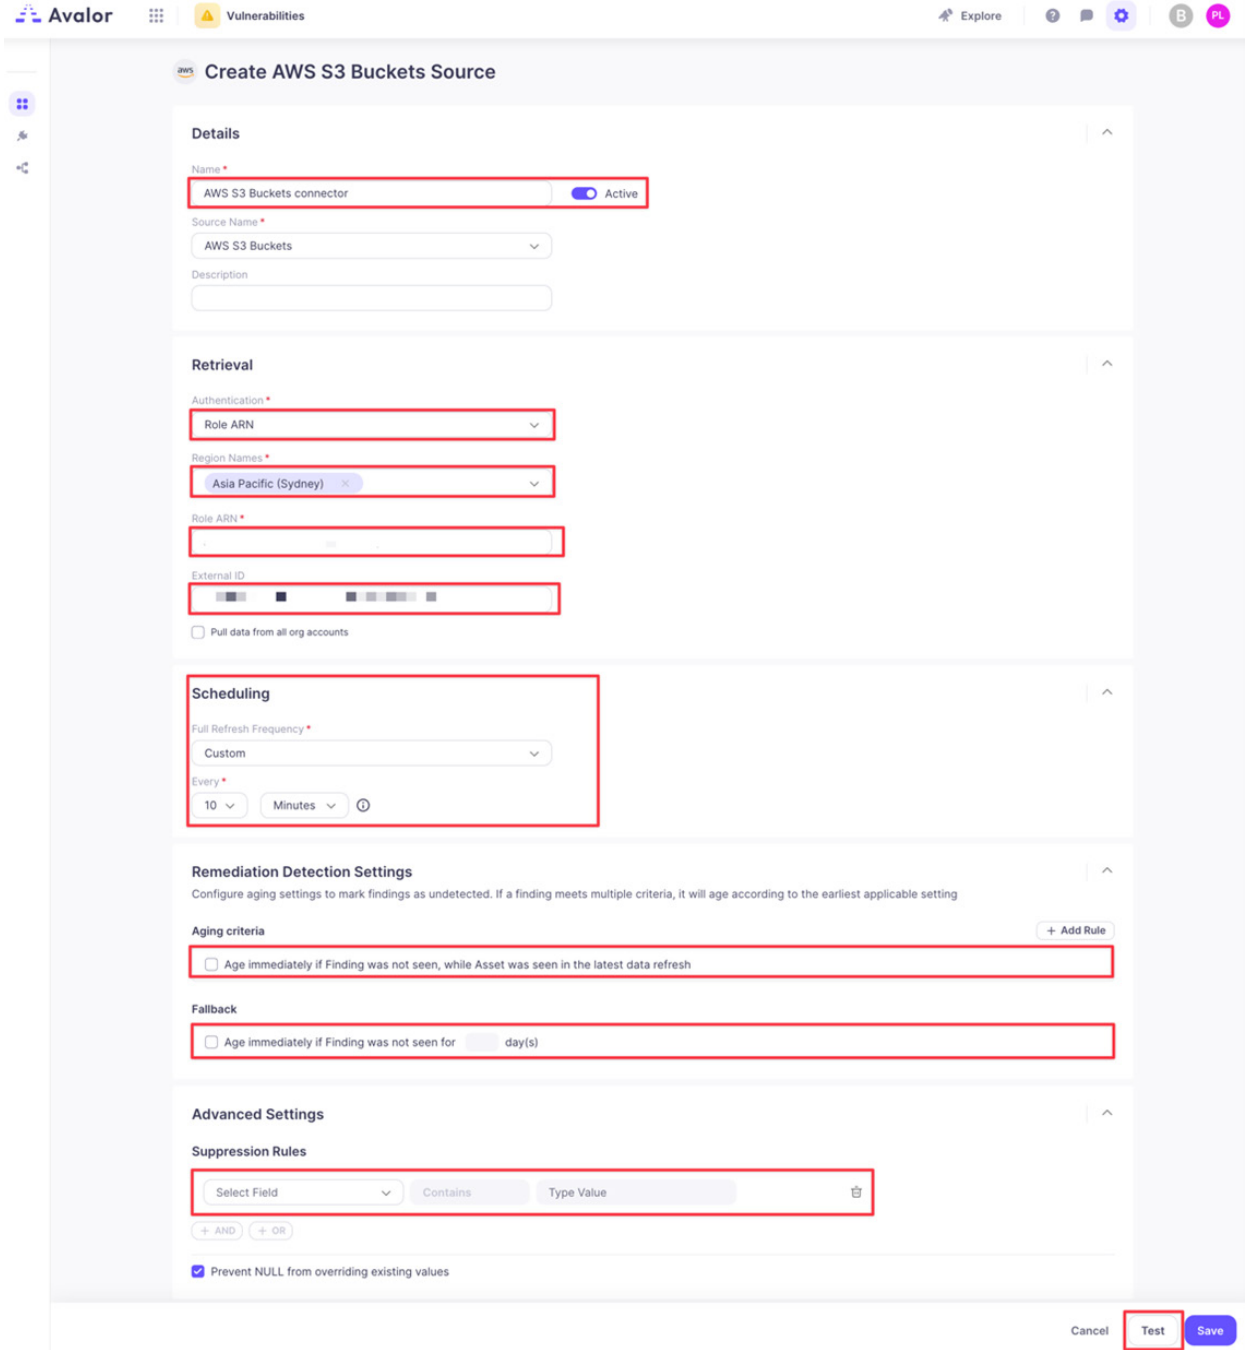


Figure 53. Create AWS S3 Bucket Source

Configure the AWS Security Hub API Data Source

To configure the AWS security hub API data source:

1. Log in to the Avalor UVM Platform.
2. Click **Configure**.




Figure 54. Avalor UVM Platform

3. Click **Create**, then search for AWS Security Hub API.




Figure 55. Connect New Data Source

4. Click the **AWS Security Hub API** application.
5. On the **Create AWS Security Hub API Source** page, complete the following:
 - a. **Name:** Enter a name for the Data Connector.
 - b. **Active:** Toggle the switch to enable the Data Connector.
 - c. **Authentication:** Enter the Role ARN.
 - d. **Region Names:** Select the Region Names to which this data source applies.
 - e. **Role ARN:** Enter the Role ARN.
 - f. **External ID:** Enter the External ID.
 - g. **Scheduling:** Set the schedule for extracting new data only. This option is more efficient because it avoids the need to retrieve all data every time.
 - h. **Remediation Detection Settings:** Select your desired option to determine when findings automatically turn undetected. To learn more, refer to the [Avalor documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - i. **Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Avalor system. To learn more, refer to the [Avalor documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with Test Passed.




Figure 56. Test Passed

7. Click **Save**.

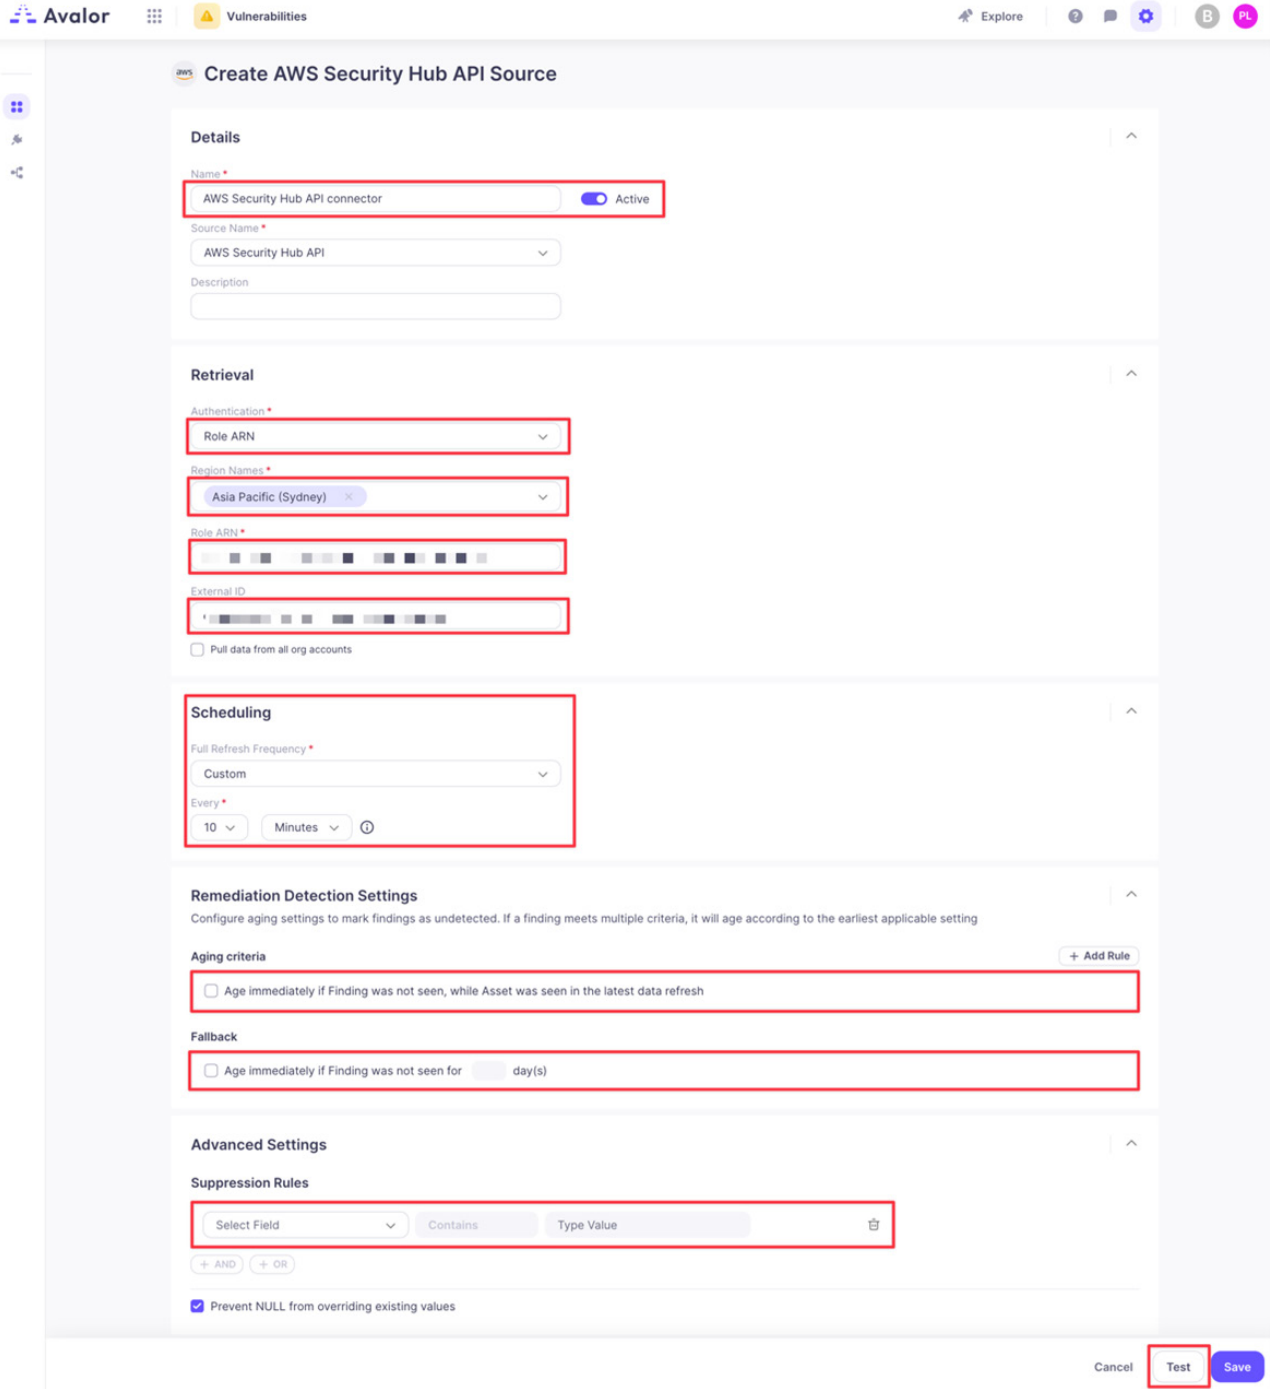


Figure 57. Create AWS Security Hub API Source

Review and Adjust Data Model Mapping

(Optional) Avalor UVM automatically maps ingested data to the default Data Model, so analysis can begin right away. However, many data sources also provide additional data points that might provide additional context to risk prioritization.

The following example shows how to leverage the *Crown Jewel* Data Model Entity based on an EC2 instance tag so that you can use that field as a Risk Factor when calculating risk for an Asset.

Create a Crown Jewel Tag for an EC2 Instance

To create a crown jewel tag for an EC2 Instance:

1. Log in to your AWS Console.
2. Select **Services** > **EC2**.

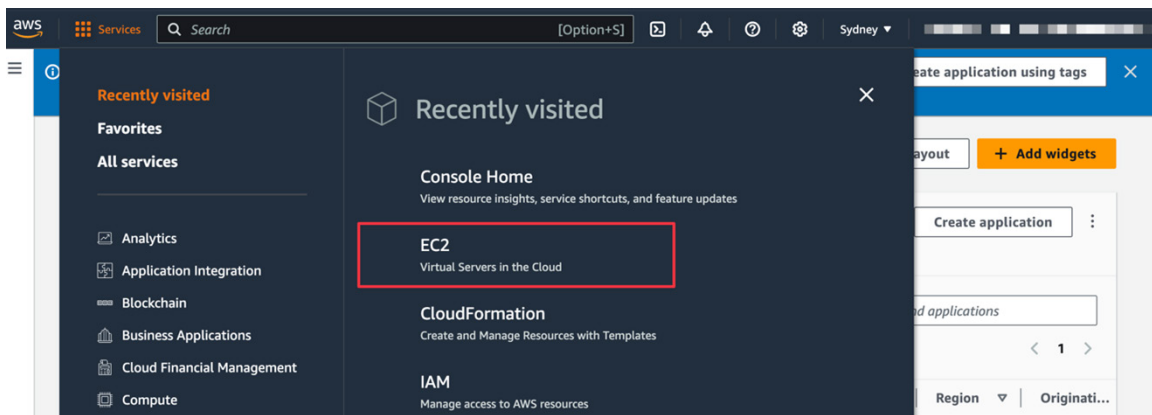


Figure 58. EC2

3. Click **Instances**.

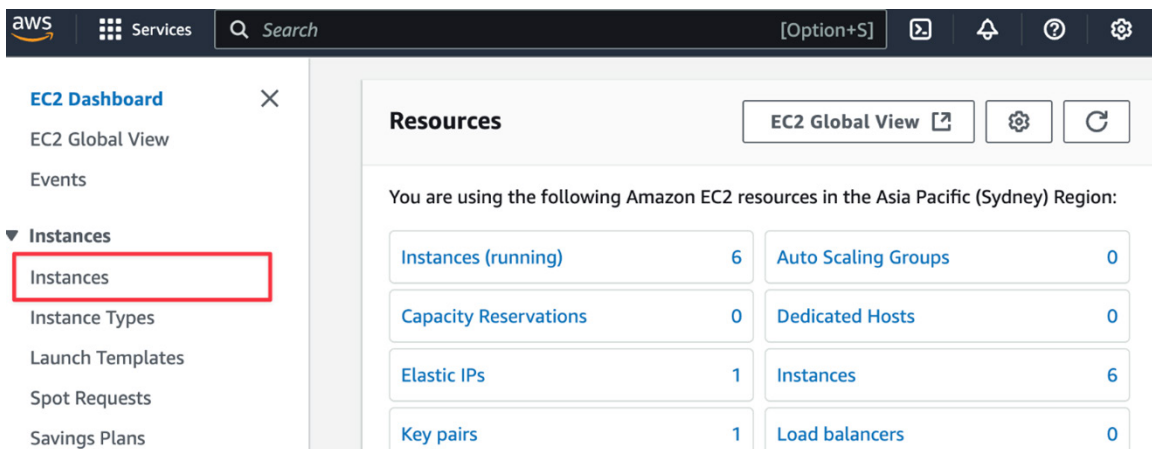


Figure 59. Instances

4. Select the instance you want to add the Crown Jewel tag to and click **Tags**.

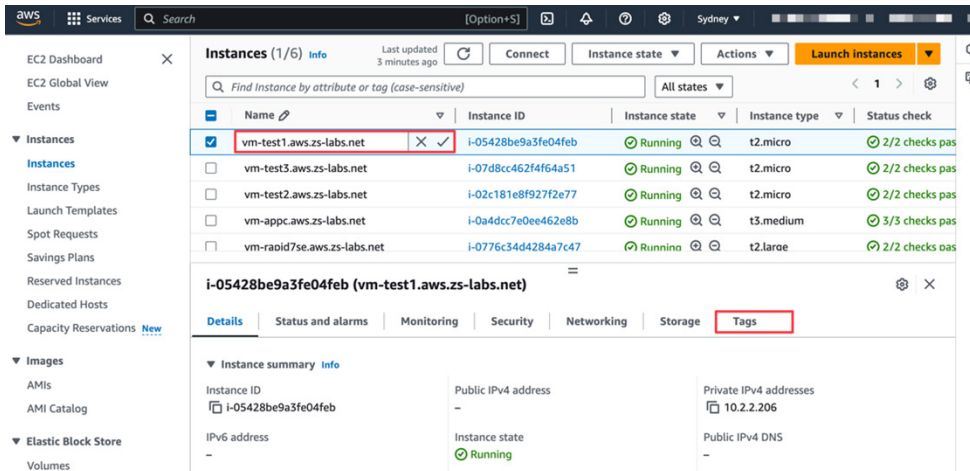


Figure 60. Tags

5. Click **Manage tags**.

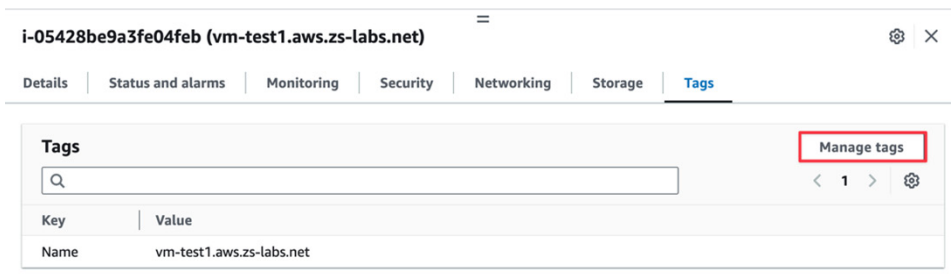


Figure 61. Manage tags

6. Click **Add new tag** and enter:
- Key:** Classification
 - Value:** Crown Jewel

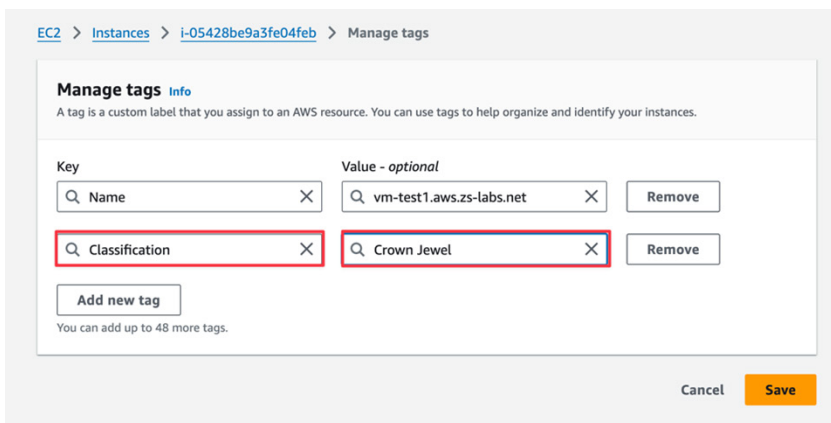


Figure 62. Manage tags

7. Click **Save**.

Map the AWS EC2 Data Source

To map the AWS EC2 data source:

1. Select **Configure** > <the newly created Zscaler Client Connector devices connector> > **Map Data**.

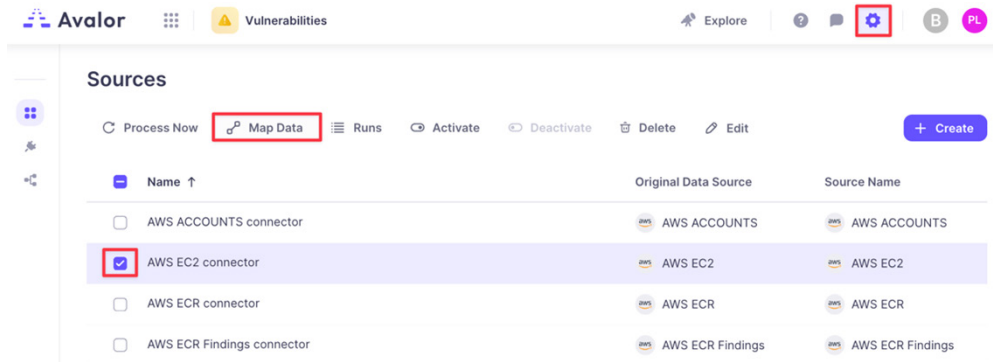


Figure 63. Map Data

2. In the **Map connector** window:
 - a. Create a new **Asset Key** with the internal DNS hostname:
 - i. On the right side, under **Asset**, drag **Key** to the **Create New Connection** element.
 - ii. On the left side, click the **Editor** element.

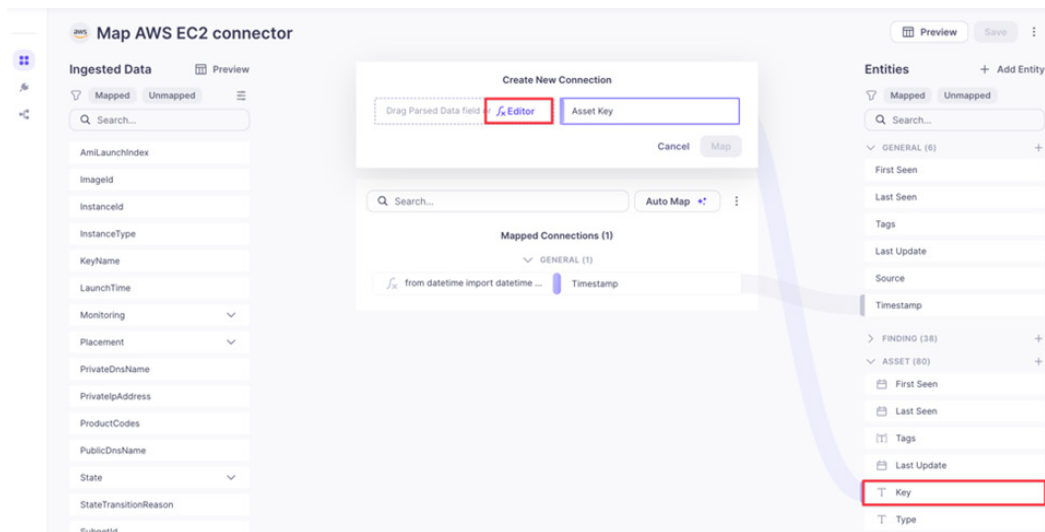


Figure 64. Asset Key

iii. Replace the text in the script field with:

```
def evaluate(row: dict) -> str:

    item = row.get("PrivateDnsName")

    clean_hostname = item.split('.')[0]

return str(clean_hostname)
```

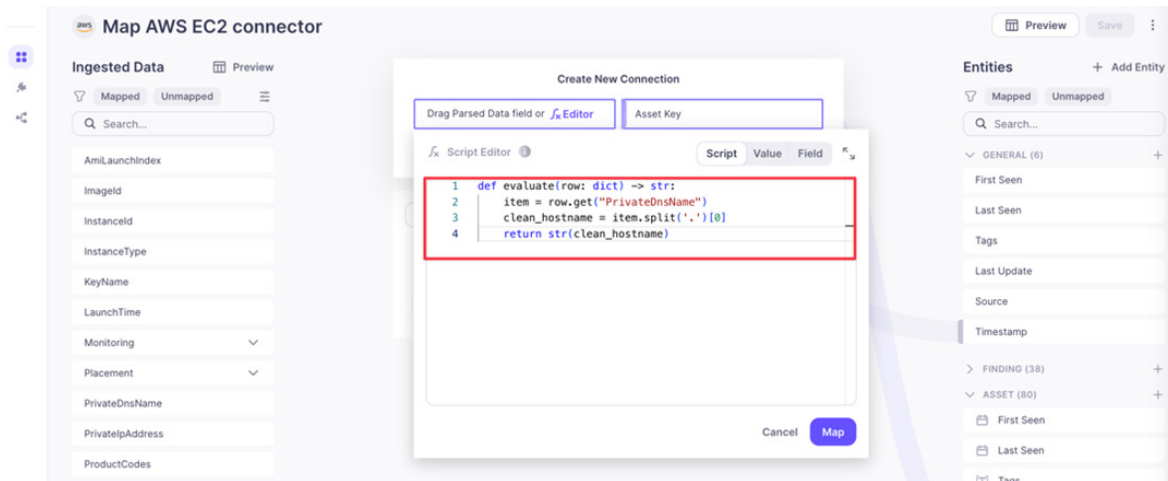


Figure 65. Script field

iv. Click **Map**, then click the **Key** icon, next to the **Asset Key** to set as a key.

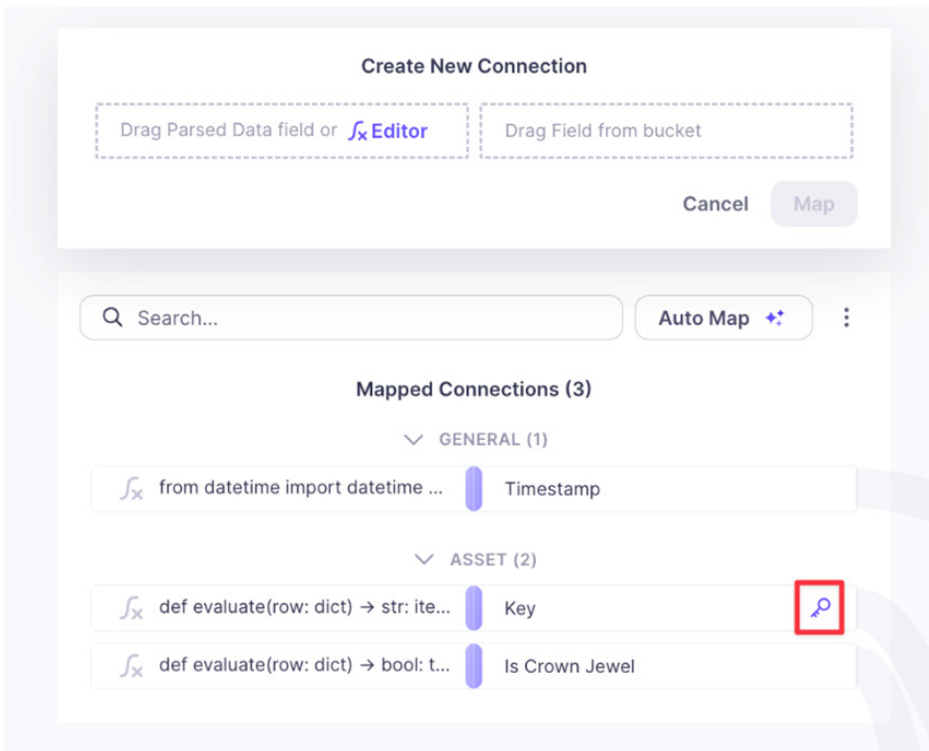


Figure 66. Asset Key

- b. Map the **Is Crown Jewel Asset** entity to the **Crown Jewel EC2** tag created earlier by:
 - i. On the right side, under **Asset**, drag **Is Crown Jewel** to the **Create New Connection** element.
 - ii. On the left side, click the **Editor** element.

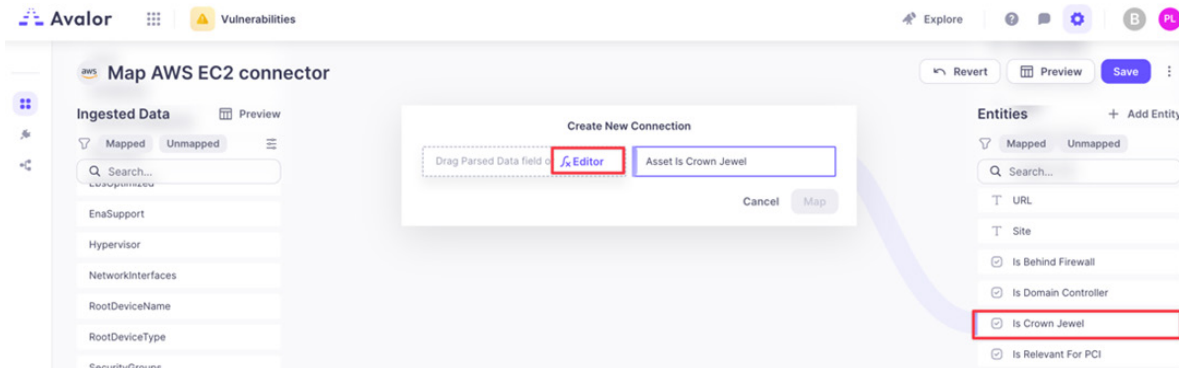


Figure 67. Editor element

- iii. Replace the text in the script field with:

```
def evaluate(row: dict) -> bool:
    tags = row.get("Tags")
    for item in tags:
```

```
    if item.get("Key") == "Classification" and item.get("Value") == "Crown Jewel":
        return True
    else:
```

```
return False
```

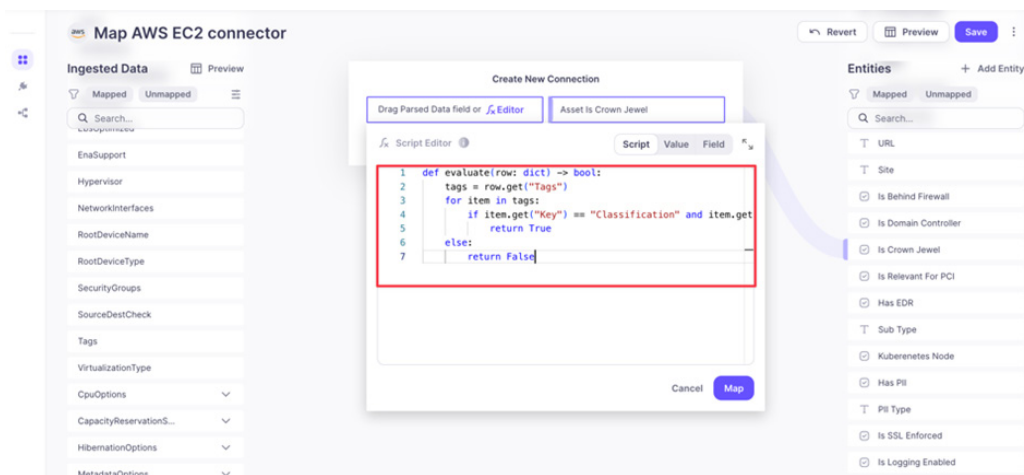
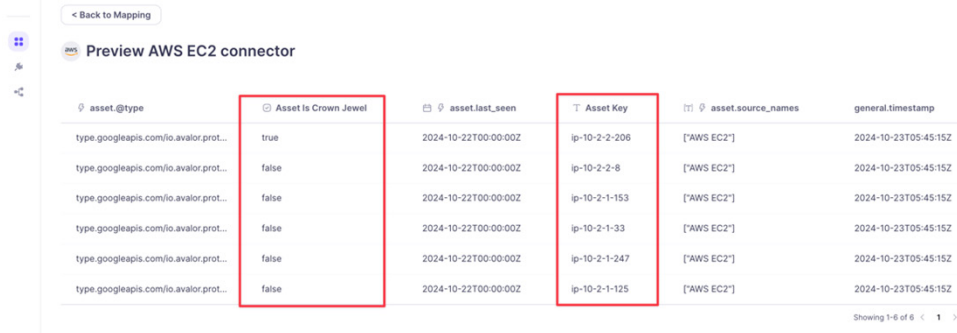


Figure 68. Script field

- iv. Click **Map**.

- c. Click **Preview**, and see the if an Asset is marked as a Crown Jewel based on its EC2 tag and its hostname is marked as its Asset Key.



< Back to Mapping

Preview AWS EC2 connector

asset.@type	Asset Is Crown Jewel	asset.last_seen	Asset Key	asset.source_names	general.timestamp
type.googleapis.com/lo.avalor.prot...	true	2024-10-22T00:00:00Z	ip-10-2-2-206	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/lo.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-2-8	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/lo.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-153	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/lo.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-33	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/lo.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-247	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/lo.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-125	["AWS EC2"]	2024-10-23T05:45:15Z

Showing 1-6 of 6 < 1 >

Figure 69. Preview

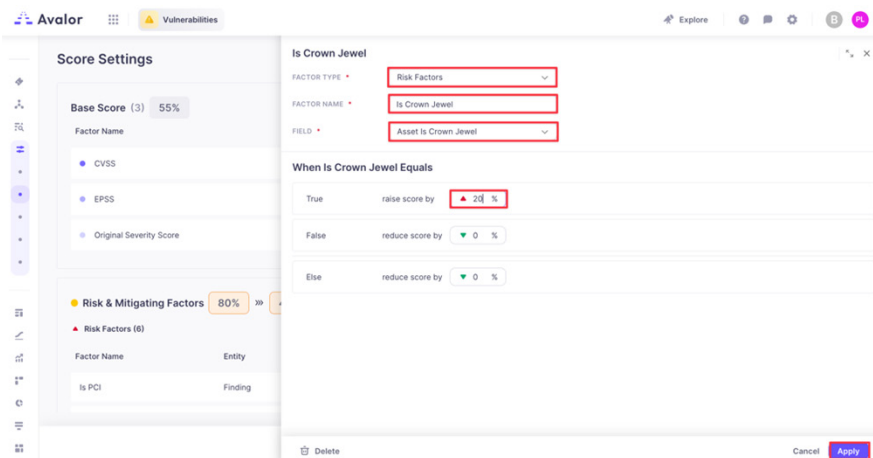
- d. Click **Back to Mapping**, then click **Save**.

Review and Adjust Risk Scoring

After the ingested data has been normalized and mapped to the Data Model, Avalor UVM can evaluate risk.

The following example shows how the *Is Crown Jewel* field is added as a Risk Factor for risk scoring. A value of True increases the risk calculation (since the asset is a Crown Jewel application).

- From the **Vulnerabilities** tab in the **Avalor dashboard (Remediation Hub)**:
 - In the left pane, select **Settings > Score**.
 - Click **Add Factor** in the **Risk & Mitigating Factors** section.
- If Crown Jewel is not already a Risk Factor, in the **Add new factor** modal:
 - Choose **Risk Factors** for **Factor Type** (**Mitigating Factors** generally lower risk scoring, while **Risk Factors** generally increase risk scoring).
 - Enter a **Name**.
 - Choose **Crown Jewel** for **Field**.
 - In the **Boolean** login section, under **True**, enter a percentage by which the risk is increased.



Avalor Vulnerabilities Explore

Score Settings

Base Score (3) 55%

Factor Name

- CVSS
- EPSS
- Original Severity Score
- Risk & Mitigating Factors 80%

Risk Factors (6)

Factor Name	Entity
Is PCI	Finding

Is Crown Jewel

FACTOR TYPE Risk Factors

FACTOR NAME Is Crown Jewel

FIELD Asset Is Crown Jewel

When Is Crown Jewel Equals

Value	raise score by	reduce score by
True	20%	0%
False	0%	0%
Else	0%	0%

Delete Cancel Apply

Figure 70. Boolean login section

- e. Click **Apply**, then **Save & Run**.

3. In the left-side pane, select the **Assets** dashboard. From the **Assets** dashboard:
 - a. Set a filter by clicking **More** and selecting **True** for **Is Crown Jewel True**.

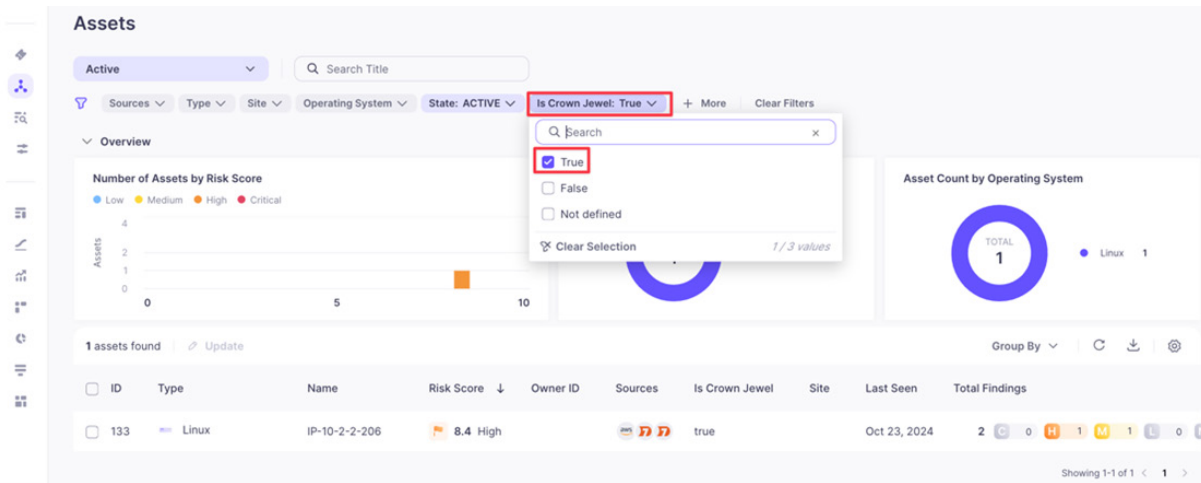


Figure 71. Assets dashboard

- b. Click one of your Assets in the filtered list.
- c. In the **Asset** modal that appears, click the **Findings** tab.
- d. Click one of the **Findings**.
- e. Review the output (notice the **Score Adjustment** section and how **Is Crown Jewel** has modified the risk scoring).

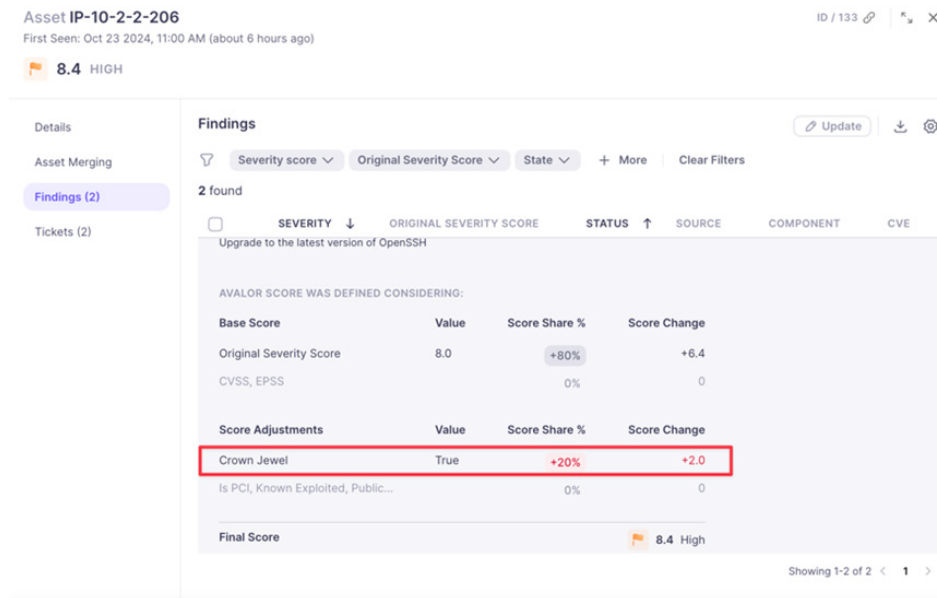


Figure 72. Findings tab

Appendix A: Requesting Zscaler Support

If you need Zscaler Support to provision certain services or to help troubleshoot configuration and service issues, it is available 24/7/365.

- 1. To contact Zscaler Support, go to **Administration > Settings > Company Profile**.

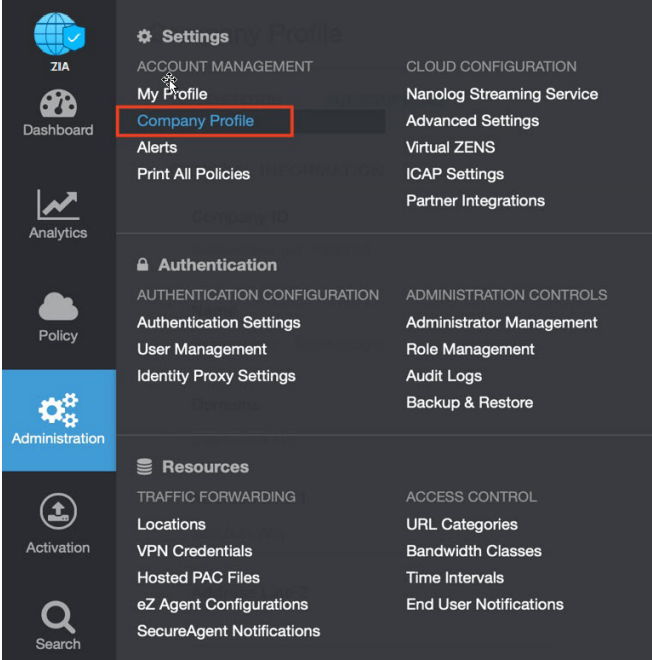


Figure 73. Collecting details to open support case with Zscaler TAC

- 2. Copy your Company ID.

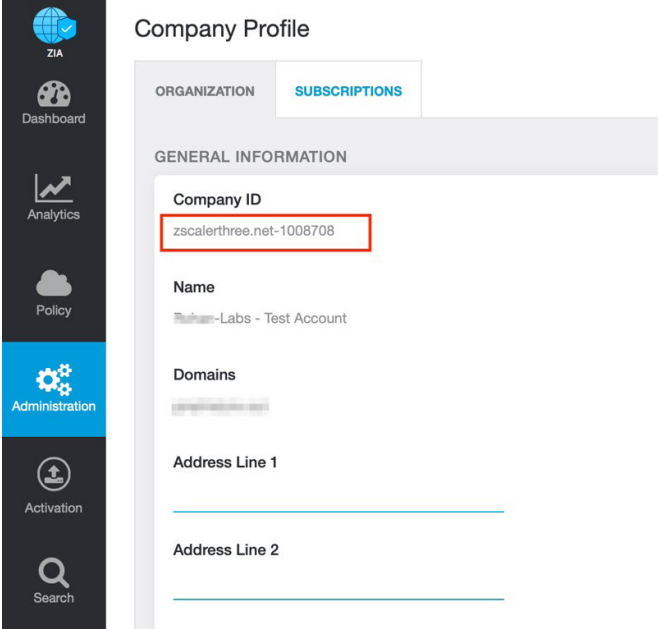


Figure 74. Company ID

3. With your company ID information, you can open a support ticket. Go to **Dashboard > Support > Submit a Ticket**.

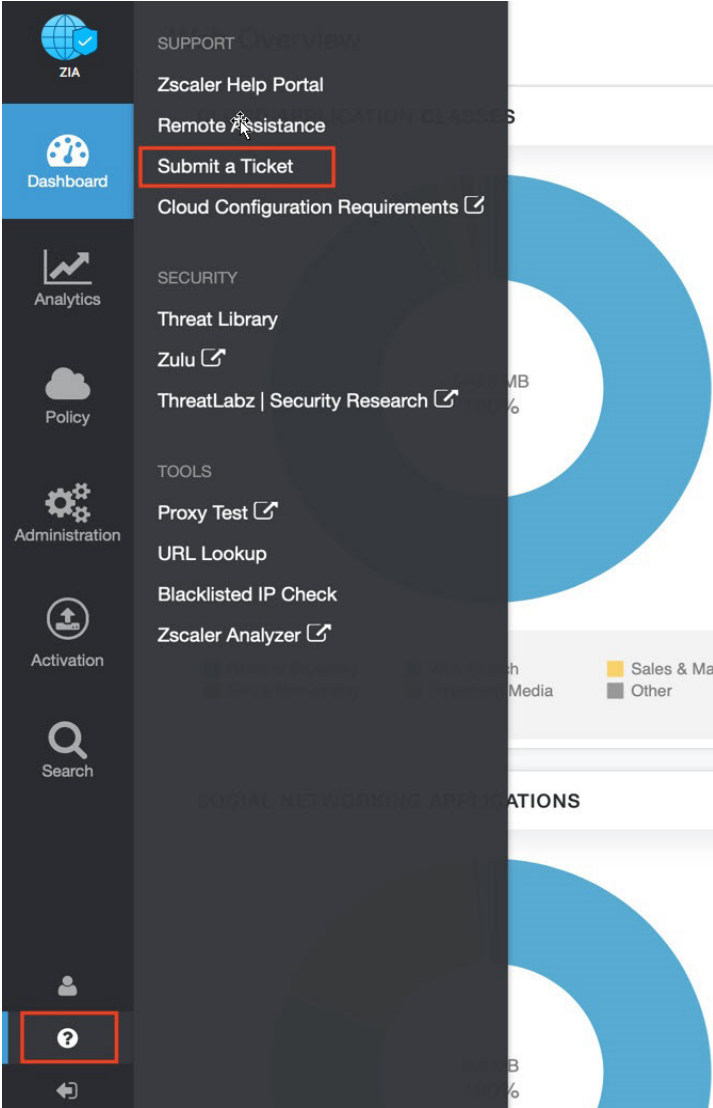


Figure 75. Submit a ticket