



ZSCALER AND AWS DEPLOYMENT GUIDE

Contents

Terms and Acronyms	7
About This Document	9
Zscaler Overview	9
AWS Overview	9
Audience	9
Software Versions	9
Request for Comments	9
Zscaler and AWS Introduction	10
ZIA Overview	10
ZPA Overview	10
Zscaler Cloud Connector Overview	10
Zscaler Zero Trust Microsegmentation Overview	11
Zscaler Deception Overview	11
Zscaler Resources	11
AWS Overview	12
AWS Workspaces Overview	12
AWS Transit Gateway Overview	12
AWS CloudWatch Overview	12
Amazon Q Overview	13
Amazon Bedrock Overview	13
Amazon GuardDuty Overview	13
AWS Snowball Overview	13
AWS Resources	14
Getting Started	15

Traffic Forwarding Options	16
Zscaler Client Connector for Users	16
PAC Files for Users or Workloads	21
Zscaler Cloud Connector for Users or Workloads	22
Deploying Zero Trust Cloud Automatically Zero Trust Gateway—Recommended	22
Deploying Zero Trust Cloud EC2 Instances Manually	27
Site-to-Site IPsec for Users or Workloads	33
Identifying the Zscaler VPN Endpoint	34
Create a Customer Gateway	36
Create a Site-to-Site VPN Connection	36
Configure ZIA	41
Configure Routing for Site-to-Site VPN Connection	45
Example Testing	49
ZIA Components that Work on AWS Infrastructure	50
Nanolog Streaming Service	51
Virtual Service Edge	52
DLP Incident Receiver	52
DLP Index Tool	52
Amazon WorkSpaces Supporting Zscaler Client Connector	53
ZPA Components that Work on AWS Infrastructure	54
App Connector	54
Private Service Edge	54
Using ZIA to Enforce Security Policy in AWS	55
Cloud App Control Policy	56
Cloud App Control Policies Available via Individual Amazon Web Services	58
File Type Control for AWS	61
Firewall Control Rules for AWS	63
DNS Control	64

Using ZPA to Enforce Security Policy in AWS	65
Deploying App Connectors	65
Prerequisites and Planning	65
Configure App Connectors in ZPA Admin Portal	66
Deploy App Connectors in AWS	67
Creating Workload Application Segments and Adjusting Policy	69
Discover and Create Applications Segments	69
Refine Policy	70
Configuring Microsegmentation Between Workloads	71
Initial Setup	71
Deploying the Agents	72
Configure AppZones	73
Configure Resource Groups	74
Configure Resource Policy	76
Extending Zscaler with AWS Service Integrations	78
Bringing Zero Trust Security to AWS Snowball Deployments	78
AWS Components	78
Zscaler Components	78
Provisioning the AWS Snowball Appliance	79
Deploy the AWS Snowball Appliance	80
Configure Policy in ZPA Admin Portal to Allow Connections	82
Deploy the Zscaler Client Connector	85
Test the Setup	87
Monitoring ZPA App Connector Health using AWS CloudWatch	90
Create an IAM Role to Use for CloudWatch Agent and Assign to App Connector	90
Log In to the App Connector and Download the CloudWatch Agent	93
Create CloudWatch Configuration File Using Wizard	93
Example Dialog	93
Start CloudWatch Agent	102

Monitor Metrics Generated by CloudWatch Agent	103
View Logs Stored in CloudWatch	104
Operationalizing Threat Intelligence with Zscaler and AWS GuardDuty	106
Feeding GuardDuty Telemetry into Zscaler Internet Access for Enforcement	106
Requirements and Components	106
Configuring ZIA	107
Feeding Zscaler Deception Telemetry into Amazon GuardDuty	114
Leveraging AWS System and User-Defined Tags for Policy	120
Requirements	120
Configuring the Partner Integration	121
Configuring Policy	124
Enhancing AWS S3 with Zscaler SaaS Security	126
Initial ZIA Configuration	126
Configure the AWS IAM Role	127
Configure the AWS Trust Relationship	129
Configure AWS CloudTrail	130
Configure the S3 Quarantine Bucket	131
Finalize Zscaler Configuration	132
Integrating Zscaler Cloud NSS with Amazon S3	133
Create a User Group in AWS IAM	133
Create a User and Access Key in AWS IAM	135
Create an S3 Bucket and Folder in S3	141
Create a Policy Granting the User Group Access to the S3 Bucket in Amazon IAM	146
Add a Cloud NSS Feed in the ZIA Admin Portal	151
Streamlining Incident Response with Workflow Automation	152
Applying Zero Trust Principles to Generative AI Workloads	153
Architecture	153
Enabling Zscaler Malware and DLP scanning of AWS S3 buckets	155
Configure Zscaler Policy Scans	156
Zscaler Data Protection Overview	159

Configure a DLP Policy for Private and Public AI Data Protection	159
DLP with Content Inspection	159
Configure DLP Dictionaries	160
Configure DLP Engine	164
Define Policy Rules	165
Configure the Zscaler Notification Framework	169
Contextualizing Risk using AWS and Zscaler UVM	172
Creating a Role ARN and an External ID in AWS	172
Configure the AWS UVM Data Connectors	175
Review and Adjust Data Model Mapping	194
Appendix A: AWS Transit Gateway Lab Environment	201
Appendix B: Creating a Trail	202
Appendix C: Testing Notes	204
Appendix D: AWS SSM Distributor	206
Prerequisites	206
Configuration	206
Upload the Installation Package	206
Configure AWS SSM Distributor	207
Deploy the Installation Package	207
Install One Time	207
Updates	207
Appendix E: Metrics Config Options	208
Appendix F: ZPA and ZIA Configuration for Private AI Data Protection	209
Appendix G: Requesting Zscaler Support	220

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
AMI	Amazon Machine Image (Amazon)
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)
EC2	Amazon Elastic Compute Cloud (Amazon)
ECS	Enhanced Cybersecurity Services
ENI	Elastic Network Interface (Amazon)
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
RDS	Remote Desktop Services
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
VLAN	Virtual Local Area Network
VPC	Virtual Private Cloud
XFF	X-Forwarded-For (RFC7239)
ZDX	Zscaler Digital Experience (Zscaler)
ZIA	Zscaler Internet Access (Zscaler)
ZPA	Zscaler Private Access (Zscaler)
ZTE	Zero Trust Exchange (Zscaler)

Trademark Notice

© 2025 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](#).

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 [AWS'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 G: 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.

ZPA Overview

ZPA is a cloud service that provides secure remote access to internal applications running on a cloud or data center using a zero trust framework. With ZPA, applications are never exposed to the internet, making them completely invisible to unauthorized users. The service enables the applications to connect to users via inside-out connectivity rather than extending the network to them.

ZPA provides a simple, secure, and effective way to access internal applications. Access is based on policies created by the IT administrator within the ZPA Admin Portal and hosted within the Zscaler cloud. On each user device, software called Zscaler Client Connector is installed. Zscaler Client Connector ensures the user's device posture and extends a secure microtunnel out to the Zscaler cloud when a user attempts to access an internal application. .

Zscaler Cloud Connector Overview

Zscaler Cloud Connector is a solution designed to secure cloud-to-cloud and cloud-to-internet communications between workloads. It enables both macro- and microsegmentation and enforces least-privileged access policies without relying on traditional network constructs like IP addresses or firewalls. By leveraging identity-based policies such as tag-based or Java Web Token (JWT), it secures east-west and outbound traffic for workloads across AWS, Azure, and other cloud platforms. The solution integrates natively with AWS as a service (Zscaler Cloud Connector) or can be installed in a DIY fashion via EC2 appliance in the customer's tenant. Like Zscaler for Users, Zero Trust Cloud offers visibility and control over workload behavior, including encrypted traffic inspection. Ultimately, it simplifies cloud security by aligning access to the identity of the workload, not the network.

Zscaler Zero Trust Microsegmentation Overview

Zscaler Zero Trust Microsegmentation is an agent-based security solution designed to prevent lateral movement within cloud and data center environments by enforcing identity-based access controls between workloads. Unlike legacy network segmentation that relies on IP addresses, firewalls, or VLANs, Zscaler's approach applies zero trust principles at the host level, allowing only verified, necessary communication between applications and services. The solution integrates natively with Zscaler Private Access (ZPA), enabling unified policy enforcement across user-to-app and app-to-app communication. Key capabilities include:

- Real-time asset discovery across public clouds and on-premises environments.
- Automated policy recommendations based on traffic flow analysis.
- Granular, host-based controls to isolate workloads and limit attack paths.
- Lightweight agents for Windows and Linux.

Zscaler Deception Overview

Zscaler Deception is a proactive security solution that uses decoys and traps to detect and contain threats inside your environment before they can cause harm. It deploys realistic, low-interaction decoys—such as fake IAM credentials, servers, or S3 buckets—across the network to lure attackers and reveal lateral movement or insider threats. Unlike traditional detection methods, Deception focuses on identifying malicious intent based on interaction with these planted assets, rather than relying solely on signatures or behavior analytics. Deception helps strengthen your zero trust strategy by turning your environment into an active defense landscape that hunts adversaries from within.

Zscaler Resources

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

Name	Definition
ZIA Help Portal	Help articles for ZIA.
ZPA Help Portal	Help articles for ZPA.
Zscaler Cloud Connector Help Portal	Help articles for Zscaler Cloud Connector.
Zscaler Deception	Help articles for Zscaler Deception.
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.
ZPA Help Portal	Help articles for ZPA.
Zscaler Cloud Connector Help Portal	Help articles for Zscaler Cloud Connector.
Zscaler Deception	Help articles for Zscaler Deception.
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 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.

AWS Workspaces Overview

Amazon WorkSpaces is a fully managed desktop virtualization service for Windows and Linux that enables you to access resources from any supported device. With Amazon WorkSpaces, you can:

- Onboard contingent workers. Easily assign and remove desktops for contractors while keeping your sensitive data secure in the cloud.
- Facilitate remote work. Enable work-from-home and remote workers to access fully functional Windows and Linux desktops from any location.
- Run powerful desktops. Provide high-performance desktops for developers and engineers to store and access proprietary models, designs, and code.
- Let contact center agents work from anywhere. Enable contact center agents to work from anywhere with a secure, easy-to-use agent experience.

AWS Transit Gateway Overview

AWS Transit Gateway is a scalable, cloud-native networking service that simplifies the way you connect Amazon VPCs, on-premises networks, and AWS services (such as WorkSpaces). It acts as a central hub for routing traffic, replacing complex peering architectures with a more manageable and efficient hub-and-spoke model. This makes it especially valuable for organizations with multi-VPC or hybrid cloud environments. Key benefits include:

- Centralized connectivity for VPCs, VPNs, and AWS Direct Connect.
- Simplified routing using route tables to control traffic flow between connected networks.
- Scalable performance, automatically handling increasing network traffic.
- Segmentation and isolation using multiple route domains.
- Hybrid cloud support, ideal for linking on-premises data centers with AWS infrastructure.

AWS CloudWatch Overview

Amazon CloudWatch is a monitoring and observability service that provides real-time insights into AWS resources, applications, and services. It collects and visualizes metrics, logs, and events to help you understand system health, detect anomalies, and take automated actions based on predefined thresholds. CloudWatch plays a critical role in performance tuning, troubleshooting, and maintaining operational excellence across AWS environments. Key benefits include:

- Metrics collection from AWS services, custom applications, and on-prem systems.
- Log aggregation and analysis with powerful filtering and search tools.
- Dashboards for visualizing performance and system health in real time.
- Alarms and notifications to alert on threshold breaches or unusual behavior.
- Automation via CloudWatch Events/Rules, enabling responses like Lambda triggers or scaling actions.
- Integration with AWS services like EC2, Lambda, RDS, and ECS for deep observability.

Amazon Q Overview

Amazon Q Business is a generative AI-powered assistant that can answer questions, provide summaries, generate content, and securely complete tasks based on data and information in your enterprise systems. It empowers employees to be more creative, data-driven, efficient, prepared, and productive. It allows end users to receive immediate, permissions-aware responses from enterprise data sources with citations, for use cases such as IT, HR, and benefits help desks.

Amazon Q supports creating its data set by using connectors to attach to AWS S3 buckets and many other supported data stores. This document covers how to protect these data stores and the Amazon Q web crawler used to populate its data sets.

Amazon Bedrock Overview

Amazon Bedrock is a fully managed service that makes high-performing foundation models (FMs) from leading AI startups and Amazon available for your use through a unified API. You can choose from a wide range of FMs to find the model that is best suited for your use case. Amazon Bedrock also offers a broad set of capabilities to build generative AI applications with security, privacy, and responsible AI. Using Amazon Bedrock, you can easily experiment with and evaluate top FMs for your use cases, privately customize them with your data using techniques such as fine-tuning and Retrieval Augmented Generation (RAG), and build agents that execute tasks using your enterprise systems and data sources.

Amazon GuardDuty Overview

Amazon GuardDuty is a threat detection service that continuously monitors your AWS accounts, workloads, and data for malicious activity and unauthorized behavior. Using machine learning, anomaly detection, and threat intelligence from AWS and third-party sources, GuardDuty helps identify potential security threats without the need for manual configuration or complex log analysis. Key benefits include:

- Continuous monitoring of AWS CloudTrail, VPC Flow Logs, and DNS logs for signs of threats.
- Detection of threats such as compromised instances, unusual API calls, reconnaissance, and data exfiltration.
- Integration with AWS Security Hub, Amazon EventBridge, and Lambda for automated responses.
- Multi-account support for centralized threat detection across an organization.
- No agent deployment required, making setup fast and maintenance minimal.
- Built-in threat intelligence from AWS, CrowdStrike, and other sources.

AWS Snowball Overview

Snowball is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS cloud. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns.

AWS Resources

The following table contains links to AWS support resources.

Name	Definition
Amazon S3 Help	Amazon Simple Storage Service documentation.
AWS CLI	AWS Command Line Interface documentation.
AWS CloudTrail Help	AWS CloudTrail documentation.
AWS IAM Help	AWS Identity and Access Management (IAM) documentation.
Amazon CloudWatch Help	Amazon CloudWatch documentation.
Amazon WorkSpaces	Amazon WorkSpaces Documentation.
AWS Site-to-Site VPN Connection	Help for configuring AWS site-to-site VPNs.
AWS Transit Gateway	Help for AWS Transit Gateways.
AWS Customer Gateway	Help for AWS Customer Gateways.
AWS Snowball Appliance	AWS Snowball Edge Developer Guide.
AWS Snowball Supported OS	A list of specific AMIs that are supported by the Snowball Edge devices.
Amazon Q Getting Started	Getting started guide for Amazon Q developers.
Set Up Amazon Bedrock	Online documentation for setting up Amazon Bedrock.

Getting Started

This guide helps AWS users deploy Zscaler Internet Access (ZIA), Zscaler Private Access (ZPA), and Zscaler Cloud Connector in their environments. After implementing Traffic Forwarding Options, traffic from both workloads and end users is routed through the Zscaler cloud, where security policies are applied and enforced. This ensures consistent zero trust protection across your AWS infrastructure. After traffic is successfully flowing through the Zscaler cloud, you can begin enabling the additional integrations detailed later in this document.

Traffic Forwarding Options

The following sections describe the traffic forwarding options.

Zscaler Client Connector for Users

Zscaler Client Connector includes Amazon's support of Microsoft Windows 10/11 Desktop in WorkSpaces using the [Bring Your Own Windows Desktop Licenses](#). Additionally, Zscaler Client Connector version 4.5.x for Windows supports Windows Server 2022 bundles in WorkSpaces.

You must get the Zscaler Client Connector installation file for Windows from your administrator (there are no publicly accessible download links). This guide uses the 64-bit EXE version of Zscaler Client Connector version 4.5.x or later for Windows.

To install the file:

1. Double-click the Zscaler Client Connector installation file to start the installation. Click **Yes** when asked if you want to allow this app to make changes to your drive.

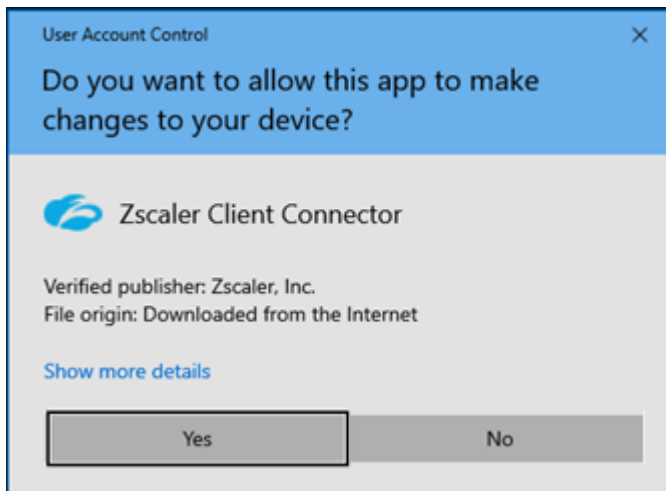


Figure 1. Allow device changes

2. Click **Next** to start the **Zscaler Setup** wizard.

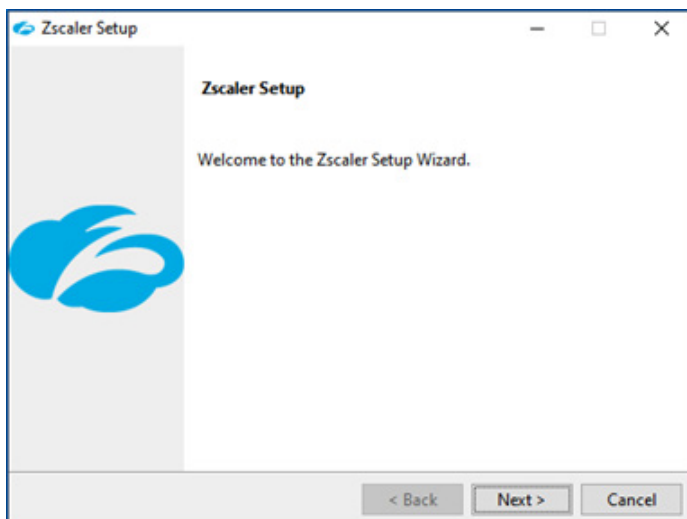


Figure 2. Zscaler Setup wizard

- Click **Next** to begin the installation.

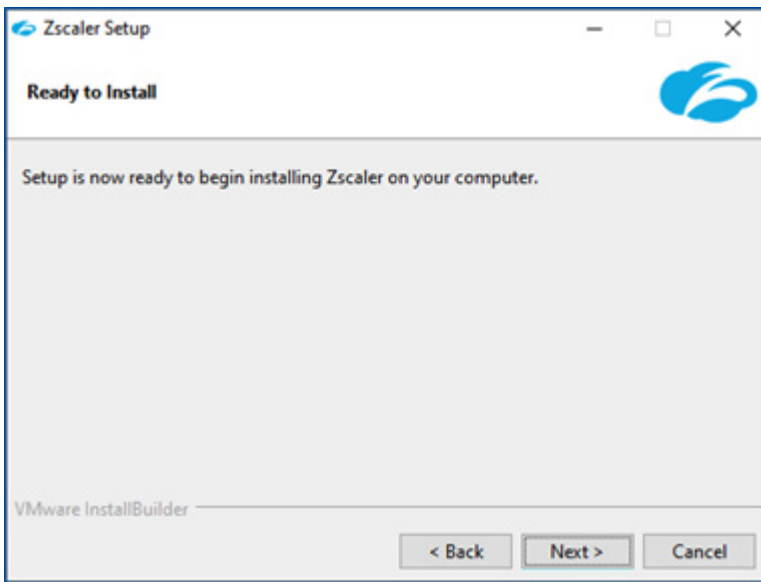


Figure 3. Install file

- Click **Finish** when the installation is completed.

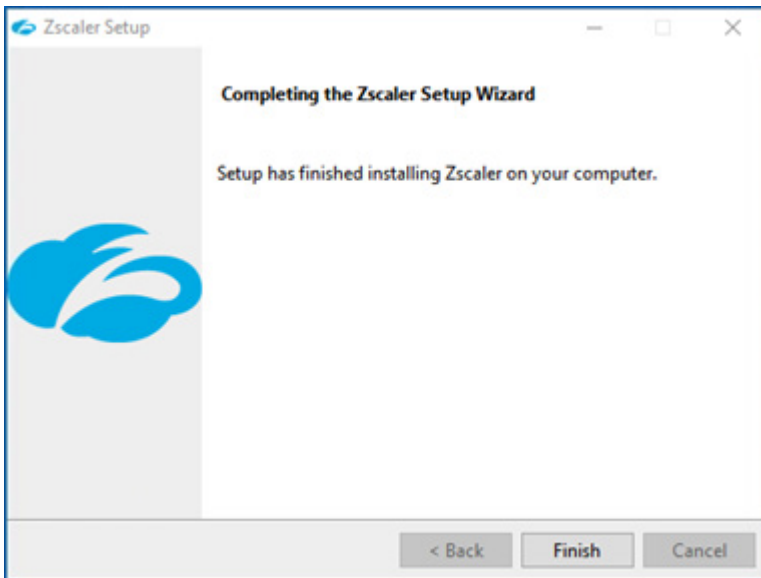


Figure 4. Installation completed

5. In the **Zscaler Client Connector** window, enter your username. After entering your username, you are redirected to your Identity Provider (IdP) to complete authentication. When authentication is successful, the Zscaler Client Connector window closes.

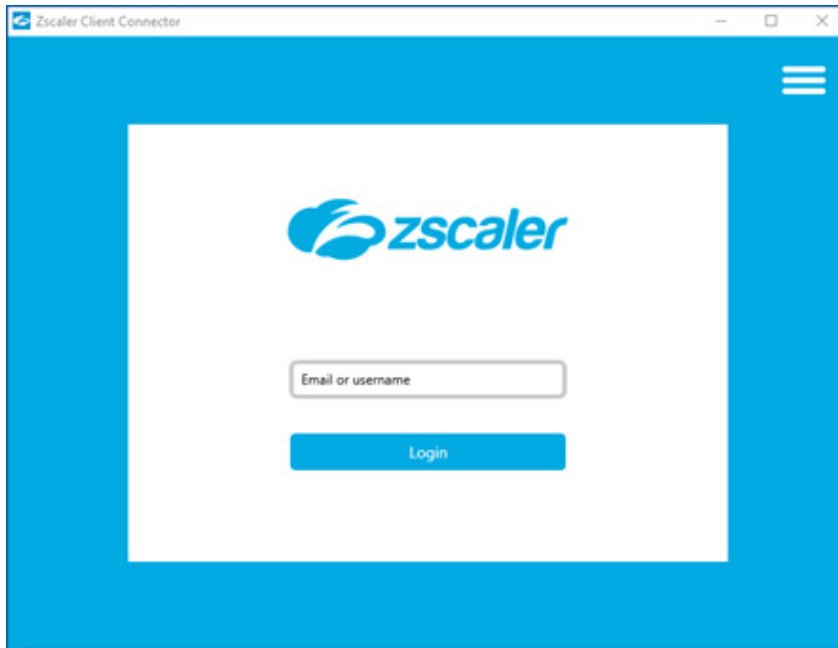


Figure 5. Zscaler Client Connector window



For detailed instructions on deploying Zscaler Client Connector, see [Zscaler Client Connector](#) (government agencies, see [Zscaler Client Connector](#)).

6. Zscaler Client Connector was tested on WorkSpaces with Tunnel (both Z-Tunnel 1.0 and Z-Tunnel 2.0) and Tunnel with Local Proxy forwarding. To learn more, see [Configuring Forwarding Profiles for Zscaler Client Connector](#) (government agencies, see [Configuring Forwarding Profiles for Zscaler Client Connector](#)).

a. Go to **App Profiles > Edit Window**.

b. In the **Destination Exclusions for IPv4** configuration, copy and paste the following list to quickly add the needed addresses into the **Destinations Exclusions** dialog box:

10.0.0.0/8, 100.64.0.0/10, 172.16.0.0/12, 192.168.0.0/16, 198.18.0.0/15,
 198.19.0.0/16, 172.31.0.0/16, 54.239.224.0/20, 54.239.236.220/32,
 127.0.0.2/32, 127.0.0.1/32, 169.254.169.123/32, 169.254.169.249/32,
 169.254.169.250/32, 169.254.169.251/32, 169.254.169.253/32,
 169.254.169.254/32:

Edit Windows Policy

Z-TUNNEL 2.0 CONFIGURATION

Application Bypass ?

None Selected

Destination Exclusions for IPv4 ? v. 2.0.0+

Use Enter to Add Multiple Items +

172.16.0.0/12	×
192.168.0.0/16	×
224.0.0.0/4	×
198.19.0.0/16	×

Figure 6. Edit Windows Policy



If you enter duplicate IP addresses, Zscaler Client Connect sends an error. Double check that the IP addresses you want to add aren't already listed in Destination Exclusions for IPv4.

For detailed information on WorkSpaces IP/port requirements, refer to the [Amazon documentation](#).



Zscaler changes the Windows firewall profile (aka network category) for the first network interface (eth0) to a domain network that can cause connection failures to the WorkSpace, or the WorkSpace to report as unhealthy.

To prevent this, add a new inbound rule to allow TCP port 8200-8250. Example using PowerShell:

```
New-NetFirewallRule -DisplayName "Allow TCP Port 8200-8250" -Direction Inbound
-LocalPort 8200-8250 -Protocol TCP -Action Allow
```

- For deploying at scale, use a dedicated Windows image with the previous changes already configured, and Zscaler Client Connector installed (but no one logged in) to create an Image/Bundle for deploying WorkSpaces.

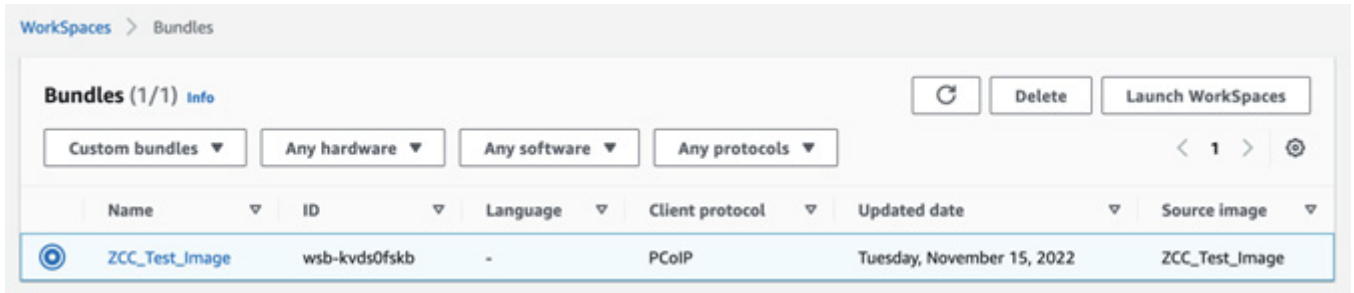


Figure 7. Bundles

- When creating a new workspace, select the custom bundle just created from the **Custom bundles** drop-down menu.
- Select **Next**.

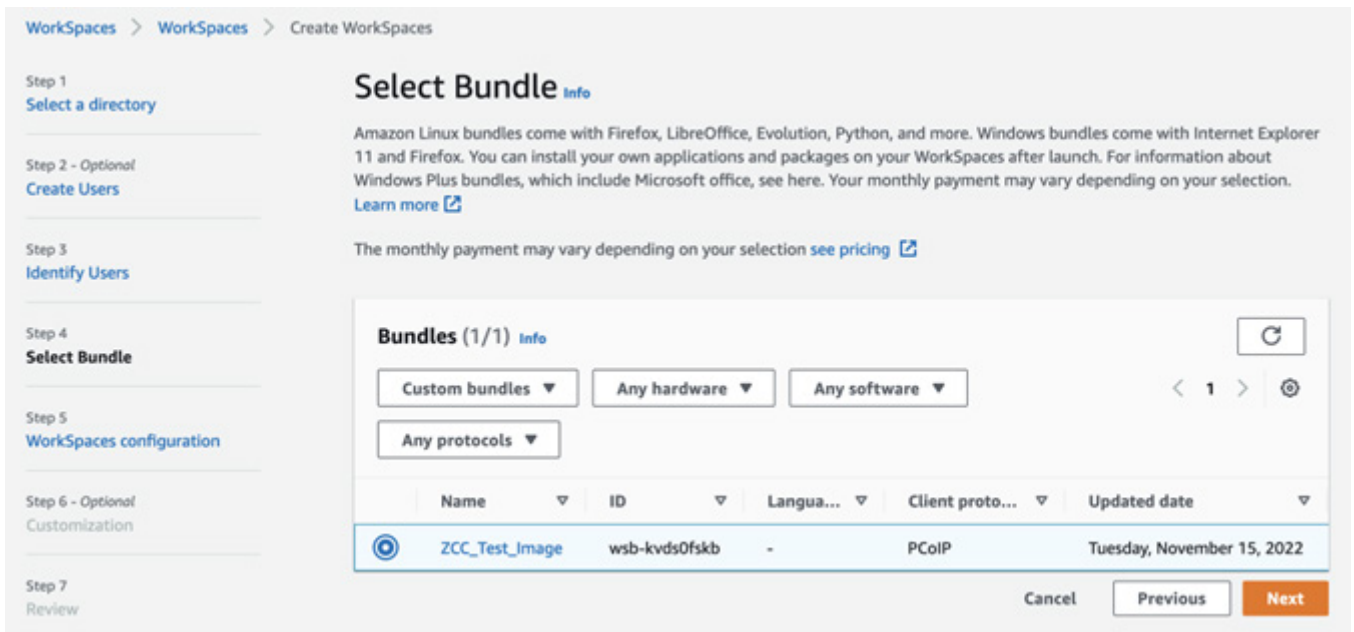


Figure 8. Select Bundle

10. Upon logging in to the new workspace, log in to Zscaler Client Connector.

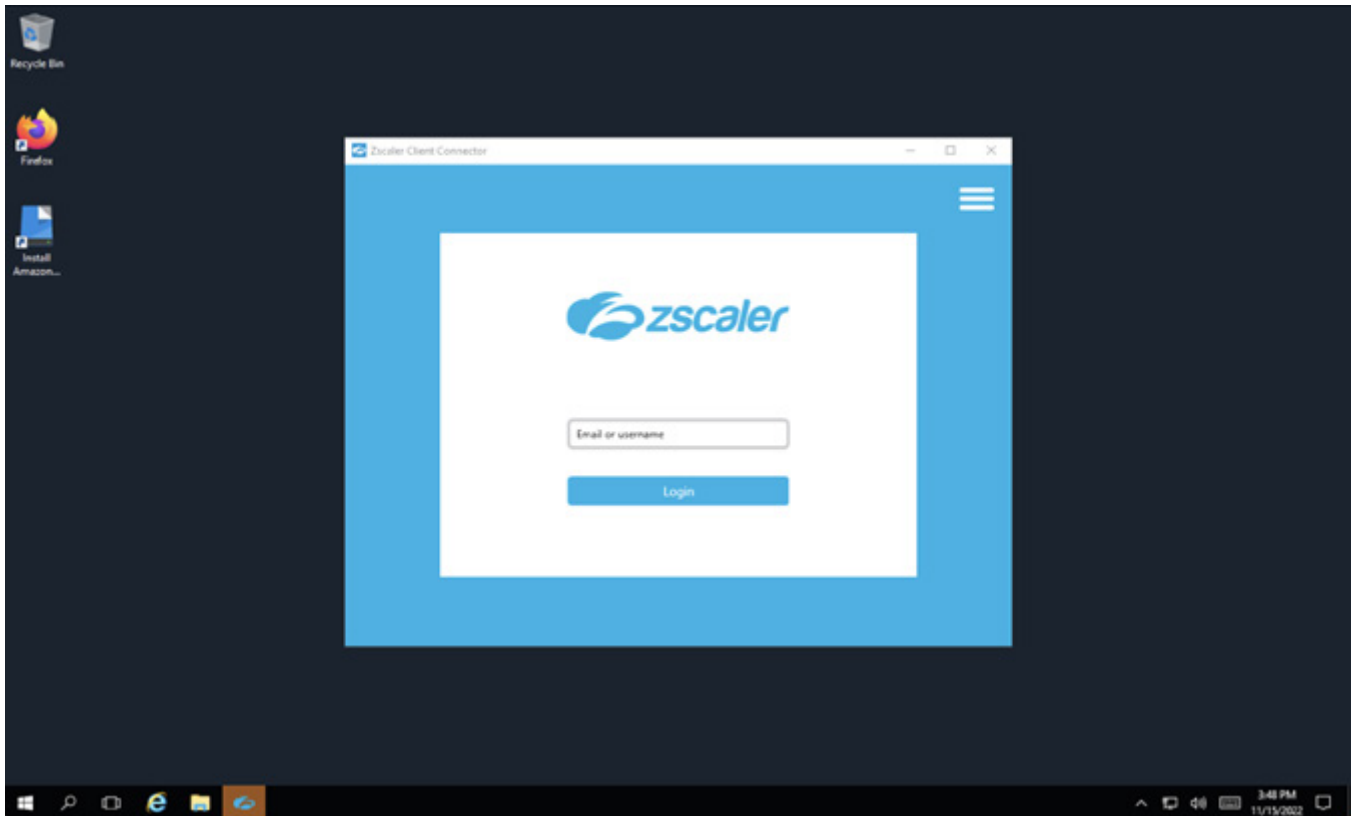


Figure 9. Zscaler Client Connector

PAC Files for Users or Workloads

You can configure any current browser to use a [PAC file](#) (government agencies, see [PAC file](#)) to forward traffic through a proxy such as ZIA. To uniquely identify WorkSpaces traffic for more granular policy control in ZIA, Zscaler recommends that you either use [a custom PAC file](#) (government agencies, see [custom PAC file](#)) with a [Dedicated Proxy Port](#) (government agencies, see [Dedicated Proxy Port](#)) and corresponding [Location](#) (government agencies, see [Location](#)), or define a location using the Elastic IP associated with WorkSpaces. You can then use this location as part of the criteria to make policy decisions for WorkSpaces web traffic.

You can find Information on configuring a browser to use a PAC File on the Zscaler Help Portal, including Chrome, Edge and Firefox. By default, Firefox on WorkSpaces is configured to use the same system proxy settings as Edge. However, the two browsers handle installing a Certificate for SSL inspection differently: Edge uses the system default certificate store, while Chrome and [Firefox use their own certificate stores](#) (government agencies, see [Firefox use their own certificate stores](#)) respectively.

SSL inspection is an option, using the [Zscaler Intermediate Certificate or a Custom Intermediate Root Certificate](#) (government agencies, see [Zscaler Intermediate Certificate or a Custom Intermediate Root Certificate](#)).

Zscaler Cloud Connector for Users or Workloads

Zero Trust Exchange (ZTE) enforces consistent security policies for cloud workloads accessing both public and private resources. It intelligently routes outbound and east-west traffic to the ZIA and ZPA platforms, ensuring all traffic is inspected and policy-enforced through the ZTE. This service supports secure, identity-based traffic handling across multi-cloud environments—without requiring any configuration changes on the source workloads.

Zero Trust Cloud is available as a fully managed Zero Trust Gateway service or as a self-managed deployment using EC2-based Cloud Connector appliances, giving organizations flexibility based on their operational model and cloud architecture.

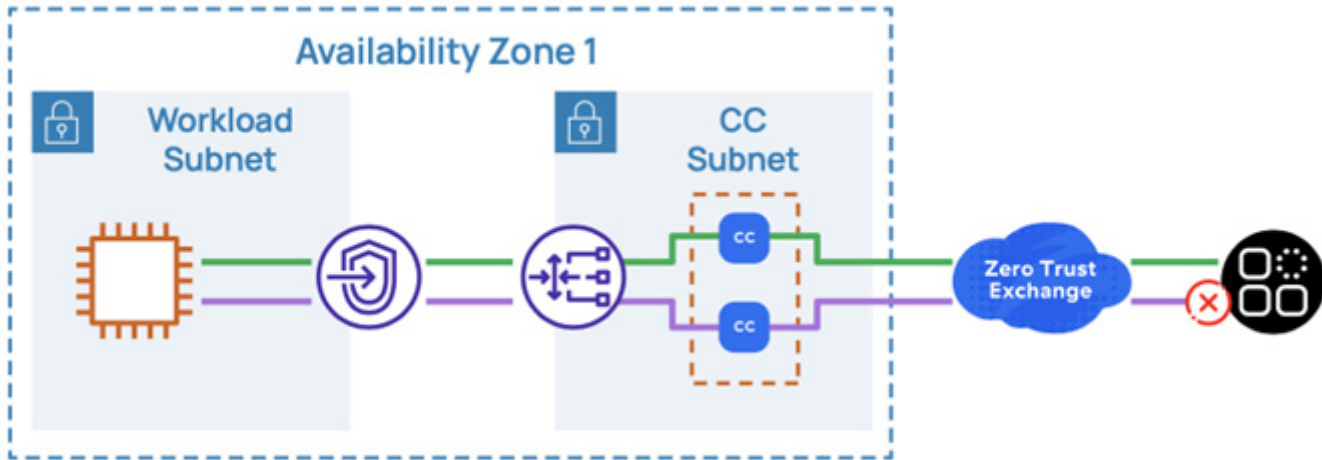


Figure 10. Cloud Connector

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](#)).

Deploying Zero Trust Cloud Automatically Zero Trust Gateway—Recommended

Zscaler Zero Trust Gateway (ZTGW) is a cloud-native SaaS offering that allows organizations to secure workload traffic in AWS without having to deploy or manage Cloud Connector infrastructure manually. Unlike the traditional Cloud Connector deployment, ZTGW is fully managed by Zscaler and provides high availability and scalability out of the box. Customers only need to configure the gateway and route workload traffic to it.

ZTGW is deployed and managed through the Zscaler Cloud & Branch Connector Portal and integrates into AWS using Gateway Load Balancer (GWLB) endpoints. It supports multiple deployment architectures, including centralized inspection with Transit Gateway, decentralized endpoints in each VPC, and hybrid combinations of both.

Deployment Steps

1. Go to **Administration > Partner Integrations**. Onboard your AWS account in the Zscaler Cloud & Branch Connector Admin Portal. This allows Zscaler to discover AWS resources and use tags in policy. To learn more, see [Adding an Amazon Web Services Account](#) (government agencies, see [Adding an Amazon Web Services Account](#)).
2. In the Zscaler Cloud & Branch Connector Admin Portal, go to **Administration > Zero Trust Gateway** and click **Add New Gateway**:

Zero Trust Gateway

AWS

[+ Add New Gateway](#)

Name	ID	Region	Availability Zone ID	Endpoint Service	Location	Endpoints	Operational Status	Service Status
us-west-1-ZTG	21705296	us-west-1 (N. California)	usw1-az3, usw1-az1	com.amazonaws	us-west-ztgw	1	Enabled	Healthy
us-west-2-ztgw	21726133	us-west-2 (Oregon)	usw2-az1, usw2-az2	com.amazonaws	us-west-2-loc	1	Enabled	Healthy

Rows per page: 10 | 1-2 of 2

Figure 11. Zero Trust Gateway

3. Fill in the required configuration:
 - a. **Name:** Logical name for the gateway.
 - b. **Region and Availability Zones:** AWS location for deployment.
 - c. **Location and Location Template:** Use default or custom as needed.

Add Zero Trust Gateway

Configuration

Enter the configuration for this gateway. This configuration will drive the AWS region and availability zones where the service will be available.

Name ⓘ

us-west-1-ZTGW

Region ⓘ

us-west-1 (N. California) ▼

Availability Zone ID ⓘ

2 items selected ▼

Location ⓘ

us-west-ztgw1

Location Template ⓘ

Default Location Template ▼

[Cancel](#) [Back](#) [Next](#)

Figure 12. Configuration

4. Assign AWS accounts:
 - a. Select discovered accounts from **Partner Integration**.
 - b. (Optional) Add accounts manually if discovery is not enabled.

Add Zero Trust Gateway

Accounts

The gateway will accept incoming endpoint requests from a list of accounts entered here. AWS accounts / account groups onboarded on the Partner Integration page can be selected from the dropdowns. For accounts not onboarded using partner integration please enter the 12 digit AWS account ID manually. Read more about partner integration [here](#).

Allowed Accounts 1 item selected ▼

Allowed Accounts Groups Select ▼

Additional AWS Accounts

Add Items Add Items

Cancel Back Next

Figure 13. Accounts

5. Review and create the gateway. Zscaler automatically provisions the infrastructure and return a healthy, ready-to-use gateway instance.

Add Zero Trust Gateway

Review

Ensure all the information below is correct before creating the gateway. Once the gateway is created it takes a few minutes to deploy the components. Please visit the Zero Trust Gateway page to view the status of the gateway.

Configuration

Name	us-west-1-ZTGW
Region	us-west-1 (N. California)
Availability Zone ID	USW1-AZ1 USW1-AZ3
Location	us-west-ztgw1
Location Template	Default Location Template

Accounts

Allowed Accounts	AS-Tag
Allowed Accounts Groups	---
Additional AWS Accounts	---

Cancel Back Create

Figure 14. Review

6. Hover over the **Activate** menu and click **Activate Now**.
7. In the AWS Management Console, go to **VPC > Endpoints** and create a new VPC Endpoint using GWLB.
 - a. Use the service name shown in the **Gateway** details.

Gateway
Status
Endpoints
Config
Analytics
Events
Traffic Test

Account Name	us-west-1-ZTGW	Zero Trust Gateway ID	21705296
Endpoint Service Name	com.amazonaws.vpce.us-west-1.vpce-svc-03b7944e6f7d60da8	Allowed Accounts	AS-Tag
Region	US_WEST_1	Allowed Account Groups	---
Availability Zones	us-west-1c, us-west-1a	Account List	---
Location	us-west-ztgw		
Public IPs	---		
usw1-az1	54.177.87.54	usw1-az3	52.9.160.88
Operational Status	Enabled		

Figure 15. Gateway details

Use the endpoints shown in **Endpoints**.

The screenshot shows the AWS VPC dashboard. In the left sidebar, under 'Virtual private cloud', the 'Endpoints' link is highlighted with a red box. The main content area, titled 'Resources by Region', shows a grid of VPC resources for 'US West 25'. The resources listed include VPCs, Subnets, Route Tables, Internet Gateways, Egress-only Internet Gateways, DHCP option sets, Endpoints, and Instance Connect Endpoints. The 'Endpoints' resource is highlighted with a blue bar and shows a count of 11. A red box is drawn around the 'Endpoints' link in the sidebar.

Figure 16. Endpoints

Select the endpoint settings.

Endpoint settings
Specify a name and select the type of endpoint.

Name tag - optional
(Name is a tag with a key of Name and a value that you specify. Tag help you find and manage your endpoints.)
ZTGW-Endpoint

Type [+/-](#)
Select a category.

- ☐ AWS services
Connect to services provided by Amazon with an interface endpoint, or a Gateway endpoint.
- ☐ PrivateLink Ready partner services
Connect to third services which have AWS Service Ready designation with an interface endpoint. Uses AWS PrivateLink.
- ☐ AWS Marketplace services
Connect to third services that you have purchased through AWS Marketplace with an interface endpoint.
- ☐ EC2 Instance Connect Endpoint
An elastic network interface that allows you to connect to resources in a private subnet.
- ☐ Resourcins - New
Connect to resources like Amazon Relational Database Services (RDS) with a Resource endpoint. Uses AWS PrivateLink.
- ☐ Service networks - New
Connect to VPC Lattice service network with a Service network endpoint. Uses AWS PrivateLink.
- ☒ Endpoint services that use NLBs and GWLBs
Find services shared with you by service name. Connect to a Network Load Balancer (NLB) service with an interface endpoint or to a Gateway Load Balancer (GWLB) service with a Gateway Load Balancer endpoint.

Service settings

Service name
com.amazonaws.vpce.us-west-1.vpce-svc-04024cc8a0ca071c [Verify service](#)

Service name verified.

Figure 17. Endpoint settings

b. Select the appropriate VPC and subnet.

Service name
com.amazonaws.vpce.us-west-1.vpce-svc-04024cc8a0ca071c [Verify service](#)

Service name verified.

Network settings
Select the VPC in which to create the endpoint.

VPC
Create the VPC endpoint in the VPC in the same AWS Region from which you will access a resource.
vpc-0f4b6d35014042c7 (AS-ZTGa5-vpc)

Subnets (1/2) [+/-](#)

Availability Zone	Subnet ID	Designate IP addresses	IPv4 address	IPv6 address
us-west-1b (us-west-1)	subnet-02a3e68dd259ec323	<input type="checkbox"/>		
us-west-1c (us-west-1)	Select a subnet	<input type="checkbox"/>		

IP address type
☒ IPv4
☐ IPv6
☐ Dualstack

Tags
No tags associated with the resource.
[Add new tag](#)
You can add 50 more tags.

[Cancel](#) [Create endpoint](#)

Figure 18. VPC and subnet

- After the endpoint is created, update the **Route Tables** in your workload subnets to route traffic (e.g., 0.0.0.0/0 or specific CIDRs) to the new VPC Endpoint.

VPC dashboard [×](#)

EC2 Global View [+](#)

Filter by VPC
vpc-0f4b6d35014042c7

Subnets (1/1) [+/-](#)

Name	Subnet ID	State	VPC	IPv4 CIDR	IPv6 CIDR
AS-ZTGa5-subnet-private1-us-west-1b	subnet-02a3e68dd259ec323	Available	vpc-0f4b6d35014042c7 (AS-ZTGa5-vpc)	10.0.128.0/20	-

Route table: rtb-027a0878991a64e / AS-ZTGa5-rb-private1-us-west-1b

[Details](#) [Flow logs](#) [Route table](#) [Network ACL](#) [CIDR reservations](#) [Sharing](#) [Tags](#)

[Edit route table association](#)

Routes (2)

Destination	Target
0.0.0.0/0	vpce-04024cc8a0ca071c
10.0.0.0/16	local

Figure 19. Route tables

- After routing is in place, all workload traffic flows through the ZTGW for inspection and policy enforcement. ZTGW also includes a built-in traffic test feature to simulate HTTP/HTTPS transactions through the gateway without creating your own endpoints. This helps verify connectivity and functionality quickly.



ZTGW also includes a built-in traffic test feature to simulate HTTP/HTTPS transactions through the gateway without creating your own endpoints. This helps verify connectivity and functionality quickly.

Deploying Zero Trust Cloud EC2 Instances Manually

The following sections describe deploying Zero Trust Cloud EC2 instances manually.

Pre-Deployment Setup

- 1. Go to **Administration > Role Management**.
- 2. Create a provisioning role with full access to Cloud Connector Provisioning and Location Management.

Add Admin Role

ADMINISTRATOR ROLE

Name

provisioning

PERMISSIONS

Dashboard

View Only

✓

None

Cloud Connector Provisioning

✓

Full

View Only

None

Template (Location & Provisioning)

Full

View Only

✓

None

Administrator Management

Full

View Only

✓

None

Location Management

✓

Full

View Only

None

Forwarding (Traffic, DNS & Logs)

Full

View Only

✓

None

API Key Management

Full

View Only

✓

None

Remote Assistance Management

Full

✓

View Only

NSS Logging

Full

✓

None

Public Cloud Config Management

Full

✓

View Only

None

Figure 20. Add Admin Role

3. Go to **Administration > Administrator Management**. Create a provisioning admin (e.g., provisioning) and assign the provisioning role from the previous step.

Add Admin [X]

ADMINISTRATOR

Login ID
provisioning @ [v]

Email
provisioning@zscaler.com

Name
Provisioning

Role
provisioning [v]

Status
Enabled [v]

Scope
Organization [v]

Comments
[Text Area]

SET PASSWORD

Figure 21. Add Admin

4. Go to **Administration > API Key Management**. Generate an API key (if necessary) or copy the existing key from the Zscaler Cloud & Branch Connector Admin Portal.
5. (Optional) Define a location template to customize naming conventions and enable features like XFF Forwarding or IPS.

6. Go to **Administration > Provisioning & Configuration**. Create a provisioning template for AWS, with or without Auto Scaling, and copy the template URL.
 - a. Enter a **Name** and click **Next**.
 - b. Select **Amazon Web Services** and click **Next**.
 - c. Select a **Location Template** and click **Next**.
 - d. Choose the **Small instance** size (Auto Scale only supports Small) and click **Next**.
 - e. Choose whether to enable **Auto Scale** and click **Next**.
 - f. Click **Save**.
 - g. Copy the **Template URL**:

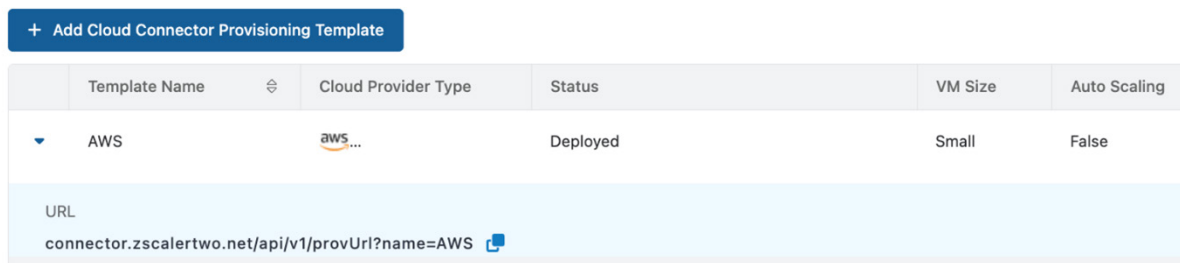


Figure 22. Template URL

7. Hover over the **Activate** menu and click **Activate Now**.

Prepare AWS Networking

Before deploying the appliance, you must ensure the VPC networking is properly configured to support Zscaler Cloud Connectors. The architecture is typically based on a hub-and-spoke model or dedicated inspection VPC, and the subnets must allow outbound access to the internet via NAT.

1. Create an **Inspection** VPC (if not already existing). This VPC hosts the Cloud Connector EC2s and related components.
 - a. Choose the **Region** in which you want to deploy.
 - b. Create the VPC with a CIDR block sufficient to support multiple EC2 instances and AWS services (Zscaler recommends /24 or greater block).
2. Create a **Service** subnet per AZ in your Inspection VPC (Zscaler recommends a minimum of /27 or greater).
3. (Optional) Create a **Management** subnet per AZ in your inspection VPC (Zscaler recommends a minimum of /27 or greater).
4. Create a **Public** subnet per AZ in your Inspection VPC (Zscaler recommends a minimum of /27 or greater).
5. Create an **Attachment** subnet per AZ in your Inspection VPC (Zscaler recommends a minimum of /27. This subnet host the Transit Gateway attachment point and Gateway Load Balancer endpoint).
6. Create and assign an Internet Gateway to the Inspection VPC.
7. Create a NAT Gateway per AZ in the **Public** subnets and assign an Elastic IP.
8. Create a **Public Route Table** for each **Public** subnet and assign the respective subnets it. Each Route Table should contain a single route for 0.0.0.0/0 directed towards the Internet Gateway.

9. Create a **Service Route Table** for each **Service** subnet and assign the respective subnets to it. Each Route Table should contain a single route for 0 . 0 . 0 . 0 / 0 directed towards the NAT Gateway in the AZ.
10. Create a **Management Route Table** for each **Management** subnet and assign the respective subnets to it. Each Route Table should contain a single route for 0 . 0 . 0 . 0 / 0 directed towards the NAT Gateway in the AZ.
11. Create a **Attachment Route Table** for each TGW **Attachment** subnet and assign the respective subnets to it. Each Route Table ultimately contains a single route for 0 . 0 . 0 . 0 / 0 directed towards the GWLB endpoint (created in the following steps).

Subscribe to the Cloud Connector AMI in AWS Marketplace

To run the CloudFormation or Terraform scripts to build the appliances, go to AWS Marketplace, search for Zscaler Cloud Connector, and click Subscribe (do not configure yet).

Store Credentials in AWS Secrets Manager

When the Cloud Connector appliance boots, it contacts AWS Secrets Manager to obtain the credentials required to register with the Zscaler Cloud & Branch Connector Admin Portal (created previously).

1. Go to **AWS Secrets Manager** and create a new **Secret**.
2. Choose **Other type of secret**.
3. In the **Key/Value pairs** section, click **Add Row** (3 rows total).
4. Add the following keys and their corresponding values in the blank fields: username, password api_key.

The screenshot shows the AWS Secrets Manager console. The 'Secret type' section has four radio buttons: 'Credentials for Amazon RDS database', 'Credentials for Amazon DocumentDB database', 'Credentials for Amazon Redshift data warehouse', and 'Other type of secret'. The 'Other type of secret' option is selected. Below this, the 'Key/value pairs' section is visible, showing a table with three rows of key-value pairs. The first row has 'username' and 'provisioning@yourdomain'. The second row has 'password' and 'somepassword'. The third row has 'api_key' and 'abcd1234'. Each row has a 'Remove' button to its right. At the bottom of the table is a '+ Add row' button.

Secret type	
<input type="radio"/> Credentials for Amazon RDS database	<input type="radio"/> Credentials for Amazon DocumentDB database
<input type="radio"/> Credentials for Amazon Redshift data warehouse	<input checked="" type="radio"/> Other type of secret API key, OAuth token, other.
<input type="radio"/> Credentials for other database	

Key/value pairs	
Key/value	Plaintext
username	provisioning@yourdomain
password	somepassword
api_key	abcd1234

+ Add row

Figure 23. Secret Type

Create EC2 Key Pair

If deploying via CloudFormation (not required with Terraform), you must create an SSH Keypair to be used for authentication when accessing the appliance for troubleshooting. Go to EC2 > Key Pairs, create a new RSA .pem key, and download it for SSH access.

Deploy the Cloud Connector Appliance

Option 1: CloudFormation

1. Using the templates from the [Zscaler GitHub repo](#), in the AWS Management Console, go to **CloudFormation** > **Create New Stack**.
 - a. Click **Choose** an existing template.
 - b. Click **Upload** a template file.
 - c. Click **Choose file** and select the **CloudFormation** template you want to deploy.

Create stack

Prerequisite - Prepare template

You can also create a template by scanning your existing resources in the [IaC generator](#).

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

☒ **Choose an existing template**
Upload or choose an existing template.

☐ **Build from Infrastructure Composer**
Create a template using a visual builder.

Specify template Info

This [GitHub repository](#) contains sample CloudFormation templates that can help you get started on new infrastructure projects. [Learn more](#)

Template source
Selecting a template generates an Amazon S3 URL where it will be stored. A template is a JSON or YAML file that describes your stack's resources and properties.

☐ **Amazon S3 URL**
Provide an Amazon S3 URL to your template.

☒ **Upload a template file**
Upload your template directly to the console.

☐ **Sync from Git**
Sync a template from your Git repository.

Upload a template file

zs_cc_cf_template_simple.yaml

JSON or YAML formatted file

Figure 24. Create stack

- d. Deploy Macro stack (zs_cc_cf_template_zscc_macro.yaml) first. There are no variables required for this script.
- e. For a single instance, deploy the zs_cc_cf_template_simple.yaml stack using the values created in the previous section(s). For more information, see the [Deploying Zscaler Cloud Connector with Amazon Web Services](#) (government agencies, see [Deploying Zscaler Cloud Connector with Amazon Web Services](#)).
- f. For a GWLB deployment (without ASG), re-run the simple stack to create a second Cloud Connector appliance. Then, deploy the GWLB stack to deploy a Gateway Load Balancer (zs_cc_cf_template_gwlb.yaml). For more information, see [Deployment Templates for Zscaler Cloud Connector](#) (government agencies, see [Deployment Templates for Zscaler Cloud Connector](#)).
- g. (Optional) If you wish to enable Auto Scale, deploy the zs_cc_cf_template_asg_gwlb.yaml stack. For more information, see the help document [Deployment Templates for Zscaler Cloud Connector](#) (government agencies, see [Deployment Templates for Zscaler Cloud Connector](#)).

Option 2: Terraform

Before deploying using Terraform, ensure you have AWS CLI and Terraform CLI installed on your machine. You must have sufficient IAM permissions to deploy resources (EC2, VPC, IAM, Secrets Manager, etc.), which are noted in the [README.md](#) file of the Github repository. Also, have your Zscaler provisioning template URL, API key, and provisioning admin credentials handy as they are input as variables in the installation.

1. Clone [Zscaler's Terraform module repo](#).
2. Choose a deployment method:
 - The easiest option is to use the zsec script for interactive CLI setup, which prompts for appropriate variables and desired topology.
 - Alternatively, edit the terraform.tfvars manually with your desired variables. Then, run terraform init, terraform plan, and terraform apply.
 - As another alternative, you can integrate the Terraform directly into your CI/CD pipeline by referencing the Terraform modules directly:

```
module "cloud_connector" {
    source = "github.com/zscaler/terraform-aws-cloud-connector-modules//modules/"
    ...
}
```

Implementing Zscaler Private Access with Zero Trust Cloud

Zscaler Private Access (ZPA) enables secure, identity-based access to internal applications hosted in AWS without exposing them to the public internet. When integrated with Zero Trust Cloud (Cloud Connector), ZPA allows private workload-to-workload communication (east-west) and secure access to private services across VPCs or even across hybrid/multi-cloud environments.

Whether you choose the manual or automatic method to deploy, after the Cloud Connectors have been, they automatically integrate with both the ZIA and ZPA fabric. There is nothing additional required in the ZPA Admin Portal beyond basic ZPA administration (such as deploying App Connectors, creating Application Segments and Access Policy). For more information on administrating ZPA, see the [ZPA help documentation](#) (government agencies, see [ZPA help documentation](#)). There are several considerations around DNS that should be taken into account, however. DNS is critical for ZPA to resolve internal application domains to synthetic IPs used by Zscaler for tunneling:

- Workloads must use a DNS forwarder (such as Route 53 or a self-managed DNS server) that sends queries through the Cloud Connector. In other words, the Cloud Connector must see the DNS exchange for an application in order for it to properly proxy the request.
- Cloud Connector is automatically configured to intercept DNS requests, modify them, and forward appropriately into the ZPA fabric. You simply must ensure the DNS request flows through the Cloud Connector.
- Zscaler maps DNS queries to the synthetic IP addresses assigned by ZPA for proper routing through the TLS tunnel.
- Ensure that Route 53 Resolver or the DNS forwarder used supports conditional forwarding, especially for hybrid environments with on-premises DNS dependencies.



The Zscaler Github repository contains a CloudFormation script that you can use to implement Route 53 architecture for ZPA. This script allows you to define domains to conditionally forward to Cloud Connector for ZPA access while leaving non-ZPA traffic to resolve normally.

If you manually install Cloud Connector using Terraform, you might choose the option to deploy Route 53 as part of the zsec script. Otherwise, you might implement Route 53 using the provided modules in your CI/CD pipeline.



For more information on DNS considerations for ZPA, see [Handling DNS Resolutions for Zscaler Cloud Connector](#) (government agencies, see [Handling DNS Resolutions for Zscaler Cloud Connector](#)).

Post Deployment

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

Site-to-Site IPsec for Users or Workloads

AWS can send traffic from a virtual private cloud (VPC) to a remote gateway via a [Site-to-Site VPN Connection](#) using IPsec tunnels. This feature routes all traffic from a VPC, such as a WorkSpaces or workload VPC, to a ZIA Public Service Edge with the following caveats:

- An AWS Site-to-Site VPN provides redundant tunnels to the same destination. Zscaler recommends that redundant tunnels use two geographically disparate data centers for failover.
- An AWS Site-to-Site VPN does not support NULL encryption for Phase 2, which requires the Zscaler Encrypted VPN subscription option to allow encrypted IPsec tunnels.
- An AWS Site-to-Site VPN does not support the [Zscaler recommended](#) (government agencies, see [Zscaler recommended](#)) IPsec SA lifetime values.

An AWS Site-to-Site VPN Connection can use either a Virtual Private Gateway or a Transit Gateway. This document uses a Transit Gateway design, but the configuration for the Site-to-Site VPN Connection is the same. Refer to [Appendix A: AWS Transit Gateway Lab Environment](#) for a lab environment to use for testing.

Identifying the Zscaler VPN Endpoint

First, determine the VPN endpoint to be used in the Zscaler cloud by going to [Cloud Enforcement Node Ranges](#) and selecting your cloud at the top (e.g., zscaler.net). In the Current Data Centers list, locate the data center location closest to your AWS region and resolve the VPN Host Name to obtain the IP address to use when configuring the AWS VPN [Customer Gateway](#).

Location	IP Address (CIDR Notation)	Proxy Hostname	GRE Virtual IP	SVPN Virtual IP	VPN Host Name	Notes
Chicago	165.225.60.0/22		165.225.56.12			Multi-cluster VIP
	104.129.196.0/23	chi1.sme.zscaler.net	104.129.196.32	104.129.196.42	chi1-vpn.zscaler.net	Multi-cluster VIP
	165.225.56.0/22	chi1-2.sme.zscaler.net	165.225.56.12	165.225.56.28	chi1-2-vpn.zscaler.net	Multi-cluster VIP

Figure 25. Current Data Centers list

To resolve the hostname, use nslookup from the command line:

```
nslookup chi1-2-vpn.zscaler.net

Non-authoritative answer:

Name:      chi1-2-vpn.zscaler.net

Address: 165.225.56.14
```

Alternatively, you can use Method 2 as described in the [SD-WAN Integrations Using API](#) (government agencies, see [SD-WAN Integrations Using API](#)). Using your Elastic IP address, you can get an automated determination of the closest Zscaler Data Center location to the AWS region. Using the following URL (with your Zscaler cloud and AWS Elastic IP substituted for <Zscaler Cloud> and <Elastic IP>), the primaryIP value returned is the Zscaler VPN endpoint you are to use.

```
https://pac.<Zscaler Cloud>.net/getVpnEndpoints?srcIp=<Elastic IP>
```

To fetch the endpoints, use curl from the command line as shown as shown in the following code sample:

```
curl https://pac.zscaler.net/getVpnEndpoints?srcIp=3.20.82.111

{
  "primaryIp": "165.225.56.14",
  "primaryMeta": {
    "region": "NorthAmerica",
    "country": "United States",
    "city": "Chicago",
    "dcName": "CHI1",
    "latitude": 41.000000,
    "longitude": -87.000000
  },
  "secondaryIp": "104.129.194.33",
```

```
"secondaryMeta": {  
    "region": "NorthAmerica",  
    "country": "United States",  
    "city": "Washington, DC",  
    "dcName": "WAS1",  
    "latitude": 39.000000,  
    "longitude": -77.000000  
},  
"tertiaryIp": "165.225.208.18",  
"tertiaryMeta": {  
    "region": "NorthAmerica",  
    "country": "Canada",  
    "city": "Toronto",  
    "dcName": "YTO3",  
    "latitude": 44.000000,  
    "longitude": -79.000000  
}  
}
```

Create a Customer Gateway

After logging in to your AWS Management Console:

1. Select **VPC Service**.
2. On the **AWS portal VPC Service** page, select **Customer Gateways** under the **Virtual Private Network (VPN)** section.
3. Click **Create Customer Gateway**.

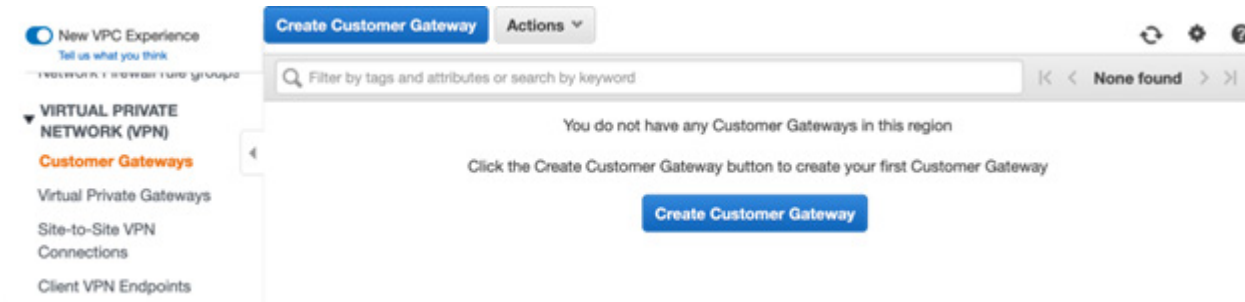


Figure 26. AWS Create Customer Gateway

4. On the **Create Customer Gateway** window:
 - a. Enter a **Name** for your **Customer Gateway**.
 - b. Select **Static** for **Routing**.
 - c. Enter the **IP Address** for your closest Zscaler VPN endpoint (determined previously).
 - d. Click **Create Customer Gateway**.

Customer Gateways > Create Customer Gateway

Create Customer Gateway

Specify the IP address for your gateway's external interface; the address must be static and may be behind a device performing network address translation (NAT). For dynamic routing, also specify your gateway's Border Gateway Protocol (BGP) Autonomous System Number (ASN); this can be either a public or private ASN (such as those in the 64512-65534 range).

VPNs can use either Pre-Shared Keys or Certificates for authentication. When using Certificate authentication, an IP address is optional. To use Certificate authentication, specify a Certificate ARN when you create your Customer Gateway. To use Pre-Shared Keys, only an IP address is required.

Name ⓘ

Routing ☐ Dynamic ☒ Static

IP Address ⓘ

Certificate ARN ⓘ ⓘ

Device ⓘ

* Required Cancel Create Customer Gateway

Figure 27. AWS Create Customer Gateway configuration

Create a Site-to-Site VPN Connection

On the VPC Service page:

1. Select **Site-to-Site VPN Connections** under the **Virtual Private Network (VPN)** section.
2. Click **Create VPN Connection**.

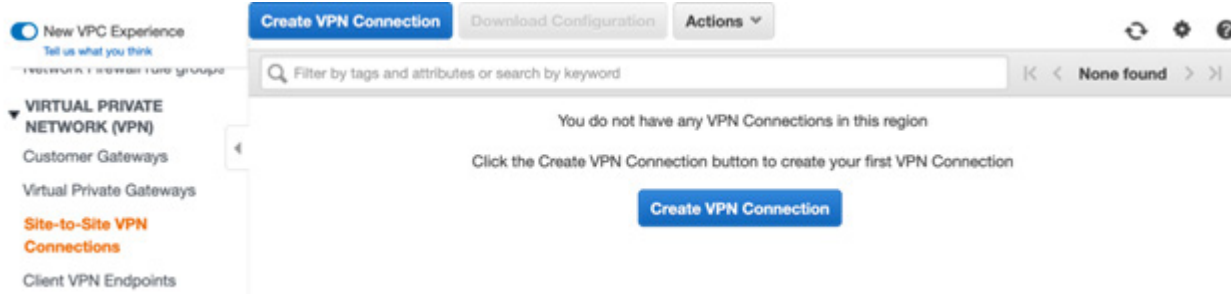


Figure 28. AWS Site-to-Site VPN Connections

3. In the **Create VPN Connection** window:
 - a. Enter a **Name** tag for your VPN Connection.
 - b. Select **Transit Gateway** for the **Target Gateway Type**.
 - c. Select your **Transit Gateway** from the drop-down menu.
 - d. Select the **Customer Gateway** you just created under **Customer Gateway ID**.
 - e. Select **Static** for **Routing Options**.
 - f. Select **IPv4** for the **Tunnel Inside Ip Version**.

Create VPN Connection

Select the target gateway and customer gateway that you would like to connect via a VPN connection. You must have entered the target gateway information already.

Name tag

Target Gateway Type
☐ Virtual Private Gateway
☒ Transit Gateway

Transit Gateway*

Customer Gateway
☒ Existing
☐ New

Customer Gateway ID*

Routing Options
☐ Dynamic (requires BGP)
☒ Static

Tunnel Inside Ip Version
☒ IPv4
☐ IPv6

Figure 29. AWS Create VPN Connection configuration

4. Scroll to **Advanced Options** for **Tunnel 1**, select **Edit Tunnel 1 Options**, then set the following options to only these values:
 - a. **Phase 1 Encryption Algorithms:** AES256
 - b. **Phase 2 Encryption Algorithms:** AES256
 - c. **Phase 1 Integrity Algorithms:** SHA2-256
 - d. **Phase 2 Integrity Algorithms:** SHA2-256
 - e. **Phase 1 DH Group Numbers:** 14
 - f. **Phase 2 DH Group Numbers:** 14
 - g. **IkeVersion:** ikev2
 - h. **DPD Timeout Action:** Restart
 - i. **Startup Action:** Start

Advanced Options for Tunnel 1 ☐ Use Default Options ☒ Edit Tunnel 1 Options

Phase 1 Encryption Algorithms ☐ AES128 ☒ AES256 ☐ AES128-GCM-16 ☐ AES256-GCM-16

Phase 2 Encryption Algorithms ☐ AES128 ☒ AES256 ☐ AES128-GCM-16 ☐ AES256-GCM-16

Phase 1 Integrity Algorithms ☐ SHA1 ☒ SHA2-256 ☐ SHA2-384 ☐ SHA2-512

Phase 2 Integrity Algorithms ☐ SHA1 ☒ SHA2-256 ☐ SHA2-384 ☐ SHA2-512

Phase 1 DH Group Numbers ☐ 2 ☒ 14 ☐ 15 ☐ 16 ☐ 17 ☐ 18 ☐ 19 ☐ 20 ☐ 21 ☐ 22 ☐ 23 ☐ 24

Phase 2 DH Group Numbers ☐ 2 ☐ 5 ☒ 14 ☐ 15 ☐ 16 ☐ 17 ☐ 18 ☐ 19 ☐ 20 ☐ 21 ☐ 22 ☐ 23 ☐ 24

IkeVersion ☐ ikev1 ☒ ikev2

Phase 1 Lifetime (seconds) ⓘ

Phase 2 Lifetime (seconds) ⓘ

Rekey Margin Time (seconds) ⓘ

Rekey Fuzz (percentage) ⓘ

Replay Window Size (packets) ⓘ

DPD Timeout (seconds) ⓘ

DPD Timeout Action ☐ Clear ☒ Restart ☐ None

Startup Action ☐ Add ⓘ ☒ Start

Figure 30. AWS advanced tunnel options (Tunnel 1)

5. Scroll to **Advanced Options for Tunnel 2** and select **Edit Tunnel 2 Options**. Select the same options and values as Tunnel 1.

Advanced Options for Tunnel 2 ☐ Use Default Options ☒ Edit Tunnel 2 Options

Phase 1 Encryption Algorithms ☐ AES128 ☒ AES256 ☐ AES128-GCM-16 ☐ AES256-GCM-16

Phase 2 Encryption Algorithms ☐ AES128 ☒ AES256 ☐ AES128-GCM-16 ☐ AES256-GCM-16

Phase 1 Integrity Algorithms ☐ SHA1 ☒ SHA2-256 ☐ SHA2-384 ☐ SHA2-512

Phase 2 Integrity Algorithms ☐ SHA1 ☒ SHA2-256 ☐ SHA2-384 ☐ SHA2-512

Phase 1 DH Group Numbers ☐ 2 ☒ 14 ☐ 15 ☐ 16 ☐ 17 ☐ 18 ☐ 19 ☐ 20 ☐ 21 ☐ 22 ☐ 23 ☐ 24

Phase 2 DH Group Numbers ☐ 2 ☐ 5 ☒ 14 ☐ 15 ☐ 16 ☐ 17 ☐ 18 ☐ 19 ☐ 20 ☐ 21 ☐ 22 ☐ 23 ☐ 24

ikeVersion ☐ ikev1 ☒ ikev2

Phase 1 Lifetime (seconds) ⓘ

Phase 2 Lifetime (seconds) ⓘ

Rekey Margin Time (seconds) ⓘ

Rekey Fuzz (percentage) ⓘ

Replay Window Size (packets) ⓘ

DPD Timeout (seconds) ⓘ

DPD Timeout Action ☐ Clear ☒ Restart ☐ None

Startup Action ☐ Add ⓘ ☒ Start

Figure 31. AWS advanced tunnel options (Tunnel 2)

6. Click **Create VPN Connection**. This automatically creates a Transit Gateway Attachment. The **Name tag** is empty, but **Resource type** is VPN.
7. Name the attachment (e.g., VPN-Attachment) for ease of identification later.
8. Select your newly created VPN Connection.
9. Click the **Tunnel Details** tab to see the Elastic IPs assigned to the tunnels in the **Outside IP Address** column. Notice that the **Status** is currently **Down** because you still must configure the Zscaler side.

VPN Connection: vpn-0f1d869b9faf85763

Details Tunnel Details Static Routes Tags

Tunnel State

Tunnel Number	Outside IP Address	Inside IPv4 CIDR	Inside IPv6 CIDR	Status
Tunnel 1	3.141.109.232	169.254.136.108/30	-	DOWN
Tunnel 2	18.218.203.101	169.254.153.40/30	-	DOWN

Figure 32. AWS Site-to-Site Connection Tunnel Details

10. Click **Download Configuration** at the top of the window.
11. Choose **Generic** for the **Vendor** and **ikev2** for the **Ike Version**.
12. Click **Download** to download the configuration.

Figure 33. Download configuration

13. Locate the **Pre-Shared Keys** for **Tunnel 1** and **Tunnel 2** in the downloaded file. The Elastic IPs for the tunnels and their corresponding Pre-Shared Keys are needed in the next section.

```

.
IPSec Tunnel #1
=====
#1: Internet Key Exchange Configuration

Configure the IKE SA as follows:
Please note, these sample configurations are for the minimum requirement of AES128, SHA1, and DH Group 2.
Category "VPN" connections in the GovCloud region have a minimum requirement of AES128, SHA2, and DH Group 14.
You will need to modify these sample configuration files to take advantage of AES256, SHA256, or other DH groups like 2, 14-18, 22, 23, and 24.
NOTE: If you customized tunnel options when creating or modifying your VPN connection, you may need to modify these sample configurations to match the custom settings for your tunnels.

Higher parameters are only available for VPNs of category "VPN," and not for "VPN-Classic".
The address of the external interface for your customer gateway must be a static address.
Your customer gateway may reside behind a device performing network address translation (NAT).
To ensure that NAT traversal (NAT-T) can function, you must adjust your firewall rules to unblock UDP port 4500.
| If not behind NAT, and you are not using an Accelerated VPN, we recommend disabling NAT-T. If you are using an Accelerated VPN, make sure that NAT-T is enabled.
- IKE version      : IKEv2
- Authentication Method : Pre-Shared Key
- Pre-Shared Key    : 1RQi1EYLBWu5LED02_3E56JYUL8Gy9Kw

.
IPSec Tunnel #2
=====
#1: Internet Key Exchange Configuration

Configure the IKE SA as follows:
Please note, these sample configurations are for the minimum requirement of AES128, SHA1, and DH Group 2.
Category "VPN" connections in the GovCloud region have a minimum requirement of AES128, SHA2, and DH Group 14.
You will need to modify these sample configuration files to take advantage of AES256, SHA256, or other DH groups like 2, 14-18, 22, 23, and 24.
NOTE: If you customized tunnel options when creating or modifying your VPN connection, you may need to modify these sample configurations to match the custom settings for your tunnels.

Higher parameters are only available for VPNs of category "VPN," and not for "VPN-Classic".
The address of the external interface for your customer gateway must be a static address.
Your customer gateway may reside behind a device performing network address translation (NAT).
To ensure that NAT traversal (NAT-T) can function, you must adjust your firewall rules to unblock UDP port 4500.
| If not behind NAT, and you are not using an Accelerated VPN, we recommend disabling NAT-T. If you are using an Accelerated VPN, make sure that NAT-T is enabled.
- IKE version      : IKEv2
- Authentication Method : Pre-Shared Key
- Pre-Shared Key    : XjnJ5ZLsv..MLDGtfo08UMr9Gp09t0n0
.

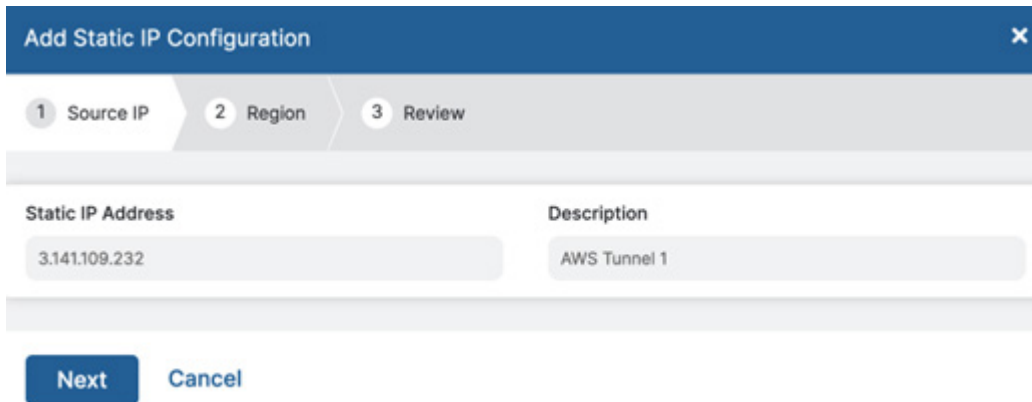
```

Figure 34. Downloaded tunnel configuration

Configure ZIA

In the ZIA Admin Portal:

1. Go to **Administration > Static IPs & GRE Tunnels**.
2. Select **Add Static IP**.
3. For **Static IP Address**, enter the outside IP address for Tunnel 1.
4. Enter a description.
5. Click **Next**.



Add Static IP Configuration [X]

1 Source IP 2 Region 3 Review

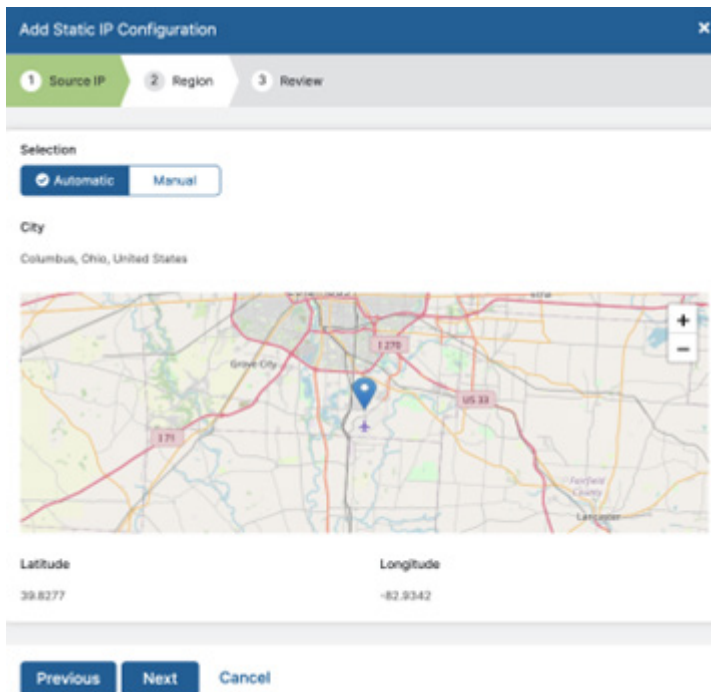
Static IP Address **Description**

3.141.109.232 AWS Tunnel 1

Next **Cancel**

Figure 35. Add Static IP Configuration page

6. Verify that the geographic location makes sense based on your AWS region.
7. Click **Next** and then **Save**. If the geographic location is not accurate, you can manually set it by **City** or **Latitude** and **Longitude**.
8. Repeat these configuration steps for the Tunnel 2 IP.



Add Static IP Configuration [X]

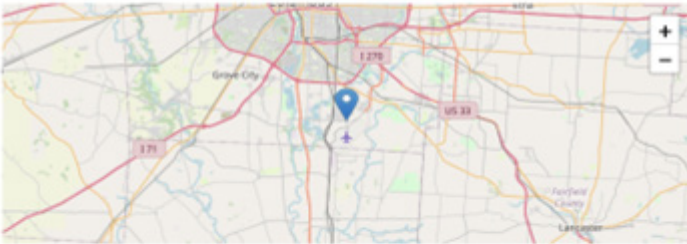
1 Source IP 2 Region 3 Review

Selection

☒ Automatic ☐ Manual

City

Columbus, Ohio, United States



Latitude **Longitude**

39.8277 -82.9342

Previous **Next** **Cancel**

Figure 36. Static IP location

9. Go to **Administration > VPN Credentials** and select **Add VPN Credentials**.
10. For **Authentication Type**, select **IP**.
11. Select your AWS Tunnel 1 **IP Address** from the drop-down menu and paste the associated **Pre-Shared Key** in the two fields.
12. Enter a comment, and click **Save**.
13. Repeat these steps for Tunnel 2 **IP Address** and associated **Pre-Shared Key**.

Add VPN Credential [X]

VPN CREDENTIAL

Authentication Type

FQDN XAUTH **IP**

IP Address

3.141.109.232

New Pre-Shared Key **Confirm New Pre-Shared Key**

Comments

AWS Tunnel 1

Save **Cancel**

Figure 37. Add VPN Credentials

14. Next, go to **Administration > Location Management** and select **Add Location**.
15. Enter a name for the location.
16. Select a **Location Type** (required).
17. Select the Tunnel 1 IP address from both the **Static IP Addresses and GRE Tunnels** and the **VPN Credentials** drop-down menus.
18. Click **Save**.
19. Repeat these steps for the Tunnel 2 IP and then **Activate** the changes.



Zscaler does not respond to tunnel initiation requests from AWS until the location configuration is activated.

Add Location

LOCATION

Name: AWS Tunnel 1

City/State/Province: Enter Text

Country: NONE

Time Zone: NONE

Manual Location Groups: None

Dynamic Location Groups: ---

Exclude from Manual Location Groups: ☒ X

Exclude from Dynamic Location Groups: ☒ X

Location Type: Corporate user traffic

Description:

ADDRESSING

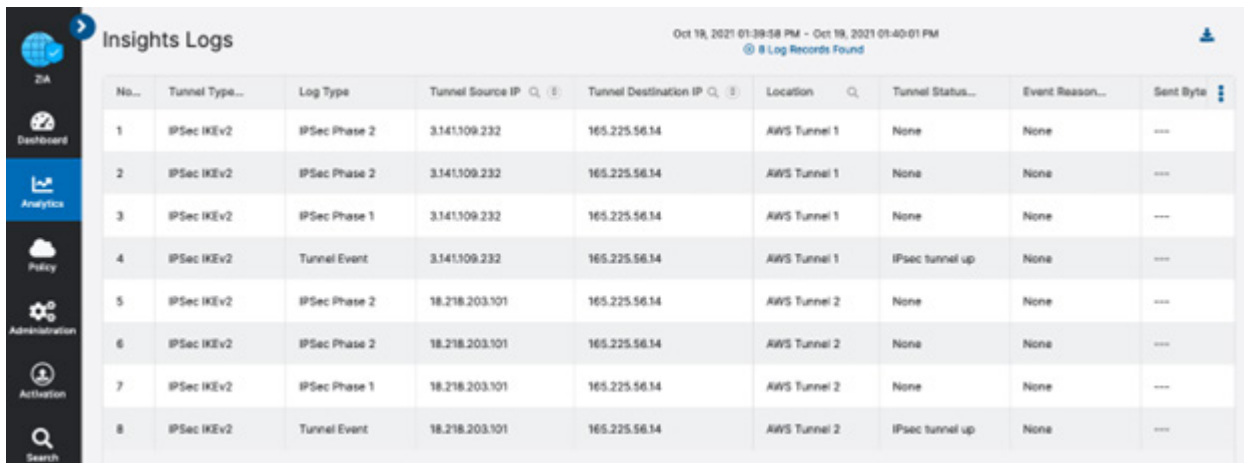
Static IP Addresses and GRE Tunnels: 3.141.109.232

Proxy Ports: None

VPN Credentials: 3.141.109.232

Figure 38. Add location configuration page

20. To verify that the tunnels are established, go to **Analytics > Tunnel Insights**.
21. Select **Logs**, and click **Apply Filters**. After a short time (you might need to refresh your view) both tunnels appear (**IPSec tunnel up**) in the **Tunnel Status** column.
22. If needed, add a filter for the AWS tunnel locations to limit the number of logs returned.



No...	Tunnel Type...	Log Type	Tunnel Source IP	Tunnel Destination IP	Location	Tunnel Status...	Event Reason...	Sent Byte
1	IPSec IKEv2	IPSec Phase 2	3.141.109.232	165.225.56.14	AWS Tunnel 1	None	None	----
2	IPSec IKEv2	IPSec Phase 2	3.141.109.232	165.225.56.14	AWS Tunnel 1	None	None	----
3	IPSec IKEv2	IPSec Phase 1	3.141.109.232	165.225.56.14	AWS Tunnel 1	None	None	----
4	IPSec IKEv2	Tunnel Event	3.141.109.232	165.225.56.14	AWS Tunnel 1	IPsec tunnel up	None	----
5	IPSec IKEv2	IPSec Phase 2	18.218.203.101	165.225.56.14	AWS Tunnel 2	None	None	----
6	IPSec IKEv2	IPSec Phase 2	18.218.203.101	165.225.56.14	AWS Tunnel 2	None	None	----
7	IPSec IKEv2	IPSec Phase 1	18.218.203.101	165.225.56.14	AWS Tunnel 2	None	None	----
8	IPSec IKEv2	Tunnel Event	18.218.203.101	165.225.56.14	AWS Tunnel 2	IPsec tunnel up	None	----

Figure 39. Tunnel Insights Logs page

23. In the AWS Management Console, in the **Site-to-Site VPN Connection** section and on the **Tunnel Details** tab for your **VPN Connection**, ensure that the **Status** is **Up**.



VPN Connection: vpn-0f1d869b9faf85763

Details Tunnel Details Static Routes Tags

Tunnel State

Tunnel Number	Outside IP Address	Inside IPv4 CIDR	Inside IPv6 CIDR	Status
Tunnel 1	3.141.109.232	169.254.136.108/30	-	UP
Tunnel 2	18.218.203.101	169.254.153.40/30	-	UP

Figure 40. AWS Site-to-Site Connection Tunnels Details

Configure Routing for Site-to-Site VPN Connection

You must route traffic to and from the active tunnels for your VPCs before traffic is sent to ZIA.

1. In the AWS Management Console **VPC Service** page, under **Transit Gateway Route Tables**, click **Create transit gateway route table**.
2. Enter an appropriate **Name**.
3. Select the Transit gateway ID from the drop-down menu and click **Create transit gateway route table**.

Figure 41. AWS Create transit gateway route table



The Transit gateway ID drop-down menu might be broken. In that case, use the [AWS CLI](#) to create a Transit Gateway Route Table for the Transit Gateway ID:

```
aws ec2 create-transit-gateway-route-table --region us-east-2 --transit-gateway-id <Your Transit gateway ID> --tag-specifications "ResourceType=transit-gateway-route-table,Tags=[{Key=Name,Value=VPN-RouteTable}]"
```

4. When the state of the newly created VPN Transit Gateway Route Table is **Available**, select the table and select the **Associations** tab.
5. Click **Create association**.
6. Select your VPN attachment from the drop-down menu under **Choose attachment to associate**.
7. Click **Create association**.



If you named the attachment earlier, look for that name.

Create association [Info](#)

Associating an attachment to a route table allows traffic to be sent from the attachment to the target route table. An attachment can only be associated to one route table.

Details

Transit gateway ID
 tgw-0ef6d0cc59759eee7

Transit gateway route table ID
 tgw-rtb-0786796cb3e790a80

Choose attachment to associate
 tgw-attach-01cdb24b94d2acecd (VPN-Attachment) ▼

Cancel **Create association**

Figure 42. AWS Create association page

8. Click the **Routes** tab.
9. Add a static route for the VPC subnet Classless Inter-Domain Routing (CIDR) range you want to send through the VPN tunnels to ZIA. Choose the appropriate Transit Gateway attachment for that VPC subnet from the **Choose attachment** drop-down menu.
10. Click **Create static route**.
11. Repeat these steps for any other VPCs that send their traffic through the VPN tunnels to ZIA. The association allows the traffic returning from the VPN tunnels to flow back to the subnet that initiated the traffic via the associated attachment.

As an example, if you are using the lab environment from [Appendix A: AWS Transit Gateway Lab Environment](#), add the following routes for the App1 and App2 VPCs.

Create static route [Info](#)

Add a static route to your transit gateway route table.

Details

Transit gateway ID
 tgw-0ef6d0cc59759eee7

Transit gateway route table ID
 tgw-rtb-0786796cb3e790a80

CIDR [Info](#)

Type [Info](#)
☒ Active
☐ Blackhole

Choose attachment
 tgw-attach-03b4601591fafc001 (App1-Attachment) ▼

Cancel **Create static route**

Figure 43. Create static route

The following image shows the static route.

Create static route Info

Add a static route to your transit gateway route table.

Details

Transit gateway ID

tgw-0ef6d0cc59759eee7

Transit gateway route table ID

tgw-rtb-0786796cb3e790a80

CIDR Info

10.1.0.0/16

X

Type Info

☒ Active

☐ Blackhole

Choose attachment

tgw-attach-082ccb79a697b6e23 (App2-Attachment)

Cancel

Create static route

Figure 44. AWS Create static route details

12. Next, change your App VPC route table's default route to point to the VPN Attachment instead of the Egress VPC. As an example, if you are using the lab environment from [Appendix A: AWS Transit Gateway Lab Environment](#), replace the following default route attachment to point to the VPN attachment.

Routes (1/3)

Filter routes

Actions

Create static route

Replace static route

Delete static route

	CIDR	Attachment ID	Resource ID	Resource type	Route type
<input checked="" type="checkbox"/>	0.0.0.0/0	tgw-attach-02f554deb7fe7a773	vpc-016c808f7890a73e2	VPC	Static
<input type="checkbox"/>	10.0.0.0/8	-	-	-	Static
<input type="checkbox"/>	172.16.0.0/12	-	-	-	Static

Figure 45. Static Routes

The following image shows the replaced static routes.

Replace static route [Info](#)

Replace a static route in your transit gateway route table.

Details

Transit gateway ID
tgw-Def6d0cc59759eee7

Transit gateway route table ID
tgw-rtb-02545278d6d03023d

CIDR [Info](#)
0.0.0.0/0

Type [Info](#)
☒ Active
☐ Blackhole

Choose attachment
 tgw-attach-01cdb24b94d2acecd (VPN-Attachment)

Cancel **Replace static route**

Figure 46. AWS Replace static route details

- Test the route, from an EC2 instance in the App VPC, through the Site-to-Site VPN Connection, to ZIA.



If you are using the lab environment from [Appendix A: AWS Transit Gateway Lab Environment](#), you must add a route for 192.168.0.0/16 in your App VPC route table pointing to the Egress-Attachment. This route allows traffic back to the Bastion host before you can connect to the App EC2 instances from the Bastion host (otherwise, the default route sends it through the VPN tunnels).

Example Testing

Use the following tests to determine if the source IP can be done using curl and the JSON output from ip.zscaler.com.

The following example shows testing from EC2 instance in App VPC with the default route pointing to the Egress attachment. The clientip shown is egress Elastic IP:

```
curl http://ip.zscaler.com?json  
  
{"srcip":"3.20.82.111","clientip":"3.20.82.111"}
```

The following example shows testing from the EC2 instance in App VPC with the default route pointing to the VPN attachment. The clientip is Tunnel 1 Outside IP address:

```
curl http://ip.zscaler.com?json  
  
{"srcip":"165.225.58.247","vip":"165.225.56.19","nodename":"zsn-chi1-4e1-sme","-  
cloud":"zscaler.net","datacenter":"Chicago","xff":"3.141.109.232","clien-  
tip":"3.141.109.232"}
```

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.

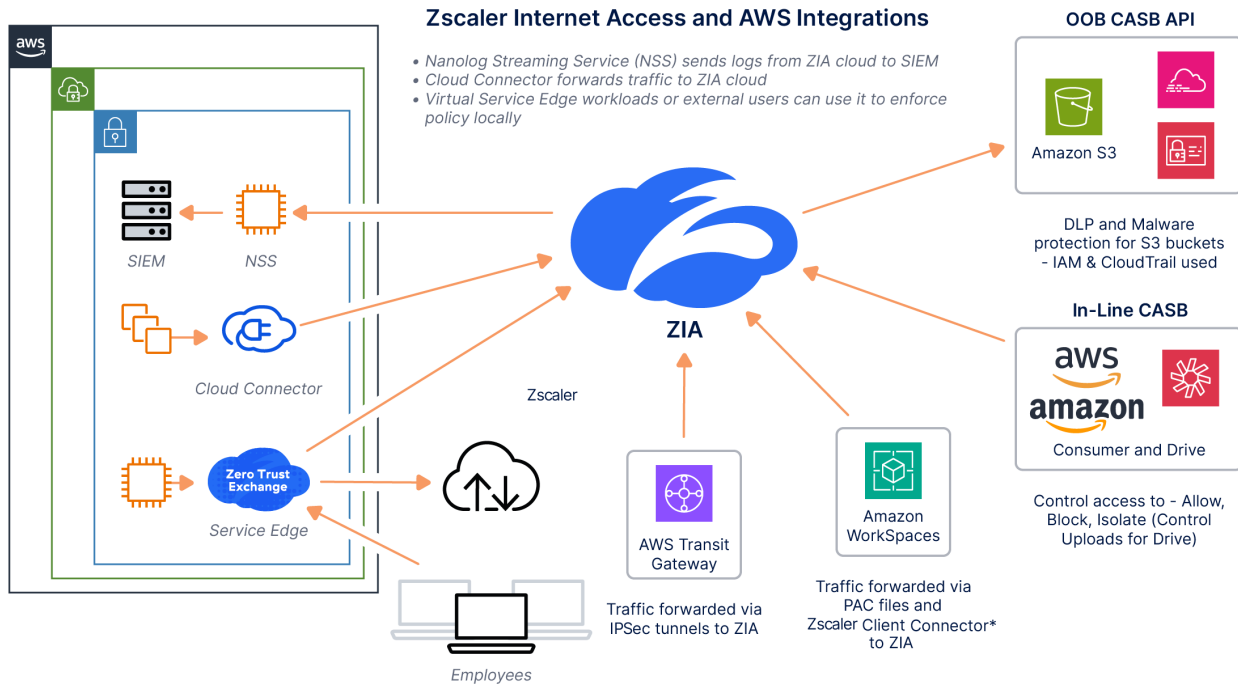


Figure 47. ZIA and AWS integration

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](#)).

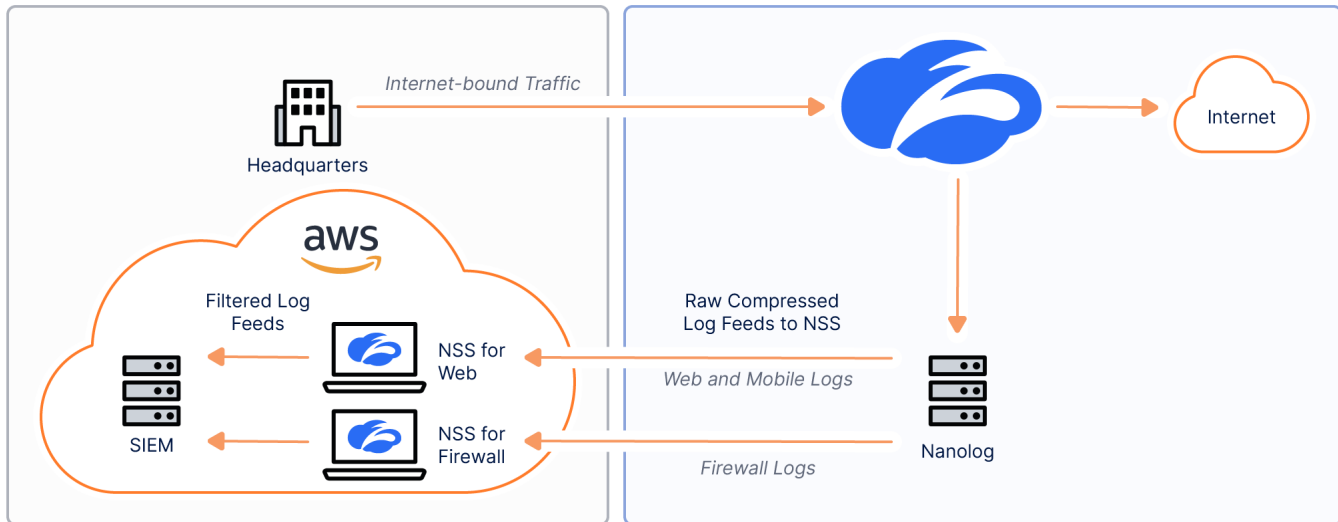


Figure 48. Zscaler NSS and AWS integration

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](#)).

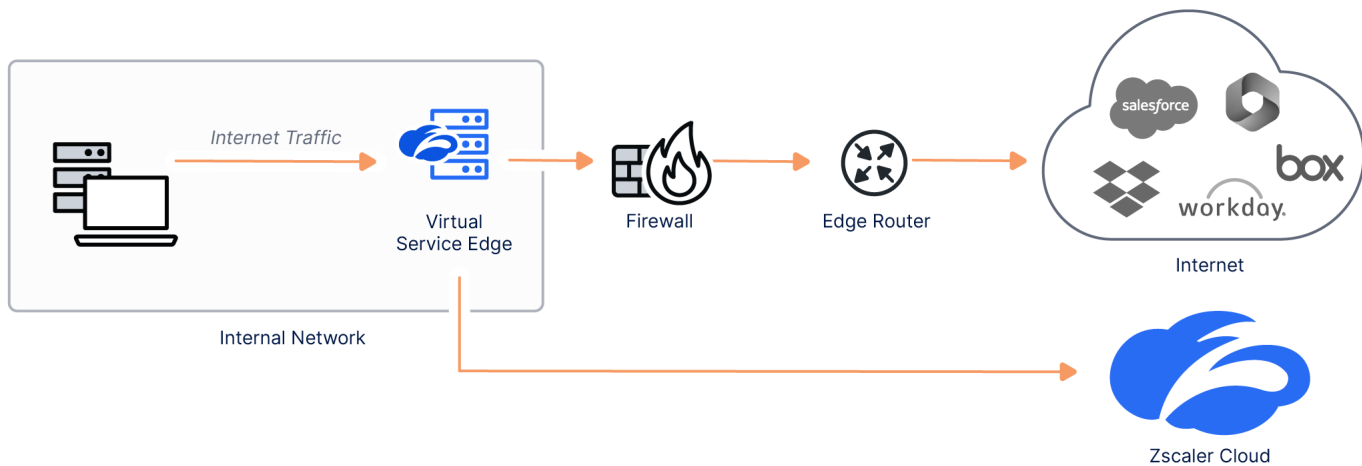


Figure 49. Virtual service edge

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](#) (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 Linux. 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 Linux for AWS.

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

ZPA Components that Work on AWS Infrastructure

The following sections describe ZPA components that work on AWS Infrastructure.

App Connector

ZPA App Connectors are lightweight virtual appliances that enable secure, zero trust access to private applications without exposing them to the internet. Deployed close to the applications—either in AWS, on-premises, or other cloud environments—they act as outbound-only proxies that establish secure TLS tunnels to the Zscaler cloud. App Connectors never accept inbound connections, which eliminates attack surface and supports the zero trust model. They broker user-to-app and app-to-app communication based on identity and policy defined in the ZPA Admin Portal. Traffic is routed to the appropriate App Connector dynamically, depending on application location, load, and availability. App Connectors are easy to deploy and scale, and they support high availability out-of-the-box. For more information, see [Zscaler Resources](#).

Private Service Edge

ZPA Private Service Edge (PSE) is a virtual appliance that brings ZPA's core functionality on-premises or into customer-controlled cloud environments. It enables organizations to enforce ZPA access policies locally while keeping traffic, authentication, and policy enforcement entirely within their environment. PSEs are ideal for scenarios with strict data residency requirements, low-latency needs, or limited internet connectivity. Like Zscaler's cloud-hosted Public Service Edges, they broker secure, outbound-only TLS tunnels from users and workloads to private applications via App Connectors. PSEs provide high availability, load balancing, and scale horizontally to support large deployments. They extend the zero trust model closer to the application, offering greater control without compromising security. For more information, see [Zscaler Resources](#).

Using ZIA to Enforce Security Policy in AWS

ZIA is a cloud-delivered security platform that protects outbound traffic from users and workloads by inspecting and enforcing security policies inline, without the need for traditional on-premises appliances. For AWS workloads, ZIA provides comprehensive protection by applying advanced controls like Cloud App Control to manage application usage, File Type Control to prevent unauthorized data transfers, and Firewall Control to enforce granular traffic filtering at the application and protocol level. It also includes DNS Control, which helps block malicious destinations and prevent command-and-control callbacks at the earliest stage. With integrated threat intelligence, SSL inspection, and Data Loss Prevention (DLP), ZIA ensures that cloud workloads maintain consistent security posture, whether accessing public internet services or third-party SaaS applications.

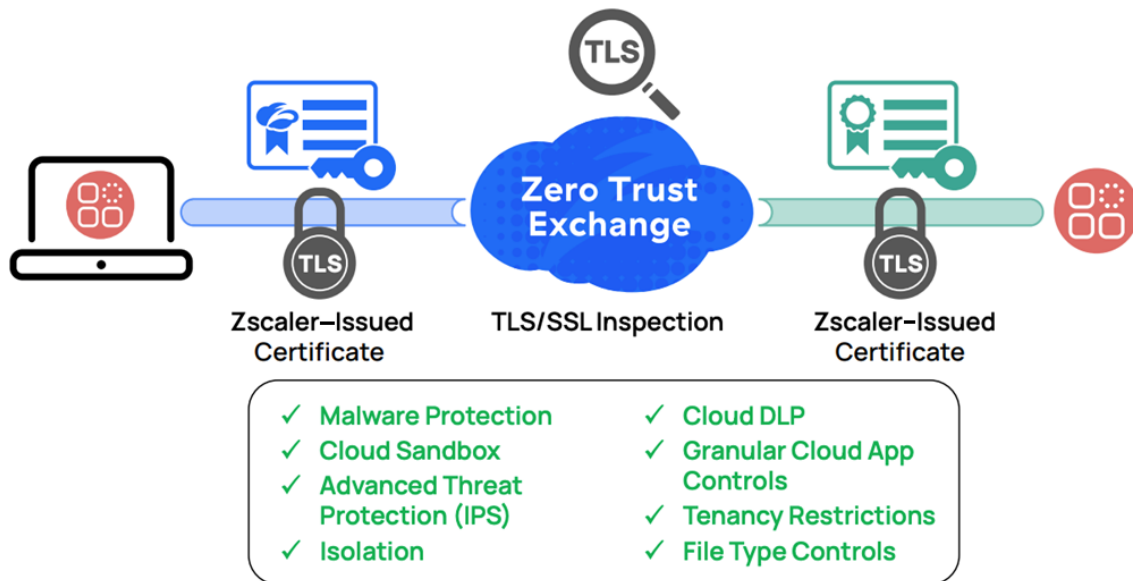


Figure 50. ZIA and AWS security policy

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](#) (government agencies, see [categories](#)) for easy reference and to define rules for similar apps.

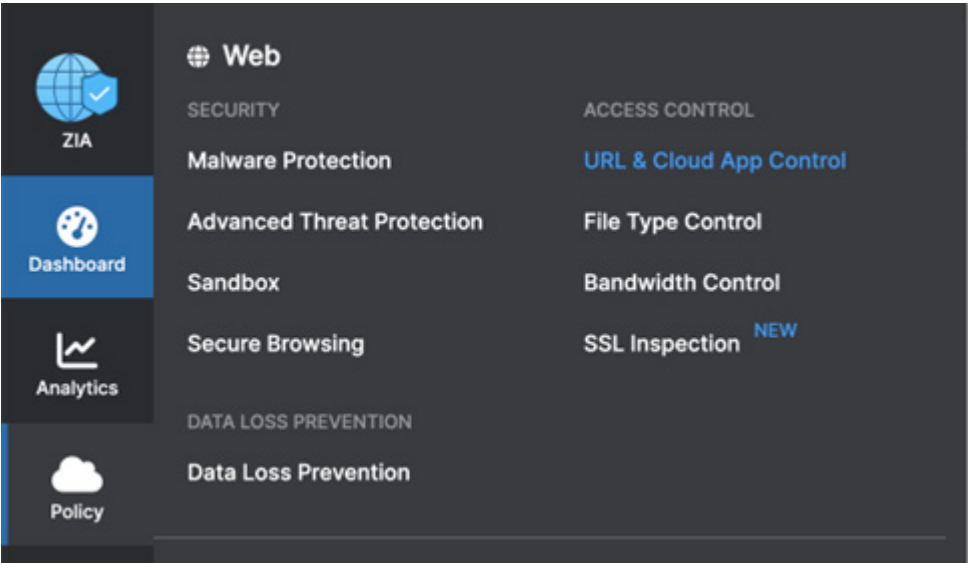


Figure 51. Cloud App Control

All policies 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).

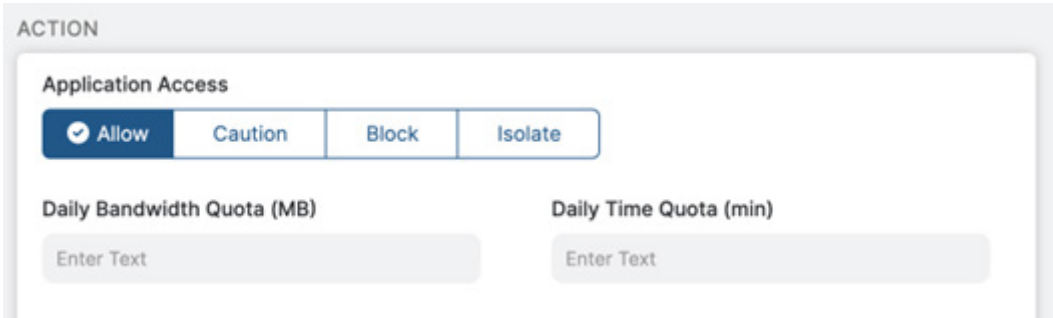


Figure 52. Policy Actions

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.

Add Collaboration and Online Meetings Rule

CLOUD APP CONTROL RULE

Rule Order: 1

Rule Name: Amazon Chime

Rule Status: Enabled

Rule Label: ---

CRITERIA

Cloud Applications: Amazon Chime

Cloud Application Risk Profile: None

Unselected Items	Selected Items (1)
Search... <input checked="" type="checkbox"/> Amazon Chime <input type="checkbox"/> WeChat Work	Amazon Chime

Figure 53. 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.

Sorry, you don't have permission to visit this site.

Probably shouldn't be going to this website.

Not allowed the use of this enterprise site

Amazon Chime

[See our internet use policy.](#)

Need help? Contact our support team at +91-9000000000, support@spaisley.com D20

Figure 54. Blocked access message

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

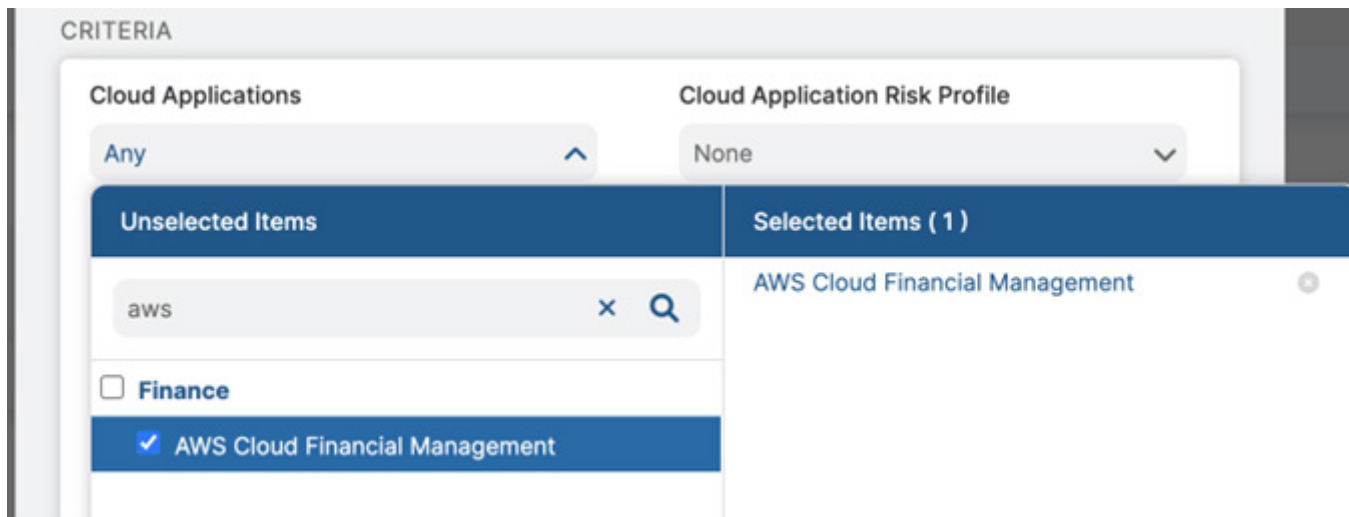


Figure 55. Policy criteria

This removes access from a user who 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 Elastic Container Service (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 Elastic Kubernetes Service (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 Amazon 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).
Amazon Trust Services	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 Resource Access Manager (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 Snowball	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 are 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. Note, however, that there is an additional restriction called Tenant Restriction, which you can enable for the Hosting Providers section that 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.

For more information, see [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 [Cloud App Control Policies Available via Individual Amazon Web Services](#).

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

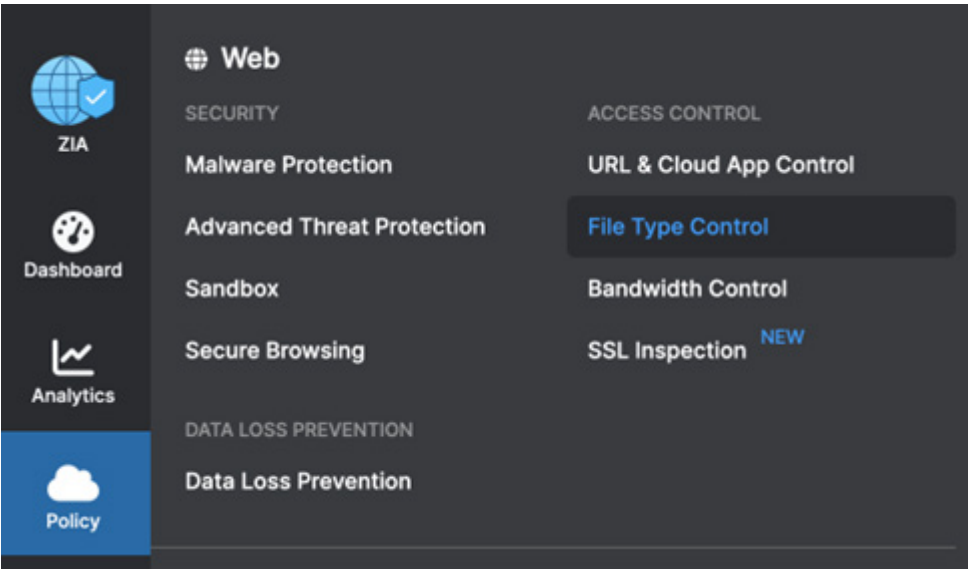


Figure 56. File Type Control

Here is an example of such a policy to prevent and block any ZIP files from being uploaded to any Amazon services:

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 57. File Type control policy

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

The screenshot shows the 'Add DLP Rule' dialog box with a close button (X) in the top right corner. The dialog is divided into two main sections: 'CRITERIA' and 'DLP INCIDENT RECEIVER'.

CRITERIA

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

DLP INCIDENT RECEIVER

Incident Receiver

ICAP ☒ Zscaler Incident Receiver

Figure 58. 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 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 specific AWS services. You can create a specific firewall rule that combines the Who, Where, When, Services, Applications, Source IP, and Destination IP.

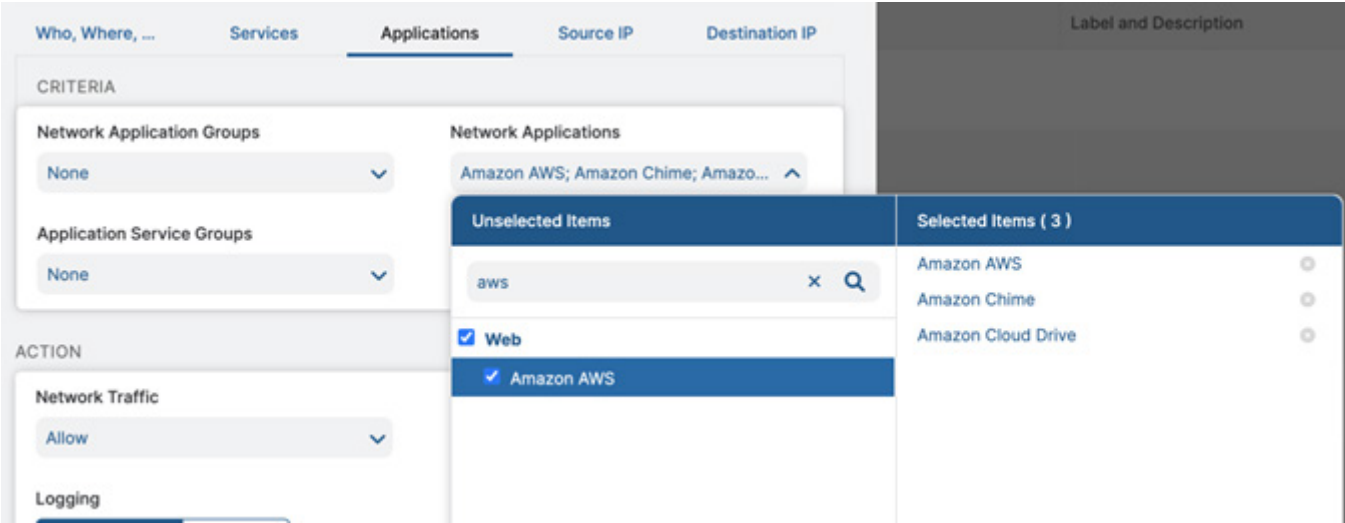


Figure 59. Firewall policy

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 a 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.

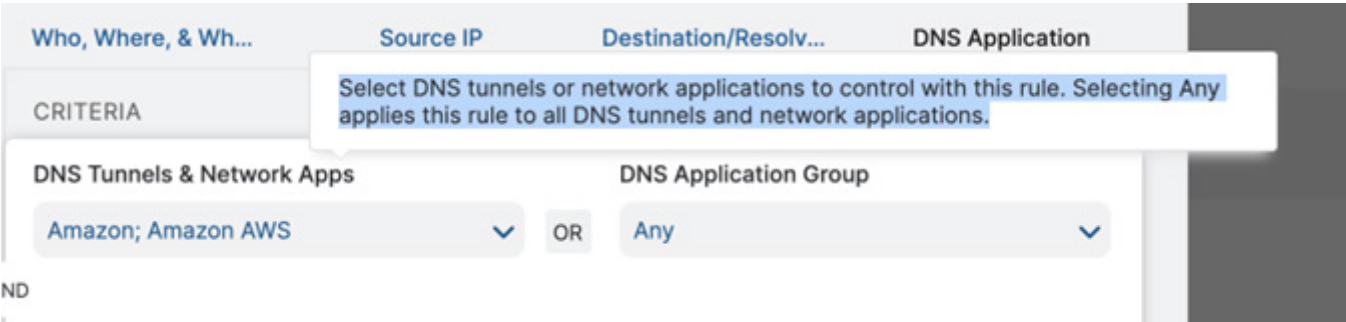


Figure 60. DNS policy

Using ZPA to Enforce Security Policy in AWS

Zscaler Private Access (ZPA) enables zero trust access to private applications hosted in AWS by eliminating traditional VPNs and network-level exposure. Instead of extending the network to users or workloads, ZPA connects users or workloads to specific applications based on identity, context, and granular policy — without ever placing them on the network. This ensures consistent and secure access whether the user is remote, in the office, or accessing applications hosted in AWS, other clouds, or on-premises environments.

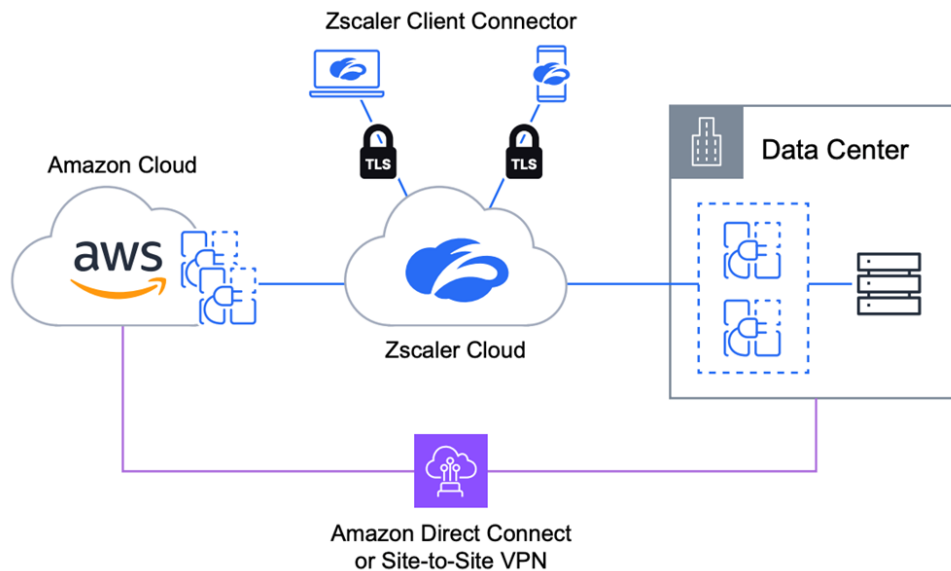


Figure 61. ZPA to enforce security policies in AWS

Deploying App Connectors

The first step to deploying ZPA within AWS is to deploy App Connectors. Zscaler App Connectors are virtual appliances that enable secure, outbound-only connections between private applications and the Zscaler cloud, and you can deploy them across multiple environments, including AWS. In AWS, deployment is streamlined through a prebuilt, regularly updated Amazon Machine Image (AMI) provided by Zscaler, allowing you to launch App Connectors much like any other EC2 instance. Successful deployment involves selecting the appropriate instance type, ensuring required network connectivity and security group configurations, and leveraging AWS tools like Auto Scaling for scalability and high availability.

Prerequisites and Planning

Make sure the following are met:

- **Instance Placement Recommendation:** Deploy App Connectors as close to the application or workload as possible. This facilitates a proper zero trust deployment as it prevents users from gaining access to applications or workloads outside of the proximity of the App Connector. App Connectors require outbound connectivity to register to the Zscaler cloud, so ensure a path to the internet exists.
- **Instance Type Recommendation:** For production environments, deploy App Connectors on AWS EC2 instances of type m5a.large or m5a.xlarge to ensure adequate performance. For non-production or low-traffic scenarios, t3.xlarge instances might suffice. For more information, see [App Connector Deployment Prerequisites](#) (government agencies, see [App Connector Deployment Prerequisites](#)).
- **Security Group Configuration:** Configure security groups to allow outbound traffic on TCP ports 80 and 443 to enable communication with Zscaler services. Ensure that inbound traffic is restricted according to your organization's security policies. Remember, App Connectors do not require inbound connectivity.

Configure App Connectors in ZPA Admin Portal

To configure an App Connector:

1. From the ZPA Admin Portal, go to **Configuration & Control > Private Infrastructure**.
2. Click **App Connectors**.
3. Click **Add**.
4. If you have an existing **Provisioning Key**, you can select it or choose the option to **Create** a new provisioning key.
5. Click **Next**.
6. Select the Connector certificate and click **Next**.
7. Select **Add App Connector Group**.
8. Provide the following details (leave items not listed at default or modified as per your preference):
 - a. **Name**: Enter a name for the App Connector Group.
 - b. **Status**: Select **Enabled**.
 - c. **App Connector Location**: Search for and select a location that aligns with where the App Connector is deployed (e.g., Northern California / US-West-1, Oregon / US-West-2, Virginia / US-East-1, Northern Ohio / US-East-2).

The screenshot shows the 'Add App Connector' form in the ZPA Admin Portal. The form is titled 'Add App Connector' and has a close button (X) in the top right corner. Below the title is a progress bar with six steps: 2. Signing Certificate, 3. App Connector Group (current step), 4. Create Provisioning Key, 5. Review, and 6. Review Documentation. The main form area is divided into two sections: 'Select App Connector Group' and 'Add App Connector Group'. The 'Add App Connector Group' section contains the following fields and options:

- Name**: A text input field with the value 'AWS West App Connectors'.
- Status**: A radio button group with 'Enabled' selected and 'Disabled' unselected.
- Description**: A large text area for additional information.
- DNS Resolution Option**: A radio button group with 'IPv4', 'IPv6', and 'IPv4 and IPv6' (selected).
- TCP Quick Acknowledgement**: A radio button group with 'Enabled' and 'Disabled' (selected).
- Disaster Recovery**: A radio button group with 'Enabled' and 'Disabled' (selected).
- Disable AppProtection**: A radio button group with 'Yes' and 'No' (selected).

At the bottom of the form, there is a 'VERSION PROFILE CONFIGURATION' section. Below the form are three buttons: 'Next', 'Previous', and 'Cancel'.

Figure 62. Add App Connector

- d. Click **Next**.
- e. Enter a name for the **Provisioning Key** as well as a reuse limit (how many App Connectors can use the Provisioning Key).
- f. Click **Next**.
- g. Click **Save**.

Add App Connector

2 Signing Certificate 3 App Connector Group 4 Create Provisioning Key **5 Review** 6 Review Documentation

Certificate Name
Connector
App Connector Group
AWS West App Connectors
Provisioning Key
AWS West Prov Key

Review all of the information before clicking Save

Save Previous Cancel

Figure 63. Review

- h. Copy the Provisioning Key presented on the screen into a notepad. This is needed when deploying the App Connector AMI in AWS.

Deploy App Connectors in AWS

You can deploy App Connectors programmatically (i.e., Terraform), or manually using the AWS EC2 wizard. The following steps depict the manual method. For information on programmatic deployment, see the [README.md](#) file within Zscaler Github repository for Terraform.

To deploy an App Connector:

1. (Optional) Before beginning, you might want to [create a new SSH Key pair](#) so that you can access the new App Connector appliance.
2. From the **AWS** dashboard, ensure you are in the appropriate region and go to **EC2** (using the search bar, or by clicking **EC2** from the home screen).
3. Click **Launch instance(s)**.
4. Enter a **Name** for the instance.
5. In the search field, search for **Zscaler App Connector**.
6. Click the **AWS Marketplace AMIs** tab.

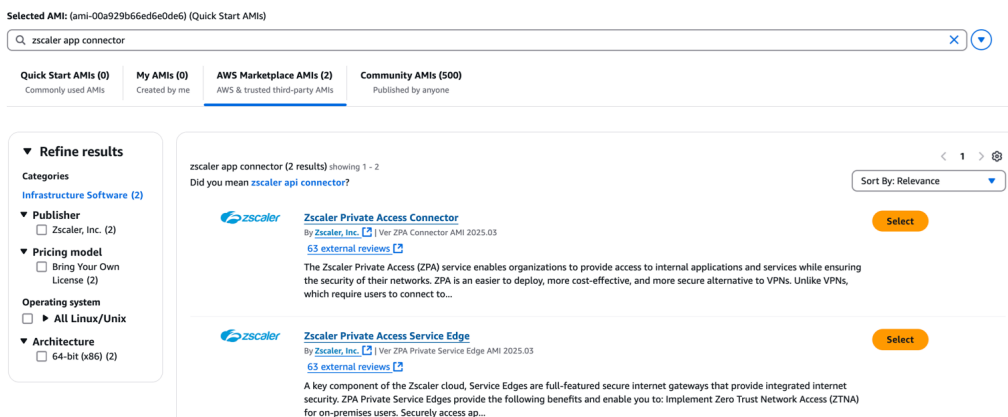


Figure 64. AWS Marketplace AMIs

7. Click **Select** next to Zscaler Private Access Connector.
8. (Optional) If asked to subscribe to the AMI, click **Subscribe Now**.
9. In the **Key pair (login)** section, choose the SSH Key pair you want to use for SSH access to the appliance.
10. In the **Network settings** section, click **Edit**.
 - a. **VPC:** Select the VPC to instantiate the App Connector within (as a best practice, instantiate the App Connector within the VPC closest to the workloads it services).
 - b. **Subnet:** Select the subnet to instantiate the App Connector within.
 - c. **Auto-assign public IP:** The App Connector is not designed to be publicly accessible. Furthermore, it is likely to be instantiated in a private subnet. Hence, a public IP is not required.

▼ Network settings [Info](#)

VPC - required [Info](#)

vpc-02a15af16aff9b40c (MyVPC)
10.0.0.0/16



Subnet [Info](#)

subnet-00b60ed8932038060

PrivateSubnet

VPC: vpc-02a15af16aff9b40c Owner: 084828604313 Availability Zone: us-east-1a
Zone type: Availability Zone IP addresses available: 249 CIDR: 10.0.2.0/24



[Create new subnet](#)

Auto-assign public IP [Info](#)

Disable



Figure 65. Networks settings



Ensure that the App Connector has proper outbound connectivity—such as routing through a NAT Gateway or Internet Gateway—to reach the Zscaler ZPA cloud. Additionally, verify that the necessary ports are open for outbound traffic, including TCP/443 for ZPA communication and UDP/53 for DNS resolution, along with any required ports to access the applications the App Connector supports. For detailed information on security and connectivity requirements, see [Zscaler Private Access Firewall Allowlist](#).

11. In the **Firewall (security groups)** section, select **Create security group**.
12. Select **Allow SSH traffic** and a valid source is selected in the drop-down menu.

Firewall (security groups) [Info](#)

A security group is a set of firewall rules that control the traffic for your instance. Add rules to allow specific traffic to reach your instance.

☒ Create security group

☐ Select existing security group

We'll create a new security group called 'Zscaler Private Access Connector-ZPA Connector AMI 2025.03-AutogenByAWSMP--1' with the following rules:

☒ Allow SSH traffic from
Helps you connect to your instance

Custom


Add CIDR, prefix list or security group

10.0.0.0/8



Figure 66. Allow SSH traffic

13. Click **Launch instance**.
14. After the instance has booted, log in to the appliance with the username `admin` using the SSH Key pair assigned to the instance during the creation process.
15. Stop the ZPA service by executing the following command: `sudo systemctl stop zpa-connector`
16. Create a provisioning key file by executing the command: `sudo touch /opt/zscaler/var/provision_key`
17. Update the permissions of the file with the command: `sudo chmod 644 /opt/zscaler/var/provision_key`
18. Open the `provision_key` file in the vi editor: `sudo vi /opt/zscaler/var/provision_key`
19. Paste in the provisioning key copied from the previous section.
20. Verify the provisioning key by executing the command: `sudo cat /opt/zscaler/var/provision_key`

 After the ZPA service is started, this file is encrypted and no longer readable in cleartext.

21. Start the ZPA service via: `sudo systemctl start zpa-connector`.

Creating Workload Application Segments and Adjusting Policy

The following sections describe creating workload application segments and adjusting policy.

Discover and Create Applications Segments

After App Connectors are deployed, you can create a discovery Application Segment passively observes which applications are accessed in AWS, helping you map user-to-app relationships without enforcing restrictions. This allows you to identify the full scope of accessed apps and begin organizing them into Application Segments based on function, department, or access needs. A discovery Application Segment simply uses a wildcard in the application definition, such as `*.safemarch.local`. For more information, see [About Application Discovery](#) (government agencies, see [About Application Discovery](#)).

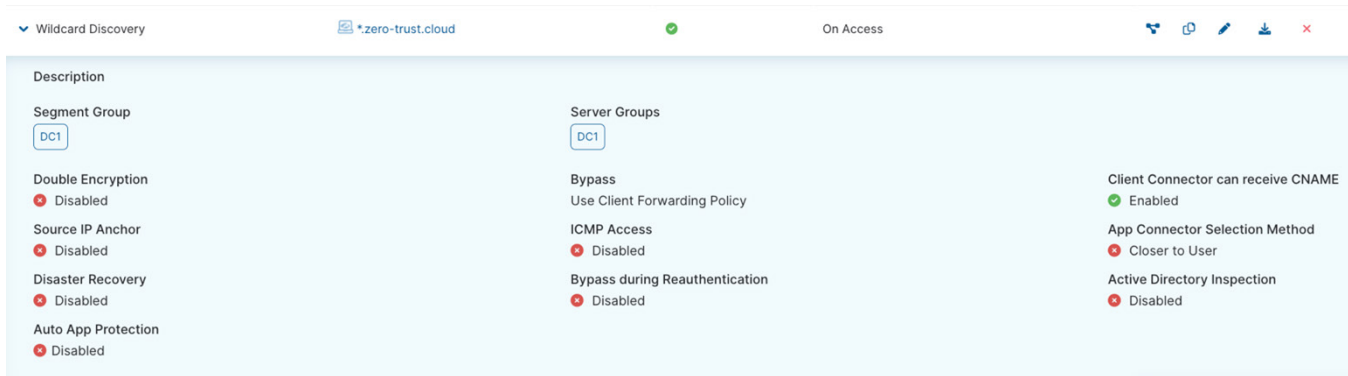


Figure 67. Wildcard discovery

If licensed, ZPA also provides AI-Powered Recommendations to help you organize Application Segments and drive towards a zero trust. The AI-Powered Recommendations tab is located at the top of the Application Segment screen. For more information, see [About AI-Powered Recommendations for Application Segments](#).

Application Segments Browser Access AppProtection Privileged Remote Access Segment Groups

Defined Application Segments **AI-Powered Recommendations** Application Segment Import

Next update in 8 days Download as CSV


No filters have been applied

Name	Applications	Grouping Reasons	Confidence	Attack Surface Reduction	Actions
20231201-5	webserver.zero-trust.cloud	USER ACCESS SIMILARITY Applications grouped by access pattern for similar user types.	70%	100.0%	+ > ⋮
20231201-1	musfrgscwks1pr.test5.zero-trust.cloud mushqexst1dv.test5.zero-trust.cloud +1 Applications	PORTS AND PROTOCOLS SIMILARITY Consists of RDP applications.	86%	99.98%	+ > ⋮
20231201-0	1jhvmv2.test5.zero-trust.cloud 2hqxhk2.test5.zero-trust.cloud +2 Applications	PORTS AND PROTOCOLS SIMILARITY Consists of Microsoft Endpoint Configuration Manager applications.	99%	99.96%	+ > ⋮
20231201-9	cgvgmr03.consumer.zero-trust.cloud cgvgmr04.consumer.zero-trust.cloud	USER ACCESS SIMILARITY Applications grouped by access pattern for similar user types. DOMAIN NAME SIMILARITY Applications grouped based on domain name	41%	99.84%	+ > ⋮

Figure 68. Application Segments

Refine Policy

After the Application Segment discovery and creation phase is complete, gradually lock down access to known applications, starting with critical or sensitive systems. Define and enforce Access Policies based on workload, tag, user role, group, and business context—restricting unnecessary access (e.g., developers accessing finance tools). As policies mature, you can continue to refine segmentation and phase out broad discovery policies, keeping them only where continuous visibility is needed (e.g., for employee access to newly launched services).

 ZPA relies heavily on DNS. As such there are a few DNS and routing considerations to be aware of:

- **Internal DNS Resolution:** ZPA relies on the Zscaler Client Connector to intercept DNS queries for internal apps. These queries are redirected through ZPA, where the application resolves to a synthetic IP.
- **Avoid DNS Conflicts:** ZPA uses UDP 53 for DNS functionality—avoid using it for app traffic. Use split DNS or configure DNS search domains in the Zscaler Client Connector settings to prioritize internal resolution.
- **Domain Bypass:** Place public-facing apps using the same domain as internal apps in *Always Bypass* segments to ensure they're resolved externally and not intercepted by ZPA.
- **Support for FQDNs and IPs:** ZPA works best with fully qualified domain names (FQDNs). For legacy apps using IPs, you can define IP-based segments temporarily, but FQDN usage is strongly recommended for long-term policy management.

Configuring Microsegmentation Between Workloads

Zscaler Microsegmentation is an agent-based solution within the Zscaler Private Access (ZPA) platform that helps protect workload environments by monitoring and controlling server-to-server communications across public clouds and data centers. It supports common Windows and Linux operating systems, collects network flow data, and visualizes it in the ZPA Admin Portal. Security teams can define and enforce granular, port-level inbound and outbound rules using the native security frameworks of the OS. Typical use cases include preventing lateral movement of threats, monitoring app-to-app traffic, and meeting compliance requirements like HIPAA, PCI, and cyber insurance standards.

Initial Setup

To set up:

1. From the ZPA Admin Portal, go to **Configuration & Control > Certificate Management > Enrollment Certificates**.
2. Create a new ZPA Enrollment Certificate to be used for Microsegmentation. Click **Generate Certificate**:
 - a. **Name**: Enter **Microsegmentation**.
 - b. **Type**: Select **Intermediate**.
 - c. **Parent**: Select **Root**.
 - d. **Client Type**: Select **None**.
3. Ensure the Agent can communicate outbound with the Zscaler cloud (*.prod.zpath.net and *.private.zscaler.com) on TCP/UDP port 443 by adjusting Security Groups or upstream firewalls as necessary.



If the AWS environment includes Zscaler Cloud Connector, adjust the [Cloud Connector Forwarding Policy](#) (government agencies, see [Cloud Connector Forwarding Policy](#)) to send Zscaler Microsegmentation traffic Direct:

- Direction: Select **Outbound**
- Protocol: Select **TCP/UDP**.
- Port: Enter 443.
- Source: Enter the workload networks where Agents are deployed.
- Destination: Enter *.prod.zpath.net, *.private.zscaler.com.

Deploying the Agents

To deploy the agents:

1. From the ZPA Admin Portal, go to **Microsegmentation > Agent Management > Agent Groups**.
2. Click the **Add Agent Group**.
3. Complete the wizard, ensuring to set the environment type details and version profile configuration:
 - a. **Name:** Enter a name.
 - b. **Cloud:** Select **AWS**.
 - c. **Auto Update:** Select whether to auto-update the agent software and confirm the frequency.

The screenshot shows the 'Agent Group' configuration page. At the top, there's a header 'Agent Group'. Below it, the 'Name' field is filled with 'agent-aws-vpc123' and has an 'Admin Status' toggle set to 'Enabled'. The 'Description' field is optional and empty. The 'Cloud' section has three radio buttons: 'AWS' (selected), 'Azure', and 'On Premises'. Below this, a note says 'Region, VPC ID and Subnet ID will be discovered by IMDS.' The 'Version Profile & Configurations' section is expanded, showing 'Auto Update' as 'Enabled' and 'Version Profile' as 'Latest'. The 'Agent Version' is '1.1.1'. The 'Schedule Agent Upgrade On' is set to 'Sunday', 'At' is '00:00', and 'Time zone' is 'America/Chicago'. The 'Update Sequence' is 'Serial' and 'In case of upgrade failure' is 'Halt next agent upgrade'.

Figure 69. Agent Group

- d. Click **Next**.
- e. Create a new **Provisioning Key** by providing a **Name** and **Maximum Reuse of Key** value.
- f. Select the signing certificate that was created during initial tenant configuration:

The screenshot shows the 'Provisioning Key' configuration page. The 'Name' field is filled with 'prov-key'. The 'Maximum Reuse of Key' field is filled with '100' and has an information icon. The 'Signing Certificate' dropdown menu is set to 'Microsegmentation'.

Figure 70. Provisioning Key

4. Click **Next**.

- Click **Save**.
- Download the installation media and save the provisioning key by clicking the **Download** icon to the right of the **Agent Group**.
- Deploy the Agent. There are many ways to do this. To learn more, see [Configuring Agent Groups](#). Alternatively, Zscaler recommends using a software distribution framework like AWS SSM. See [Appendix E: Metrics Config Options](#).
- Verify the Agent has registered with ZPA by going to **Microsegmentation > Agent Management > Agents**.

Name	Current Software Version	Connection Status	Admin Status	Policy Status	Actions
admineast1-4ad11085c0ef170001e5	1.1.1	Connected	Enabled	Enabled	Edit Delete
admineast1-f56ce01219b24e04e57f	1.1.1	Connected	Enabled	Enabled	Edit Delete
adminwest1-5cef285d8c2a87e562e	1.1.1	Connected	Enabled	Enabled	Edit Delete
adminwest1-72b38b187333e8dd3d	1.1.1	Connected	Enabled	Enabled	Edit Delete
app1.sg8jyuga.uat.cxlab.local-6aaa4	1.1.1	Connected	Enabled	Enabled	Edit Delete
db1.sg8jyuga.uat.cxlab.local-d1b69f	1.1.1	Connected	Enabled	Enabled	Edit Delete
dceast1-579583063134486e2a6b9	1.1.1	Connected	Enabled	Enabled	Edit Delete
dceast2-d128d513c2dc1963793412	1.1.1	Connected	Enabled	Enabled	Edit Delete
dcwest1-36306c03ef2a686717bcef	1.1.1	Connected	Enabled	Enabled	Edit Delete
dcwest2-53ada4387e724b9b95e4	1.1.1	Connected	Enabled	Enabled	Edit Delete

Figure 71. Agent management

Configure AppZones

AppZones in Microsegmentation are applications grouped together into zones based on the applications' topology and their underlying network connectivity with each other. Unless otherwise required, create AppZones that mirror the customer's environments. For example, AWS us-east-1 and AWS us-west-2 could be two separate AppZones.

Guidelines and Considerations

The following are guidelines and considerations for configuring AppZones:

- Map an AppZone to a network routing boundary.
- Do not create multiple AppZones for a single routable network (except in advanced cases).
- Log flow matching is performed within an AppZone. If a connection is logged between a host in an AppZone and a host outside of that AppZone, even if both hosts are running the Microsegmentation agent, the logs aren't flow matched and it show as two events in the ZPA Admin Portal.
- Policy is applied within an AppZone.
- A *managed-to-managed* connection can only occur within an AppZone. Inter-AppZone communications are always unmanaged.
- Flow logs are not used for Dashboard until a host is a member of an AppZone.

Configuration

To configure an AppZone:

- 1. From the ZPA Admin Portal, go to **Microsegmentation > Resource Management > AppZones**.
- 2. Click **Add AppZone**.
- 3. Specify the **Name** and match **Criteria**, or optionally configure the VPCs/VNETs to be included (or use the Include All option).

GENERAL

Name

appzone1

Description

CRITERIA

Region

us-east-1

☒ Include all VPCs/VNETs for selected regions

Figure 72. AppZones

Configure Resource Groups

Resource Groups are collections of Resources, where each Resource refers to a server running the Microsegmentation Agent.

Defining Resource Group Membership

To define a Resource Group membership:

- Static Membership: Manually select one or more Resources from the available list.
- Dynamic Membership: Define group membership based on metadata attributes, such as cloud tags.

App1 Backend

General Information

Member Resources

General Information

Name: App1 Backend

Type: MANAGED

Description:

Static Membership:

Dynamic Membership: [(ApplicationTier EQ Back-End) AND (Environment EQ production)]

Figure 73. General Information

The following image shows the Member Resources.

App1 Backend

General Information		Member Resources	
Resource Name		AppZone	
back-end-2.us-west-2.compute.internal		Production	
back-end-1.us-west-2.compute.internal		Production	

Rows per page: 20 ▾ 1-2 of 2 < 1 / 1 >

Figure 74. Member Resources

Guidelines and Considerations

Make sure the following guidelines and considerations are followed:

- As a reminder, policies are applied between Resource Groups
- Consider what level of security is appropriate for a given application stack. Also, take into account the number of Resources in the application stack:
 - If the application stack must be very secure, it might require a Resource Group per tier or collection of similar Resources. This allows for more policies to be applied between related Resource Groups.
 - If the application stack has a lot of Resources, it might require a Resource Group per tier or collection of similar Resources in order to avoid having a Resource Group with a lot of Resources.
 - If an application does not require very strict enforcement (Dev/Test, for example), it might be acceptable to combine Resources into a single Resource Group in order to minimize the amount of policy to create and manage.
- Do not put unrelated Resources in a Resource Group.
- Do not put Resources in separate physical locations/networks in a Resource Group.
- More Resource Groups provide more granular control after policies are applied inbound, outbound, and internal to these groups.
- Fewer Resource Groups allow for quicker setup and less ongoing management overhead.

Configuring a Resource Group

To configure a Resource Group:

1. From the ZPA Admin Portal, go to **Microsegmentation > Resource Management > Resource Groups**.
2. Click **Add Resource Group**.
3. Enter a **Name**.
4. Specify the match criteria and membership.

The screenshot shows a configuration form titled 'Criteria'. It includes a 'Resource Group Type' section with 'Managed' and 'Unmanaged' buttons, where 'Managed' is selected. Below is a 'Membership' section with 'Static Membership' and 'Dynamic Membership' options. Under 'Static Membership', a dropdown menu shows 'admineast1'. The 'Dynamic Membership' section has a '+ Add Row' button and a small icon.

Figure 75. Criteria

5. Click **Next**.
6. Click **Save**. Depending on the membership option chosen, resources are automatically assigned to the Resource Group created.

Configure Resource Policy

Resource Policies are policy rules that are applied between two Resource Groups or between a Resource Group and a special object (**ANY**, for example). Policies are programmed to the OS security frameworks (WFP, nftables) via the Microsegmentation agent.



The **ANY** object is equivalent to 0.0.0.0/0.

Policies are evaluated from the most significant priority (lowest priority number) to the least (highest priority number). You can create Resource Policies to enforce connectivity inbound, outbound, and internal, relative to a Resource Group.

Configuration

From the ZPA Admin Portal:

1. Go to **Microsegmentation > Policy > Resource Policy**.
2. Click **Add Rule**.
3. Enter a **Name**.
4. Click **Next**.
5. Identify allowed connections. Be sure to create **Allow** rules relative to the target Resource Group:
 - a. Inbound
 - b. Outbound
 - c. RG to RG (host-to-host in an RG as well as app-to-app on a host in an RG)

6. If installing a block policy, Zscaler recommends creating a Sim Block with a less significant priority (higher priority number) than the **Allow** rules.

Rule Configuration

Actions

Priority: 101

Action: ☐ Allow ☒ Sim Block ☐ Block

Source (maximum of 10): admin-aws-east

Destination (maximum of 10): ANY

Destination Ports

Default Port Ranges: Select

TCP Port Ranges

22 To 22

+ Add TCP Port Range

UDP Port Ranges

From To

+ Add UDP Port Range

Scope

Figure 76. Rule Configuration

This is one or more policies, depending on if you want to have a default block inbound, outbound, and/or internal to the Resource Group.

- a. Run with this configuration for a period of time (weeks to months) to build confidence in the policy set.
 - b. Review connections that triggered the **Simulated Block** rule. Add these items to existing or new **Allow** rules.
7. Change the **Action** of the **Simulated Block rule(s)** to **Block** to enable full enforcement.

Extending Zscaler with AWS Service Integrations

The following sections describe extending Zscaler with AWS Service Integrations.

Bringing Zero Trust Security to AWS Snowball Deployments

This integration allows you to establish a zero trust environment by deploying ZPA directly on the AWS Snowball device. By running an App Connector on Snowball, authorized users—defined by policy—can securely access private applications without any exposure to the internet.

You can think of this user-to-application connection like a virtual dark fiber: a dedicated, isolated path between the user and the app, accessible only to authorized users. No ports are open to the public internet, ensuring that the environment remains completely private.

Applications configured by policy are accessible to external users from anywhere. This works by having the internal App Connector on Snowball initiate an outbound connection to the Zscaler cloud, which then brokers a secure, inbound connection for the user—eliminating the need for any inbound ports. Additionally, if the Snowball App Connector can reach networks beyond the Snowball appliance, users can also securely access those connected applications and workloads.

AWS Components

- Snowball Edge
- AWS AMIs compatible with Zscaler App Connector

Zscaler Components

- Zscaler Client Connector
- ZPA
 - App Connector Deployment
 - Application Segment Configuration
 - Access Policy

The following is a diagram outlining the ZPA and AWS Snowball integration:

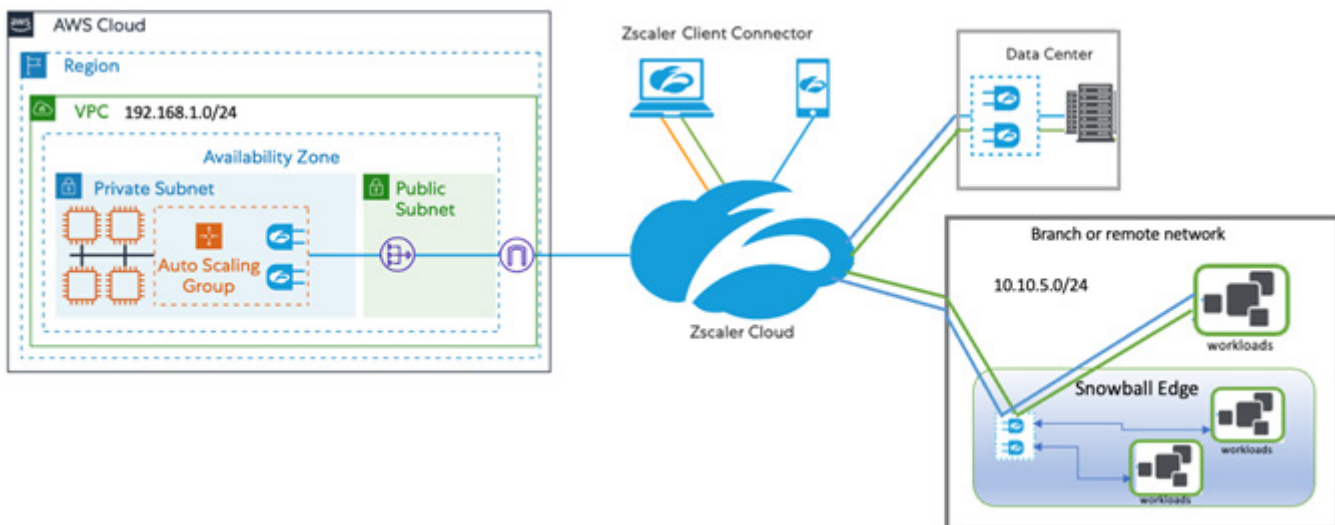


Figure 77. Zscaler and AWS Snowball integration

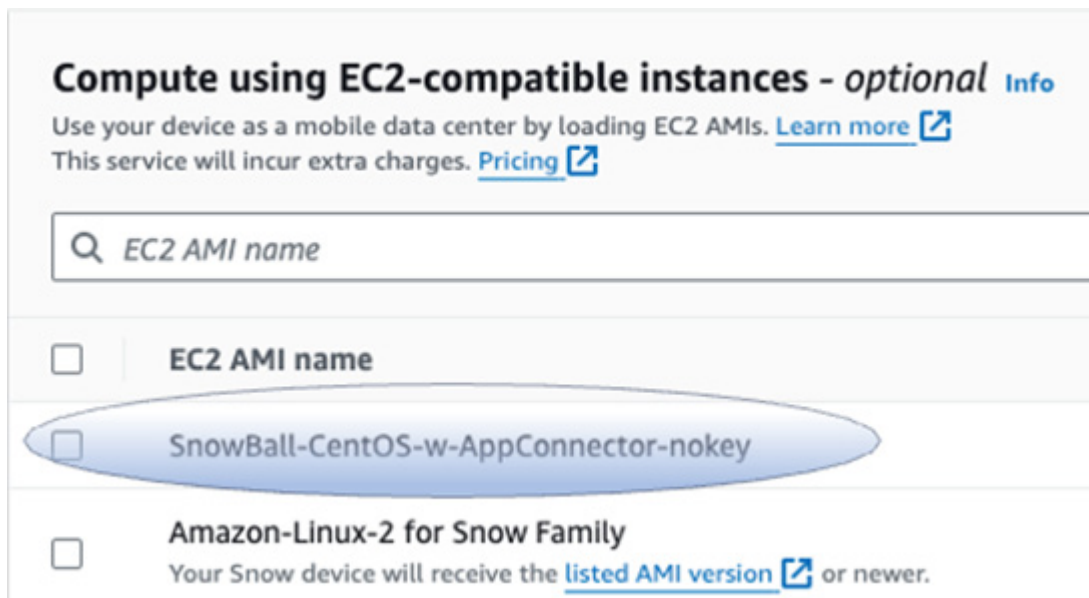
A fundamental principle of zero trust is ensuring that private applications are completely invisible to the outside world, and accessible only by authorized users. In the diagram, the two small Zscaler icons on the left represent App Connectors, which establish secure connections between users and internal applications. No matter where the user is located, as long as they are authorized, they can securely access resources in the private subnet. The only requirement is that the App Connector must be able to reach the target application, typically hosted within the VPC.

On the right side of the diagram, you see a remote branch network using AWS Snowball Edge. If the App Connectors deployed on the Snowball Edge have network access to applications within the branch network, authorized users can also securely connect to those resources. Any application that is reachable by the App Connector—whether in the cloud, on Snowball, or on-premises—can be made available to users through Zscaler’s secure access model.

Provisioning the AWS Snowball Appliance

To run App Connectors on the AWS Snowball appliance, ensure that you have provisioned the Snowball with AMIs that are compatible with the ZPA App Connector. While App Connector AMIs do not natively run on the specific Snowball hardware, the Zscaler App Connector can run on OSs supported by the Snowball appliance. To learn more about App Connector compatibility, see [App Connector Deployment Guide for CentOS, Oracle, and Red Hat](#) (government agencies, see [App Connector Deployment Guide for CentOS, Oracle, and Red Hat](#)). To learn more about AMI compatibility, refer to the [AWS Snowball documentation](#).

When ordering your Snowball device, Zscaler recommends creating an EC2 instance and creating an AMI that you can have pre-installed on your Snowball appliance. If the AMI you created in your AWS service account is not listed when you order a Snowball, submit a support ticket to have it added to the list.



Compute using EC2-compatible instances - optional [Info](#)

Use your device as a mobile data center by loading EC2 AMIs. [Learn more](#)

This service will incur extra charges. [Pricing](#)

Q *EC2 AMI name*

- ☐ **EC2 AMI name**
- ☒ **SnowBall-CentOS-w-AppConnector-nokey**
- ☐ **Amazon-Linux-2 for Snow Family**
Your Snow device will receive the [listed AMI version](#) or newer.

Figure 78. Compatible AMI list

This guide uses an AMI called *SnowBall-CentOS-w-AppConnector-nokey*, which was added to the list by requesting AWS support add the AMI to the Snowball request form. When selected, it is preloaded on the Snowball appliance.

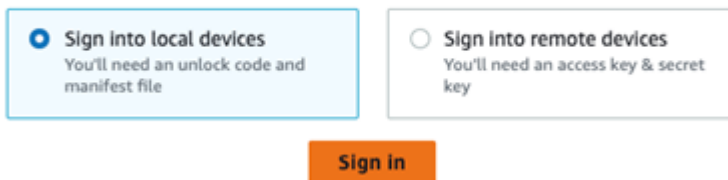
Deploy the AWS Snowball Appliance

After you have received the appliance, you can deploy it in your local network and deploy the App Connector. Follow the Snowball deployment guide by going to the Snow Family > Jobs section of your AWS Management Console and locating the AWS OpsHub link to download. The AWS Management Console has the unlock code and the manifest file that enables you to connect to the Snowball appliance on your network.

The steps are:

1. Connect the Snowball appliance to the network.
2. Use the AWS OpsHub to unlock and connect to the appliance.
3. Launch an EC2 instance for the ZPA App Connector.

Get started with OpsHub



☒ **Sign into local devices**
You'll need an unlock code and manifest file

☐ **Sign into remote devices**
You'll need an access key & secret key

Sign in

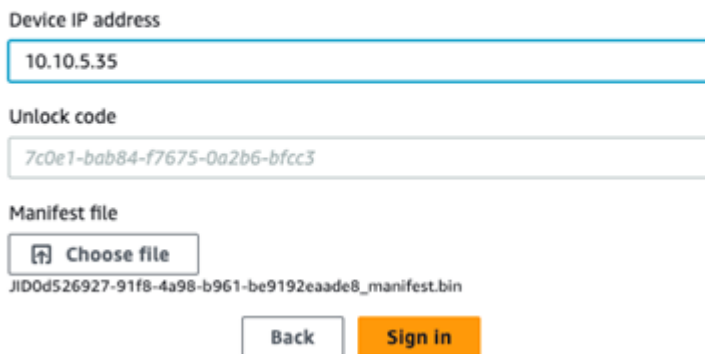
Figure 79. OpsHub

You can sign in remotely as well. After you have the App Connector running and ZPA configured, however, you won't need to sign in using the remote option. You can access any local device (with an enabled policy) securely from anywhere.

Sign in with your credentials.

Sign into your Snowball Edge

Sign in with an unlock code and manifest file



Device IP address
10.10.5.35

Unlock code
7c0e1-bab84-f7675-0a2b6-bfcc3

Manifest file
JID0d526927-91f8-4a98-b961-be9192eaade8_manifest.bin

Choose file

Back **Sign in**

Figure 80. Sign in to Snowball Edge

After logging in, launch an EC2 instance that becomes the App Connector.

Launch instance X

Device
10.10.5.35 ▼

Image (AMI)
SnowBall-CentOS-w-AppConnector-nokey ▼

Instance type
sbe-c.large ▼

☒ Create public IP address (VNI) ☐ Use existing IP address (VNI) ☐ Do not attach IP address

Physical network interface
RJ45: s.ni-8b0dd4979be942965 ▼

IP Address assignment
DHCP ▼

Key pair
☒ Create key pair ☐ Use existing key pair ☐ Do not attach key pair

Name
SSHkey
The name can include up to 255 ASCII characters. It can't include leading or trailing spaces.

Create key pair

Cancel Launch

Figure 81. Launch instance

Notice Create public IP address (VNI) is selected. While there is a routable IP in the AWS IP space, that IP is not public. You must use the local network address (which, in this case, is 10.10.5.35).

The Snowball appliance uses a public address space of 34.223.x.x on the Snowball appliance, and the 10.10.5.x/24 network is NAT'd.

The following shows the ZPA instance. It demonstrates an SSH connection to the 10.10.5.19 address. When you look at the interface, it displays 34.223.14.194. This is important as the interface in the ZPA Admin Portal displays 34.223.14.194, not 10.10.5.19 as expected. For more information, refer to [Understanding Virtual Network Interfaces on AWS Snowball Edge](#).

```
ssh: connect to host 10.10.5.34 port 22: Connection refused
scott@Scotts-MBP AWS % ssh -i snowballkey.pem centos@10.10.5.19
Last login: Mon Mar 25 21:51:35 2024 from scotts-mbp.local.tld
[centos@ip-34-223-14-194 ~]$ ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
      inet 34.223.14.194 netmask 255.255.255.128 broadcast 34.223.14.255
```

Figure 82. Example connecting to the ec2 snowball instance

Next, deploy the ZPA software on the EC2 instance on the Snowball appliance. To learn more, see [App Connector Deployment Guide for CentOS, Oracle, and RedHat](#) (government agencies, see [App Connector Deployment Guide for CentOS, Oracle, and RedHat](#)). Make sure you create a provisioning key in ZPA as you need it to start the ZPA App Connector.

Configure Policy in ZPA Admin Portal to Allow Connections

After the App Connector is running on the Snowball EC2 instance, verify if it's working by logging in to the ZPA Admin Portal.

From the ZPA Admin Portal, navigate to Configuration & Control > Private Infrastructure > App Connectors. A screen similar to the following is displayed:

Name	Manager Version	Current Software Version	Connection Status	Upgrade Status	Status	Actions
> AWS Knotted...	23.120.1	24.63.1	Connected	Success	✓	
✓ Snowball Provisioning Key-1709161911475	23.374.1	24.63.1	Connected	Success	✓	

Description:

App Connector Group: Snowball App Connector Group	App Connector Host Platform: AWS	App Connector Host OS: CentOS Linux 7
App Connector Package OS: Enterprise Linux 7	Last Software Update: Mar 25th, 01:01 AM (EDT)	Public Service Edge: US-VA-9418
Last Connection to Zscaler: Mar 27th, 03:48 (EDT)	Last Disconnect from Zscaler: Mar 27th, 03:29 (EDT)	Location: Ashburn, VA, USA
Public IP: 70.106.210.13	Private IP: 34.223.14.194	Uptime: 2day(s) 12hrs 1mins
Enrollment Certificate: Connector		

Figure 83. App Connectors

The Private IP is the NAT'd IP address from the Snowball appliance and not the network with which you configured it.

Now that you have the App Connector connected to the Zscaler cloud, you can configure policies and begin to move private traffic.

When defining a new application within an Application Segment, you can:

- [Define applications individually](#) or [enable Application Discovery](#) (government agencies, see [Define applications individually](#) or [enable Application Discovery](#)). This includes defining applications for [Browser Access](#), and for use with ZIA as part of [Source IP Anchoring](#) (government agencies, see [Browser Access](#) and [Source IP Anchoring](#)).
- Specify the Server Groups hosting the applications.
- Specify the [App Connector groups](#) (government agencies, see [App Connector groups](#)) that have access to those Server Groups.

To add an Application Segment:

1. Go to **Resource Management > Application Management > Application Segments > Defined Application Segments**.
2. Click **Add Application Segment**. The **Add Application Segment** window appears.

The following example shows the Application Segment as a subnet. In this case, the Branch network is 10.10.5.0/24

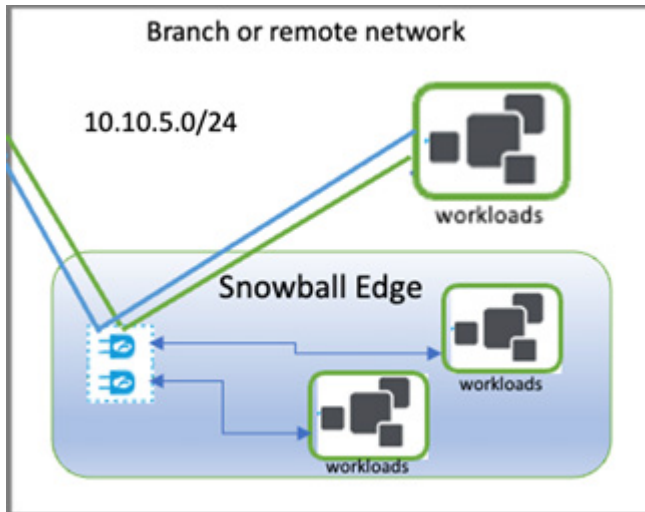


Figure 84. Branch network

The Application Segment 10.10.5.0/24 is defined in ZPA:

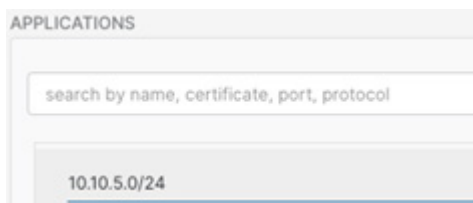


Figure 85. Applications

The Application Segment uses all ports. You can restrict the ports and protocols to an exact list. The following example enables the entire subnet to work for just one individual.

The screenshot shows the 'Default Port Ranges' configuration page. At the top, there is a 'Default Port Ranges' dropdown menu set to 'Select'. To the right, there is a 'TCP Keepalive' section with 'Enabled' and 'Disabled' buttons, where 'Disabled' is selected. Below this, there are two sections: 'TCP Port Ranges' and 'UDP Port Ranges'. Each section has two input fields for port ranges. For TCP, the first range is '1' to '52' and the second is '54' to '65535'. For UDP, the first range is '1' to '52' and the second is '54' to '65535'.

Figure 86. Default Port Ranges

3. After the App Connector is configured, enable an authorized user to connect to anything on the network, or just one application, etc. The following example adds App Connectors in a VPC in the AWS cloud on the subnet 192.168.1.0/24.

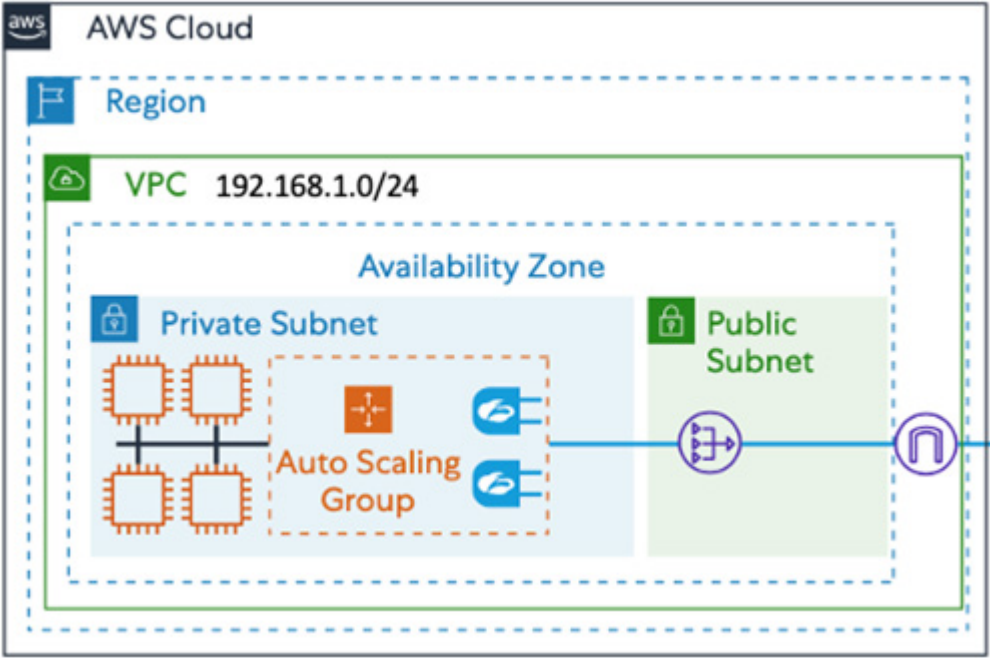


Figure 87. AWS VPC

If an additional Application Segment is created that includes the same subnet, users can seamlessly access resources in either subnet without needing to worry about the specific path. The App Connectors automatically identify the destination resource and securely route the connection accordingly.

Name	Applications	Status
> AWS Knotted Cloud Segment	192.168.1.0/24	
> AWS Snowball Cloud Segment	10.10.5.0/24	

Figure 88. Application segments in ZPA

[Access policy](#) (government agencies, see [Access policy](#)) rules enable you to implement role-based access control. To configure an access policy rule, you must first define the users and then define which applications or Segment Groups they can access. For example, you would specify the users first (i.e., Sales Staff), then specify which Application Segments or Segment Groups they can access (i.e., Sales App and Intranet Group).

If you want to configure application-based access control, you must create an access policy rule for specific Application Segments or Segment Groups. When you need to apply different policies to individual applications, create an Access Policy rule that includes one or more Application Segments. However, if you want all applications within a group of users who need a similar level of access across those applications, create an access policy rule that includes one or more Segment Groups.

Here's the example of the policy for access.

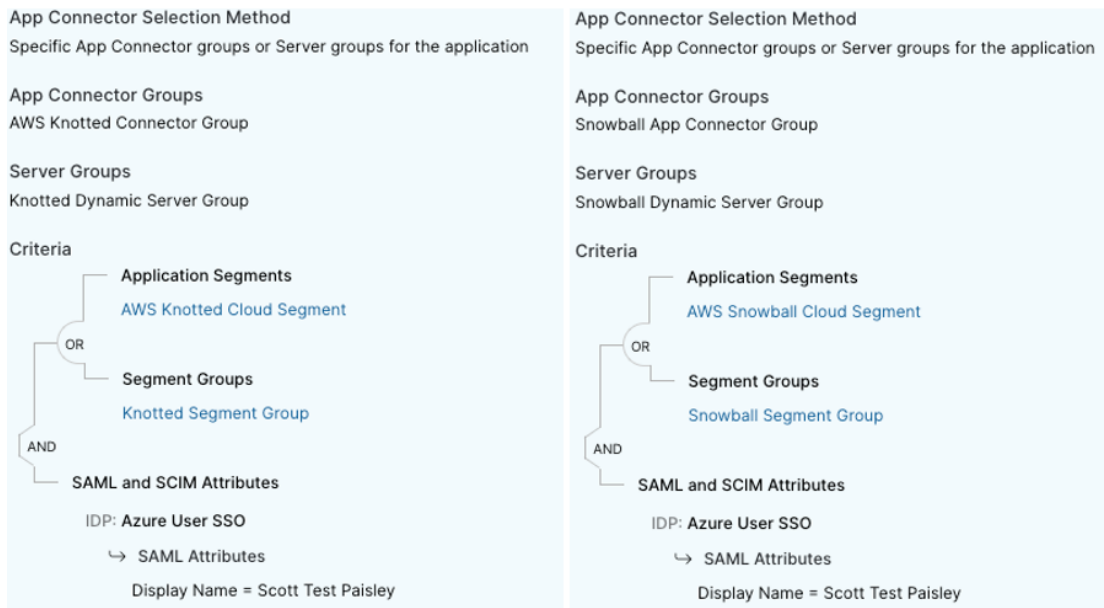


Figure 89. Access policies

The Knotted Cloud Segment allows Scott Test Paisley access to only the 192.168.1.0/24 network hosted in the AWS cloud by the App Connector AML from the Amazon Marketplace. The Snowball Cloud Segment allows Scott Test Paisley access to only the 10.10.5.0/24 network hosted at the branch office, enabled by an App Connector running on an AWS Snowball appliance.

After these policies are configured, you are ready to deploy the Zscaler Client Connector.

Deploy the Zscaler Client Connector

The Zscaler Client Connector is a lightweight traffic director agent that resides on the endpoints for personal entities (laptops, tablets, mobile devices, etc.). You can deploy the Zscaler Client Connector in production via an endpoint management platform such as Microsoft Intune or via Group Policy. You can download the Zscaler Client Connector from the Zscaler Client Connector Portal for laptops or via the applicable app store for mobile devices. For more information on downloading software for Zscaler Client Connector, see [Understanding Zscaler Client Connector App Downloads](#) (government agencies, see [Understanding Zscaler Client Connector App Downloads](#)).

The following figure shows downloading the Zscaler Client Connector software via the Zscaler Client Connector Portal.

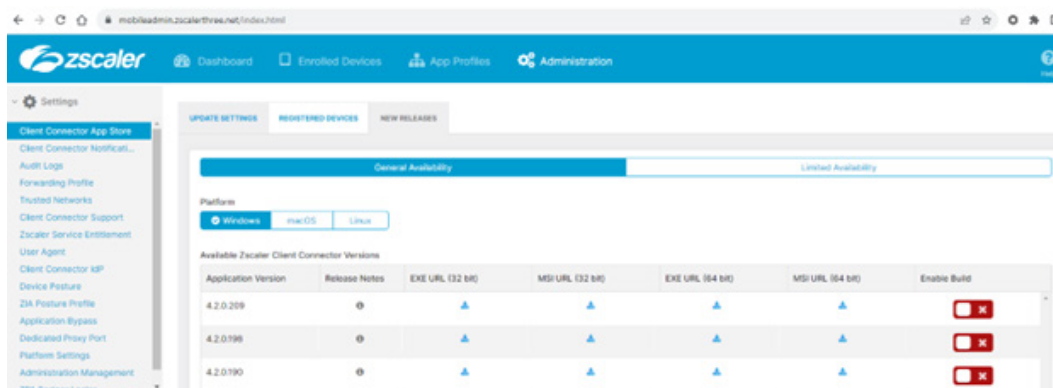


Figure 90. Zscaler Client Connector

You can find instructions for installing the software on any specific platform on the Zscaler Help site.

You can configure the Trusted Network. In this case, the Trusted Network refers to the Local Environment Network. The purpose is to recognize when the user is operating on-premises. It changes the behavior when you are on what you determine is the local network or off the local network.

The following is the criteria used to define a Trusted Network for Zscaler Client Connector:

- **DNS Server:** The DNS servers to which your corporate network sends DNS requests. Enter the DNS servers, separated by commas. IPv6 addresses are supported if you're using Zscaler Client Connector version 3.4 or later. The app verifies at least one DNS server.
- **DNS Search Domains:** The search domains configured as the primary domains for the network adapter used for connecting to Zscaler. Enter the search domains, separated by commas. The app only verifies the primary domains assigned to the active network adapter.
- **Hostname and IP:** A hostname and the IP addresses to which the hostname resolves when users are on the corporate network. For Hostname, enter the hostname. For Resolved IPs for Hostname, enter the IP addresses that the hostnames resolve to, separated by commas. IPv6 addresses are supported if you're using Zscaler Client Connector version 3.4 or later. The app verifies at least one IP address.

Zscaler recommends selecting DNS Server and DNS Search Domains for Trusted Network Criteria, because they are static properties on the network interface.

Add Trusted Network

NETWORK DEFINITION

Network Name ?
Local Network

TRUSTED NETWORK CRITERIA

Add Condition ?
Select Add Condition

Condition Match
Any

DNS Servers ?
10.10.5.254

Save Cancel

Figure 91. Add Trusted Network

The Zscaler Client Connector configuration looks similar to the following image:

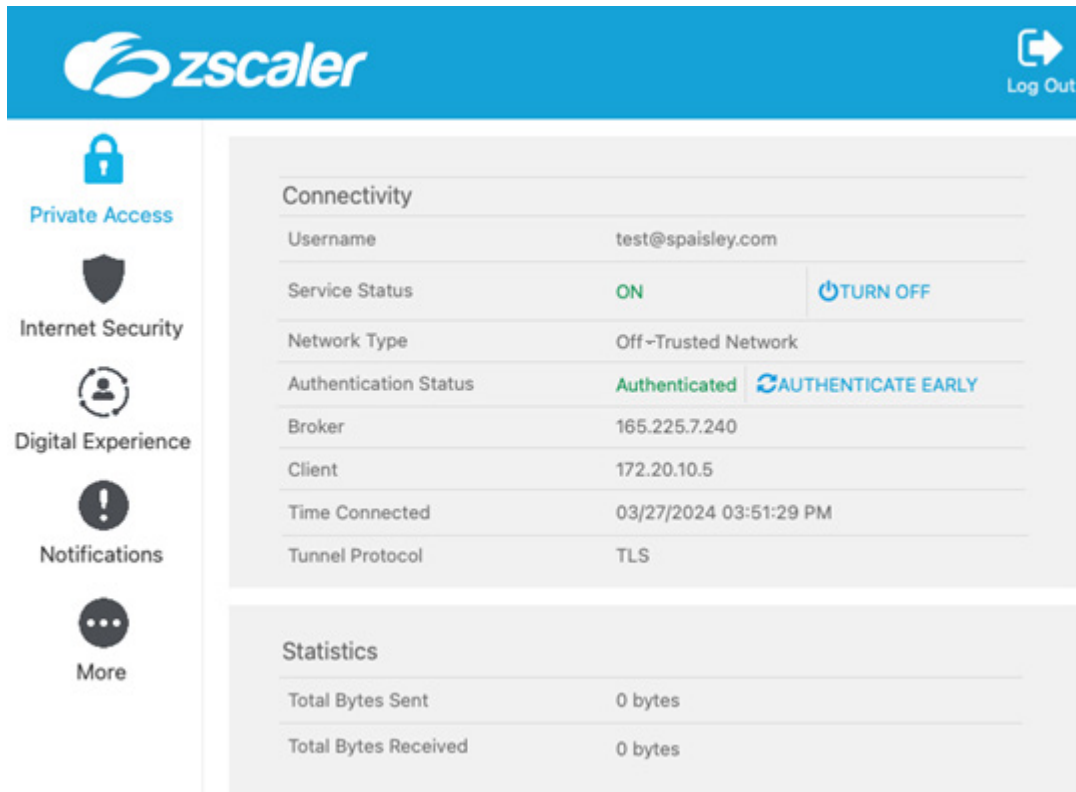


Figure 92. Zscaler Client Connector connectivity

Test the Setup

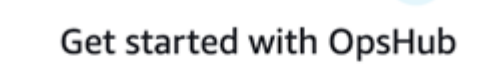
In the following example, the local IP is set to 172.20.10.5, and the user is off the trusted (local) network. Since Zscaler Client Connector is enabled and has a policy to allow connections to the 192.168.1.0/24 and 10.10.5.0/24 networks, the user can connect to any of those IPs just as if they were local to that network.

1. To test, attempt to go to a website in the AWS Knotted Cloud VPC.



Figure 93. Test webpage

You can also connect to the Snowball appliance with a client as if it is on the local network.



Get started with OpsHub

- ☒ **Sign into local devices**
You'll need an unlock code and manifest file
- ☐ **Sign into remote devices**
You'll need an access key & secret key

[Sign in](#)

Figure 94. OpsHub

The following shows signing in to Snowcone.


Sign into your Snowcone

Sign in with an unlock code and manifest file

Device IP address

Unlock code

Manifest file

 Choose file

JID0d526927-91f8-4a98-b961-be9192eaade8_manifest.bin

Back

Sign in

Figure 95. Sign in to your Snowcone

The following shows the local devices.

aws

OpnHub

Remote devices

Local devices

« Back

Local devices

Local devices

Devices (0)

Q

Filter devices by device id

	Device id	IP address	Unlock status	Network status	Services
<input type="radio"/>	JID0d526927-91f8-4a98-b961-be9192eaade8	10.10.5.35	Unlocked	-	6

Figure 96. Local devices

The following is a sample website on a computer in the 10. network as if the user is local to this network.

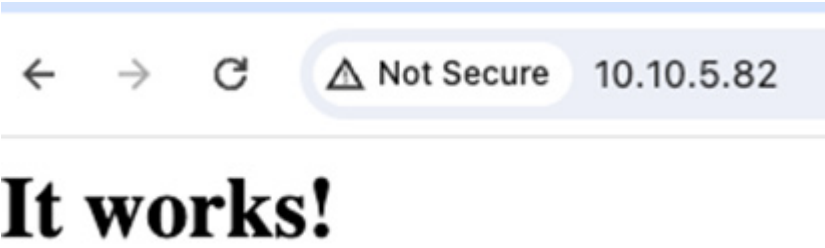


Figure 97. Sample webpage on the 10. network

- 2. To see all the connections, the following are the logs from the connections in your ZPA Admin Portal:

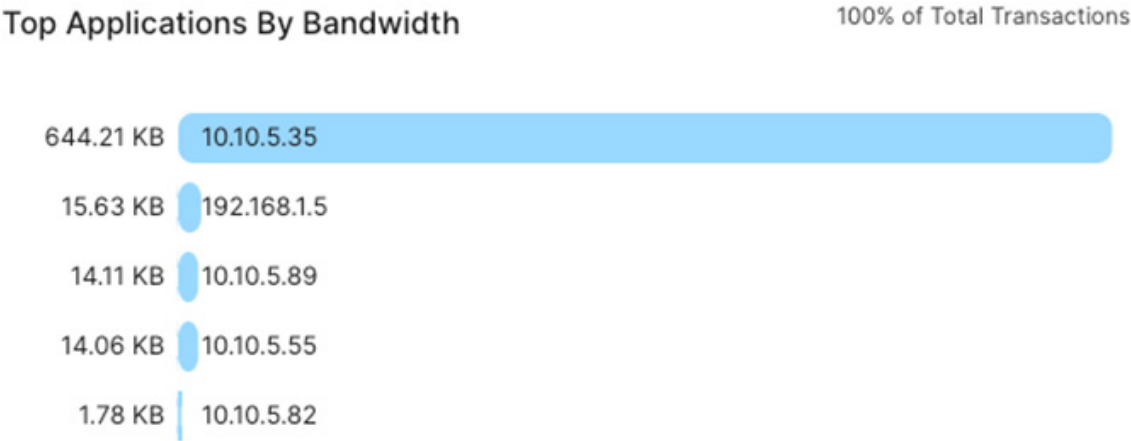


Figure 98. Top Applications by Bandwidth

The following image shows Application Segments by bandwidth.

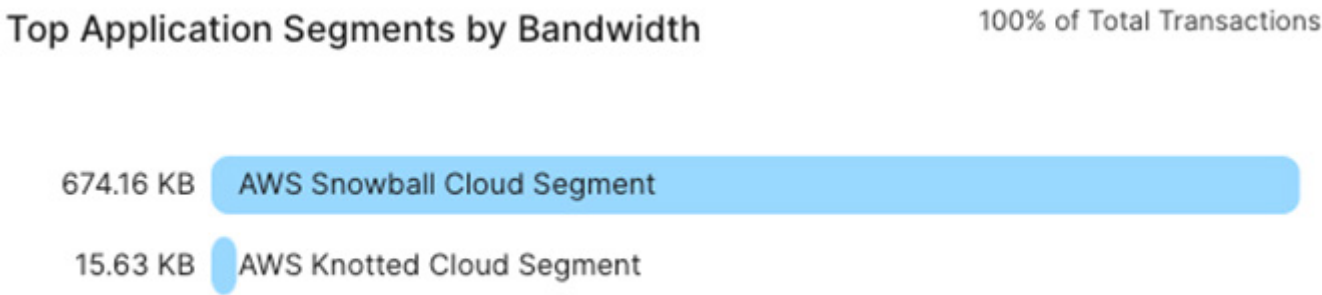


Figure 99. Top Application Segments by Bandwidth

Monitoring ZPA App Connector Health using AWS CloudWatch

AWS CloudWatch agent is installed on an App Connector to send various App Connector EC2 metrics and system log messages to CloudWatch. This agent enables admins to access a central repository of health and performance data rather than logging into App Connectors individually. See the following Zscaler community post for [Instrumenting Deployed App Connectors](#). The following steps are based on procedures documented on the Amazon AWS website.

Create an IAM Role to Use for CloudWatch Agent and Assign to App Connector

To learn more about creating an IAM role for CloudWatch, see the [AWS documentation](#).



If you plan to use AWS Systems Manager to download and install the CloudWatch agent, refer to the [AWS documentation](#) instead of the linked steps in this guide. Systems Manager is not used in this guide.

1. In the AWS Management Console, go to **IAM > Roles**.

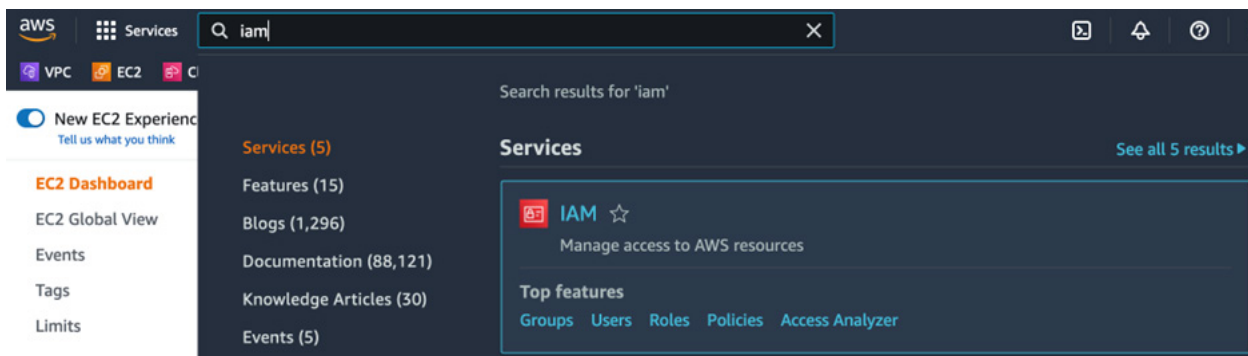


Figure 100. IAM services

2. Click **Create role**.

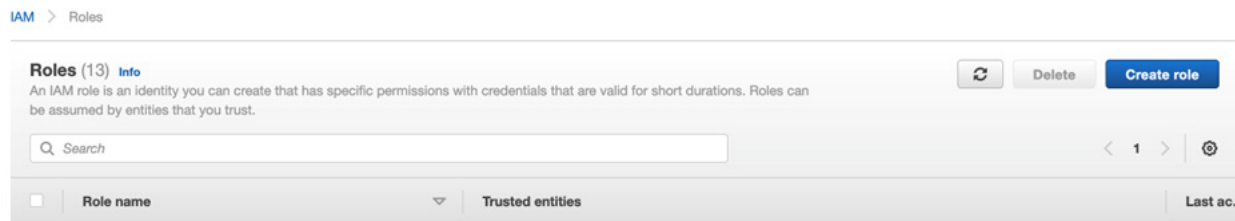


Figure 101. IAM role creation

3. Select **trusted entity**, ensure that **AWS service** is selected under **Trusted entity type** and **EC2** is selected under **Common use cases**.
4. Click **Next**.

Step 1
Select trusted entity

Step 2
Add permissions

Step 3
Name, review, and create

Select trusted entity

Trusted entity type

- ☒ **AWS service**
Allow AWS services like EC2, Lambda, or others to perform actions in this account.
- ☐ **AWS account**
Allow entities in other AWS accounts belonging to you or a 3rd party to perform actions in this account.
- ☐ **Web identity**
Allows users federated by the specified external web identity provider to assume this role to perform actions in this account.
- ☐ **SAML 2.0 federation**
Allow users federated with SAML 2.0 from a corporate directory to perform actions in this account.
- ☐ **Custom trust policy**
Create a custom trust policy to enable others to perform actions in this account.

Use case
Allow an AWS service like EC2, Lambda, or others to perform actions in this account.

Common use cases

- ☒ **EC2**
Allows EC2 instances to call AWS services on your behalf.
- ☐ **Lambda**
Allows Lambda functions to call AWS services on your behalf.

Use cases for other AWS services:
Choose a service to view use case

Cancel Next

Figure 102. IAM role Trusted entity type selection

5. Add permissions, select the **CloudWatchAgentServerPolicy** policy. If necessary, you can use the search box to filter the policy names.
6. Click **Next**.

Step 1
Select trusted entity

Step 2
Add permissions

Step 3
Name, review, and create

Add permissions

Permissions policies (Selected 1/746)
Choose one or more policies to attach to your new role.

Filter policies by property or policy name and press enter 1 match

"CloudWatchAgentServerPolicy" X Clear filters

<input checked="" type="checkbox"/>	Policy name	Type	Description
<input checked="" type="checkbox"/>	CloudWatchAgentServerPolicy	AWS m...	Permissions required to use AmazonCloudWatchAgent on servers

Set permissions boundary - optional
Set a permissions boundary to control the maximum permissions this role can have. This is not a common setting, but you can use it to delegate permission management to others.

Cancel Previous Next

Figure 103. IAM role permissions

- Name, review, create, name the role **CloudWatchAgentServerRole**, scroll to the bottom, and click **Create role**. Optionally you can change the default description.

Step 1
[Select trusted entity](#)

Step 2
[Add permissions](#)

Step 3
Name, review, and create

Name, review, and create

Role details

Role name
Enter a meaningful name to identify this role.

Maximum 128 characters. Use alphanumeric and '+', '=', '@', '-' characters.

Description
Add a short explanation for this policy.

Maximum 1000 characters. Use alphanumeric and '+', '=', '@', '-' characters.

Figure 104. IAM role naming

- Go to the EC2 instance of your App Connector in **EC2 > Instances** and select the **App Connector** (you can apply a filter, if needed, to find it).
- Under **Actions**, select **Security > Modify IAM Role**.

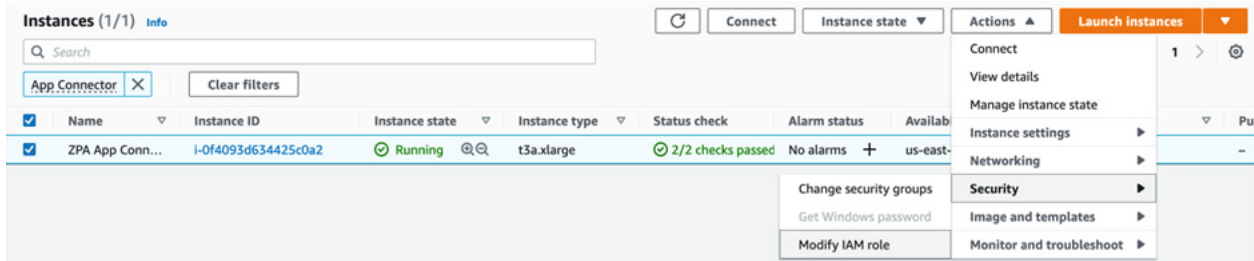


Figure 105. IAM EC2 modify role

- Select the **CloudWatchAgentServerRole** from the drop-down menu to attach the IAM role to this EC2 instance.

EC2 > Instances > i-Of4093d634425c0a2 > Modify IAM role

Modify IAM role [Info](#)

Attach an IAM role to your instance.

Instance ID
i-Of4093d634425c0a2 (ZPA App Connector)

IAM role
Select an IAM role to attach to your instance or create a new role if you haven't created any. The role you select replaces any roles that are currently attached to your instance.

[Create new IAM role](#)

Cancel **Save**

Figure 106. IAM role assignment

- Click **Save**.

Log In to the App Connector and Download the CloudWatch Agent

The following describes how to log in to the App Connector and Download the CloudWatch agent.

1. SSH into the App Connector (to the assigned public IP, or through ZPA). The Zscaler App Connector is currently based on CentOS. You can find the CentOS/x86-64 download link for the CloudWatch agent in the table on the Download the Cloudwatch Agent page linked from this guide. Use `wget` to download the RPM file to the current directory:

```
wget https://s3.amazonaws.com/amazoncloudwatch-agent/centos/amd64/latest/ama-
zon-cloudwatch-agent.rpm
```

2. Install the CloudWatch agent on the App Connector using the command:

```
sudo rpm -U ./amazon-cloudwatch-agent.rpm
```
3. The installation creates a user and group named `cwagent`. Although you can install the agent to run as the `cwagent`, the agent doesn't have access to the `/var/log/messages` file. The CloudWatch agent typically runs as the root user.

Create CloudWatch Configuration File Using Wizard

You can create the CloudWatch configuration file manually or by using a wizard. To create the configuration file manually, refer to the [AWS documentation](#). There are additional options available if you manually create the file. This guide shows how to use the wizard to create the configuration file. Run the following command: `sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-config-wizard`.

Example Dialog

```
=====
= Welcome to the Amazon CloudWatch Agent Configuration Manager =
=
= CloudWatch Agent allows you to collect metrics and logs from =
= your host and send them to CloudWatch. Additional CloudWatch =
= charges may apply. =
=====
```

On which OS are you planning to use the agent?

1. linux
2. windows
3. darwin

default choice: [1]:

Trying to fetch the default region based on ec2 metadata...

Are you using EC2 or On-Premises hosts?

1. EC2

2. On-Premises

default choice: [1]:

Which user are you planning to run the agent?

1. root

2. cwagent

3. others

default choice: [1]:

Do you want to turn on StatsD daemon?

1. yes

2. no

default choice: [1]:

2

Do you want to monitor metrics from CollectD? WARNING: CollectD must be installed or the Agent will fail to start

1. yes

2. no

default choice: [1]:

2

Do you want to monitor any host metrics? e.g. CPU, memory, etc.

1. yes

2. no

default choice: [1]:

Do you want to monitor cpu metrics per core?

1. yes

2. no

default choice: [1]:

Do you want to add ec2 dimensions (ImageId, InstanceId, InstanceType,

AutoScalingGroupName) into all of your metrics if the info is available?

1. yes

2. no

default choice: [1]:

Do you want to aggregate ec2 dimensions (InstanceId)?

1. yes

2. no

default choice: [1]:

Would you like to collect your metrics at high resolution (sub-minute resolution)? This enables sub-minute resolution for all metrics, but you can customize for specific metrics in the output json file.

1. 1s

2. 10s

3. 30s

4. 60s

default choice: [4]:

Which default metrics config do you want?

1. Basic

2. Standard

3. Advanced

4. None

default choice: [1]:

Current config as follows:

```
{
  "agent": {
    "metrics_collection_interval": 60,
    "run_as_user": "cwagent"
  },
```

```

"metrics": {
  "aggregation_dimensions": [
    [
      "InstanceId"
    ]
  ],
  "append_dimensions": {
    "AutoScalingGroupName": "${aws:AutoScalingGroupName}",
    "ImageId": "${aws:ImageId}",
    "InstanceId": "${aws:InstanceId}",
    "InstanceType": "${aws:InstanceType}"
  },
  "metrics_collected": {
    "cpu": {
      "measurement": [
        "cpu_usage_idle",
        "cpu_usage_iowait",
        "cpu_usage_user",
        "cpu_usage_system"
      ],
      "metrics_collection_interval": 60,
      "resources": [
        "*"
      ],
      "totalcpu": false
    },
    "disk": {
      "measurement": [
        "used_percent",
        "inodes_free"

```

```

    ],
    "metrics_collection_interval": 60,
    "resources": [
        "*"
    ]
},
"diskio": {
    "measurement": [
        "io_time"
    ],
    "metrics_collection_interval": 60,
    "resources": [
        "*"
    ]
},
"mem": {
    "measurement": [
        "mem_used_percent"
    ],
    "metrics_collection_interval": 60
},
"swap": {
    "measurement": [
        "swap_used_percent"
    ],
    "metrics_collection_interval": 60
}
}
}
}

```

Are you satisfied with the above config? Note: it can be manually customized after the wizard completes to add additional items.

1. yes

2. no

default choice: [1]:

Do you have any existing CloudWatch Log Agent (<http://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AgentReference.html>) configuration file to import for migration?

1. yes

2. no

default choice: [2]:

Do you want to monitor any log files?

1. yes

2. no

default choice: [1]:

Log file path:

/var/log/messages

Log group name:

default choice: [messages]

Log stream name:

default choice: [{instance_id}]

Log Group Retention in days

1. -1

2. 1

3. 3

4. 5

5. 7

- 6. 14
- 7. 30
- 8. 60
- 9. 90
- 10. 120
- 11. 150
- 12. 180
- 13. 365
- 14. 400
- 15. 545
- 16. 731
- 17. 1827
- 18. 3653

default choice: [1]:

Do you want to specify any additional log files to monitor?

- 1. yes
- 2. no

default choice: [1]:

2

Saved config file to /opt/aws/amazon-cloudwatch-agent/bin/config.json successfully.

Current config as follows:

```
{
  "agent": {
    "metrics_collection_interval": 60,
    "run_as_user": "cwagent"
  },
  "logs": {
    "logs_collected": {
      "files": {
```

```

        "collect_list": [
            {
                "file_path": "/var/log/messages",
                "log_group_name": "messages",
                "log_stream_name": "{instance_id}",
                "retention_in_days": -1
            }
        ]
    }
}

},
"metrics": {
    "aggregation_dimensions": [
        [
            "InstanceId"
        ]
    ],
    "append_dimensions": {
        "AutoScalingGroupName": "${aws:AutoScalingGroupName}",
        "ImageId": "${aws:ImageId}",
        "InstanceId": "${aws:InstanceId}",
        "InstanceType": "${aws:InstanceType}"
    },
    "metrics_collected": {
        "cpu": {
            "measurement": [
                "cpu_usage_idle",
                "cpu_usage_iowait",
                "cpu_usage_user",
                "cpu_usage_system"
            ]
        }
    }
}

```

```

    ],
    "metrics_collection_interval": 60,
    "resources": [
        "*"
    ],
    "totalcpu": false
},
"disk": {
    "measurement": [
        "used_percent",
        "inodes_free"
    ],
    "metrics_collection_interval": 60,
    "resources": [
        "*"
    ]
},
"diskio": {
    "measurement": [
        "io_time"
    ],
    "metrics_collection_interval": 60,
    "resources": [
        "*"
    ]
},
"mem": {
    "measurement": [
        "mem_used_percent"
    ],

```

```

        "metrics_collection_interval": 60
    },
    "swap": {
        "measurement": [
            "swap_used_percent"
        ],
        "metrics_collection_interval": 60
    }
}
}
}

```

Please check the above content of the config.

The config file is also located at `/opt/aws/amazon-cloudwatch-agent/bin/config.json`.

Edit it manually if needed.

Do you want to store the config in the SSM parameter store?

1. yes

2. no

default choice: [1]:

2

Program exits now.

Start CloudWatch Agent

Now that the configuration file is built, you can start the agent. Enter the following command to start the agent and use the local config file that was just created with the wizard:

```
sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl -a fetch-config -m ec2 -s -c file:/opt/aws/amazon-cloudwatch-agent/bin/config.json
```

Example Output

```
***** processing amazon-cloudwatch-agent *****
```

```
/opt/aws/amazon-cloudwatch-agent/bin/config-downloader --output-dir /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.d --download-source file:/opt/aws/amazon-cloudwatch-agent/bin/config.json --mode ec2 --config /opt/aws/amazon-cloudwatch-agent/etc/common-config.toml --multi-config default
```

```
Successfully fetched the config and saved in /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.d/file_config.json.tmp
```

Start configuration validation...

```
/opt/aws/amazon-cloudwatch-agent/bin/config-translator --input /opt/aws/amazon-cloud-
watch-agent/etc/amazon-cloudwatch-agent.json --input-dir /opt/aws/amazon-cloud-
watch-agent/etc/amazon-cloudwatch-agent.d --output /opt/aws/amazon-cloudwatch-agent/
etc/amazon-cloudwatch-agent.toml --mode ec2 --config /opt/aws/amazon-cloudwatch-agent/
etc/common-config.toml --multi-config default
```

```
2022/04/18 18:43:52 Reading json config file path: /opt/aws/amazon-cloudwatch-agent/etc/
amazon-cloudwatch-agent.d/file_config.json.tmp ...
```

Valid Json input schema.

I! Detecting run_as_user...

No csm configuration found.

Configuration validation first phase succeeded

```
/opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent -schematest -config /opt/
aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.toml
```

Configuration validation second phase succeeded

Configuration validation succeeded

amazon-cloudwatch-agent has already been stopped

```
Created symlink from /etc/systemd/system/multi-user.target.wants/amazon-cloud-
watch-agent.service to /etc/systemd/system/amazon-cloudwatch-agent.service.
```

Redirecting to /bin/systemctl restart amazon-cloudwatch-agent.service

Monitor Metrics Generated by CloudWatch Agent

1. In the AWS Management Console, go to **CloudWatch** > **All Metrics**.
2. Select the **CWAgent** custom namespace.



CWAgent is the default name. Customize the name using the namespace field in the metrics section of the agent configuration file. If you modify the configuration file, you must restart the agent using the fetch-config option for the change to take effect.

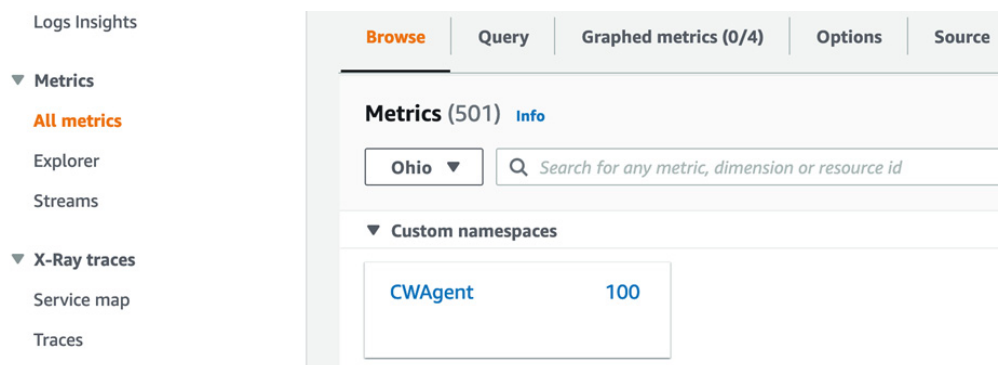


Figure 107. CloudWatch all metrics

- You can select and graph metrics.

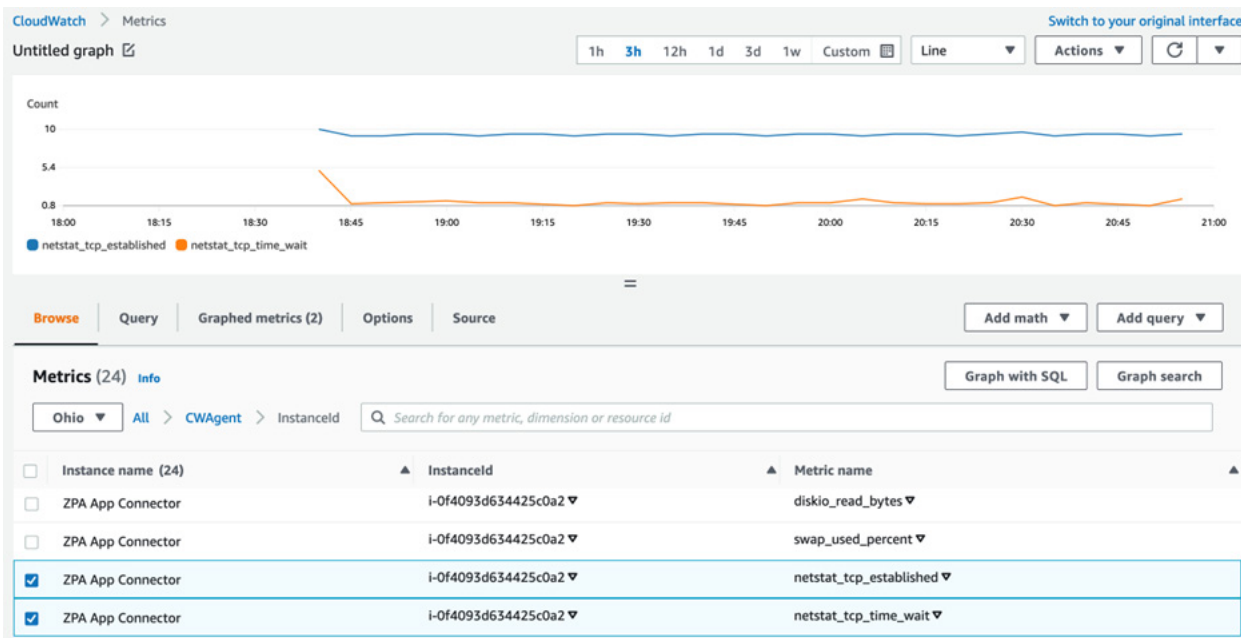


Figure 108. Metrics graph example



The metrics shown are available as part of the Advanced metric config. See [Appendix E: Metrics Config Options](#) for a description of the metrics available in each config option.

View Logs Stored in CloudWatch

To view logs stored in CloudWatch:

- In the AWS Management Console, go to **CloudWatch > Log groups**.
- Click **messages** under **Log groups**.

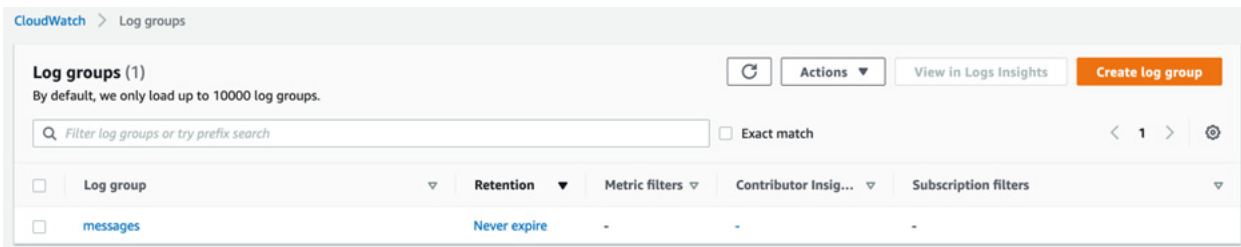


Figure 109. CloudWatch log groups



To set the retention time, click **Never expire** and select your desired time period.

3. You can see the app connectors that you are collecting messages from listed by instance ID under **Log streams**. Select the desired Instance ID to view the log file contents.

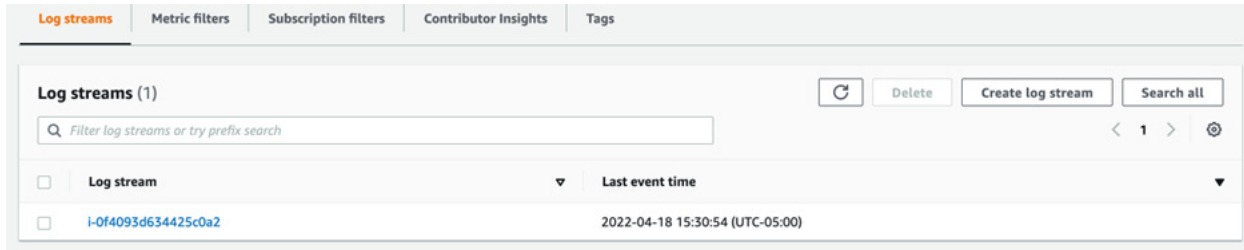


Figure 110. CloudWatch Log Stream

4. The newest lines are shown. Click **Load more** to see additional lines.

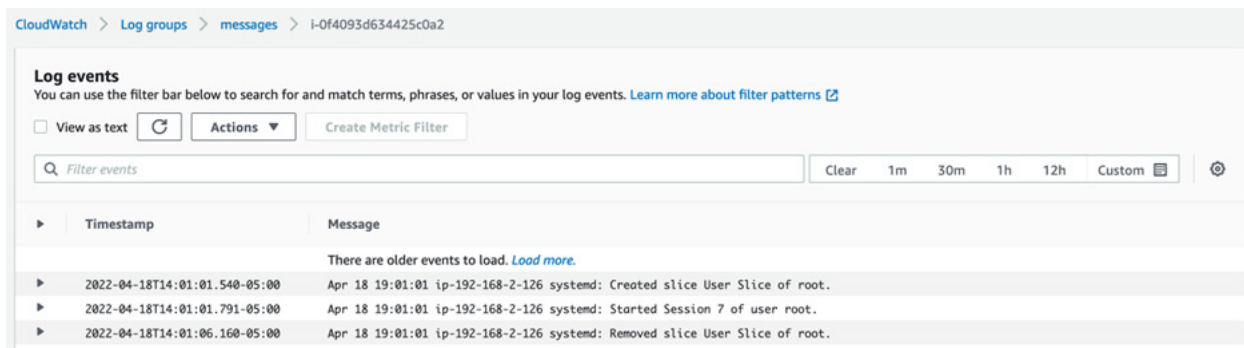


Figure 111. CloudWatch log events

5. Search messages by keyword using the Filter events input area.

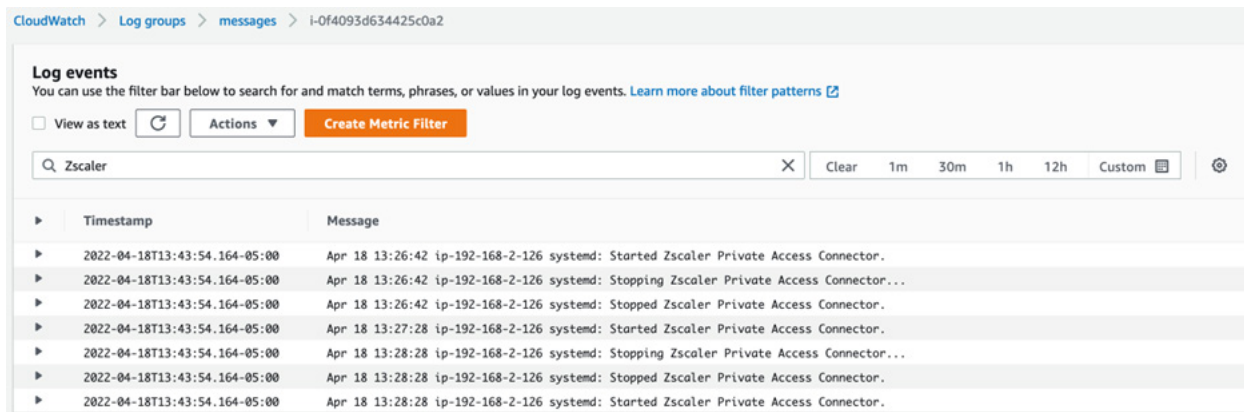


Figure 112. CloudWatch log filter



The filter ability is case sensitive.

Operationalizing Threat Intelligence with Zscaler and AWS GuardDuty

The following sections describe operationalizing threat intelligence with Zscaler and GuardDuty.

Feeding GuardDuty Telemetry into Zscaler Internet Access for Enforcement

This integration combines AWS GuardDuty's threat intelligence with Zscaler's ZTE to enhance security and visibility for traffic targeting potentially malicious destinations. Enabled via AWS Lambda, the system checks GuardDuty findings every 5 minutes (by default) and extracts suspicious FQDNs and IP addresses. These are then used to create or update IP and FQDN Destination Groups within ZIA. The updated Destination Groups are applied to a Firewall policy, which blocks and logs access to the flagged destinations. Each time the script runs, it overwrites the previous entries in the Destination Group, ensuring Destination Groups remain current. As GuardDuty automatically ages out entries over time, these changes are reflected in the ZIA Destination Group, maintaining an up-to-date defense.

Requirements and Components

Prior to deploying the integration, make sure the following requirements are met:

- A subscription to ZIA.
- A Zscaler Cloud Access API Key.
- Valid API service credentials that Lambda can use to run Read/Write operations against ZIA.
- A subscription to AWS GuardDuty with an active Detector (the Lambda function fails if an active Detector is not found).
- Access to AWS CloudFormation with permission to create a CloudWatch rule and Lambda Function with the following permissions:
 - secretsmanager:GetSecretValue
 - logs:CreateLogGroup
 - logs:CreateLogStream
 - logs:PutLogEvents
 - guardduty:ListDetectors
 - guardduty:ListFindings
 - guardduty:GetFindings
 - events:PutRule
 - events:PutTargets
 - events:RemoveTargets
 - events>DeleteRule
 - iam:PassRole



This guide assumes that ZIA and Deception have already been deployed within the environment. It does not include basic setup and configuration information. For information on configuring ZIA or Deception, see the [ZIA help documentation](#) (government agencies, see [ZIA help documentation](#)) or the [Deception help documentation](#).

This guide assumes that AWS GuardDuty has already been deployed within the environment. It does not include basic setup and configuration information. For information on configuring AWS GuardDuty, refer to the [AWS documentation](#).

Configuring ZIA

To integrate ZIA with AWS Lambda, generate or retrieve the Zscaler Cloud Services API key and base URL, then create a dedicated API service account with appropriate firewall and policy permissions. These credentials, along with the Login ID and Password, are stored in AWS Secrets Manager for secure authentication.

To generate or retrieve a Zscaler Cloud Services API key:

1. From the ZIA Admin Portal, go to **Administration > Authentication > Cloud Service API Key**.
2. Click **Add API Key** (or copy the existing key, if already present).

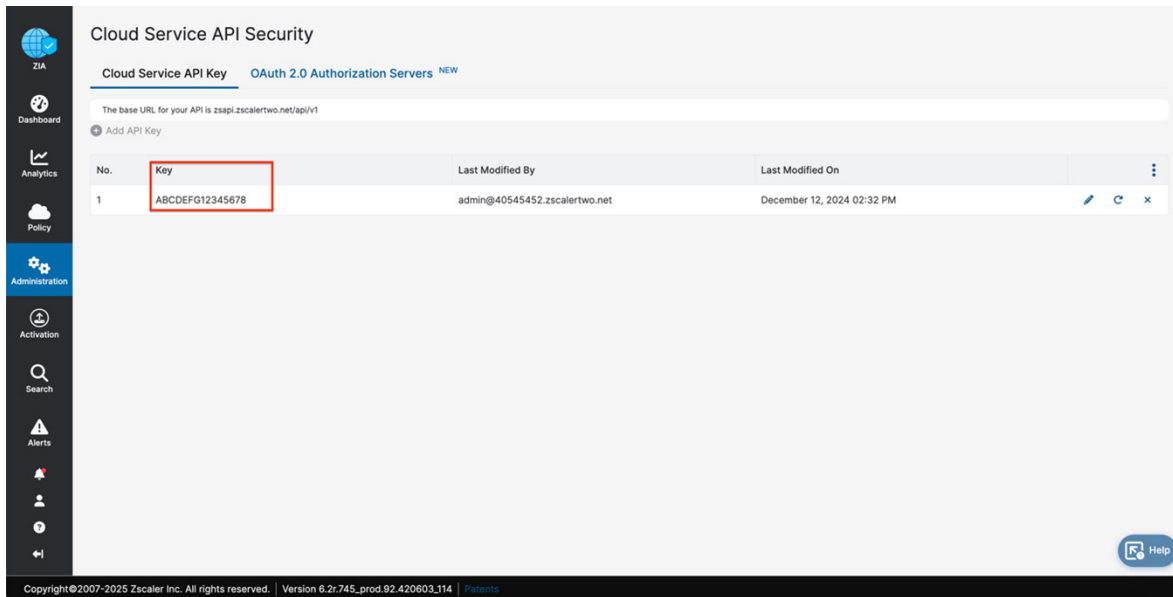


Figure 113. Cloud Service API Key

Zscaler API Service Account

Use an API Service Account to authenticate the Lambda service to the Zscaler cloud. Zscaler does not recommend using the Super Admin account to run the Lambda function:

1. Go to **Administration > Administration Controls > Role Management**.
2. Click **Add Administrator Role**.
3. Enter a **Name**.
4. Enable the **Firewall** and **Policy Access** permissions.

Add Admin Role [X]

ADMINISTRATOR ROLE

Name: provisioning

PERMISSIONS

Dashboard	Cloud Connector Provisioning
View Only <input type="radio"/> None <input checked="" type="radio"/>	<input checked="" type="radio"/> Full <input type="radio"/> View Only <input type="radio"/> None
Template (Location & Provisioning)	Administrator Management
Full <input type="radio"/> View Only <input type="radio"/> None <input checked="" type="radio"/>	Full <input type="radio"/> View Only <input type="radio"/> None <input checked="" type="radio"/>
Location Management	Forwarding (Traffic, DNS & Logs)
<input checked="" type="radio"/> Full <input type="radio"/> View Only <input type="radio"/> None	Full <input type="radio"/> View Only <input type="radio"/> None <input checked="" type="radio"/>
API Key Management	Remote Assistance Management
Full <input type="radio"/> View Only <input type="radio"/> None <input checked="" type="radio"/>	Full <input type="radio"/> View Only <input checked="" type="radio"/>
NSS Logging	Public Cloud Config Management
Full <input type="radio"/> None <input checked="" type="radio"/>	Full <input checked="" type="radio"/> View Only <input type="radio"/> None

Figure 114. Add Administrator Role

5. Click **Save**.
6. Create a new **Service Account** for Lambda:
 - (Option 1) If using the built-in ZIA database:
 - Go to **Administration > Administration Controls > Administrator Management**.
 - Click **Add Administrator**.
 - Enter a **Login ID**, **Email**, and **Name**.
 - Select the **Role**.
 - Click **Save**.
 - (Option 2) If using an Identity Provider, create the account in the IdP and ensure the role created previously is attached to this account.
7. Note the **Login ID** and **Password** for this account (these values are placed in the AWS Secrets Manager object along with the Cloud Services API Key and base URL).

Installing the Integration Script and Configuring AWS

To set up the integration, download the required files from the GitHub repository, create an S3 bucket to store the integration script, and configure an AWS Secrets Manager object to securely store ZIA credentials. Then, deploy the CloudFormation script to create the Lambda function.

Download the Integration Script

Download the `lambda_guardduty_threatfeed.zip` and `threatfeed.yaml` files from the [Github repository](#).

Creating an S3 Bucket

An S3 bucket houses the appropriate Python libraries and script necessary to provide the integration. The CloudFormation script leverages this S3 bucket to create the Lambda function.

1. Log in to your AWS account.
2. Go to the **S3** dashboard.
3. Click **Create bucket**.
4. Enter a **Bucket name** (note this name for future use).
5. Select **Block all public access**.

Create bucket [Info](#)

Buckets are containers for data stored in S3.

General configuration

AWS Region
US East (Ohio) us-east-2

Bucket type [Info](#)

☒ **General purpose**
Recommended for most use cases and access patterns. General purpose buckets are the original S3 bucket type. They allow a mix of storage classes that redundantly store objects across multiple Availability Zones.

☐ **Directory**
Recommended for low-latency use cases. These buckets use only the S3 Express One Zone storage class, which provides faster processing of data within a single Availability Zone.

Bucket name [Info](#)

Bucket name must be unique within the global namespace and follow the bucket naming rules. [See rules for bucket naming](#)

Copy settings from existing bucket - optional
Only the bucket settings in the following configuration are copied.

Format: s3://bucket/prefix

Object Ownership [Info](#)
Control ownership of objects written to this bucket from other AWS accounts and the use of access control lists (ACLs). Object ownership determines who can specify access to objects.

☒ **ACLs disabled (recommended)**
All objects in this bucket are owned by this account. Access to this bucket and its objects is specified using only policies.

☐ **ACLs enabled**
Objects in this bucket can be owned by other AWS accounts. Access to this bucket and its objects can be specified using ACLs.

Object Ownership
Bucket owner enforced

Block Public Access settings for this bucket
Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, access point policies, or all. In order to ensure that public access to this bucket and its objects is blocked, turn on Block all public access. These settings apply only to this bucket and its access points. AWS recommends that you turn on Block all public access, but before applying any of these settings, ensure that your applications will work correctly without public access. If you require some level of public access to this bucket or objects within, you can customize the individual settings below to suit your specific storage use cases. [Learn more](#)

☒ **Block all public access**
Turning this setting on is the same as turning on all four settings below. Each of the following settings are independent of one another.

☐ **Block public access to buckets and objects granted through new access control lists (ACLs)**

Figure 115. Create bucket

6. Verify remaining settings as per your organization and click **Create Bucket**.
7. Select your new S3 bucket.
8. Click **Upload**.
9. Upload the downloaded `lambda_guardduty_threatfeed.zip`.

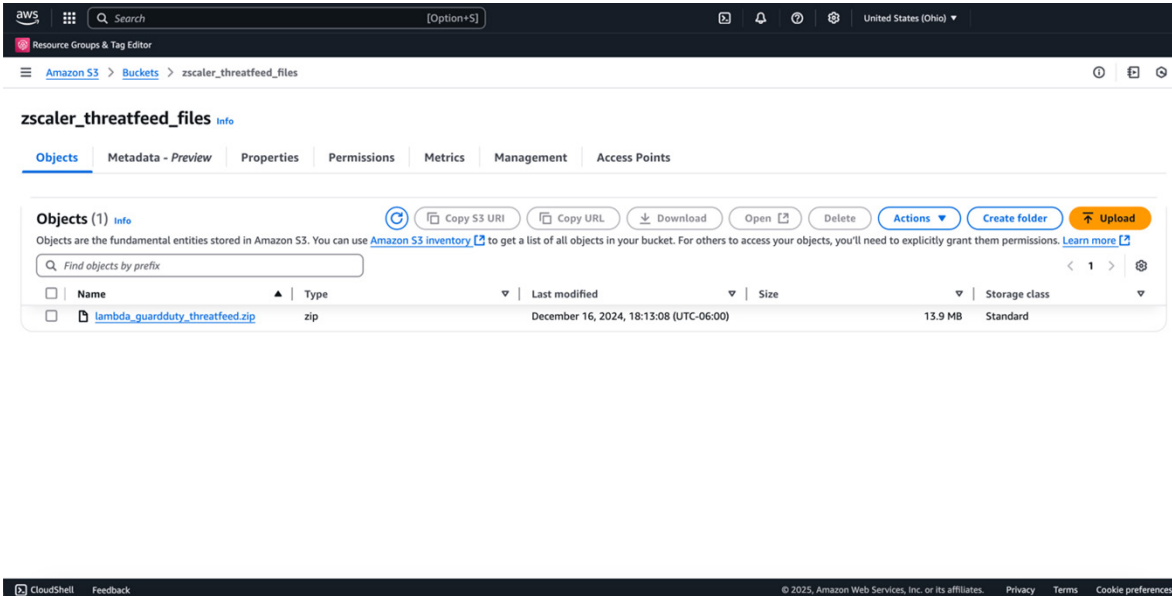


Figure 116. Resource Groups & Tag Editor

Creating a Secrets Manager Object

Create an AWS Secrets Manager object to securely house the ZIA credentials that Lambda uses to make policy updates.

1. Go to the **Secrets Manager** dashboard.
2. Click **Store a New Secret**.
3. Choose **Other Type of Secret**.
4. Enter the following Key/Value pairs:

Key	Value
zscaler_api_base_url	<Your API Base URL>
zscaler_username	<Your API Service Account username>
zscaler_password	<Your API Service Account password>
zscaler_api_key	<Your Zscaler Cloud Service API Key>

The screenshot shows the AWS Secrets Manager console. The breadcrumb trail is 'AWS Secrets Manager > Secrets > Store a new secret'. The left sidebar shows a progress bar with four steps: 'Step 1: Choose secret type' (active), 'Step 2: Configure secret', 'Step 3 - optional: Configure rotation', and 'Step 4: Review'. The main content area is titled 'Choose secret type'. Under 'Secret type', there are four radio buttons: 'Credentials for Amazon RDS database', 'Credentials for Amazon DocumentDB database', 'Credentials for Amazon Redshift data warehouse', and 'Other type of secret' (which is selected). Below this is the 'Key/value pairs' section, which has a 'Key/value' tab selected. It contains a table with four rows, each with a key, a value, and a 'Remove' button. The keys are 'zscaler_api_base_url', 'zscaler_username', 'zscaler_password', and 'zscaler_api_key'. The values are 'https://zsapi.zscaler.net/api/v1', 'username@domain.com', 'password', and 'ABCDEF12345678' respectively. There is an '+ Add row' button at the bottom of the table.

Key	Value	Action
zscaler_api_base_url	https://zsapi.zscaler.net/api/v1	Remove
zscaler_username	username@domain.com	Remove
zscaler_password	password	Remove
zscaler_api_key	ABCDEF12345678	Remove

Figure 117. Choose secret type

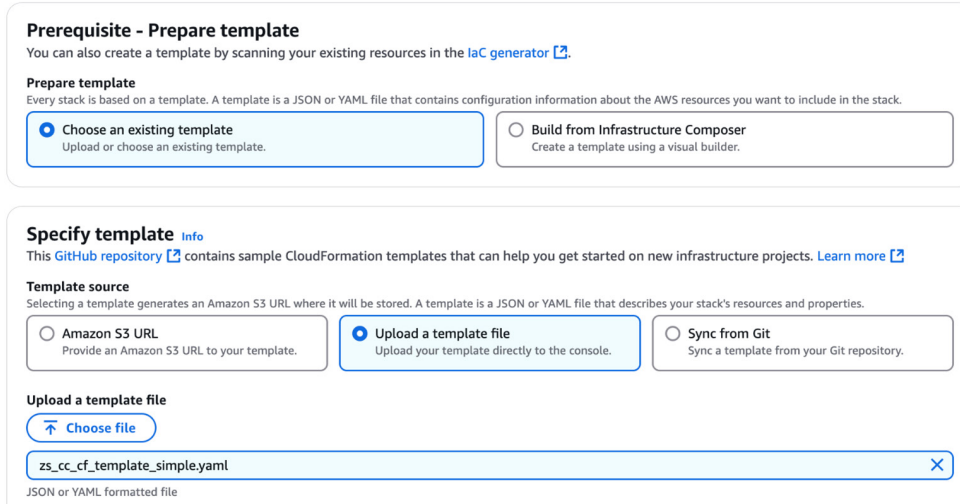
5. Click **Next**.
6. Enter a **Name**.
7. Verify remaining settings as per your organization and click **Next**.
8. Verify settings as per your organization and click **Next**.
9. Click **Store**.

Running the CloudFormation Script

Run the CloudFormation script (threatfeed.yaml) downloaded previously.

1. Browse to the **AWS CloudFormation** dashboard.
2. From the **AWS CloudFormation** dashboard, create a new stack with new resources.
3. Choose the option to **Upload a template file** and select the threatfeed.yaml file previously downloaded:

Create stack



Prerequisite - Prepare template
You can also create a template by scanning your existing resources in the [IaC generator](#).

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

☒ **Choose an existing template**
Upload or choose an existing template.

☐ **Build from Infrastructure Composer**
Create a template using a visual builder.

Specify template [Info](#)
This [GitHub repository](#) contains sample CloudFormation templates that can help you get started on new infrastructure projects. [Learn more](#)

Template source
Selecting a template generates an Amazon S3 URL where it will be stored. A template is a JSON or YAML file that describes your stack's resources and properties.

☐ **Amazon S3 URL**
Provide an Amazon S3 URL to your template.

☒ **Upload a template file**
Upload your template directly to the console.

☐ **Sync from Git**
Sync a template from your Git repository.

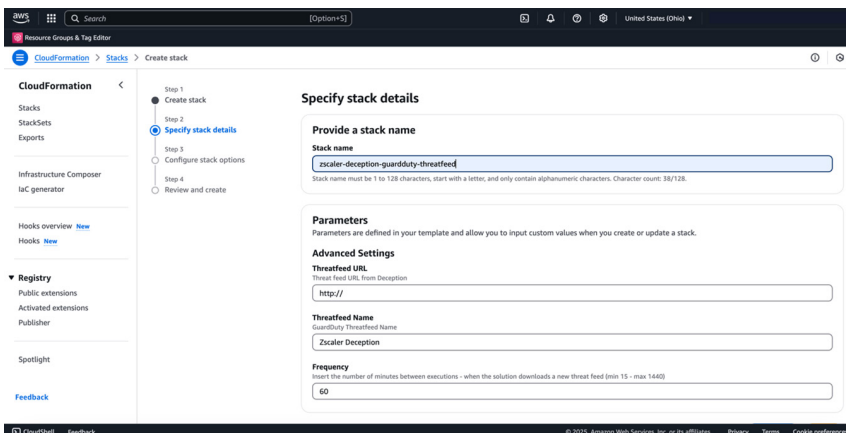
Upload a template file
[Choose file](#)

zs_cc_cf_template_simple.yaml

JSON or YAML formatted file

Figure 118. Create stack

4. Click **Next**.
5. Enter a **Stack Name** for the stack.
6. Enter a **Lambda Schedule Interval frequency** for the Lambda function to execute.
7. Enter your **S3 Bucket Name** created previously.
8. Enter the **S3 Key** (lambda_guarddduty_threatfeed.zip).
9. Enter the **Zscaler Secret Name**.



Specify stack details

Provide a stack name

Stack name
zscaler-deception-guardduty-threatfeed
Stack name must be 1 to 128 characters, start with a letter, and only contain alphanumeric characters. Character count: 38/128.

Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Advanced Settings

Threatfeed URL
Threat feed URL from Description
http://

Threatfeed Name
GuardDuty Threatfeed Name
Zscaler Deception

Frequency
Insert the number of minutes between executions - when the solution downloads a new threat feed (min 15 - max 1440)
60

Figure 119. Specify stack details

10. Click **Next**.
11. Proceed through the remainder of the wizard and update remaining default values as you see fit for your environment.

Verification

You can verify the integration in several ways:

From CloudWatch Logs:

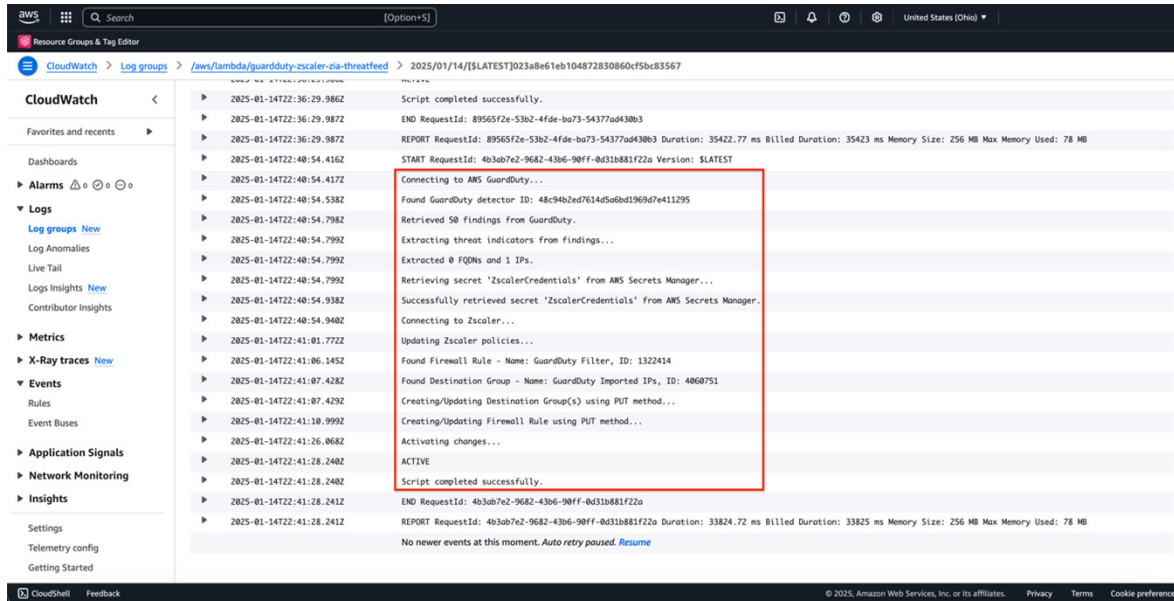


Figure 120. CloudWatch logs

From the Lambda Function itself (use the Test tab to verify the script can be manually executed). The script does not require any Event parameters to run:

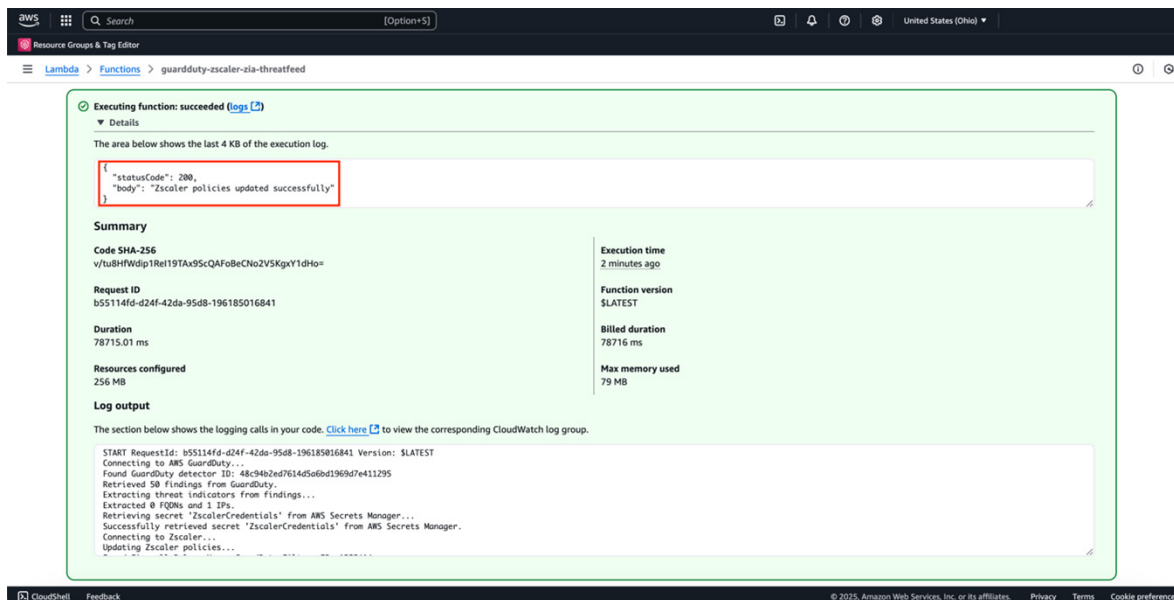


Figure 121. Lambda Function

From the ZIA Admin Portal:

1. Go to **Administration > IP & FQDN Groups**.
2. In the **Destination IPv4 Groups** tab, verify that new GuardDuty Imported groups appear (the description of these groups should also include a timestamp of when they were last updated).

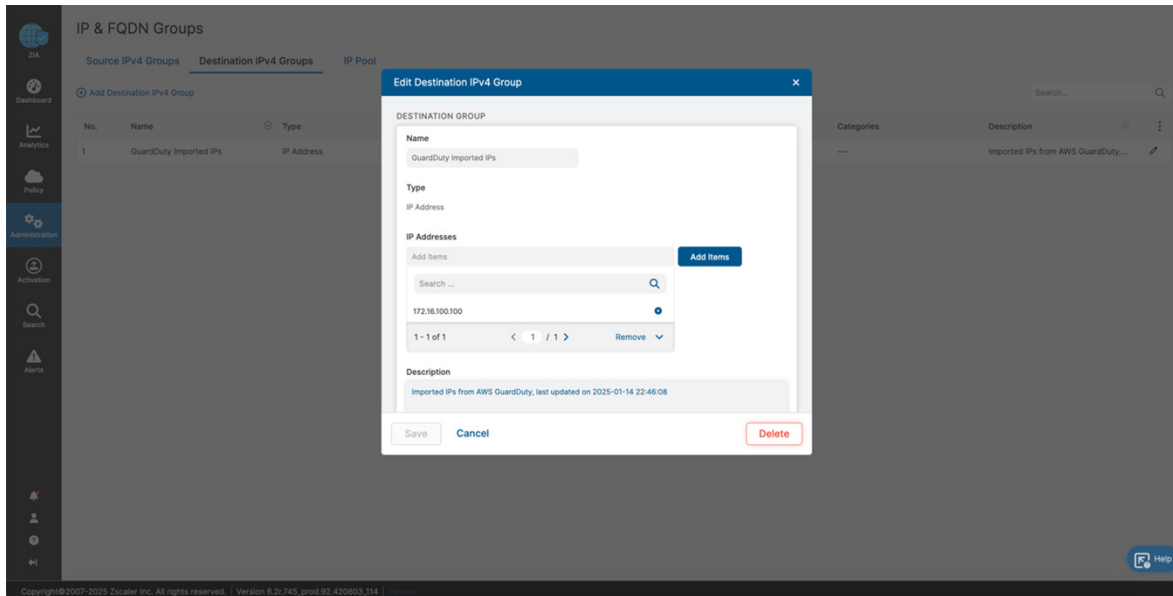


Figure 122. Destination IPv4 Groups

Feeding Zscaler Deception Telemetry into Amazon GuardDuty

Zscaler Deception enhances AWS GuardDuty by forwarding high-fidelity alerts as custom findings. This integration is powered by an AWS Lambda function that processes Deception event logs. During operation, Deception publishes its threatfeed to a web endpoint, which the Lambda function fetches and stores in an Amazon S3 bucket as a flat text file. The S3 file is updated only when the Deception feed detects changes, ensuring efficiency. GuardDuty then periodically retrieves updated findings from the S3 bucket. This setup allows Deception to deliver its advanced threat detection insights directly into GuardDuty, empowering security teams to swiftly identify and address threats.

Requirements and Components

Prior to deploying the integration, make sure the following requirements are met:

- A subscription to Zscaler Deception.
- A subscription to AWS GuardDuty with an active Detector (the Lambda function fails if no active Detectors are found).
- Access to AWS CloudFormation with permission to create a CloudWatch Rule, an S3 bucket, and Lambda Function with the following permissions:
 - logs:CreateLogGroup
 - logs:CreateLogStream
 - logs:PutLogEvents
 - guardduty:ListDetectors
 - guardduty:CreateThreatIntelSet
 - guardduty:GetThreatIntelSet
 - guardduty:ListThreatIntelSets

- guardduty:UpdateThreatIntelSet
- events:PutRule
- events:PutTargets
- events:RemoveTargets
- events>DeleteRule
- iam:PutRolePolicy
- iam>DeleteRolePolicy
- s3:getobject
- s3:putobject



This guide assumes that Deception is already deployed in the environment. It does not include basic Deception setup and configuration information. For information on configuring Deception, see the [Deception help portal](#).

This guide assumes that AWS GuardDuty is already deployed in the environment. It does not include basic setup and configuration information. For information on configuring AWS GuardDuty, refer to the [AWS documentation](#).

Configuring Deception

To configure Deception to export findings to AWS, use the following steps:

1. From the Deception Admin Portal, go to **Orchestrate > Rules**.
2. Click **Add Rule**.
3. In the window that appears, enter a **Name** for the rule and create the condition that triggers the rule (i.e., attacker.score > 1 and decoy.group is Threat Intelligence).

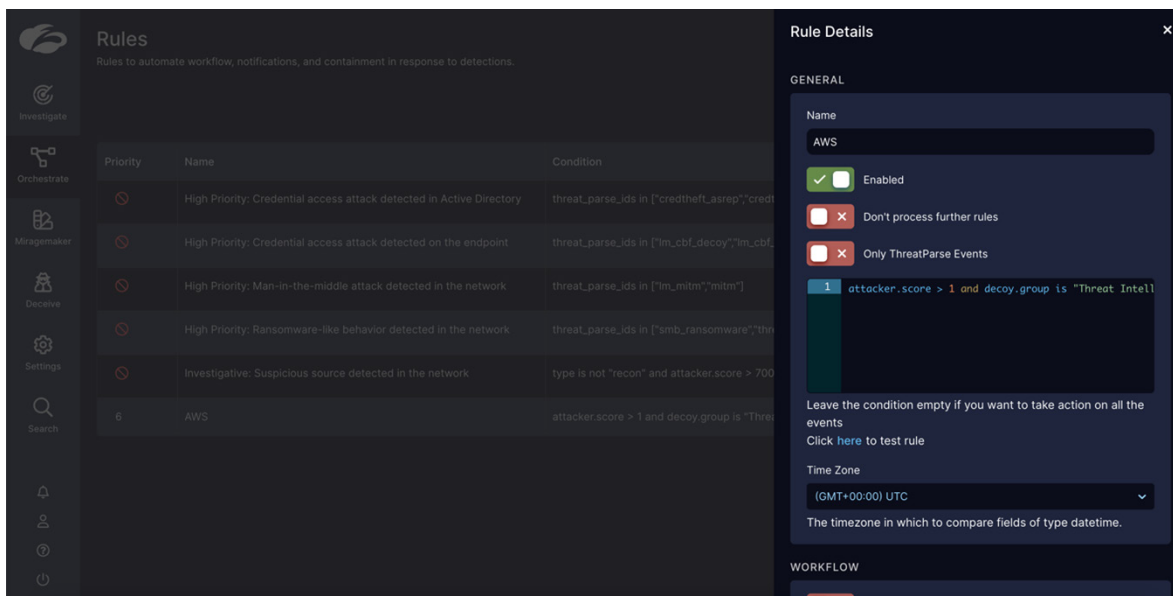


Figure 123. Rule Details

4. Select **Enabled** under **AWS**. Choose **AWS GuardDuty (Threat Feeds - Attacker IPs)** from the drop-down menu.
5. Click **Save**.
6. Go to **Orchestrate > Containment**.
7. In the list that appears, find the AWS GuardDuty integration and click **Edit**.

8. In the pane that appears, enable the integration.
9. Click **Save**.
10. Click **Edit** next to **AWS GuardDuty**.
11. In the window that appears, note the URL. This is entered into the CloudFormation template wizard in the next section.

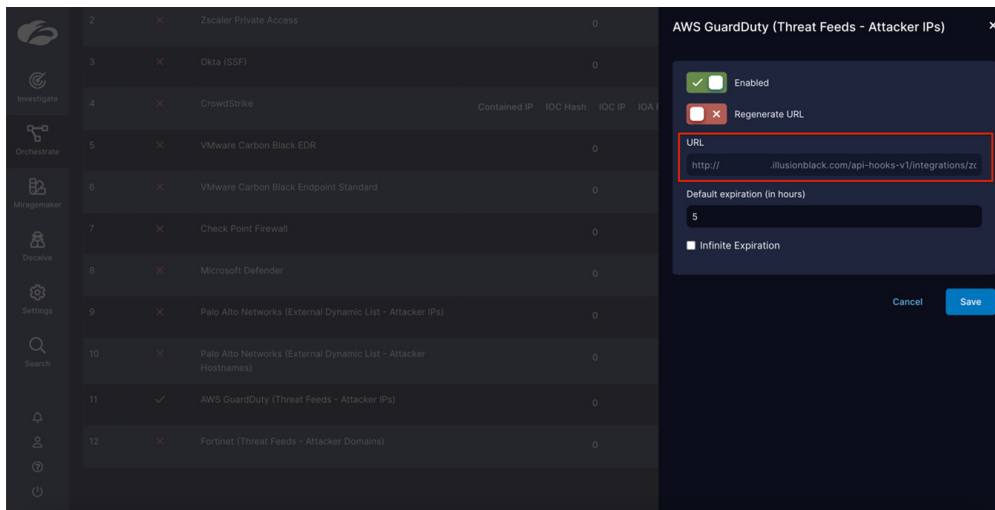


Figure 124. AWS GuardDuty (Threat Feeds – Attacker IPs)

Installing the Integration Script in AWS

To configure AWS GuardDuty to consume Deception findings, use the following steps:

1. Browse to the **Zscaler Script Samples** repository and download the `threatfeed.yaml` file from the `deception-to-AWS-guardduty` folder:

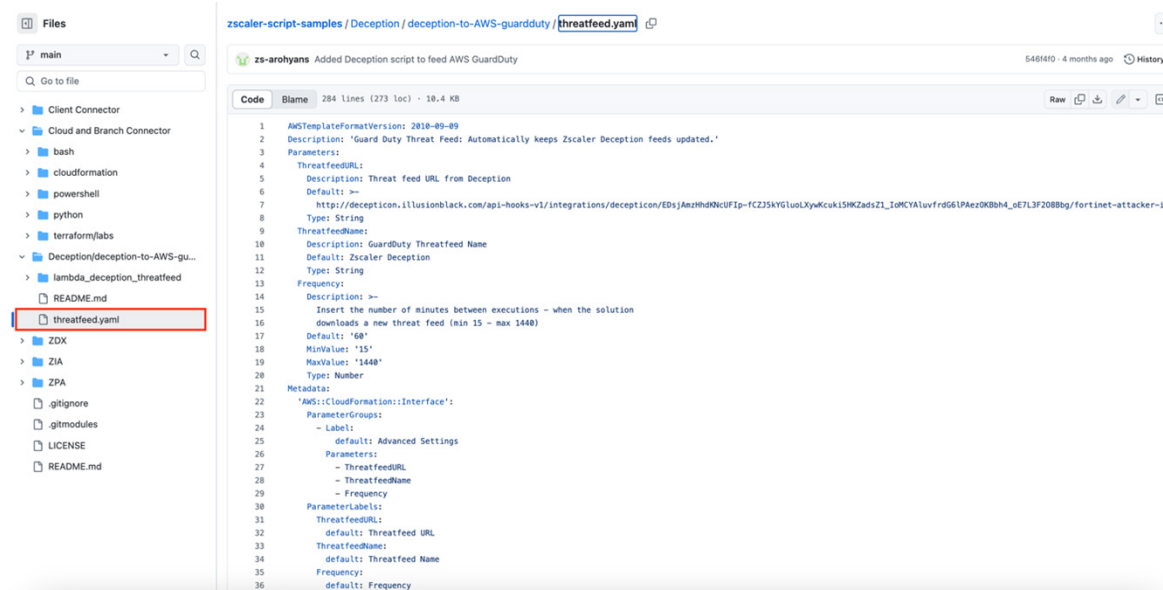


Figure 125. Zscaler Script Samples

2. Go to your AWS account and log in.
3. Go to the **AWS CloudFormation** dashboard.
4. From the **AWS CloudFormation** dashboard, create a new stack with new resources.
5. Choose the option to upload a template file and select the threatfeed.yaml file previously downloaded.
6. Click **Next**.
7. Enter a **Stack name** for the stack.
8. Enter the **Threatfeed URL** obtained from Deception (in [Configuring Deception](#)).

Specify stack details

Provide a stack name

Stack name

zscaler-guardduty-zia

Stack name must be 1 to 128 characters, start with a letter, and only contain alphanumeric characters. Character count: 21/128.

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Lambda Schedule Interval (in minutes)

The interval (in minutes) to trigger the threatfeed function (minimum 1).

5

S3 Bucket Name

The S3 bucket where the ZIP file is stored.

zscaler_threatfeed_files

S3 Key

The path (key) to the ZIP file in the S3 bucket.

lambda_guardduty_threatfeed.zip

Zscaler Secret Name

The name of the AWS Secrets Manager secret that contains Zscaler credentials. The script will use these credentials to create new FQDN/IP Destination Groups as well as Firewall policy

Zscaler_Credentials

Cancel Previous Next

Figure 126. Specify stack details

9. Enter the **Threatfeed Name** (which is used to name the **CloudWatch Rule** and **Lambda Function**).
10. Set the polling **Frequency** (how often Lambda checks Deception for new findings) in the Frequency field.
11. Click **Next**.
12. Proceed through the remainder of the wizard and update remaining default values as you see fit for your environment.

Verification

You can verify the integration in several ways, as per the following.

Directly from CloudWatch logs:

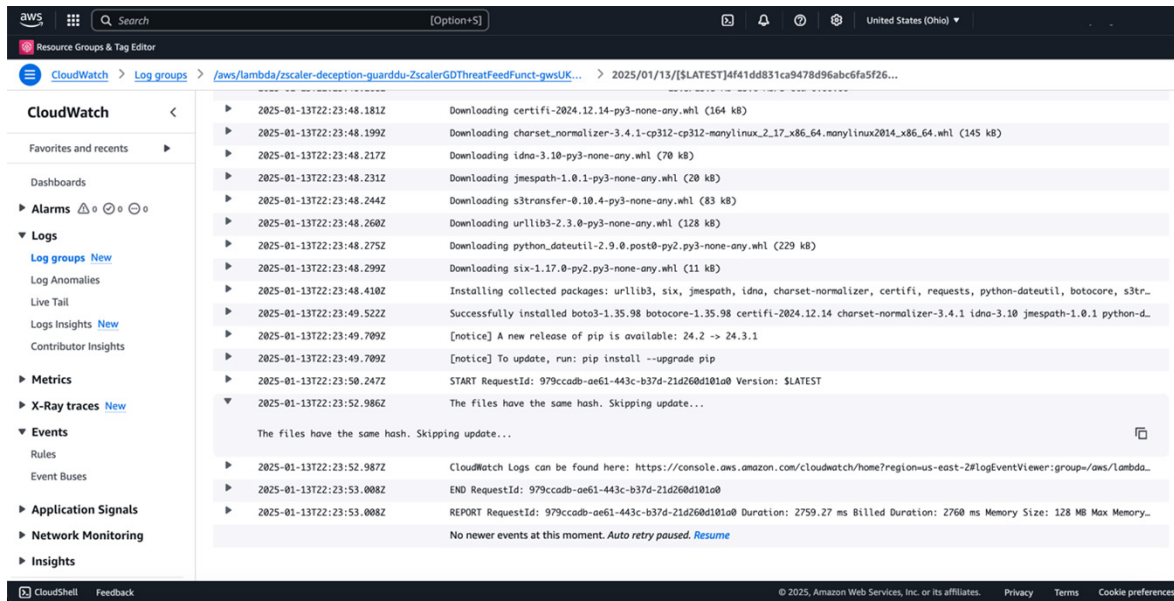


Figure 127. CloudWatch logs

From the **Lambda Function** (use the Test tab to verify that you can manually execute the script). The script does not require any Event parameters to run:

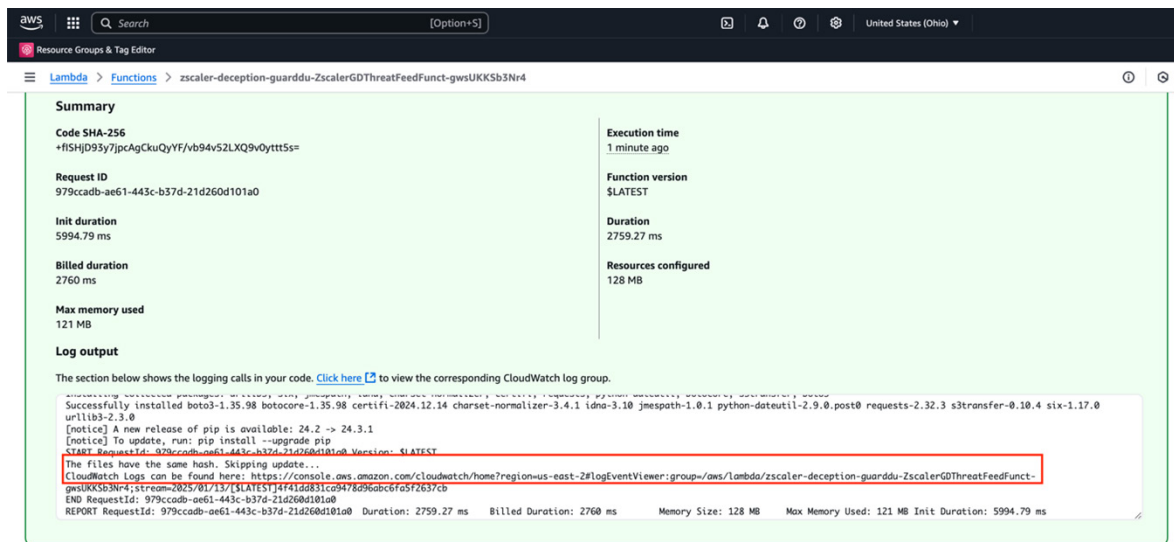


Figure 128. Lambda Function

From the **GuardDuty Overview** page (verify that new Findings appear within the dashboard):

The screenshot shows the AWS GuardDuty console. The left sidebar contains the navigation menu. The main area displays a list of findings. The right sidebar shows the details for a specific finding.

Findings (207)

Findings list (partial):

- The API GetTemplateSummary was invoked from an IP address on a custom threat list.
- The API PutObject was invoked from an IP address on a custom threat list.
- The reconnaissance API ListStacks was invoked from an IP address on a custom threat list.
- The API FilterLogEvents was invoked from an IP address on a custom threat list.
- The reconnaissance API DescribeMetricFilters was invoked from an IP address on a custom threat list.
- The reconnaissance API ListNotificationHubs was invoked from an IP address on a custom threat list.
- The API GetFunction20150331v2 was invoked from an IP address on a custom threat list.
- The reconnaissance API ListLayers20181031 was invoked from an IP address on a custom threat list.

The API PutObject was invoked from an IP address on a custom threat list.

High First seen a month ago, last seen a month ago

An API was used to access a bucket from an IP address on a custom threat list.

[Investigate with Detective](#)

This finding is

Overview

Finding ID	14c9e972a788d74f8281381484cdd64c
Type	UnauthorizedAccess:S3/MaliciousIPCaller.Cust
Severity	HIGH
Region	us-east-2
Count	1
Account ID	084828604313
Resource ID	cf-templates-1hzm754gwbre-us-east-2
Created at	12-16-2024 22:34:23 (a month ago)
Updated at	12-16-2024 22:34:23 (a month ago)

Resource affected

Figure 129. GuardDuty Overview page

From the S3 Bucket (verify that *.txt files are placed in this bucket by the Lambda function):

The screenshot shows the AWS S3 console. The left sidebar contains the navigation menu. The main area displays the details for a specific S3 bucket.

Amazon S3

zscaler-deception-guarddu-zscalergdthreathfeedoutpu-4r0ikmfjcbw

Objects (1)

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

Find objects by prefix

Name	Type	Last modified	Size	Storage class
dec-ip-list.txt	txt	December 20, 2024, 11:14:36 (UTC-06:00)	1.8 KB	Standard

Figure 130. S3 Bucket

Leveraging AWS System and User-Defined Tags for Policy

The Zscaler Workload Discovery (Tagging) feature enables organizations to apply granular security policies to groups of cloud workloads based on metadata. This metadata includes user-defined tags created by the enterprise or system-defined tags provided by AWS. By leveraging these tags, customers can create dynamic policies for traffic inspection, access control, and data protection across their cloud environments.

Requirements

Make sure the following requirements are met:

- The use of Zscaler Cloud Connectors.
- Activation of the Zscaler Workload Discovery feature via a Zscaler Support case.
- EC2 read permissions to allow Zscaler to retrieve resource tags and attributes from the AWS environment.

The workload discovery service, delivered as a service by Zscaler, continuously polls the AWS environment (about every 5 minutes) to detect new or changed workloads and updates tag information accordingly. This data flows through the Cloud Connectors to the Zscaler ZTE, allowing administrators to dynamically apply policies based on real-time workload attributes.

In addition to tags, administrators can use more static attributes (like VPC ID, Subnet ID, or ENI) to define policies. This is particularly useful when there's concern about privileged users modifying tags, as attributes remain stable over time.

Additionally, to handle environments with overlapping IP addresses, Zscaler leverages AWS Namespaces. A namespace is a unique identifier applied at the VPC level, allowing Zscaler to differentiate between conflicting IP ranges and ensure policies are applied correctly. These namespaces are visible in the Zscaler Cloud & Branch Connector Admin Portal, helping administrators manage complex, multi-VPC environments effectively.

Configuring the Partner Integration

1. On the Cloud & Branch Connector Admin Portal, go to **Administration > Partner Integration**.
2. Click **Add AWS Account**.
3. Provide the following information:
 - a. **Name**: Provide a name.
 - b. **AWS Account ID**: Provide your AWS Account ID.
 - c. **AWS Role Name**: Provide the name you want to assign to the IAM Role created with this service.
 - d. **Account Group**: Assign the account to a group (optional).

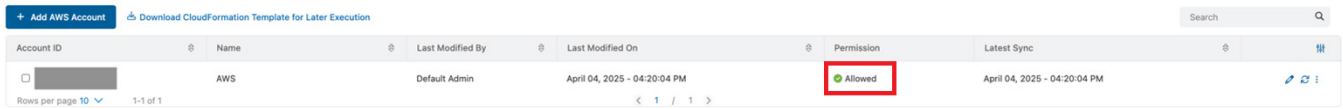
The screenshot shows a form titled 'Add AWS Account' with the following fields and values:

- Name**: AWS
- AWS Account ID**: 084828601234
- AWS Role Name**: ZscalerTagRole
- Account Group (optional)**: Select (dropdown menu)
- External ID**: 2feb42b91269f3c224fdf195013dc952
- Event Bus Name**: [Redacted] Copy
- Trusted Account ID**: [Redacted] Copy
- Trusted Role**: [Redacted] Copy

Figure 131. Add AWS account

4. Make note of the **External ID**.
5. Click **Next**.
6. Select a **Region**.
7. Click **Next**.
8. Click **Save** and **Next**.
9. Hover over the **Activation** menu and click **Activate**.
10. Click the **Launch Cloudformation** template in the AWS Management Console link.
11. From the **AWS Cloudformation** console, select the checkbox to **Acknowledge** the changes.
12. Click **Create Stack**.

13. From the Cloud & Branch Connector Admin Portal, click **Finish**.
14. In the **Partner Integrations** window, verify the integration is successful (Allowed).



Account ID	Name	Last Modified By	Last Modified On	Permission	Latest Sync
[REDACTED]	AWS	Default Admin	April 04, 2025 - 04:20:04 PM	Allowed	April 04, 2025 - 04:20:04 PM

Figure 132. Partner Integrations


15. Click the AWS account created.
16. Click the region.
17. Verify tag values are flowing by clicking one of the IP addresses present in the list.

Workloads (258)

Private IP Address	VPC ID	Namespace	User Defined Tags
10.0.1.4	vpc-[REDACTED]	default	9
10.0.1.5	vpc-[REDACTED]	default	9
10.0.1.6	vpc-[REDACTED]	default	9
10.0.1.7	vpc-[REDACTED]	default	9
10.0.1.8	vpc-[REDACTED]	default	9
10.0.1.9	vpc-[REDACTED]	default	9
10.0.1.10	vpc-[REDACTED]	default	9
10.0.1.11	vpc-[REDACTED]	default	9
10.0.1.12	vpc-[REDACTED]	default	9
10.0.1.13	vpc-[REDACTED]	default	9

Figure 133. Workloads

The following image shows the Attributes.

 10.0.1.4

Attributes

Instance Role

AMI ID

Platform

VM ID

Security Group ID

VPC ID

vpc-

Subnet ID

subnet-

Network Interface ID

eni-

Security Group Name

AllowAll

User Defined Tags (9)

Resource Type	Key	Value
VPC		
ENI		
SUBNET	aws:cloudformation:logical-id	Subnet
VPC	aws:cloudformation:stack-name	zscaler-bedrock-ui
SUBNET	aws:cloudformation:stack-id	arn:aws:cloudformation:us-east-1:08482...
VPC	aws:cloudformation:logical-id	VPC
VPC	aws:cloudformation:stack-id	arn:aws:cloudformation:us-east-1:08482...
SUBNET	Name	
SUBNET	aws:cloudformation:stack-name	

Rows per page 10 1-9 of 9 < 1 / 1 >

Figure 134. Attributes

Configuring Policy

Using AWS resource tags collected through the Workload Discovery/Tagging Service, administrators can build Workload Groups in ZIA. These groups act as logical collections of cloud workloads that share common identifiers—such as environment type.

After Workload Groups are defined, you can use them across ZIA policies such as URL filtering, SSL inspection, firewall rules, and DLP policies. This enables cloud traffic to be protected with the same level of visibility and control as user traffic, without the need for manual IP-based policy maintenance. The process is dynamic: if new instances are spun up or tags change, the Cloud Connector and workload discovery service ensure Workload Groups are automatically updated to reflect these changes in near real time.

1. Go to the ZIA Admin Portal.
2. From the **Administration** menu, click **Workload Groups**.
3. Click **Add Workload Group**.
4. Enter a **Name** and **Description**.
5. In the **Criteria** section, identify the metadata that classifies what resources to include in the group. For example, select a **Tag Type** of **VPC**, then click the **Add (+)** symbol to supply the value. Select **Vpc-id** as the key, then provide your VPC ID.

Name

Prod-Workloads

Description

Criteria

A maximum of 8 tags can be added

Tag Type

VPC

Vpc-id =vpc-03bbc483b5588b6f-X



Add More

Figure 135. Criteria

6. Add a maximum of 8 tags to each Workload Group. Evaluate tags using a Boolean AND/OR expression.
7. As an example, enable a policy to block outbound TCP/22 from the Prod-Workloads group created previously. Go to **Policy > Firewall Control**.
8. Click **Add Firewall Filtering Rule**.
9. Assign an **Order** and enter a **Name** for the rule.
10. From the **Workload Groups** drop-down menu, select **Prod-Workloads**.

Add Firewall Filtering Rule

Rule Order

5

Rule Name

BlockSSH-Prod-Workloads

Rule Status

Enabled

Rule Label

Who, Where, ...
Services
Applications
Source IP
Destination IP

CRITERIA

Users

Any

Groups

Any

Departments

Any

Locations

Any

Location Groups

Any

Time

Always

Devices

Device Groups

Device Trust Level

Workload Groups

Prod-Workloads

ACTION

Network Traffic

Block/Drop

Save

Cancel

Figure 136. Add Firewall Filtering Rule

Enhancing AWS S3 with Zscaler SaaS Security

Zscaler's SaaS security for AWS S3 provides continuous monitoring and policy enforcement for sensitive data stored in S3 buckets, without requiring inline traffic inspection. By integrating directly with AWS through APIs, Zscaler scans S3 buckets for sensitive content, policy violations, and misconfigurations such as public exposure or overly permissive access controls. This enables organizations to detect and respond to risks like data leakage, malware uploads, or shadow IT use of S3—all without interrupting the flow of data or requiring network changes.

The solution leverages Zscaler's advanced Data Loss Prevention (DLP) engine primarily to classify sensitive data at rest. It also supports policy-driven remediation actions, such as alerting, quarantining files, or adjusting access permissions. Because it operates out-of-band, Zscaler can monitor both managed and unmanaged S3 instances, providing full visibility into cloud storage use across the organization.

Initial ZIA Configuration

1. From the ZIA Admin Portal, go to **Administration > SaaS Application Tenants** and select **Add SaaS Application Tenant**.
2. Select the **Amazon S3** tile.

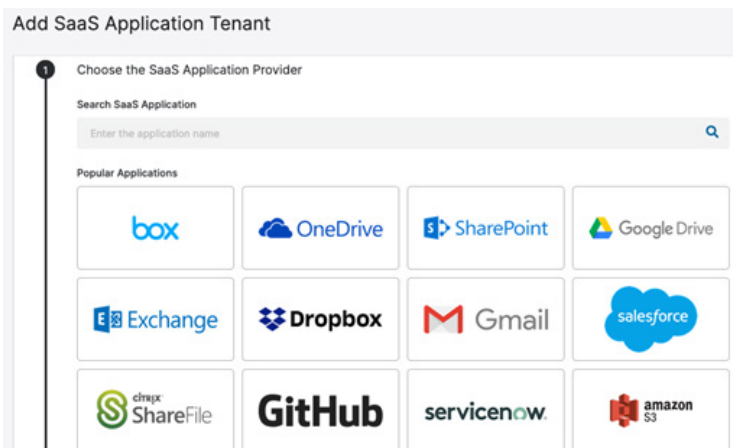


Figure 137. Add SaaS Application Tenant

3. Enter a name to use for this S3 tenant, and then copy the Zscaler Connector Account Number and Zscaler Connector User ARN that are created for later use.

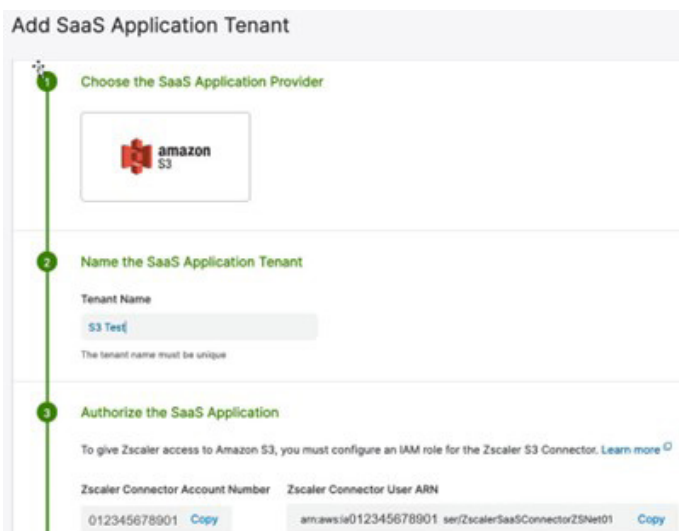


Figure 138. Tenant name

Configure the AWS IAM Role

The next steps are also documented in [Adding SaaS Application Tenants](#) (government agencies, see [Adding SaaS Application Tenants](#)).

1. Log in to the AWS Management Console.
2. Go to **Services** > **IAM**, then click **Access Management** > **Roles** in the left-side navigation.
3. Click **Create Role**.

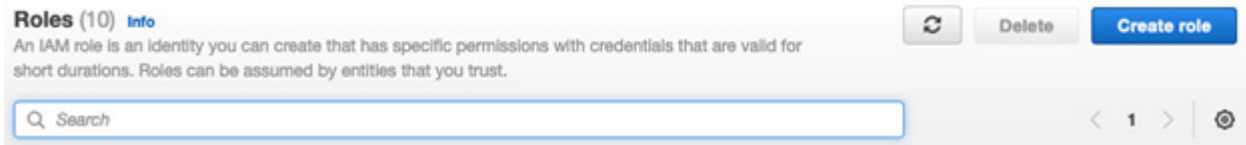
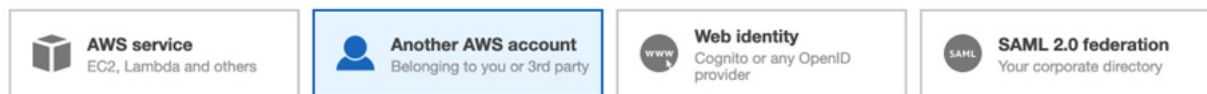


Figure 139. Create Role

4. Click the **Another AWS account** tile as the type of trusted entity.

Select type of trusted entity



Allows entities in other accounts to perform actions in this account. [Learn more](#)

Specify accounts that can use this role

Account ID* ⓘ

Options ☐ Require external ID (Best practice when a third party will assume this role)
☐ Require MFA ⓘ

Figure 140. Trusted entity type

5. Enter the Zscaler Connector Account Number that you copied earlier in the **Account ID** field, and make sure both **Options** are deselected.
6. Click **Next: Permissions** located at the bottom of the screen.

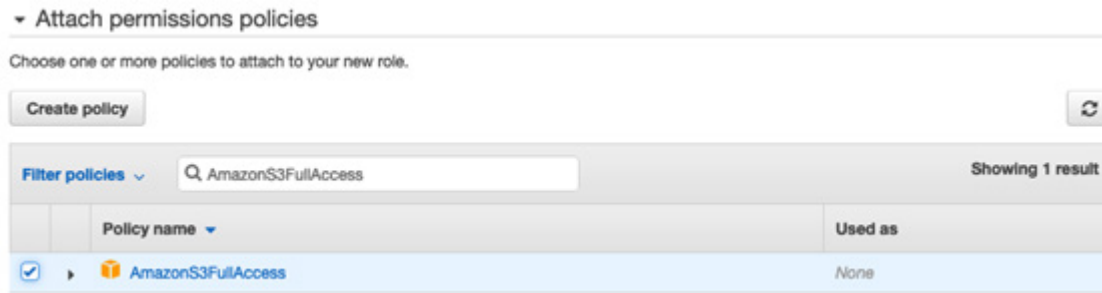
Specify accounts that can use this role

Account ID* ⓘ

Options ☐ Require external ID (Best practice when a third party will assume this role)
☐ Require MFA ⓘ

Figure 141. Which accounts can use the role

- Search for `AmazonS3FullAccess`, and select the policy name when found. Click **Next: Tags** located at the bottom of the window.



Attach permissions policies

Choose one or more policies to attach to your new role.

Create policy ↻

Filter policies Showing 1 result

	Policy name	Used as
<input checked="" type="checkbox"/>	AmazonS3FullAccess	None

Figure 142. Attach permissions policies

- Add tags if needed, and then click **Next: Review** located at the bottom of the window.

Add tags (optional)

IAM tags are key-value pairs you can add to your role. Tags can include user information, such as an email address, or can be descriptive, such as a job title. You can use the tags to organize, track, or control access for this role. [Learn more](#)

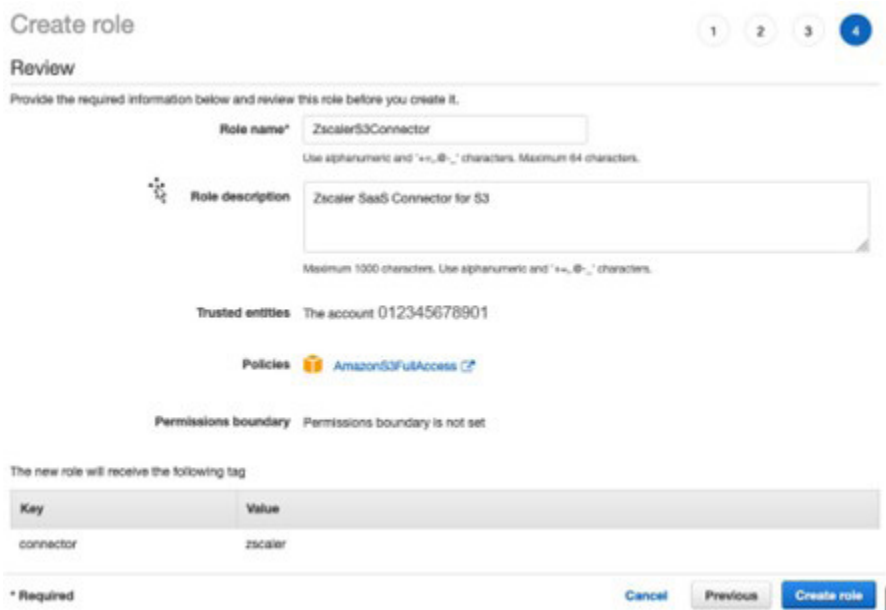


Key	Value (optional)	Remove
connector	zscaler	✕
Add new key		

You can add 49 more tags.

Figure 143. Add tags

- Enter a **Role name** to use for this role.
- (Optional) Provide a description.
- Click **Create role** at the bottom.



Create role 1 2 3 4

Review

Provide the required information below and review this role before you create it.

Role name*
Use alphanumeric and "+, @, _" characters. Maximum 64 characters.

Role description
Maximum 1000 characters. Use alphanumeric and "+, @, _" characters.

Trusted entities The account 012345678901

Policies AmazonS3FullAccess [↗](#)

Permissions boundary Permissions boundary is not set

The new role will receive the following tag

Key	Value
connector	zscaler

* Required Cancel Previous Create role

Figure 144. Create role

Configure the AWS Trust Relationship

The next steps are also documented in the [Adding SaaS Application Tenants](#) (government agencies, see [Adding SaaS Application Tenants](#)).

1. Search for the newly created role by entering `Zscaler`, and selecting the role name when found.

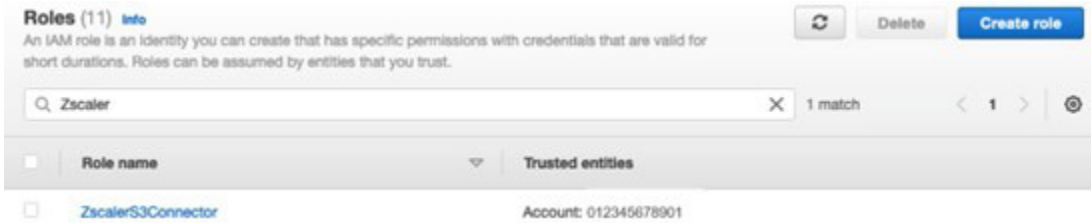


Figure 145. Roles

2. Click the **Trust relationships** tab, and then click **Edit trust relationship**.

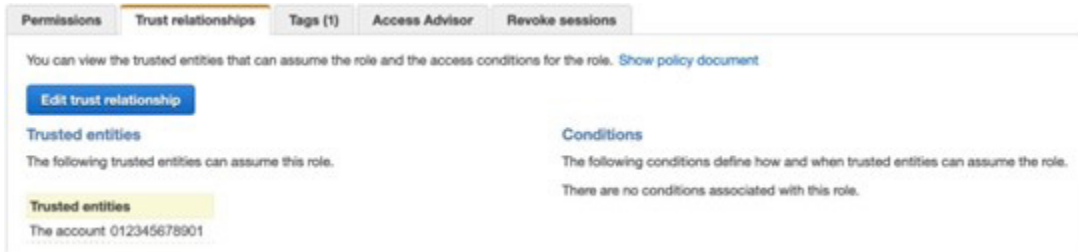


Figure 146. Trust relationships

3. Under **Policy Document**, replace the default AWS value with the Zscaler Connector User ARN that you copied earlier, and click **Update Trust Policy**.



Figure 147. Policy document

4. Under the **Summary**, copy the **Role ARN** for later use (as the IAM Role ARN).



Figure 148. Role ARN

Configure AWS CloudTrail

The next steps are also documented in [Adding SaaS Application Tenants](#) (government agencies, see [Adding SaaS Application Tenants](#)) in the Obtain the CloudTrail Bucket ARN section.



The S3 bucket selected for the trail won't be available to scan in the SaaS Security Scan Configuration, as it is marked Internal.

1. Go to **Services > CloudTrail**, and click **Trails** located in the left-side navigation.

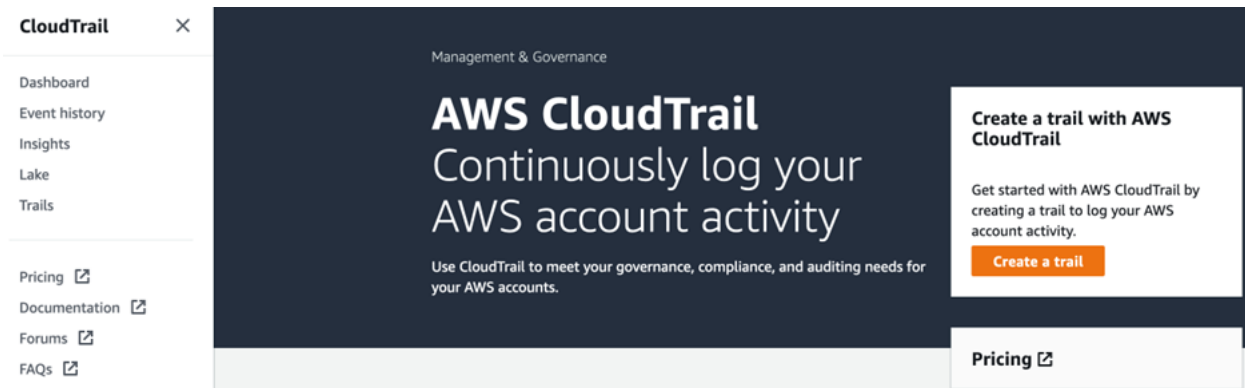


Figure 149. Create a CloudTrail



In the Adding SaaS Application Tenants help, step iii under section c shows an existing trail. If you don't already have a trail, you must create one. See [Appendix B: Creating a Trail](#) on how to create a trail before proceeding.

2. Select the trail name to use from the list.

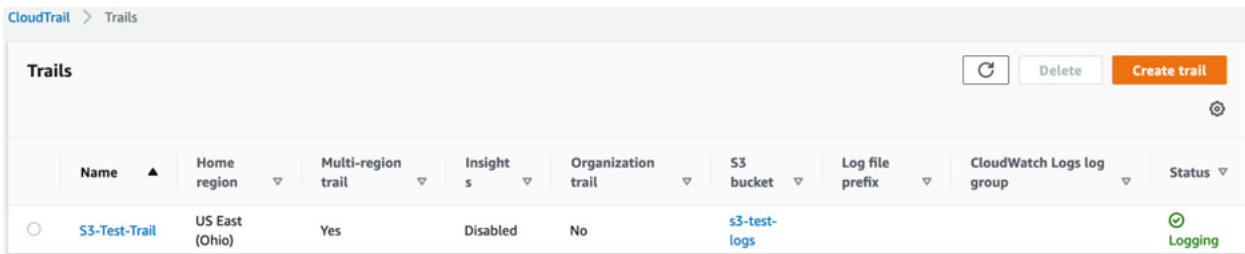


Figure 150. Select CloudTrail name

3. Click the **Traill log location** (in blue) in **General details**.

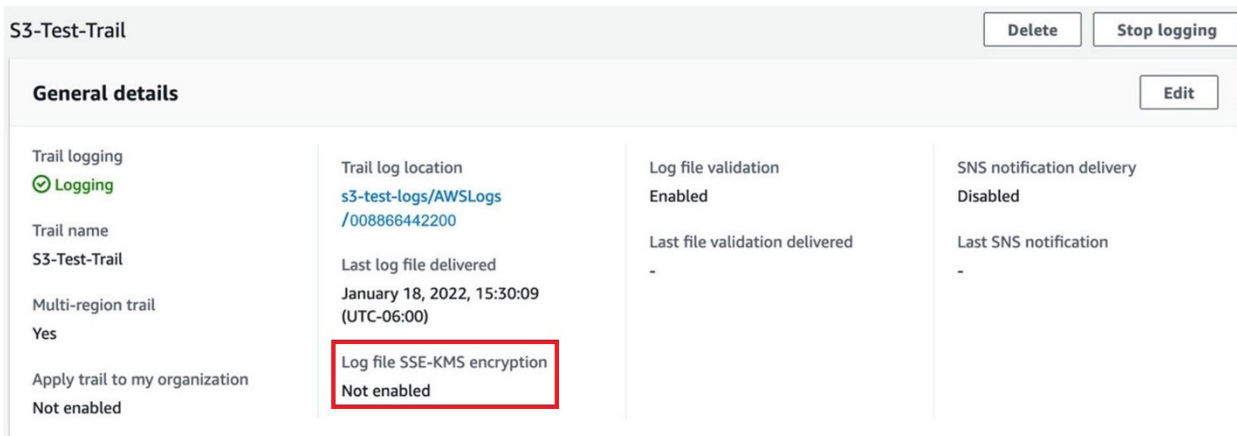


Figure 151. CloudTrail general details

- On the **Objects** tab, click **CloudTrail/**.

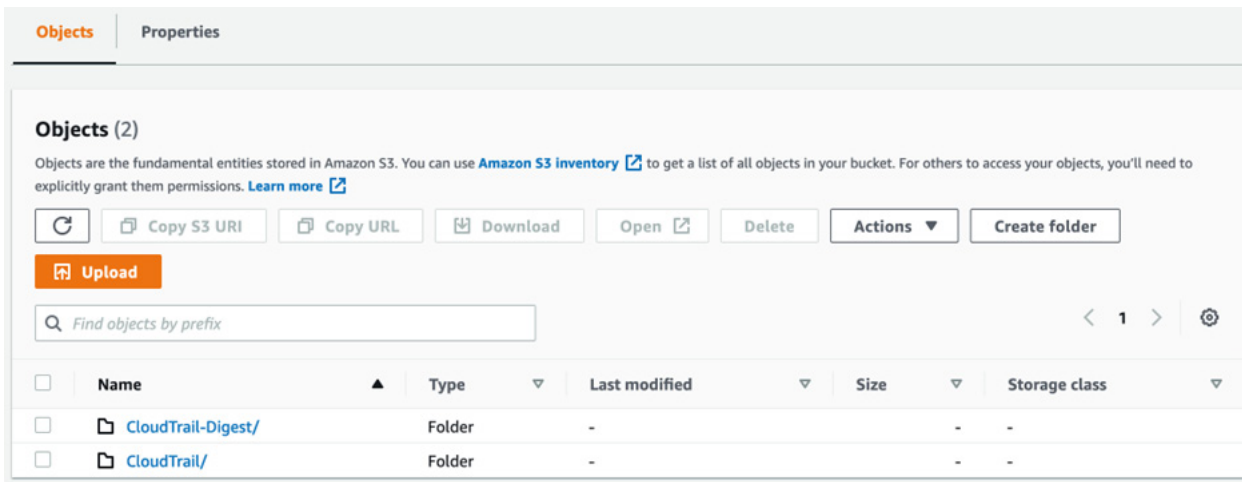


Figure 152. CloudTrail objects

- Click the **Properties** tab, and copy the **Amazon Resource Name (ARN)** to use later (as the CloudTrail Bucket ARN).



Figure 153. CloudTrail properties

Configure the S3 Quarantine Bucket

The next steps are documented in [Adding SaaS Application Tenants](#) (government agencies, see [Adding SaaS Application Tenants](#)) in the Create a Quarantine Bucket section.



If you already have a bucket, you don't need to create one. However, verify that the following settings match step iii of the procedure described in the online documentation. A directory called Zscaler_Quarantine is created in this bucket, but only when malware files are quarantined.

- Block all public access: Select.
- Bucket Versioning: Disable.
- Server-side encryption: Disable.

The S3 bucket selected for use with the quarantined files is not available in the SaaS Security Scan Configuration and is marked Internal.

1. Go to **Services > S3** and click **Buckets** in the left-side navigation.
2. Record the name of the S3 bucket you identified as the Quarantine bucket (either existing or newly created). You must refer to this name later.

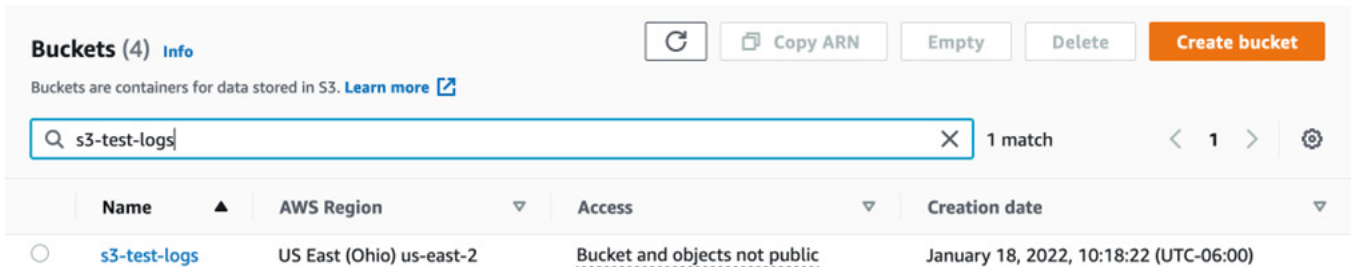


Figure 154. Buckets configuration

Finalize Zscaler Configuration

To complete the Zscaler configuration:

1. From the ZIA Admin Portal, go to **Administration > SaaS Application Tenants**.
2. Click **Add SaaS Application Tenant**.
3. Enter the details in each field.
4. Click **Save**.



- You can find your AWS Account ID in the user details in the upper right-hand corner of the AWS Management Console. You can find information on how to obtain your AWS Account ID in the [AWS documentation](#).
- The Quarantine Bucket Name is the name that you copied in the Quarantine Bucket configuration.
- The IAM Role ARN is the role ARN that you copied earlier during the Trust Relationship configuration.
- The CloudTrail Bucket ARN is the Amazon Resource Name (ARN) that you copied earlier during the CloudTrail configuration.

The screenshot shows the '4 Register the SaaS Application' step in the Zscaler Admin Portal. It prompts the user to enter Zscaler S3 Connector details. The form contains the following fields:

Field	Value
AWS Account ID	0088-6644-2200
IAM Role ARN	iam::008866442200:role/ZscalerS3Connector
Quarantine Bucket Name	s3-test-logs
CloudTrail Bucket ARN	s3-test-logs/AWSLogs/008866442200/CloudTrail/

Figure 155. Register the SaaS Application

- Save and activate so the status is **Validating**.

No.	Application	Tenant Name	Status
1	Amazon S3	S3 Tenant	● Validating

Figure 156. S3 Tenant validating

- After a short period, when access is successful, the status is **Active**. Proceed with [Understanding the Data at Rest Scanning Policy](#) (government agencies, see [Understanding the Data at Rest Scanning Policy](#)).

No.	Application	Tenant Name	Status
1	Amazon S3	S3 Tenant	● Active

Figure 157. S3 Tenant active

Integrating Zscaler Cloud NSS with Amazon S3

This section covers the integration of Zscaler Cloud Nanolog Streaming Service (NSS) with Amazon Web Services' (AWS) Simple Storage Service (Amazon S3). With a Cloud NSS subscription, you can enable direct cloud-to-cloud streaming of ZIA traffic logs into Amazon S3, where log data is stored in containers known as buckets. This integration supports long-term log retention, allows for preprocessing of log data before ingestion, and ensures compatibility with analytics solutions that can easily access and analyze data from S3 buckets. For more information about geo-availability and subscription qualifications for Cloud NSS, contact Zscaler Support.

Create a User Group in AWS IAM

- Log in to the AWS Management Console.
- In the search bar, enter **IAM** and select **IAM**.

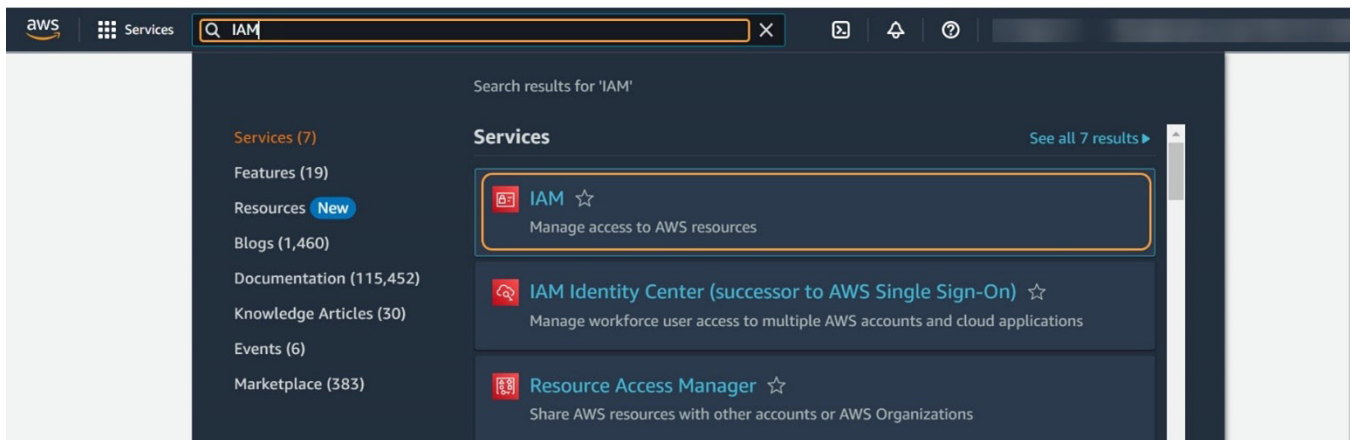


Figure 158. Search for Identity and Access Management (IAM) in AWS Management Console

3. In the left-side navigation, go to **Access management** > **User groups**.

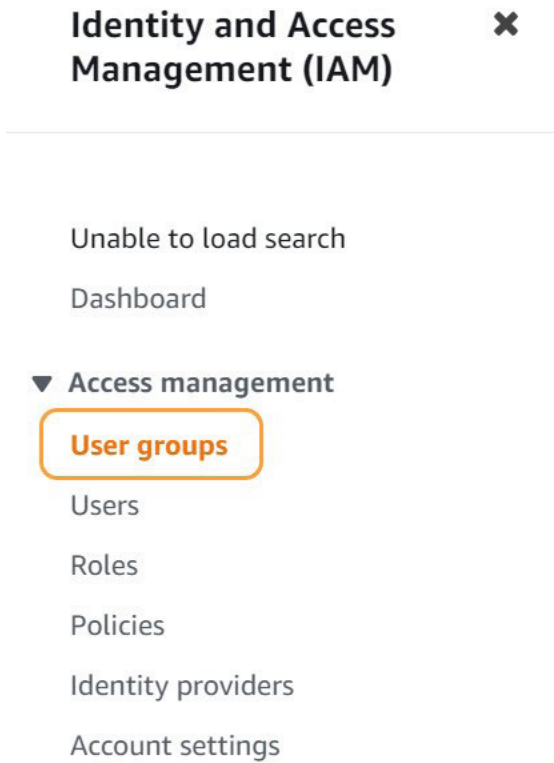


Figure 159. AWS IAM menu with User groups selected

4. Click **Create group**. The **Create user group** page appears.

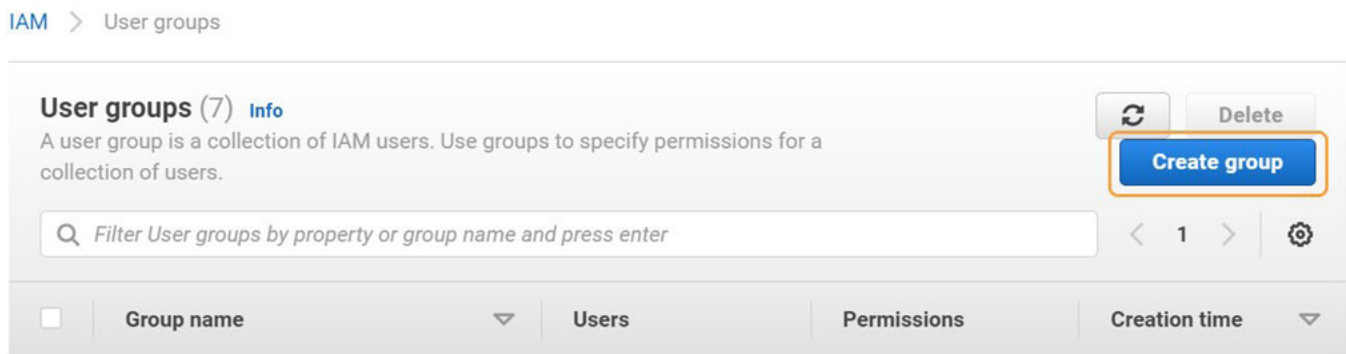


Figure 160. User groups page in AWS IAM with Create group button selected

- On the **Create user group** page, create a user group. Enter a name for the user group (e.g., `Zscaler_Group_Test`).

[IAM](#) > [User groups](#) > Create user group

Create user group

Name the group

User group name
Enter a meaningful name to identify this group.

Maximum 128 characters. Use alphanumeric and '+=,@>' characters.

Figure 161. Create user group wizard in AWS IAM showing Name the group field

- Skip the options to add users and attach permissions policies.
- Click **Create group**. You are redirected to the **User groups** page and a success message appears.

✓ Zscaler_Group_Test user group created.

[View group](#)



[IAM](#) > [User groups](#)

Figure 162. Success message in AWS IAM after a user group was created

Create a User and Access Key in AWS IAM

- In the left-side navigation of IAM, go to **Access management** > **Users**.

Identity and Access Management (IAM)



Unable to load search

[Dashboard](#)

▼ Access management

[User groups](#)

[Users](#)

[Roles](#)

[Policies](#)

[Identity providers](#)

[Account settings](#)

Figure 163. AWS IAM menu with Users selected

- Click **Add users**. The **Create user** wizard appears.

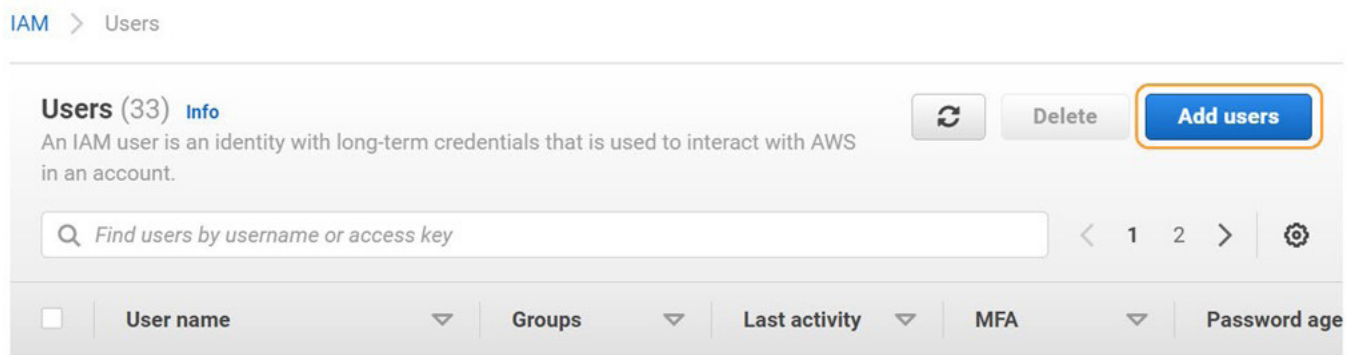


Figure 164. Users page in AWS IAM with Add users button selected

- In the **Create user** wizard, create a user. Enter a user name (e.g., `Zscaler_User_Test`), then click **Next**.

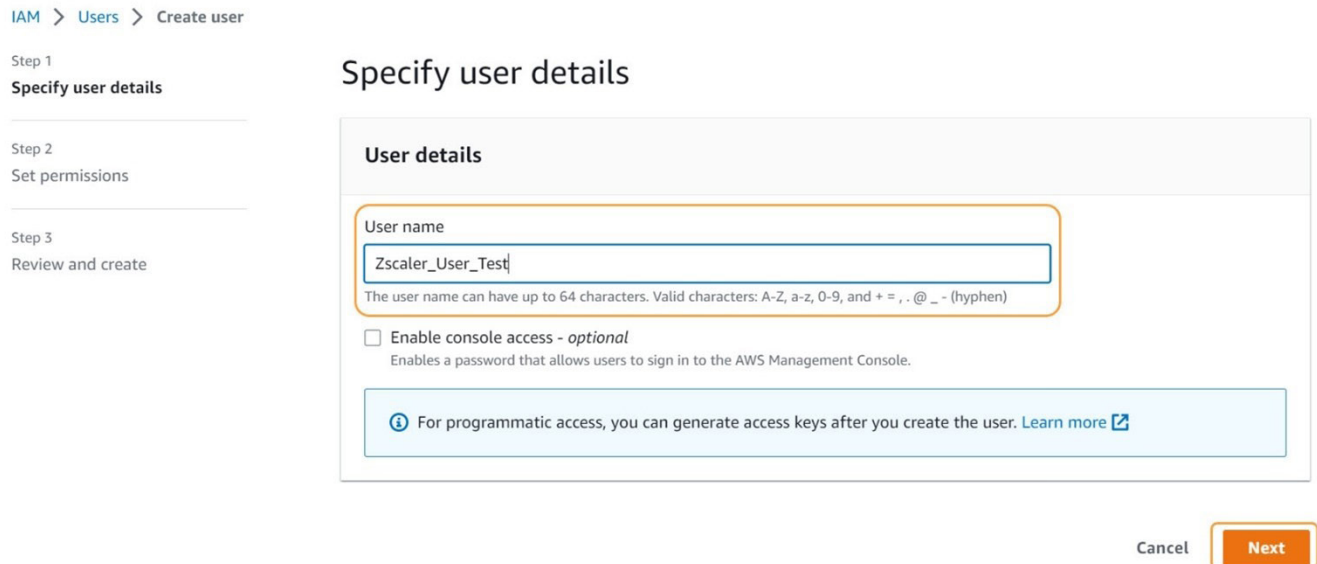


Figure 165. Set user details field in AWS IAM

4. Add the user to the newly created user group (e.g., `Zscaler_Group_Test`), then click **Next**.

IAM > Users > Create user

Step 1
Specify user details

Step 2
Set permissions

Step 3
Review and create

Set permissions

Add user to an existing group or create a new one. Using groups is a best-practice way to manage user's permissions by job functions. [Learn more](#)

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.

User groups (1/10) Refresh Create group

Search: Zscaler_Group X 1 match < 1 > Settings

<input checked="" type="checkbox"/>	Group name ?	Users	Attached policies ?	Created
<input checked="" type="checkbox"/>	Zscaler_Group_Test	0	None	2022-12-13 (1 mont...

► **Permissions boundary - optional**
Set a permissions boundary to control the maximum permissions for this user. Use this advanced feature used to delegate permission management to others. [Learn more](#)

Cancel Previous **Next**

Figure 166. Adding a user to a user group in AWS IAM

5. Review your choices, then click **Create user**.

Permissions summary

< 1 >

Name ?	Type	Used as
Zscaler_Group_Test	Group	Permissions group

Tags - optional

Tags are key-value pairs you can add to AWS resources to help identify, organize, or search for resources. Choose any tags you want to associate with this user.

No tags associated with the resource.

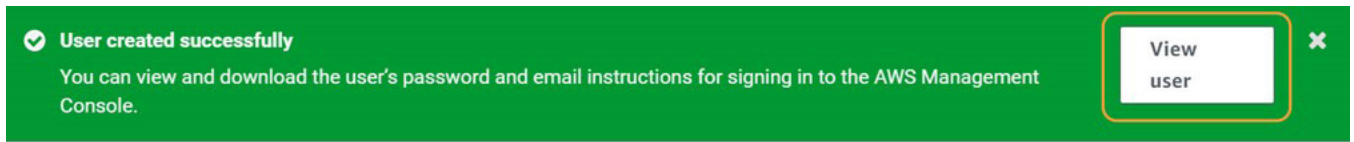
Add new tag

You can add up to 50 more tags.

Cancel Previous **Create user**

Figure 167. Create user wizard in AWS IAM

You are redirected to the Users page and a success message appears.



IAM > Users

Figure 168. Success message in AWS IAM after a user was created

- Click **View user** in the success message, or use the search bar to find the user by name, then select the new user.

IAM > Users

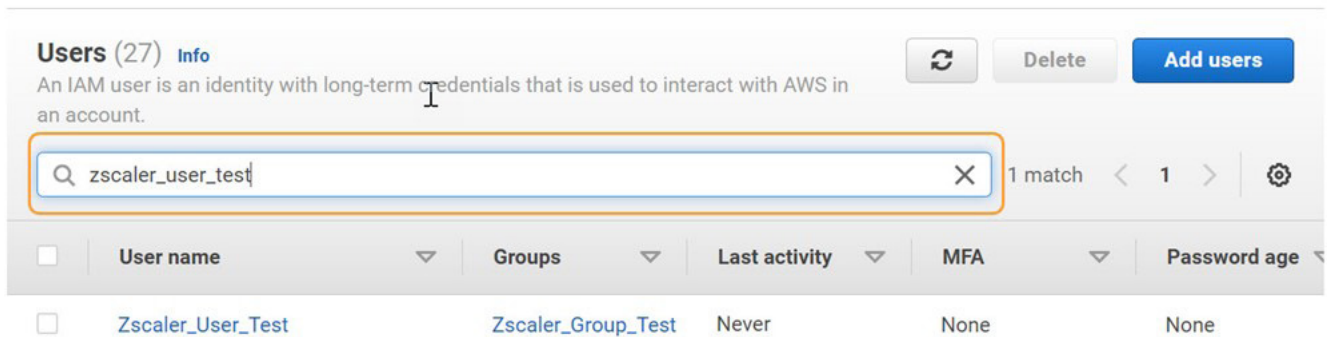


Figure 169. Users page in AWS IAM with search bar

- On the **Summary** page for the newly created user, scroll down and click the **Security credentials** tab.



Figure 170. Security credentials tab in user Summary page in AWS IAM

- On the **Security credentials** tab, scroll down to the **Access keys** section and click **Create access key**. The **Create access key** wizard appears.

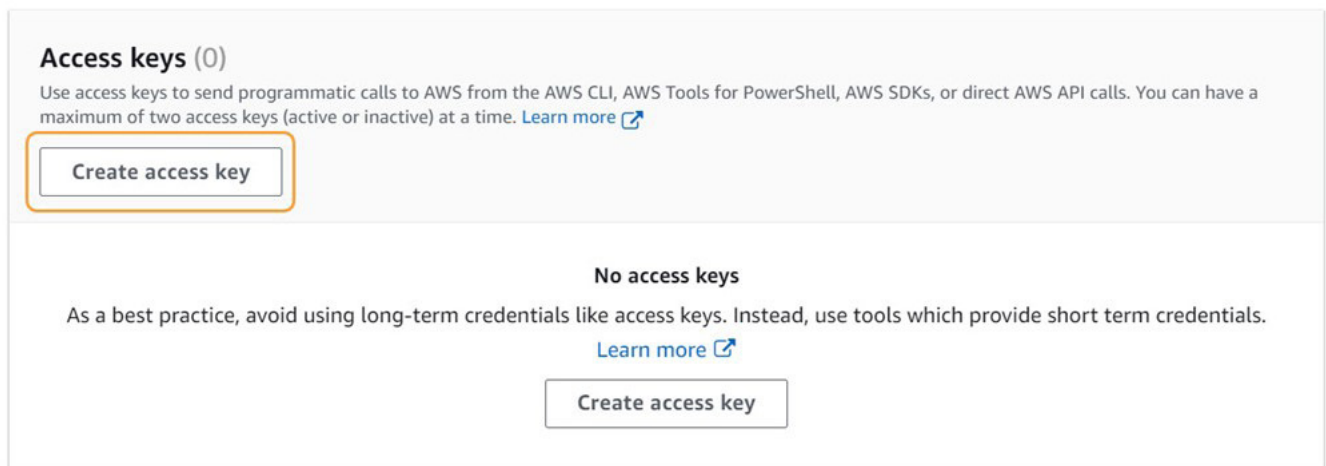


Figure 171. Access keys section in the Security credentials tab on the user Summary page in AWS IAM

9. In the **Create access key** wizard, create an access key. Select a use case, then click **Next**.

Step 1

Access key best practices & alternatives

Step 2 - optional

Set description tag

Step 3

Retrieve access keys

Access key best practices & alternatives

Avoid using long-term credentials like access keys to improve your security. Consider the following use cases and alternatives.

☐ **Command Line Interface (CLI)**
You plan to use this access key to enable the AWS CLI to access your AWS account.

☐ **Local code**
You plan to use this access key to enable application code in a local development environment to access your AWS account.

☐ **Application running on an AWS compute service**
You plan to use this access key to enable application code running on an AWS compute service like Amazon EC2, Amazon ECS, or AWS Lambda to access your AWS account.

☐ **Third-party service**
You plan to use this access key to enable access for a third-party application or service that monitors or manages your AWS resources.

☐ **Application running outside AWS**
You plan to use this access key to enable an application running on an on-premises host, or to use a local AWS client or third-party AWS plugin.

☒ **Other**
Your use case is not listed here.

It's okay to use an access key for this use case, but follow the best practices:

- Never store your access key in plain text, in a code repository, or in code.
- Disable or delete access keys when no longer needed.
- Enable least-privilege permissions.
- Rotate access keys regularly.

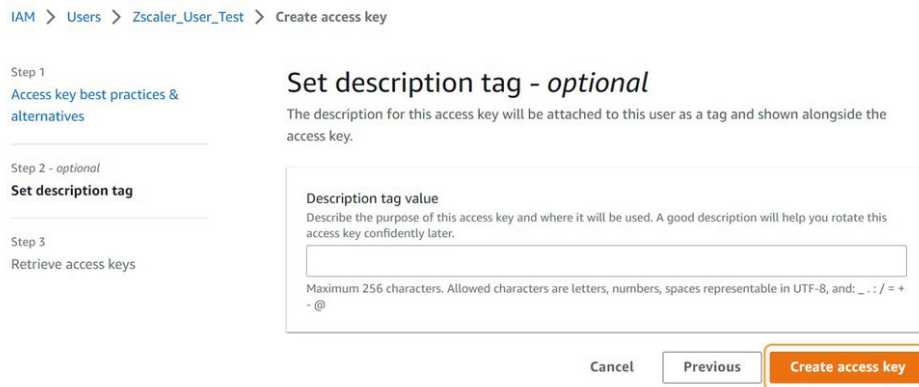
For more details about managing access keys, see the [Best practices for managing AWS access keys](#).

Cancel
 Next

Figure 172. Create access key wizard in AWS IAM



This is the only time that you can view or download the secret access key.

10. Click **Create access key**.


IAM > Users > Zscaler_User_Test > Create access key

Step 1
Access key best practices & alternatives

Step 2 - optional
Set description tag

Step 3
Retrieve access keys

Set description tag - optional

The description for this access key will be attached to this user as a tag and shown alongside the access key.

Description tag value
Describe the purpose of this access key and where it will be used. A good description will help you rotate this access key confidently later.

Maximum 256 characters. Allowed characters are letters, numbers, spaces representable in UTF-8, and: _ . : / = + - @

Cancel Previous **Create access key**

Figure 173. Create access key wizard in AWS IAM with Create access key button selected

The following image shows the success message in AWS IAM.

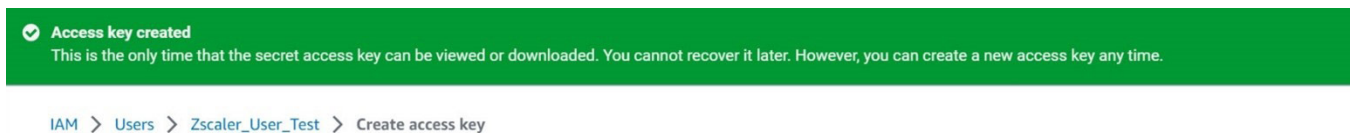
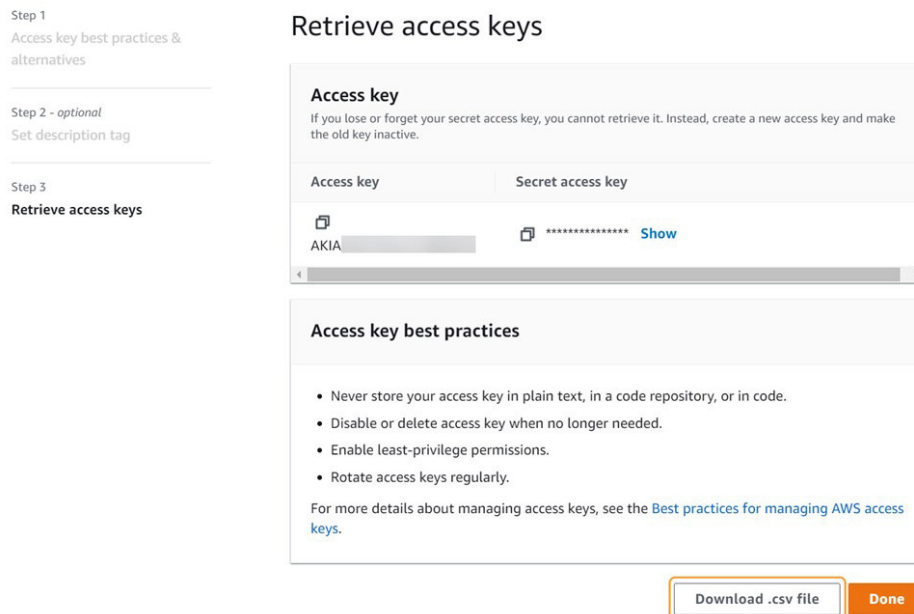


Figure 174. Success message in AWS IAM after an access key was created

11. Click **Download .csv file** to download and save a CSV file containing the access key ID and secret access key required for creating a Cloud NSS feed in the ZIA Admin Portal.


Step 1
Access key best practices & alternatives

Step 2 - optional
Set description tag

Step 3
Retrieve access keys

Retrieve access keys

Access key
If you lose or forget your secret access key, you cannot retrieve it. Instead, create a new access key and make the old key inactive.

Access key	Secret access key
AKIA [redacted]	[redacted] Show

Access key best practices

- Never store your access key in plain text, in a code repository, or in code.
- Disable or delete access key when no longer needed.
- Enable least-privilege permissions.
- Rotate access keys regularly.

For more details about managing access keys, see the [Best practices for managing AWS access keys](#).

Download .csv file Done

Figure 175. Download .csv file button in Create user wizard in AWS IAM

12. Click **Done**.

Create an S3 Bucket and Folder in S3

1. In the search bar at the top of the screen, enter S3 and select **S3**.

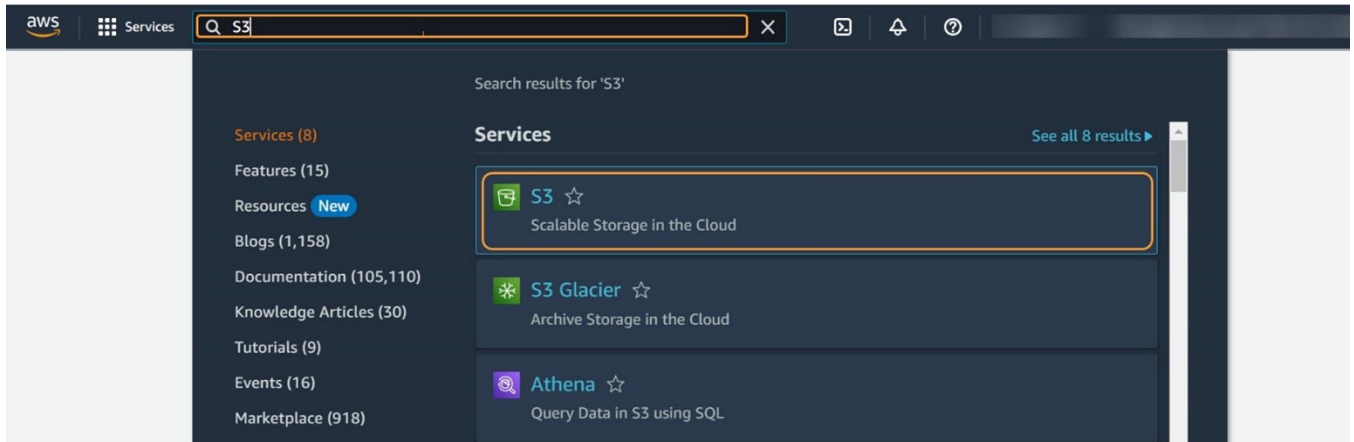


Figure 176. Search for S3 in AWS Management Console

2. In the left-side navigation, go to **Buckets**.



Figure 177. Amazon S3 menu with Buckets selected

3. Click **Create bucket**. The **Create bucket** page appears.

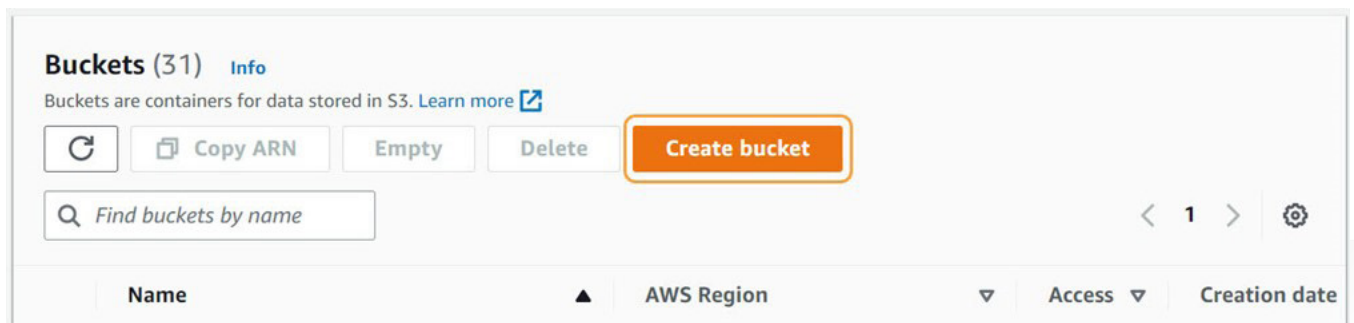


Figure 178. Buckets page in Amazon S3 with Create bucket button selected

- On the **Create bucket** page, create a bucket. Enter a name for the bucket (e.g., `zscaler-bucket-test`). The bucket name is part of its Amazon Resource Name (ARN), which is required for creating a policy in AWS.

Figure 179. Create bucket wizard in Amazon S3 with Bucket name field



S3 buckets cannot have an underscore (`_`) due to the Amazon S3 bucket naming convention. Use a hyphen (`-`) if you want a separation in the characters for the S3 name.

- Select your **AWS Region**. The region is part of the URL required for creating a Cloud NSS feed in the ZIA Admin Portal.

Figure 180. Create bucket wizard in Amazon S3 with AWS Region field

- (Optional) Maintain the default configurations for the remaining settings (e.g., Bucket Versioning, Default encryption, etc.).
- Click **Create bucket**.

Figure 181. Create bucket button in Amazon S3

You are redirected to the **Buckets** page and a success message appears.

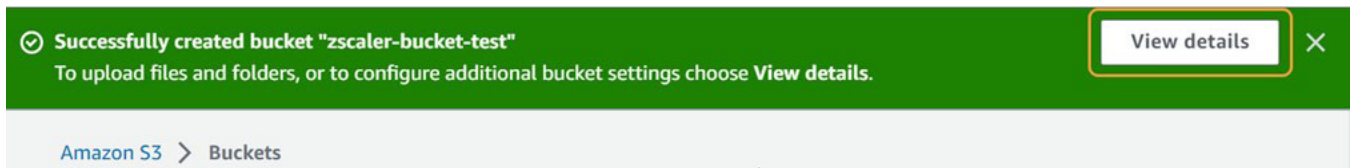


Figure 182. Success message in Amazon S3 after a bucket was created

8. Click **View details** in the success message, or use the search bar to find the bucket by name, then select the new bucket.

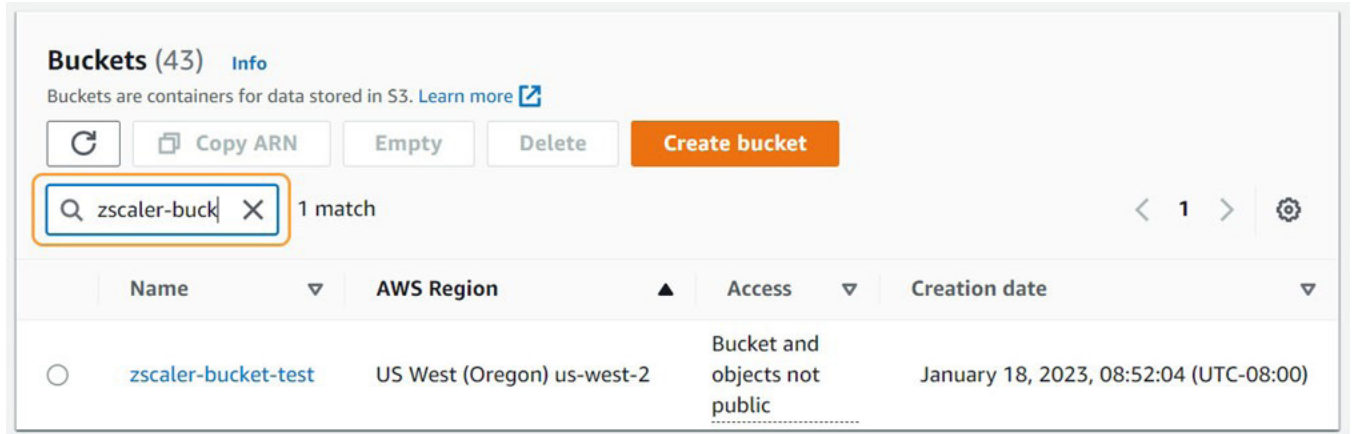


Figure 183. Search for bucket in Amazon S3

9. On the **Objects** tab of the bucket page, click **Create folder**. The **Create folder** page appears.

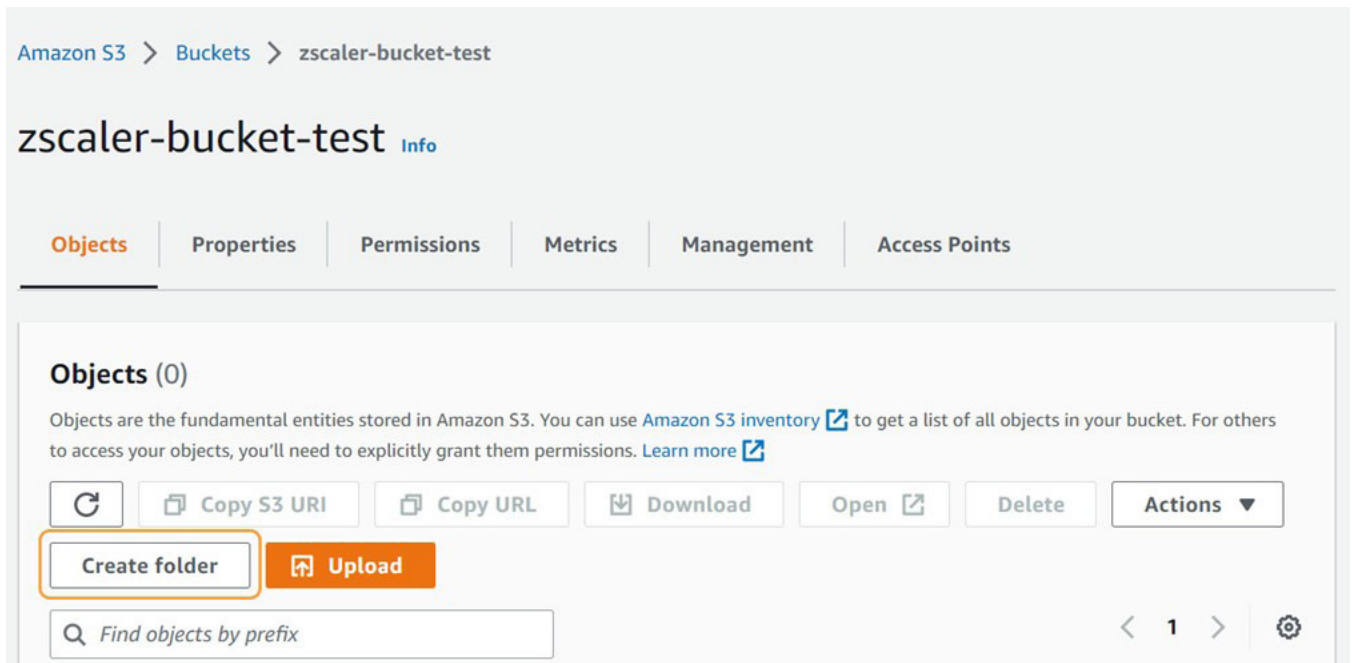


Figure 184. Bucket overview page in Amazon S3 with Create folder button selected

10. On the **Create folder** page, create a folder. Enter a **Folder name** (e.g., logs-test).

Figure 185. Create folder page in Amazon S3 showing the Folder name field

11. Maintain the default Server-side encryption settings and click **Create folder**.

Figure 186. Create folder page in Amazon S3 with Create folder button selected

You are redirected to the **Bucket** page and a success message appears.

Figure 187. Success message in Amazon S3 after a folder was created

12. Select the folder and click **Copy URL**. Save the URL (e.g., `https://zscaler-bucket-test.s3.us-west-2.`

amazonaws.com/logs-test/) required for creating a Cloud NSS feed in the ZIA Admin Portal. The name of your region (e.g., us-west-2) must be present in the URL.

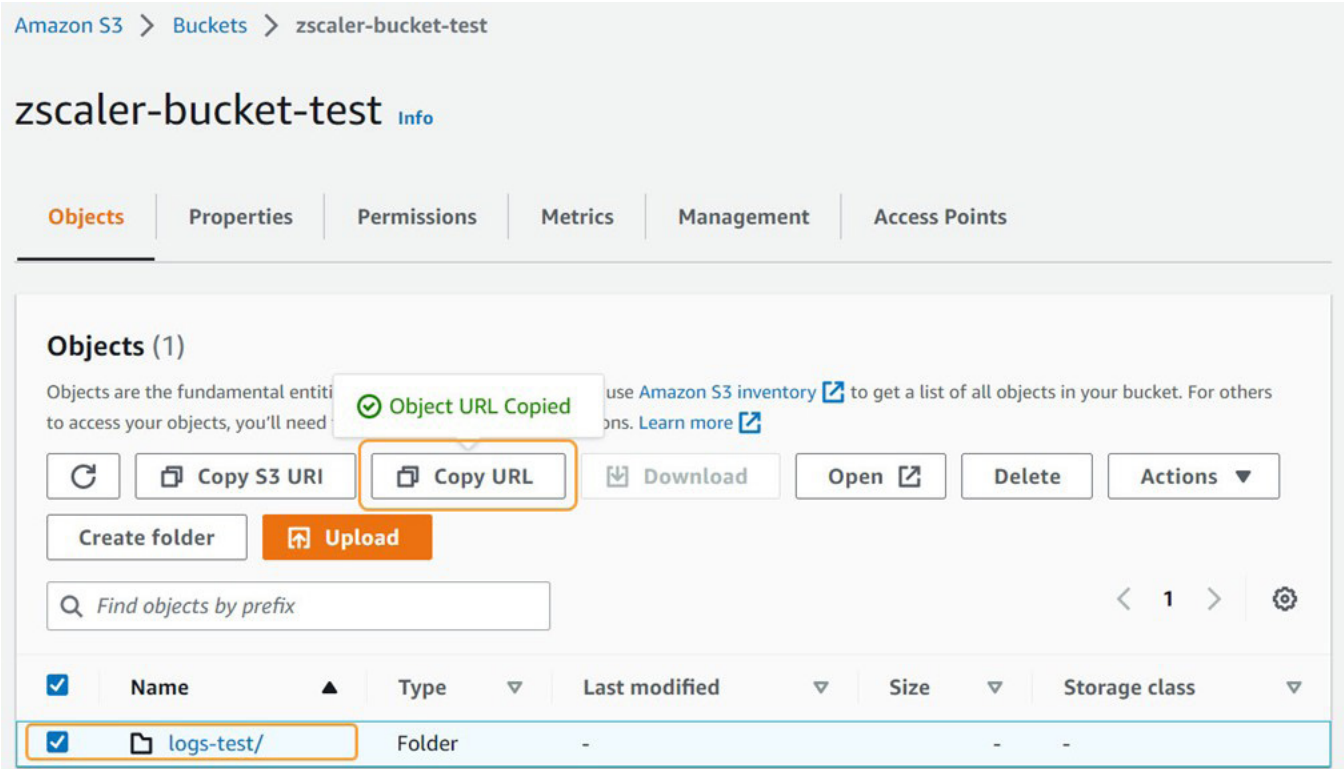


Figure 188. Bucket overview page in Amazon S3 with folder and Copy URL button selected

- 13. Click the **Properties** tab, then copy and save the **ARN** (e.g., `arn:aws:s3:::zscaler-bucket-test`) required for creating a policy in AWS.

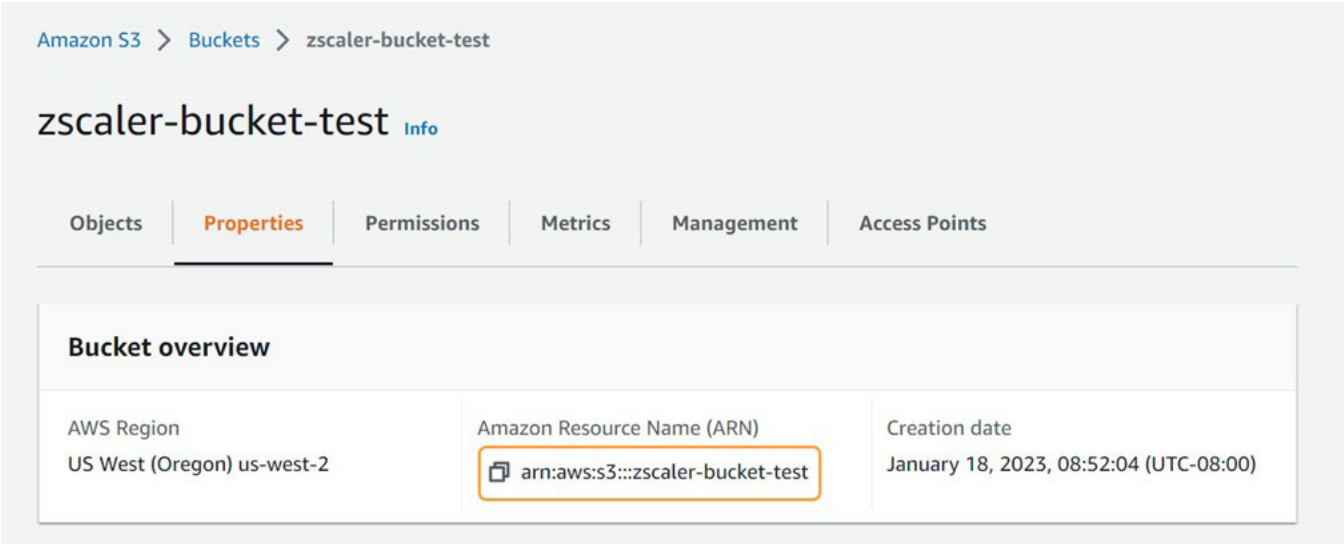


Figure 189. Bucket Properties and ARN in Amazon S3

Create a Policy Granting the User Group Access to the S3 Bucket in Amazon IAM

A policy is a JSON document in AWS that specifies who has access to AWS resources and what actions they can perform on those resources. You can attach a policy to an identity (e.g., user group) or resource (e.g., S3 bucket) to define its permissions. To learn more, refer to the [AWS documentation](#).

To integrate with Cloud NSS, the user group (e.g., Zscaler_Group_Test) needs permission to perform the PutObject action on the S3 bucket (e.g., zscaler-bucket-test). The PutObject action adds an object to a bucket. The user must have Write permissions to perform the PutObject action. To learn more, refer to the [AWS API Reference documentation](#).

To create a policy granting the user group PutObject access to the S3 bucket:

1. Go to the IAM Management Console.
2. In the left-side navigation, go to **Access management** > **Policies**.

Identity and Access Management (IAM)

Unable to load search

Dashboard

▼ Access management

User groups

Users

Roles

Policies

Identity providers

Figure 190. AWS IAM menu with Policies selected

3. Click **Create policy**. The **Create policy** wizard appears.

IAM > Policies

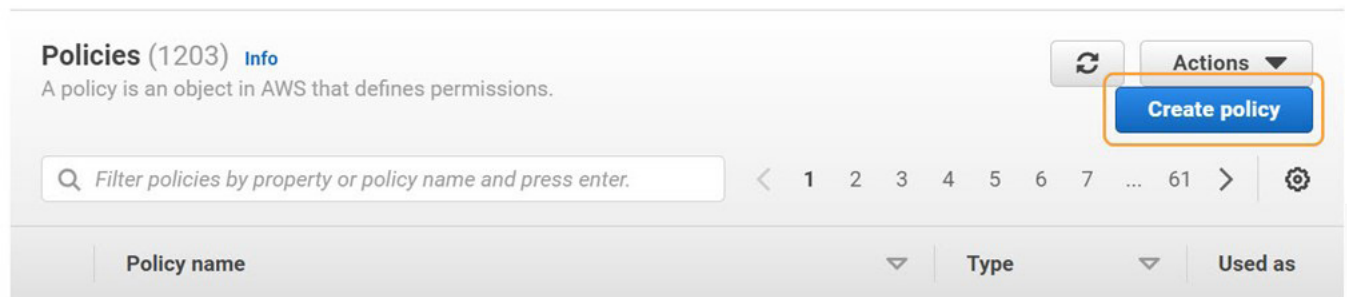


Figure 191. Policies page in AWS IAM with Create policy button selected

4. In the **Create policy** wizard, create a policy.
5. Click the **JSON** tab.

Create policy

1 2 3

A policy defines the AWS permissions that you can assign to a user, group, or role. You can create and edit a policy in the visual editor and using JSON. [Learn more](#)

Visual editor

JSON

[Import managed policy](#)

Figure 192. Create policy wizard in AWS IAM showing JSON tab

6. In the **JSON editor**, write a policy that allows PutObject access to the S3 bucket (e.g., `zscaler-bucket-test`). See the following example.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PutObjectActions",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::zscaler-bucket-test/*"
      ]
    }
  ]
}
```

7. Click **Next: Tags**.
8. Click **Next: Review**.

Create policy

1 2 3

Add tags - optional
Tags are key-value pairs that you can add to AWS resources to help identify, organize, or search for resources.

No tags associated with the resource.

Add tag

You can add up to 50 more tags.

Cancel Previous **Next: Review**

Figure 193. Create policy wizard in AWS IAM showing Tags screen

9. Enter a name for the policy (e.g., `zscaler_policy_test`).

Review policy

Name* zscaler_policy_test

Use alphanumeric and '+,_,@,-' characters. Maximum 128 characters.

Description

Maximum 1000 characters. Use alphanumeric and '+,_,@,-' characters.

Figure 194. Review policy page in AWS IAM showing Name field

10. Review the policy **Summary information** and click **Create policy**.

Summary

Filter

Service	Access level	Resource	Request condition
Allow (1 of 363 services) Show remaining 362			
S3	Limited: Write	BucketName string like zscaler-nss-test, ObjectPath string like All	None

Tags

Key	Value
No tags associated with the resource.	

Cancel Previous **Create policy**

Figure 195. Policy Summary page in AWS IAM with Create policy button selected

You are redirected to the **Policies** page and a success message appears.



IAM > Policies

Figure 196. Success message in AWS IAM after a policy was created

- 11. Attach the policy to the newly created user group. Click the link in the success message, or use the search bar to filter the policies by name, then select the new policy (e.g., `zscaler_policy_test`). The policy **Summary** page appears.

IAM > Policies

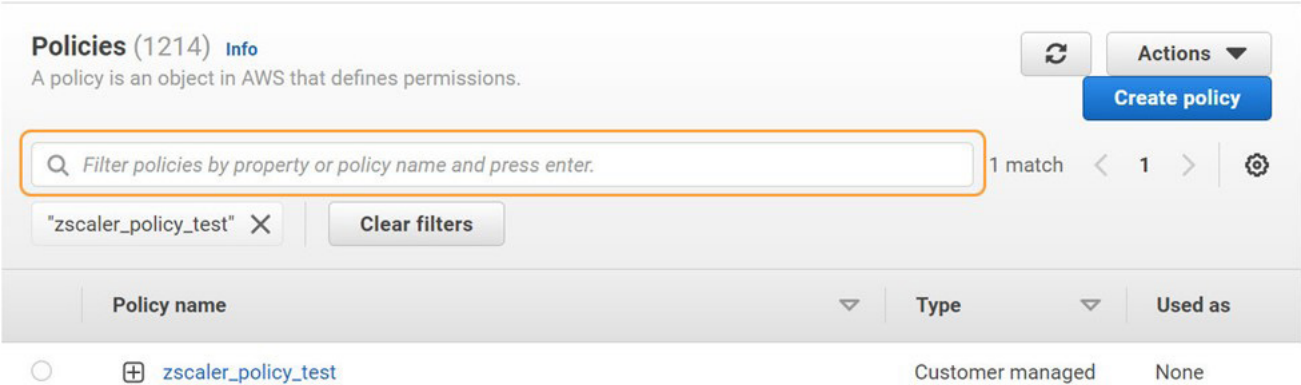


Figure 197. Search for policy in AWS IAM

- 12. On the policy **Summary** page, click the **Policy usage** tab, then click **Attach**. The **Attach policy** page appears.

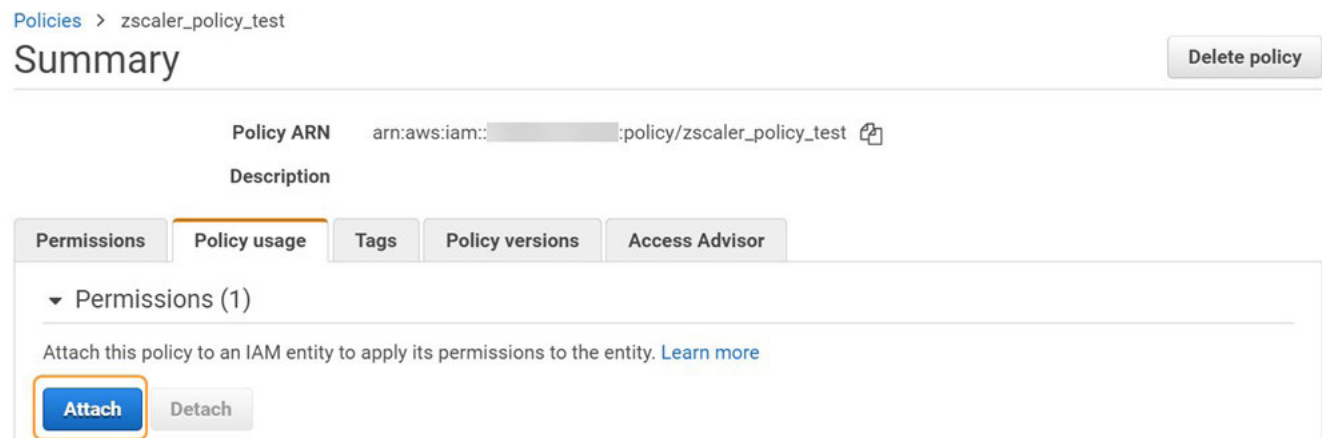


Figure 198. Attach button in Policy usage tab in AWS IAM

13. On the **Attach policy** page, search for and select the newly created user group (e.g., `zscaler_group_test`), then click **Attach policy**.

Attach policy

Attach the policy to users, groups, or roles in your account

Filter: Filter

Showing 1 result

<input checked="" type="checkbox"/>	Name	Type
<input checked="" type="checkbox"/>	Zscaler_Group_Test	Group

Cancel

Attach policy

Figure 199. Attach policy page with user group selected for attachment in AWS IAM

You are redirected to the **Summary** page, which shows the user group (e.g., `zscaler_group_test`) under **Permissions**.

Policies > zscaler_policy_test

Summary

Delete policy

Policy ARN

arn:aws:iam:::policy/zscaler_policy_test

Description

Permissions

Policy usage

Tags

Policy versions

Access Advisor

Permissions (1)

Attach this policy to an IAM entity to apply its permissions to the entity. [Learn more](#)

Attach

Detach

Filter: Filter

Showing 1 result

<input type="checkbox"/>	Name	Type
<input type="checkbox"/>	Zscaler_Group_Test	Group

Figure 200. Policy Summary page in AWS IAM with user group attached

Add a Cloud NSS Feed in the ZIA Admin Portal

See [Adding Cloud NSS Feeds](#) (government agencies, see [Adding Cloud NSS Feeds](#)) and select the type of feed (e.g., Web Logs) that you want to configure.

1. The following fields require specific inputs:
 - a. **SIEM Type:** Select **S3**.
 - b. **AWS Access Id:** Enter the access key ID for the user created in AWS.
 - c. **AWS Secret Key:** Enter the secret access key for the user created in AWS.
 - d. **Max Batch Size:** Enter the recommended maximum batch size based on the log type. For Web and Firewall log types, the recommended maximum batch size is 8 MB. For DNS, Tunnel, and all other log types (e.g., SaaS Security), it is 1 MB.
 - e. **S3 Folder URL:** Enter the URL of the folder created in the S3 bucket (e.g., `https://zscaler-bucket-test.s3.us-west-2.amazonaws.com/logs-test/`).
 - f. **Feed Output Type:** Select **JSON**.
 - g. **Feed Escape Character:** Enter `,` `\` (comma, backslash, quote).
 - h. **Feed Output Format:** Zscaler recommends adding `"time": "%d{epochtime}"` to the Feed Output Format. See the following feed output formats by log type. For [Cloud NSS Feeds for Web Logs](#) (government agencies, see [Cloud NSS Feeds for Web Logs](#)), copy and paste the prepopulated Feed Output Format with the following:

```
{
  "time": "%d{epochtime}",
  "act": "%s{action}",
  "reason": "%s{reason}",
  "app": "%s{proto}",
  "dhost": "%s{ehost}",
  "dst": "%s{sip}",
  "src": "%s{cintip}",
  "sourceTranslatedAddress": "%s{cip}",
  "in": "%d{respsize}",
  "out": "%d{reqsize}",
  "request": "%s{eurl}",
  "requestContext": "%s{ereferer}",
  "outcome": "%s{respcode}",
  "requestClientApplication": "%s{ua}",
  "requestMethod": "%s{reqmethod}",
  "user": "%s{elogin}",
  "spriv": "%s{elocation}",
  "externalId": "%d{recordid}",
  "fileType": "%s{filetype}",
  "destinationServiceName": "%s{appname}",
  "cat": "%s{urlcat}",
  "deviceDirection": "1",
  "cn1": "%d{riskscore}",
  "cn1Label": "riskscore",
  "cs1": "%s{dept}",
  "cs1Label": "dept",
  "cs2": "%s{urlcat}",
  "cs2Label": "urlcat",
  "cs3": "%s{malwareclass}",
  "cs3Label": "malwareclass",
  "cs4": "%s{malwarecat}",
  "cs4Label": "malwarecat",
  "cs5": "%s{threatname}",
  "cs5Label": "threatname",
  "cs6": "%s{band5}",
  "cs6Label": "md5hash",
  "rulelabel": "%s{rulelabel}",
  "ruletype": "%s{ruletype}",
  "urlclass": "%s{urlclass}",
  "DeviceVendor": "Zscaler",
  "DeviceProduct": "NSSWeblog",
  "devicemodel": "%s{devicemodel}"
}
```

For [Cloud NSS Feeds for Firewall Logs](#) (government agencies, see [Cloud NSS Feeds for Firewall Logs](#)), copy and paste the prepopulated Feed Output Format with the following:

```
{
  "datetime": "%s{time}",
  "user": "%s{elogin}",
  "department": "%s{edepartment}",
  "locationname": "%s{elocation}",
  "cdport": "%d{cdport}",
  "csport": "%d{csport}",
  "sdport": "%d{sdport}",
  "ssport": "%d{ssport}",
  "csip": "%s{csip}",
  "cdip": "%s{cdip}",
  "ssip": "%s{ssip}",
  "sdip": "%s{sdip}",
  "tsip": "%s{tsip}",
  "tunsport": "%d{tsport}",
  "tuntype": "%s{ttype}",
  "action": "%s{action}",
  "dnat": "%s{dnat}",
  "stateful": "%s{stateful}",
  "aggregate": "%s{aggregate}",
  "nwsvc": "%s{nwsvc}",
  "nwapp": "%s{nwapp}",
  "proto": "%s{ipproto}",
  "ipcat": "%s{ipcat}",
  "destcountry": "%s{destcountry}",
  "avgduration": "%d{avgduration}",
  "rulelabel": "%s{erulelabel}",
  "inbytes": "%ld{inbytes}",
  "outbytes": "%ld{outbytes}",
  "duration": "%d{duration}",
  "durationms": "%d{durationms}",
  "numsessions": "%d{numsessions}",
  "ipsrulelabel": "%s{ipsrulelabel}",
  "threatcat": "%s{threatcat}",
  "threatname": "%s{ethreatname}",
  "deviceowner": "%s{deviceowner}",
  "devicehostname": "%s{devicehostname}"
}
```

For [Cloud NSS Feeds for DNS Logs](#) (government agencies, see [Cloud NSS Feeds for DNS Logs](#)), copy and paste the prepopulated Feed Output Format with the following:

```
\{"datetime":"%s{time}","user":"%s{ellogin}","department":"%s{edepartment}","location":"%s{elocation}","reagation":"%s{reagation}","resaction":"%s{resaction}","regrulelabel":"%s{regrulelabel}","resrulelabel":"%s{resrulelabel}","dns_reqtype":"%s{reqtype}","dns_req":"%s{req}","dns_resp":"%s{res}","srv_dport":"%d{sport}","durationms":"%d{durationms}","clt_sip":"%s{cip}","srv_dip":"%s{sip}","category":"%s{domcat}","respipcategory":"%s{respipcat}","deviceowner":"%s{deviceowner}","devicehostname":"%s{devicehostname}"\}
```



When logs are streaming, Zscaler creates a file for every batch of logs with the following path (id1 and id2 represent internal IDs): S3bucket/feedtype/feedname=feed_name/year=YYYY/month=MM/day=DD/epochtime_id1_id2_samesecondcount

See the following example: zscaler-bucket-test/_weblog/feedname=s3test_feed/year=2023/month=01/day=23/1674506076_40960_24_2

If you do not include %d{epochtime} in the Feed Output Format, the file path substitutes ingestiontime for epochtime. Ingestion time is when the NSS uploads the feed to the S3 bucket.

You can specify the file extension (e.g., GZIP) of the log data stored in your configured S3 bucket, according to your integration requirements. To enable and set the file extension, contact Zscaler Support.

Streamlining Incident Response with 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](#) (government agencies, see [Configuring the DLP Application Integration Using Amazon Web Services](#)).

AWS integration requires three AWS resources:

1. S3 Bucket: The S3 bucket names share a common prefix.
2. SNS Topic: The metadata S3 bucket pushes notifications to the SNS topic that is subscribed by the Workflow Automation SQS Queue.
3. 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.

Applying Zero Trust Principles to Generative AI Workloads

The following sections describe applying zero trust principles to generative AI (Gen AI) workloads.

Architecture

This integration focuses on securing the data sources that Amazon Q uses to generate responses to user prompts. There are two key components to enable full protection for Amazon Q data flows:

- Deploy Zscaler Cloud Connectors to ensure all traffic consumed by Amazon Q is inspected by the Zscaler ZTE.
- Enable Zscaler's out-of-band CASB scanner to monitor and secure data stored in Amazon S3 buckets.

When setting up Amazon Q for business, follow the official Amazon Q setup guide. Typically, Amazon Q aggregates data from multiple sources to build its database. To secure these data sources, deploy Zscaler Cloud Connectors within a Virtual Private Cloud (VPC) and route Amazon Q's egress traffic through this environment. This ensures that all data accessed by Amazon Q is scanned for malware and potential data loss, providing comprehensive protection for your AI workflows.

The following is an example of data sources that an Amazon Q solution could use:

Data sources (3) [Info](#)

	Name	Source	Data source state
<input type="radio"/>	AWS-Q-Internal-Wiki	WEBCRAWLER	Active
<input type="radio"/>	Global-Cycling-Net...	WEBCRAWLER	Active
<input type="radio"/>	qgenaitest	S3	Active

Figure 201. Data sources

You must ensure the web crawler sites go through a VPC that egresses through a Cloud Connector (Zero Trust Gateway):

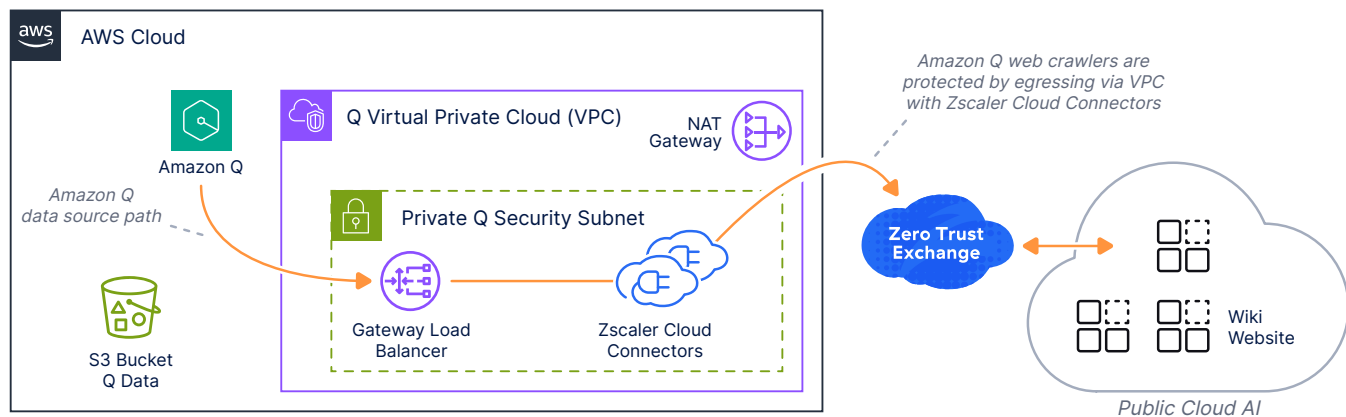


Figure 202. Zscaler and Amazon Q architecture

To ensure that all Amazon Q web crawler traffic is inspected and secured, deploy a Cloud Connector within a dedicated VPC and route the crawler traffic through it. By directing traffic to this VPC, you enable full inspection through the Zscaler ZTE, ensuring that all data collected is free of malware and safeguarded against data loss or exposure of sensitive company information.

For detailed guidance, see [Deploying a Zscaler Cloud Connector for Amazon Web Services](#) (government agencies, see [Deploying a Zscaler Cloud Connector for Amazon Web Services](#)). After your VPC with Cloud Connector is in place, follow the steps outlined in this section to configure Amazon Q's web crawlers to route traffic through your secured VPC environment.

1. In the AWS Management Console under your Amazon Q for Business, click **Add data source**.

Data sources (3) [Info](#)

Sync now

Stop sync

Actions ▼

Add data source

Figure 203. Add data source

2. Select **Web crawler**.



Figure 204. Web crawler

3. Provide the source URLs from where you want to collect data, and then click **Configure VPC and security group**.

Configure VPC and security group - optional [Info](#)

Virtual Private Cloud (VPC)

Select a VPC that defines the virtual networking environment for this repository instance. [Manage VPCs](#)



Figure 205. Configure VPC and security group

4. Select the VPC you configured with your Cloud Connectors.
5. Select the subnet that is configured to forward traffic to the Cloud Connectors. This allows the Amazon Q web crawler to egress out the cloud connectors.
6. Select the security group you configured to allow traffic to flow from the Amazon Q web crawler out to the sites from which you want to collect data.

After you configure the web crawler to scan regularly, you see traffic from the user associated with your cloud connector instance. It might have a default name such as east-1-vpc-01faf2c78bdbd9cbc.

Enabling Zscaler Malware and DLP scanning of AWS S3 buckets

To learn more about setting up S3 bucket scanning, see [Enhancing AWS S3 with Zscaler SaaS Security](#).

For example, you must scan the Amazon Q dataset stored in the AWS S3 bucket, which is configured in the Amazon Q data sources. The following setup enables the Zscaler ZTE to scan and ensure there are no DLP violations or malware stored on the S3 bucket.

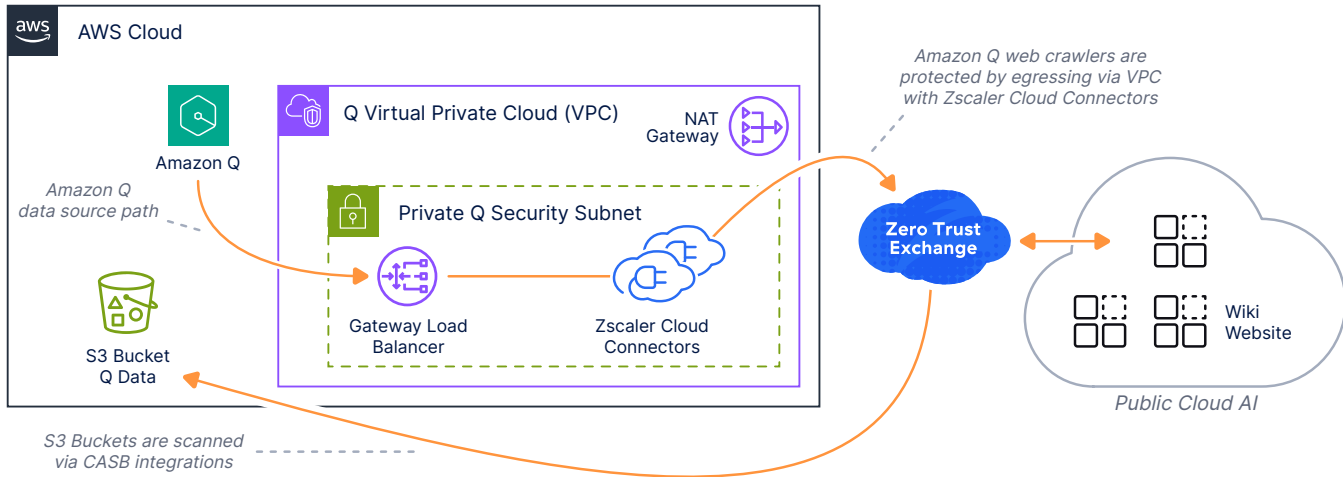


Figure 206. Zscaler and AWS S3 architecture

To configure your Zscaler Tenant to scan the S3 bucket:

1. In the ZIA Admin Portal, go to **Administration** > **Cloud Configuration** > **SaaS Application Tenants**.

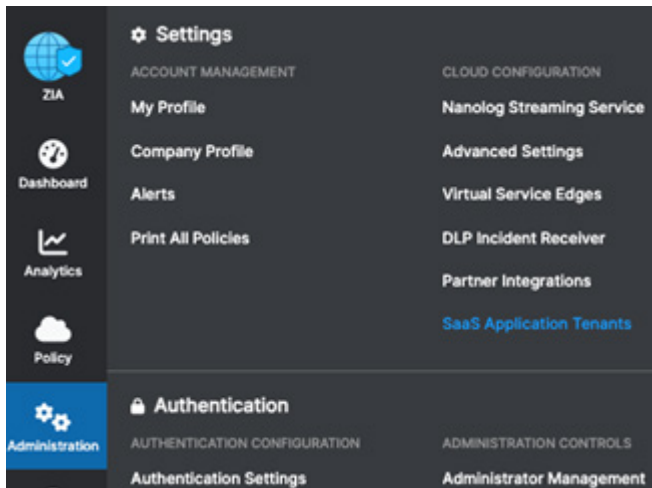


Figure 207. SaaS Application Tenants

2. Click **Add SaaS Application Tenant** and then click the **Amazon S3** bucket.

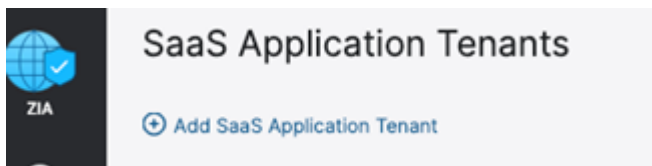


Figure 208. Add SaaS Application Tenant

- Configure the tenant to match the fields shown in the following image (with your information). To learn more about configuring the IAM roles, see [Adding SaaS Application Tenants](#) (government agencies, see [Adding SaaS Application Tenants](#)).

Edit SaaS Application Tenant

- Choose the SaaS Application Provider**
- Name the SaaS Application Tenant**
 Tenant Name: Status: Active
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. [Learn more](#)
 Zscaler Connector Account Number: [Copy](#)
 Zscaler Connector User ARN: [Copy](#) External ID: [Copy](#)
[Reauthorize](#)
- Register the SaaS Application**
To give Zscaler access to Amazon S3, you must configure an IAM role for the Zscaler S3 Connector. [Learn more](#)
 AWS Account ID: IAM Role ARN: [Copy](#)
 Quarantine Bucket Name: CloudTrail Bucket ARN: [Copy](#)
[Validate](#)

Figure 209. Add SaaS Application Tenant

Configure Zscaler Policy Scans

After authorizing the SaaS Application Tenant for S3, configure the scans in your Zscaler policy:

- In the ZIA Admin Portal, go to the **Policy > SaaS Security API > SaaS Security API Control**.

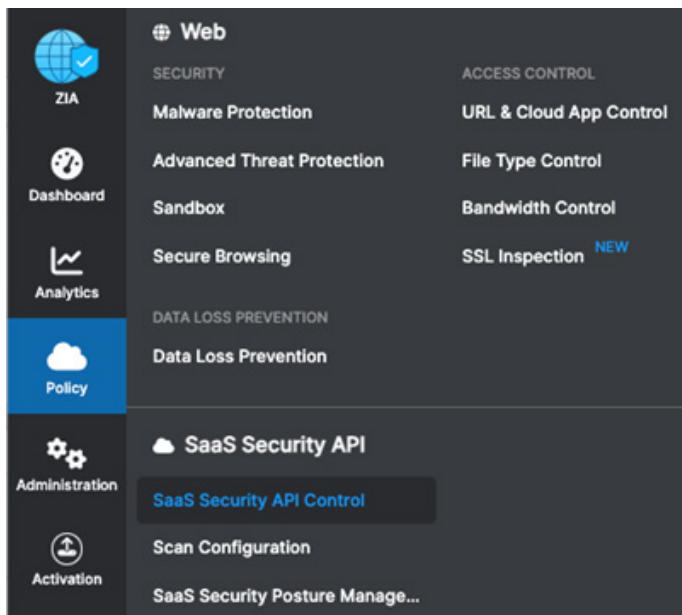


Figure 210. SaaS Security API Control

- Click **Add DLP** rule and configure the settings to match the fields shown in the following image. After you have enabled the DLP rule, set up scans.

Add DLP Rule

DLP RULE

Rule Order

1

Rule Name

SaaS_Storage_App_Rule_1

Rule Status

Enabled

Rule Label

CRITERIA

SaaS Application Tenant

AWSq Gen AI

Buckets

All Buckets Selected in the Scan Schedule

Bucket Owner

Select Bucket Owner

DLP Engines

ClassificationConfidential; Source Code

Collaboration Scope

Any - Any

DLP INCIDENT RECEIVER

Zscaler Incident Receiver

None

ACTION

Action

Report Incident Only

Severity

High

NOTIFICATION

Auditor Type

Hosted

☒ External

Auditor Email Address

spaisley@zscaler.com

Notification Template

DLP Email Template

Figure 211. Add DLP Rule

3. After you have the DLP rule enabled you can setup scans. Go to **Policy > SaaS Security API > Scan Configuration**.

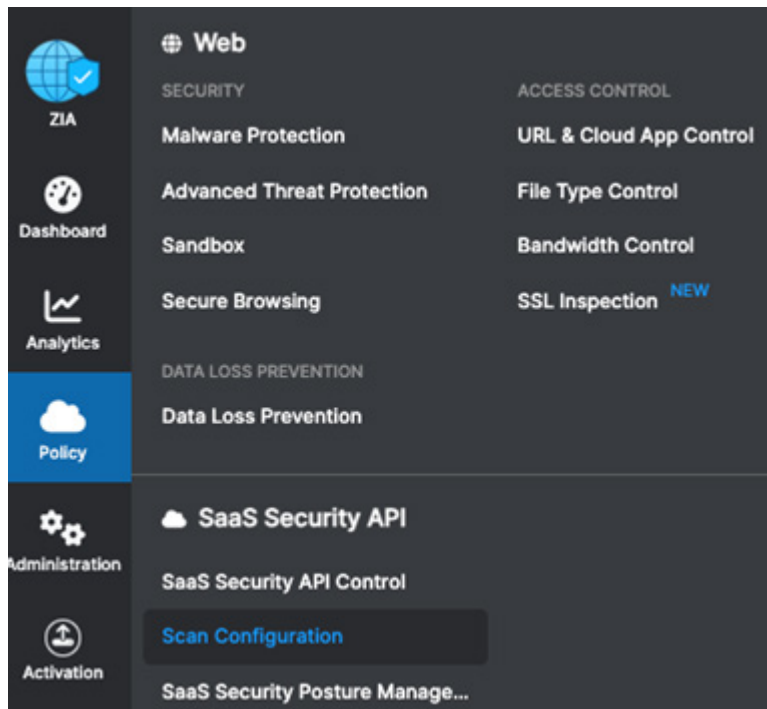


Figure 212. Scan Configuration

4. Select the S3 bucket that is configured as the **Amazon Q Business Data Source**.

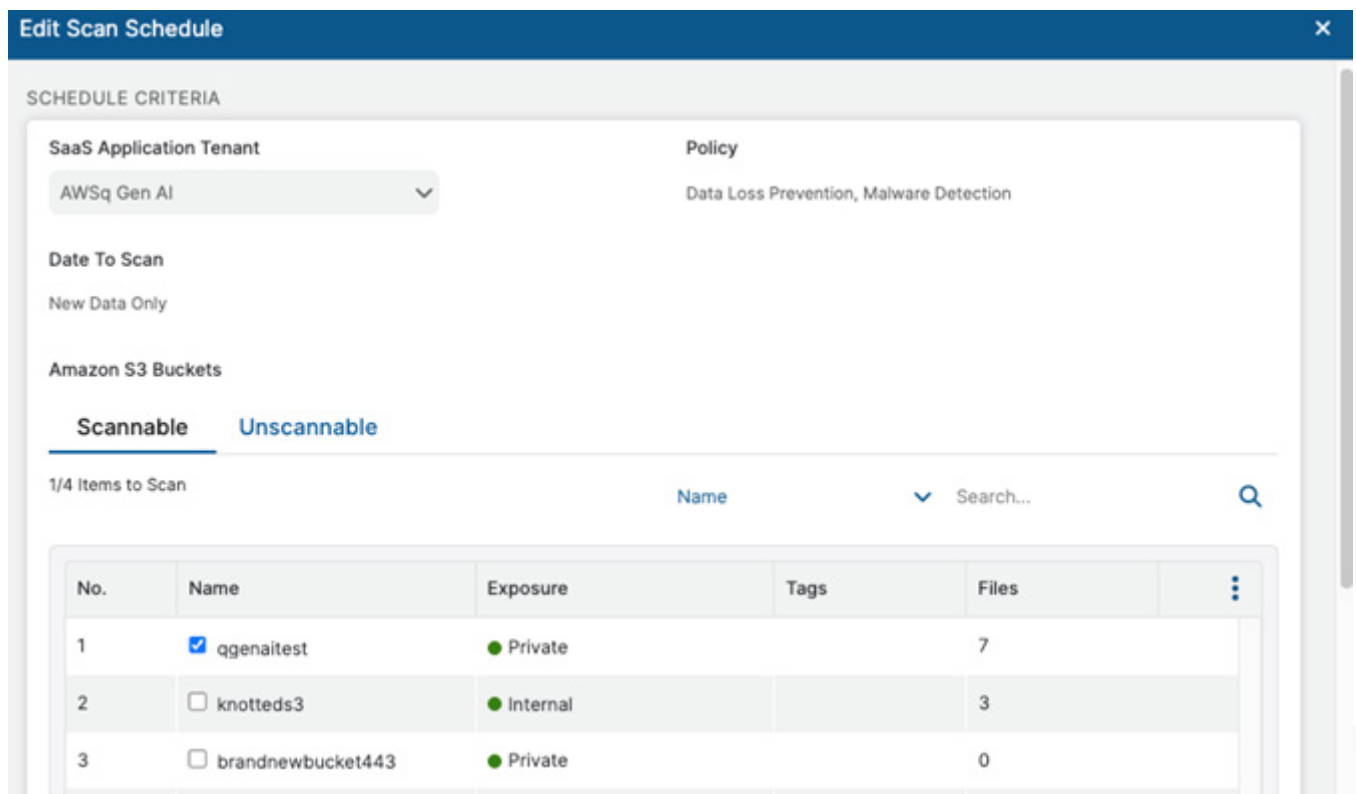


Figure 213. Edit Scan Schedule

5. Click **Save** and **Activate** your policy. Amazon Q business data sources are now protected.

DLP policies within the ZTE help safeguard Gen AI platforms like Amazon Bedrock and Amazon Q, while also ensuring that users only access approved, corporate-sanctioned AI tools.

In the diagram, you can see both end users and corporate users included in the protected environment. After DLP and security policies are configured, they are immediately enforced across the organization. Additionally, you can block access to unsanctioned public AI tools and seamlessly redirect users to approved corporate AI services.

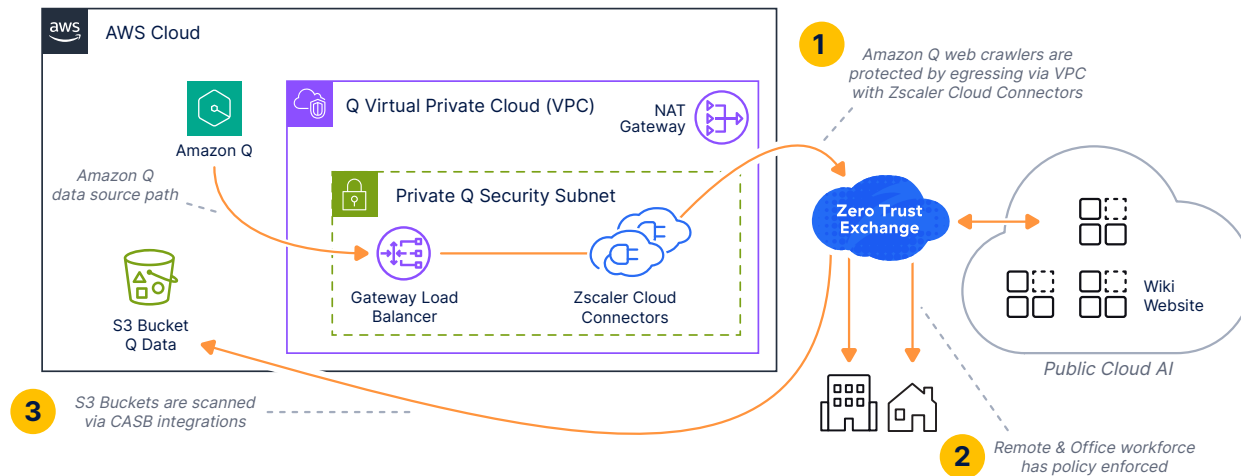


Figure 214. Implemented policies for end users and corporate users

All of the policies protect the Amazon web crawler data collection, the S3 buckets, and the end users from using unsanctioned Gen AI websites for example.

Zscaler Data Protection Overview

Zscaler's DLP engines enable you to detect sensitive data, control transactions by allowing or blocking them, and automatically notify your organization's auditor whenever a transaction triggers a DLP rule. If you're using a third-party DLP solution, Zscaler can forward details of triggered transactions to it securely via ICAP. However, Zscaler operates independently of ICAP responses—enforcement decisions are based solely on the policies configured within Zscaler. The external DLP system can then process the forwarded information and take any necessary remediation actions.

By default, Zscaler evaluates inline DLP rules sequentially and stops at the first match. If you enable Evaluate All Rules mode, Zscaler instead assesses all applicable rules and enforce the action from the most restrictive policy.



You must first delete all existing inline DLP rules before you change the way the Zscaler service evaluates rules. To learn more, see [Configuring DLP Policy Rules with Evaluate All Rules Mode Enabled](#) (government agencies, see [Configuring DLP Policy Rules with Evaluate All Rules Mode Enabled](#)).

Configure a DLP Policy for Private and Public AI Data Protection

The following sections describe configuring DLP policies for private and public AI data protection.

DLP with Content Inspection

AI DLP with content inspection can use predefined dictionaries, custom dictionaries, or exact data match algorithms to detect specific kinds of information in your users' traffic and activities to AI sites. The Zscaler service provides predefined dictionaries that you can modify and, in some cases, clone. You can also create custom dictionaries for content not covered by predefined dictionaries.

Configure DLP Dictionaries

Use DLP dictionaries and engines as defined, or modify them to suit your needs. You can also create custom dictionaries or engines. Skip this procedure if you don't want to modify or create custom DLP dictionaries and engines for AI Data Protection.

To add a custom DLP dictionary:

1. Go to **Administration > DLP Dictionaries & Engines**.
2. Click **Add DLP Dictionary**.
3. In the **Add DLP Dictionary** window:
 - a. **Name**: Enter a name for the dictionary.
 - b. **Dictionary Type**: Select a type from the drop-down menu.
 - **Patterns & Phrases**: If selected, the **Patterns & Phrases** sections appear, where you can add patterns, phrases, and apply actions to them. To learn more, see [Defining Patterns for Custom DLP Dictionaries](#) and [Defining Phrases for Custom DLP Dictionaries](#) (government agencies, see [Defining Patterns for Custom DLP Dictionaries](#) and [Defining Phrases for Custom DLP Dictionaries](#)).

The screenshot shows the 'Add DLP Dictionary' window. The 'Dictionary Type' is set to 'Patterns & Phrases'. The 'Patterns' section is highlighted with an orange box and contains a table with columns 'Pattern' and 'Action'. The 'Phrases' section is also highlighted with an orange box and contains a table with columns 'Phrase' and 'Action'. A tooltip 'Patterns & Phrases Dictionary Type' is visible over the 'Add Pattern' button.

Figure 215. Add DLP Dictionary Patterns & Phrases

- **Microsoft Information Protection (MIP):** If selected, the MIP labels appear, where you can select the MIP labels. To learn more, see [Adding an MIP Account](#) and [Defining Microsoft Information Protection Labels for Custom DLP Dictionaries](#) (government agencies, see [Adding an MIP Account](#) and [Defining Microsoft Information Protection Labels for Custom DLP Dictionaries](#)).

Add DLP Dictionary

DLP DICTIONARY

Name

Dictionary Type

Microsoft Information Protection (MIP)

Match On

<input type="checkbox"/>	Label Name	Label Value
<input type="checkbox"/>	AccessRestriction	e4129fe8-719c-408e-8bd6-183231fd25de
<input type="checkbox"/>	Confidential	04e197b9-da12-45e9-a208-88ff2b1ab2f3
<input type="checkbox"/>	Confidential:All Employee...	59cc00fc-f0c1-48bf-a940-61d8bacf7b81
<input type="checkbox"/>	Confidential:Anyone (no...	017bcd9-d7b-46ce-b018-899f1a673625
<input type="checkbox"/>	Confidential:Finance	c9d1f7ec-1f52-4ab2-82d8-d2ed61e8ced9
<input type="checkbox"/>	Confidential:Recipients ...	299d9dce-432b-44d1-be90-90e205df1c10
<input type="checkbox"/>	Confidential:TestConfid...	3a426b94-e393-4033-af03-c2bcb1bebe5f
<input type="checkbox"/>	General	3ba90ed3-9f2a-40ac-830b-12d5a2e411af
<input type="checkbox"/>	Highly Confidential	e1deef7f-80b4-4ebd-b720-323020a9146c
<input type="checkbox"/>	Highly Confidential:All E...	e1cb6ae8-b499-454c-9776-8d929ad18a7d
<input type="checkbox"/>	Highly Confidential:Any...	1100943e-c605-4117-ab72-88d042357632
<input type="checkbox"/>	Highly Confidential:Reci...	68fda39a-f3e4-4be5-9ab4-ffb82de16cbc
<input type="checkbox"/>	labelTest	474d3452-31ec-40e9-a0a7-04980200fa89
<input type="checkbox"/>	Personal	e792bc4a-6adc-4eda-b3a3-026511cc985e
<input type="checkbox"/>	Public	fb02abc-25e3-4ce4-bcd6-eed8957ba70d
<input type="checkbox"/>	Secret	ea6d7a48-b5f4-46f6-861d-b7be1437dad3
<input type="checkbox"/>	TestLabel	d6f4e25c-f5cc-4a64-9530-9a4851176029

Save
Cancel

Figure 216. Add DLP Dictionary Microsoft Information Protection

- **Indexed Document Match:** If selected, the Indexed Document Match section appears, where you can select existing IDM templates and choose match accuracy levels for those templates. To learn more, see [Creating an Indexed Document Match Template](#) and [Defining IDM Match Accuracy for Custom DLP Dictionaries](#) (government agencies, see [Creating an Indexed Document Match Template](#) and [Defining IDM Match Accuracy for Custom DLP Dictionaries](#)).

Add DLP Dictionary

DLP DICTIONARY

Name
Enter Text

Dictionary Type
Indexed Document Match

Exclude 100% Match
☐ Yes ☒ No

Description

INDEXED DOCUMENT MATCH

Index Template
None

Save Cancel

Figure 217. Add DLP Dictionary Index Document Match

- **Exact Data Match:** If selected, the Exact Data Match (EDM) section appears, where you can select existing EDM templates and add data fields from those templates. To learn more, see [Creating an Exact Data Match Template](#) and [Defining Exact Data Match Fields for Custom DLP Dictionaries](#) (government agencies, see [Creating an Exact Data Match Template](#) and [Defining Exact Data Match Fields for Custom DLP Dictionaries](#)).

The screenshot shows a modal window titled "Add DLP Dictionary". It contains two main sections. The first section, "DLP DICTIONARY", includes a "Name" input field, a "Dictionary Type" dropdown menu currently showing "Exact Data Match", and a "Description" text area. The second section, "EXACT DATA MATCH", is highlighted with an orange border and contains a "Data Template" dropdown menu set to "NONE" and an "Add Template" button. At the bottom of the modal are "Save" and "Cancel" buttons.

Figure 218. Add DLP Dictionary Exact Data Match

- c. **Match Type:** This is only applicable if you are configuring a Patterns & Phrases type dictionary. Select a Match Type from the drop-down menu to configure how the dictionary triggers when matching patterns and phrases.
 - **Match Any:** This is the default setting. If selected, the dictionary triggers when a transaction matches any one of the dictionary's patterns or phrases
 - **Match All:** If selected, the dictionary triggers when a transaction matches all of the dictionary's patterns and phrases
 - d. **Description:** (Optional) Enter a description for the dictionary.
4. Click **Save** and **Activate** the change.

Configure DLP Engine

Adding a custom DLP engine is one of the tasks you can complete when configuring DLP policy rules.



You can add a custom DLP engine on the Add DLP Engine window or through the [Cloud Service API](#) (government agencies, see [Cloud Service API](#)).

To add a custom DLP Engine:

1. Go to **Administration > DLP Dictionaries & Engines**.
2. In the **DLP Engines** tab, click **Add DLP Engine**.
3. In the **Add DLP Engine** window, enter the **Name** for the custom DLP engine.
4. In the **Engine Builder** section, add operators and DLP dictionaries to build an expression. You can see your expression in the **Expression Preview**.

Figure 219. ZIA Engine Builder

5. Under **Expression**:
 - a. Select an operator to build your expression. The operators include **All (AND)**, **Any (OR)**, **Exclude (AND NOT)**, and **Sum**. The Sum operator is available for count-based DLP dictionaries (e.g., Credit Cards, Social Security Numbers, etc.) and allows you to specify the sum total of matches that trigger a group of dictionaries specified in the DLP engine.
 - b. Select a dictionary from the drop-down menu, then specify a [match count](#) (government agencies, see [match count](#)) as needed.
 - c. Click **Add** to add a Dictionary or a Subexpression. Click **Remove** to delete dictionaries or subexpressions.
 - If you use the **Sum** operator, select two or more predefined or custom DLP dictionaries. You must set a value for the match count. You can enter any value less than 1000.
 - If you use the **All**, **Any**, or **Exclude** operators, you must select a predefined or custom DLP dictionary. Certain dictionaries require you to set a value for the match count. You can enter any value less than 1000.

- If you use **Subexpression**, you must select an operator. The operators include **All (AND)**, **Any (OR)**, **Exclude (AND NOT)**, and **Sum**. The **Sum** operator is available for count-based DLP dictionaries (e.g., Credit Cards, Social Security Numbers, etc.) and allows you to specify the sum total of matches that trigger a group of dictionaries specified in the subexpression
- 6. Continue adding dictionaries and operators to the expression as needed. At each level, you can create up to 4 subexpressions, use up to 4 operators, and add up to 16 dictionaries per operator.
- 7. (Optional) For **Description**, enter any additional notes or information. The description cannot exceed 255 characters.
- 8. Click **Save** and **Activate** the change.

Define Policy Rules

To define your policy rules:

1. Go to **Policy > Data Loss Prevention**.
2. Click **Add** and select **Rule With Content Inspection**.
3. In the **Add DLP Rule** window:
 - a. **Rule Order**: Policy rules are evaluated in ascending numerical order (Rule 1 before Rule 2, and so on), and the Rule Order reflects this rule's place in the order. You can change the value, but if you've enabled **Admin Ranking**, the assigned Admin Rank determines the Rule Order values you can select.
 - b. **Admin Rank**: Enter a value from 0–7 (0 is the highest rank). Your assigned admin rank determines the values you can select. You cannot select a rank that is higher than your own. The rule's Admin Rank determines the value you can select in Rule Order so that a rule with a higher Admin Rank always precedes a rule with a lower Admin Rank.
 - c. **Rule Name**: Enter a unique name for the DLP rule or use the default name.
 - d. **Rule Status**: An enabled rule is actively enforced. A disabled rule is not actively enforced, but does not lose its place in the Rule Order; the service skips it and moves to the next rule.
 - e. **Rule Label**: Select a rule label to associate it with the rule. To learn more, see [About Rule Labels](#) (government agencies, see [About Rule Labels](#)).
4. Define the following **Criteria**:
 - a. **DLP Engines**: Select **Any** to choose all DLP engines for this rule, or select up to 4 engines. You can search for DLP engines or click the Add icon to create a new DLP engine.
 - b. The **Match Only** option takes effect for both **Allow** and **Block** rule actions. You can select **Match Only** to configure how engines must trigger in order for the service to take action. To learn more, see [DLP Policy Configuration Example: Match Only](#) (government agencies, see [DLP Policy Configuration Example: Match Only](#)).
5. For Public AI Sites, select the in-scope AI site under the **URL Categories or Cloud Applications**.
6. For **Private AI Sites**:
 - a. **ZPA Application Segment**: Select **Any** to apply the rule to all [ZPA application segments](#) (government agencies, see [ZPA application segments](#)), or select up to 255 ZPA application segments. You can also search for ZPA application segments.
 - b. **File Type**: From the drop-down menu, choose the file types for the rule. You can create DLP policy rules that apply just to content sent via specific file types. Policies that reference Zscaler DLP engines support different file types than [policies that reference external DLP engines](#) (government agencies, see [policies that reference external DLP engines](#)). Zscaler DLP engines can scan files of up to 100 MB. For an archived file, the size of individual files when decompressed can also be a maximum of 100 MB.

- c. **Minimum Data Size:** Enter the minimum size requirement that data must meet before the DLP rule applies. The default minimum data size, 0 KB, means there is no minimum data size requirement.
- d. **Users:** You can specify how the DLP rule applies to your users.
 - Choose **Include** to apply the rule to selected users and no other users. From the drop-down menu, choose **Any** to apply the rule to all users or select up to 4 users.
 - Choose **Exclude** to apply the rule to all other users and not selected users. You can select up to 256 users.
- e. **Groups:** You can specify how the DLP rule applies to your groups.
 - Choose **Include** to apply the rule to selected groups and no other groups. From the drop-down menu, choose **Any** to apply the rule to all groups or select up to 8 groups.
 - Choose **Exclude** to apply the rule to all other groups and not selected groups. You can select up to 256 groups.
- f. **Departments:** You can specify how the DLP rule applies to your departments.
 - Choose **Include** to apply the rule to selected departments and no other departments. From the drop-down menu, choose **Any** to apply the rule to all departments or select up to 8 departments.
 - Choose **Exclude** to apply the rule to all other departments and not selected departments. You can select up to 256 departments.
- g. **User Risk Profile:** Select the user risk score levels to which the rule applies. Selecting no value ignores this criterion in the policy evaluation. Users are assigned a risk score based on their browsing activities. A range of risk scores is grouped as a risk score level:
 - **Low:** Level with user risk scores ranging from 0 to 29.
 - **Medium:** Level with user risk scores ranging from 30 to 59.
 - **High:** Level with user risk scores ranging from 60 to 79.
 - **Critical:** Level with user risk scores ranging from 80 to 100.
- h. **Locations:** Select **Any** to apply the rule to all [locations](#) (government agencies, see [locations](#)), or select up to 8 locations. You can also search for a location or click **Add** to add a new location.
- i. **Location Groups:** Select **Any** to apply the rule to all [location groups](#) (government agencies, see [location groups](#)), or select up to 32 location groups. You can also search for a location group.
- j. **Time:** Select **Always** to apply this rule to all [time intervals](#) (government agencies, see [time intervals](#)), or select up to two time intervals. You can also search for a time interval or click **Add** to add a new time interval.
- k. **Protocols:** Select the protocols to which the rule applies.
 - **HTTP:** Data transactions and file uploads from HTTP websites.
 - **HTTPS:** Data transactions and file uploads from HTTPS websites encrypted by TLS/SSL.
 - **Native FTP:** Data transactions and file uploads from native FTP servers.
- l. **Inspect Downloads:** Enable this option to allow DLP inspection for content downloaded from specific AI apps. If this option is enabled, you must choose at least one application segment. If disabled, the DLP rule only applies to content sent to cloud apps.

- m. (Optional) For **DLP Incident Receiver**, complete one of the following tasks:
- If you don't have a third-party DLP solution or don't want to forward content, leave the following **Zscaler Incident Receiver** or **ICAP Receiver** field as **None**.
 - If you want to forward the transactions captured by this policy rule to a DLP incident receiver:
 - For **Incident Receiver**, select whether the DLP incident receiver is an [ICAP](#) (government agencies, see [ICAP](#)) receiver (government agencies, see) or a [Zscaler Incident Receiver](#) (government agencies, see [Zscaler Incident Receiver](#)).
 - Select the applicable **ICAP Receiver** or **Zscaler Incident Receiver** from the drop-down menu. You must configure your ICAP receivers or Zscaler Incident Receivers in order to complete this step.
- n. Select the **Action** for the rule. You can **Allow** or **Block** transactions that match the rule. If you select **Allow**, the service allows and logs the transaction. If you select **Block**, the service blocks and logs the transaction.
- o. (Optional) Configure an email notification for the rule. If you do not select an auditor and notification template, a notification is not sent for this rule.
- For **Auditor Type**, select whether the auditor is from a **Hosted** database or **External** to your organization.
 - Select the **Auditor**:
 - If the auditor is from a hosted database, select or search for the auditor.
 - If the auditor is external, enter the auditor's email address.
 - Select a **Notification Template**, if you [configured one](#) (government agencies, see [configured one](#)). You can also search for a notification template or click the Add icon to add a new notification template.
- p. (Optional) Configure **Client Connector Notification**. You can **Enable** or **Disable** Zscaler Client Connector notifications for the rule when violations occur. The field is only available if you enable the Web DLP Violations option for your organization on the End User Notifications page in the ZIA Admin Portal and you select the Action as Block for the rule. See [Using the Zscaler Notification Framework](#) (government agencies, see [Using the Zscaler Notification Framework](#)).
- q. (Optional) Enter a **Description** including additional notes or information. The description cannot exceed 10,240 characters.
- r. Click **Save** and **Activate** the change.



You can combine public and private AI sites under a single policy if required.

Edit DLP Rule

Content Matching

Select DLP Engines

None

DLP Engines

AWS BedRock SageMaker AI Data Pro...

URL Categories

Any

Cloud Applications

Any

Cloud Application Instances

Any

ZPA Application Segment

AWS EAST IP ONLY

File Type

Any

Minimum Data Size (KB)

0

Users

Any

Groups

Any

Departments

Any

User Risk Profile

Any

Locations

Any

Location Groups

Any

Time

Always

Protocols

HTTP; HTTPS; Native FTP

Inspect Downloads

Enable

Disable

Save

Cancel

Figure 220. Zscaler Internet Access DLP Rule

Configure the Zscaler Notification Framework

You can optionally configure various user notification settings in Zscaler Client Connector. Some options become available after the user's device is enrolled in the Zscaler service, and can be adjusted directly within the Zscaler Client Connector settings.

There are two types of notification systems: the standard Windows notification system and Zscaler's built-in Notification Framework. While both deliver the same messages, users cannot disable Zscaler's Notification Framework through Windows settings.

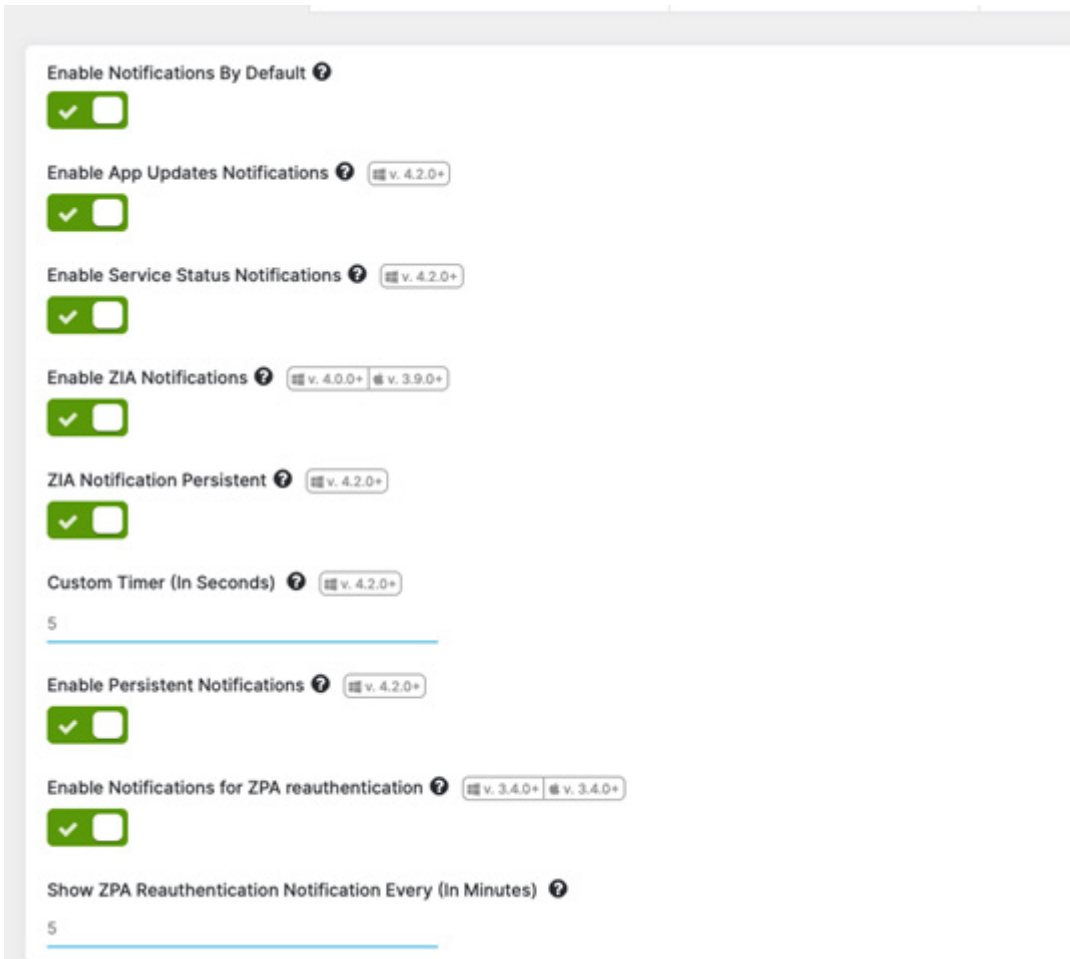
Zscaler notifications appear in the bottom right corner of the window. Up to five notifications can be shown at once, each automatically disappearing after five seconds. Users can also move or dismiss these notifications by clicking anywhere within the notification window.

To configure the Zscaler Notification Framework, follow these steps:

1. In the Zscaler Client Connector Portal, go to **Administration > Client Connector Notifications**. Click the **End User Notifications** tab and select from the following options:
2. **Enable Notifications by Default**: This setting is enabled when a user is enrolled. Users can turn this option off from the Zscaler Client Connector.
3. **Enable App Updates Notifications**: Select this option to have users receive app upgrade notifications.
4. **Enable Service Status Notifications**: Select this option to have users receive status notifications for Zscaler services, such as when a service is in Disaster Relief (DR) mode.
5. **Enable ZIA Notifications**: Select this option to have users receive notifications from ZIA, such as DLP notifications.
6. **Enable Notifications for ZPA reauthentication**: Select this option to prompt users for authentication. This option is enabled after a user is enrolled. Users can turn off this option from Zscaler Client Connector.
7. **Show ZPA Reauthentication Notifications Every (In Minutes)**: Select this option to show ZPA reauthentication notifications at a specific time interval. This setting is enabled by default. Enter a value from 2 to 1440 to set the interval in minutes.
8. **Custom Timer (In Seconds)**: Use this option to set the time the notification displays for the user. Enter a time between 5 and 60 seconds.
9. **Enable Persistent Notifications**: This setting is enabled by default and displays critical notifications until the user dismisses them. Critical notifications include ZPA reauthentication, captive portal detection, request for a system reboot, and packet capture.
10. **ZIA Notification Persistent**: When enabled, this option overrides the custom timer and makes notifications persistent.
11. Click **Save**.



Use ZIA Notification Persistence carefully. Based on ZIA policy, users might have several DLP block notifications.



The screenshot displays a configuration panel for Zscaler notifications. It includes several toggle switches, all of which are turned on (indicated by a green checkmark in a box). Each toggle is accompanied by a help icon (a question mark in a circle). Below the toggles, there are two input fields for numerical values, both set to '5'. The settings are as follows:

- Enable Notifications By Default** (Help icon): ☒
- Enable App Updates Notifications** (Help icon): ☒ (Requires v. 4.2.0+)
- Enable Service Status Notifications** (Help icon): ☒ (Requires v. 4.2.0+)
- Enable ZIA Notifications** (Help icon): ☒ (Requires v. 4.0.0+ and v. 3.9.0+)
- ZIA Notification Persistent** (Help icon): ☒ (Requires v. 4.2.0+)
- Custom Timer (In Seconds)** (Help icon): (Requires v. 4.2.0+)
- Enable Persistent Notifications** (Help icon): ☒ (Requires v. 4.2.0+)
- Enable Notifications for ZPA reauthentication** (Help icon): ☒ (Requires v. 3.4.0+ and v. 3.4.0+)
- Show ZPA Reauthentication Notification Every (In Minutes)** (Help icon):

Figure 221. Enable ZIA Notifications

You must also enable the notification framework in the Application Profile. To enable the Zscaler Notification Framework on Windows devices:

1. In the Zscaler Client Connector Portal, go to **App Profiles**.
2. Click **Add Windows Policy**.
3. Enable **Use Zscaler Notification Framework**.
4. Click **Save**.

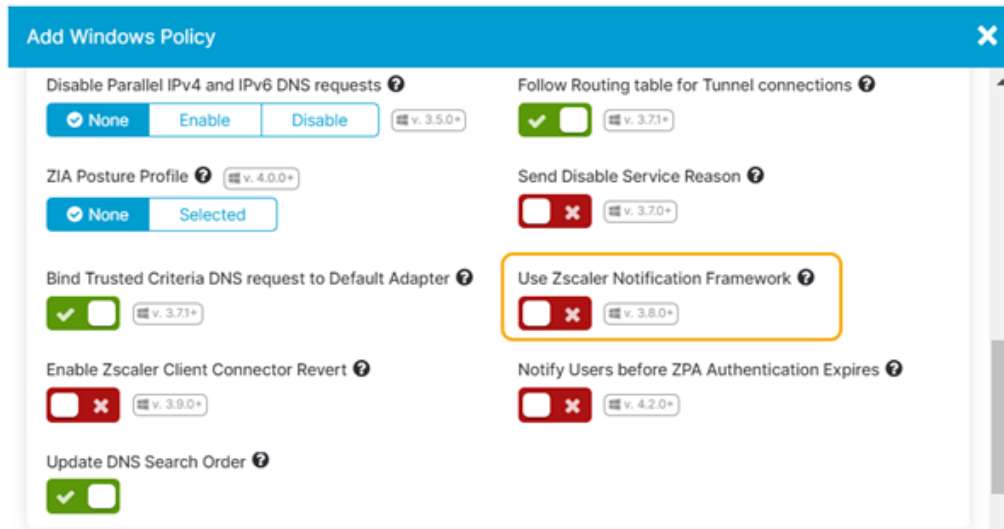


Figure 222. Enable Windows Notification Framework

To enable the Zscaler Notification Framework on macOS devices:

1. In the Zscaler Client Connector Portal, go to **App Profiles**.
2. Click **Add macOS Policy**.
3. Enable **Use Zscaler Notification Framework**.
4. Click **Save**.

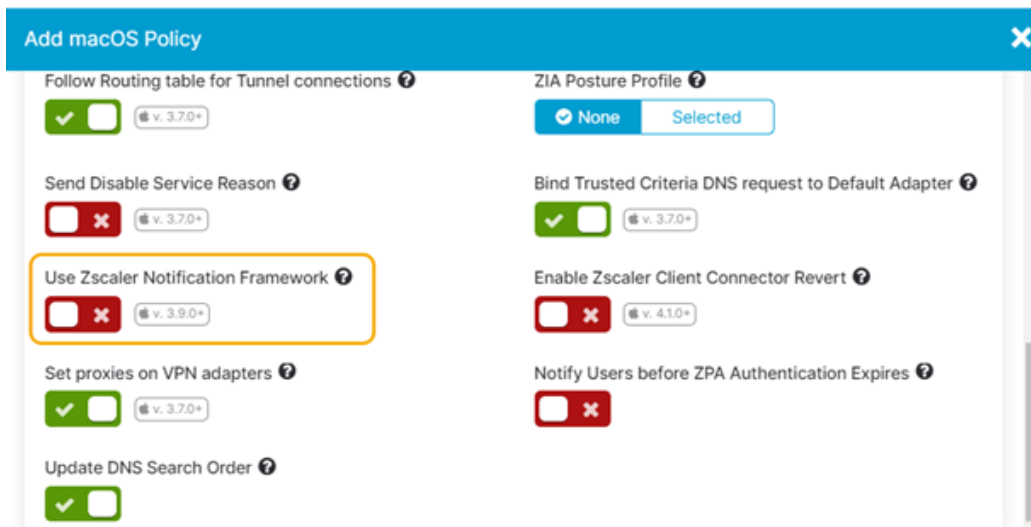


Figure 223. Enable macOS Notification Framework

Contextualizing Risk using AWS and Zscaler UVM

Zscaler'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.

Zscaler 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 Zscaler 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 [Zscaler 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
Accounts	Retrieves your organization's accounts details.	organizations:DescribeAccount organizations:ListAccounts organizations:ListTagsForResource

Note: Attach this permission to the root/organization account.

- Under the second Action in ZscalerPolicy, edit the permissions list to cover those necessary for the data you want to retrieve

```

"ZscalerPolicy": {
  "Properties": {
    "PolicyDocument": {
      "Statement": [
        {
          "Sid": "AllowSQSReceiveMessage",
          "Effect": "Allow",
          "Action": [
            "sqs:ReceiveMessage",
            "sqs:DeleteMessage",
            "sqs:ChangeMessageVisibility"
          ],
          "Resource": "arn:aws:sqs:*:*:*Zscaler*"
        },
        {
          "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 `Zscaler-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 Resources](#).
7. Run the following CloudFormation Role Stack command.

```

aws cloudformation create-stack \

--region <REGION> \

--stack-name ZscalerStackIntegration \

--capabilities CAPABILITY_NAMED_IAM \

--template-body file://Zscaler-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 `Zscaler-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 has been created with a response of a StackID, such as:

```

{

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

}

```

Output for the RoleARNID and ExternalID

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

```

"Outputs": [

  {

    "OutputKey": "RoleARNID",

    "OutputValue": "arn:aws:iam:: :*****459973:role/
    ZscalerAccess-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 Zscaler UVM Platform.
2. Click **Configure**.

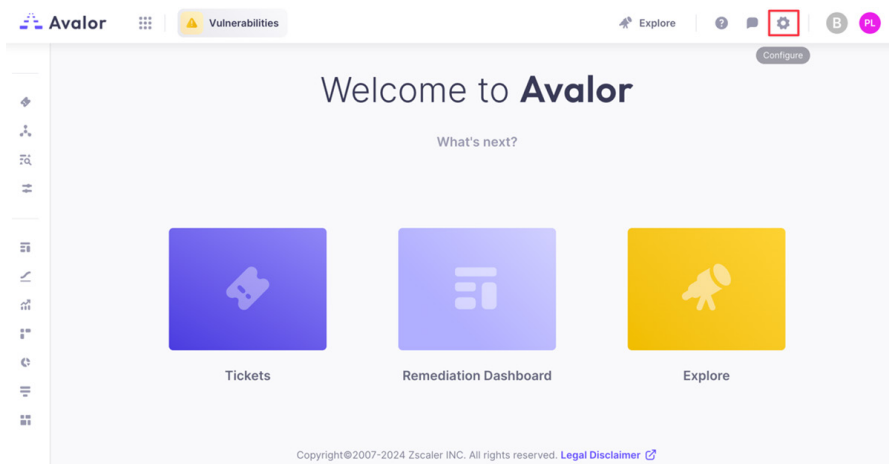


Figure 224. Zscaler UVM Platform

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

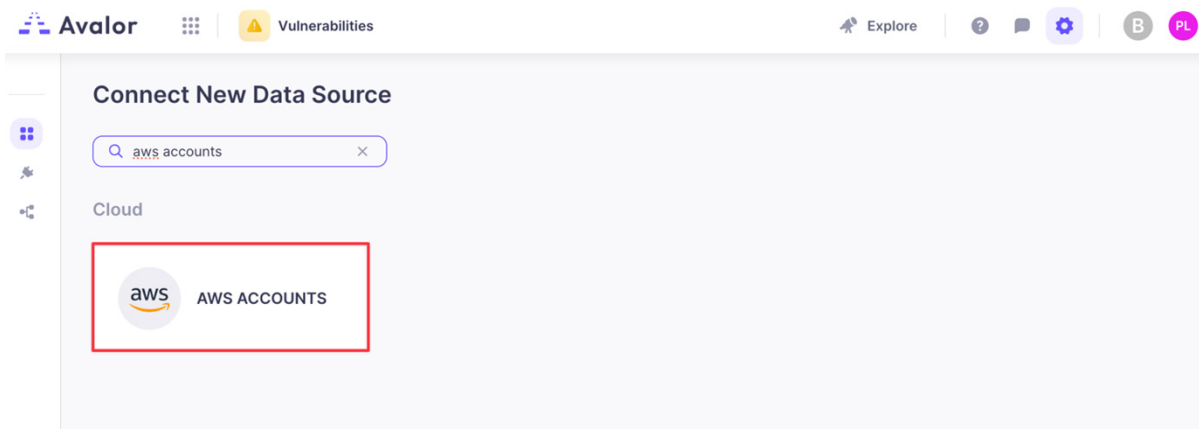


Figure 225. Connect New Data Source

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 this data source applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).
6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

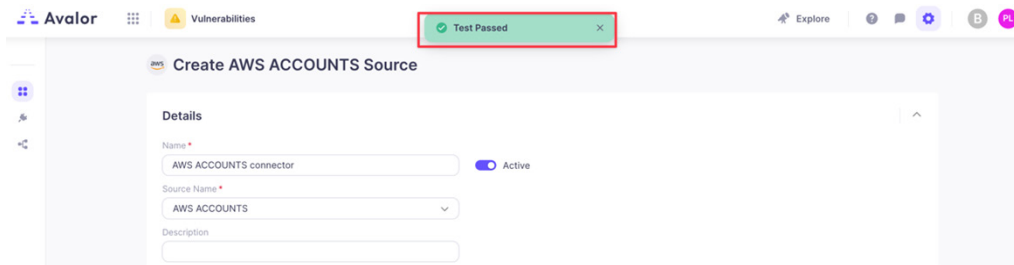


Figure 226. Test Passed

- Click **Save**.

Create AWS ACCOUNTS Source

Details

Name* ☒ Active

Source Name*

Description

Retrieval

Authentication*

Region Names*

Role ARN*

External ID

Scheduling

Full Refresh Frequency*

Time (UTC)*

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

Advanced Settings

Suppression Rules

[+ AND](#) [+ OR](#)

☒ Prevent NULL from overriding existing values

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

Figure 227. Create AWS Accounts Source

Configure the AWS EC2 Data Source

- Log in to the Zscaler UVM Platform.
- Click **Configure**.

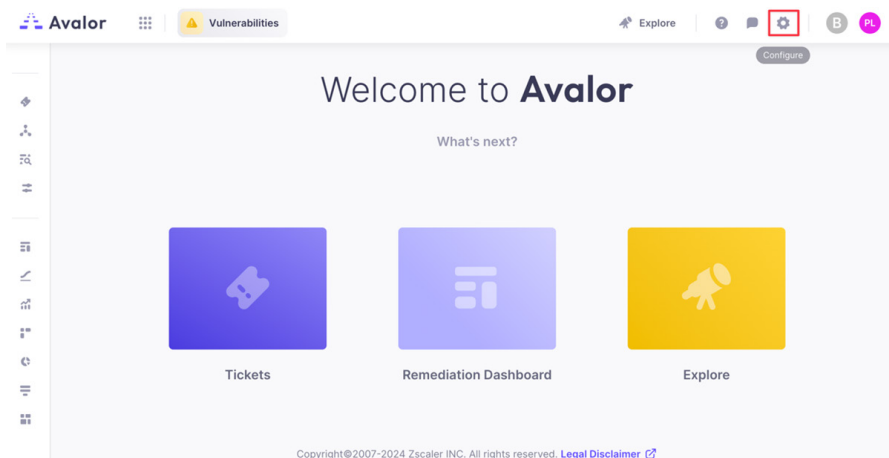


Figure 228. Zscaler UVM Platform

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

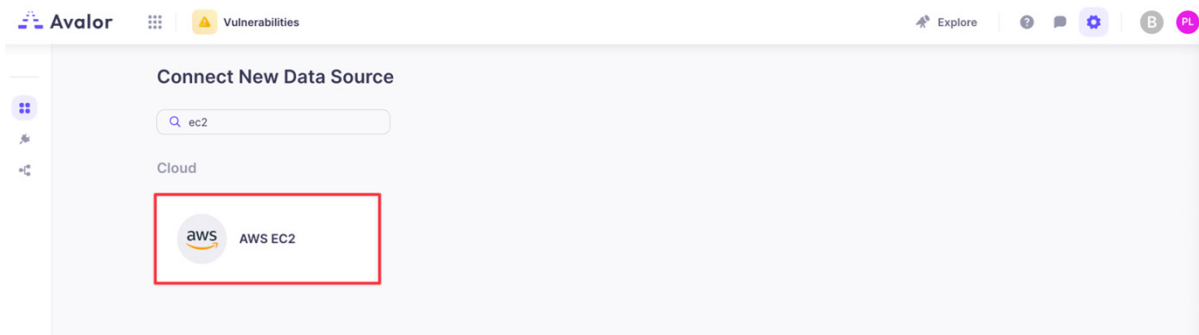


Figure 229. Connect New Data Source

- Click the **AWS EC2** application.
- On the **Create AWS EC2 Source** page, complete the following:
 - Name:** Enter a name for the Data Connector.
 - Active:** Toggle the switch to enable the Data Connector.
 - Authentication:** Enter the Role ARN.
 - Region Names:** Select the Region Names this data source applies to.
 - Role ARN:** Enter the Role ARN.
 - External ID:** Enter the External ID.
 - 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.
 - Remediation Detection Settings:** Select your desired option to determine when findings automatically turn to undetected. To learn more, see the [Zscaler documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Zscaler system. To learn more, see the [Zscaler documentation](#).
- Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

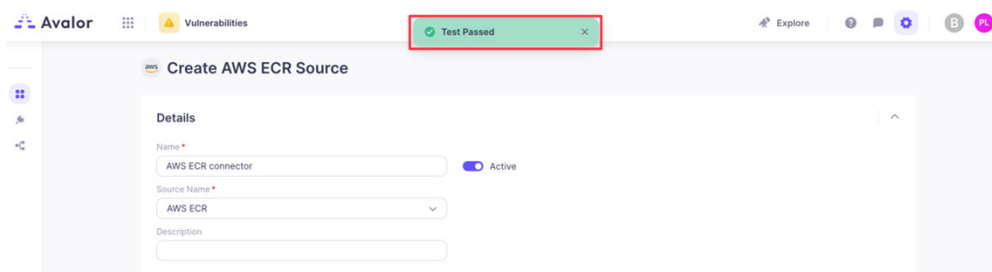


Figure 230. Test Passed

7. Click **Save**.

Create AWS EC2 Source

Details

Name * AWS EC2 connector Active

Source Name * AWS EC2

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 231. Create AWS EC2 Source

Configure the AWS ECR Data Source

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

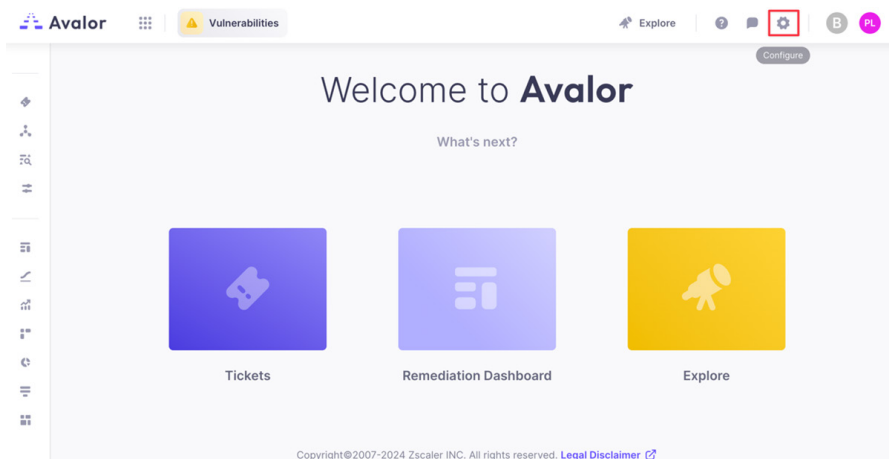


Figure 232. Zscaler UVM Platform

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

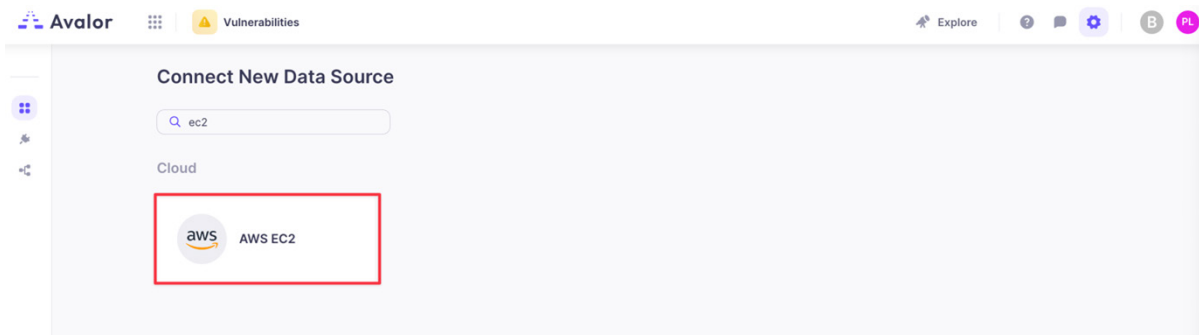


Figure 233. Connect New Data Source

- Click the **AWS ECR** application.
- On the **Create AWS ECR Source** page, complete the following:
 - Name:** Enter a name for the Data Connector.
 - Active:** Toggle the switch to enable the Data Connector.
 - Authentication:** Enter the Role ARN.
 - Region Names:** Select the Region Names this data source applies to.
 - Role ARN:** Enter the Role ARN.
 - External ID:** Enter the External ID.
 - 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.
 - Remediation Detection Settings:** Select your desired option to determine when findings automatically turn to undetected. To learn more, see the [Zscaler documentation](#). Automatic remediation detection only applies when data is refreshed fully, not incrementally.
 - Advanced Settings > Suppression Rules:** Define rules and conditions to remove specific data before it enters the Zscaler system. To learn more, see the [Zscaler documentation](#).
- Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

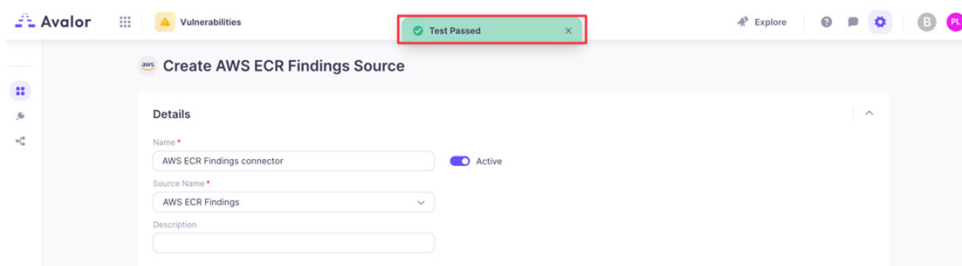


Figure 234. Test Passed

7. Click **Save**.

Avalor Vulnerabilities Explore ? ⚙️ 👤 🔒

Create AWS ECR Source

Details

Name *

AWS ECR connector Active

Source Name *

AWS ECR

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 235. Create AWS ECR Source

Configure the AWS ECR Findings Data Source

To configure the AWS ECR findings data source:

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

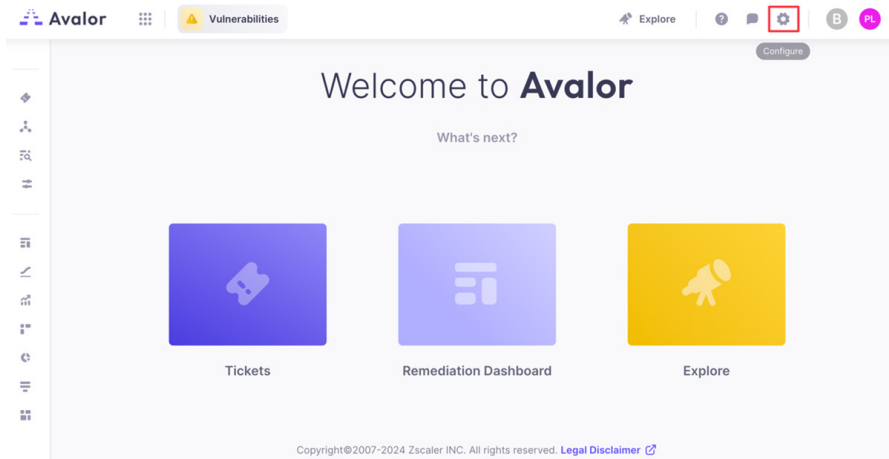


Figure 236. Zscaler UVM Platform

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

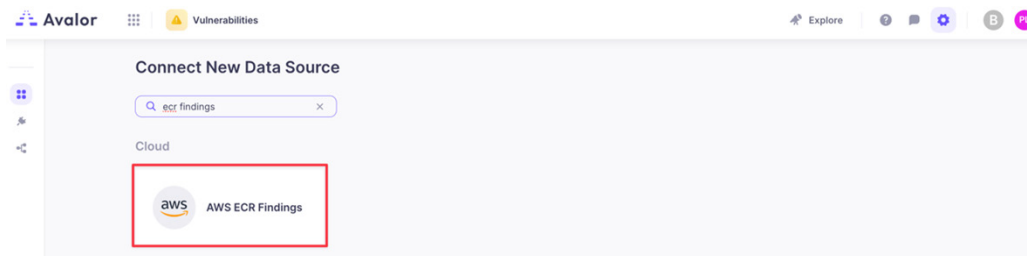


Figure 237. Connect New Data Source

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 applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

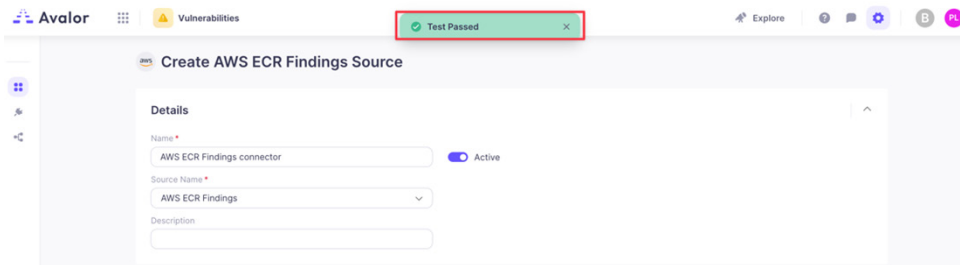


Figure 238. Test Passed

7. Click **Save**.

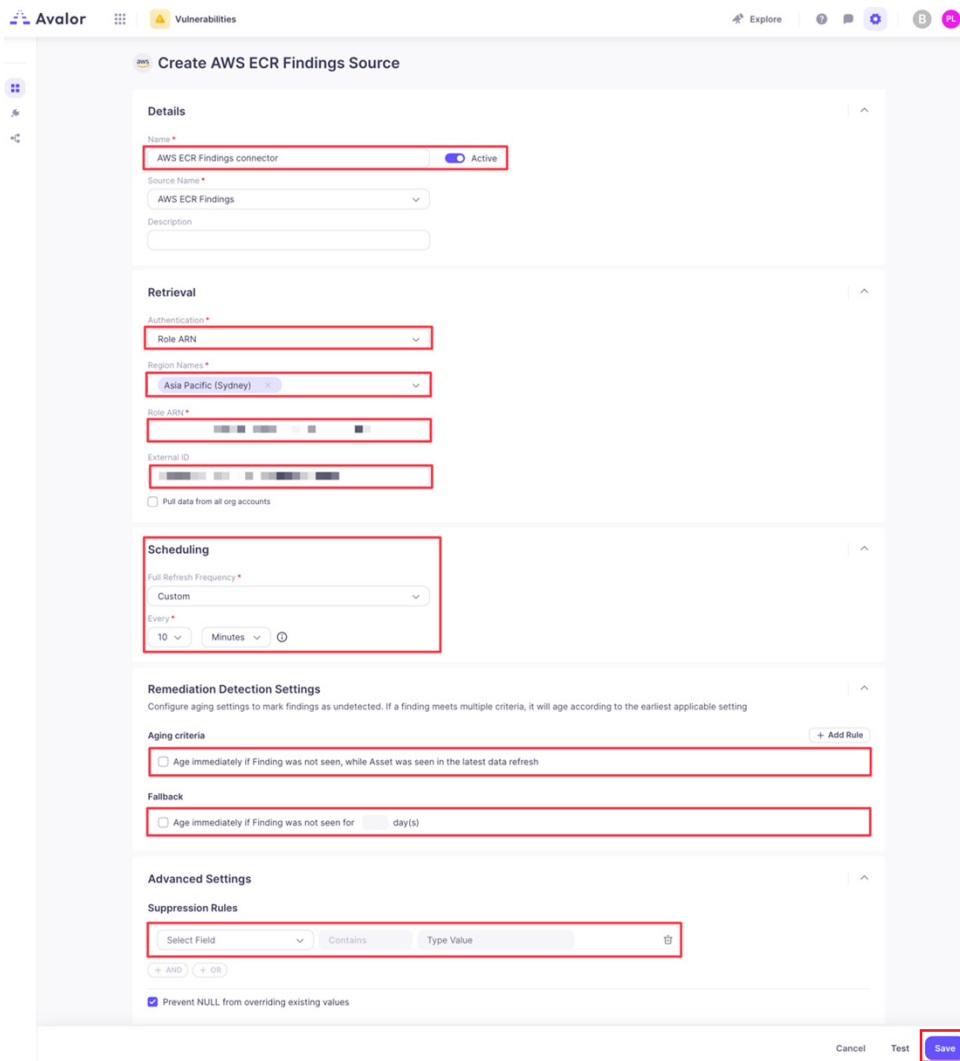


Figure 239. Create AWS ECR Findings Source

Configure the AWS EKS Clusters Data Source

To configure the AWS EKS clusters data source:

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

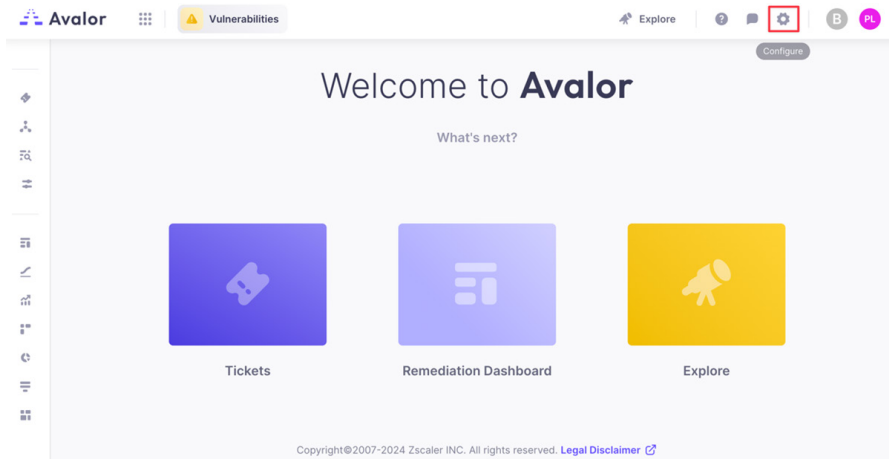


Figure 240. Zscaler UVM Platform

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

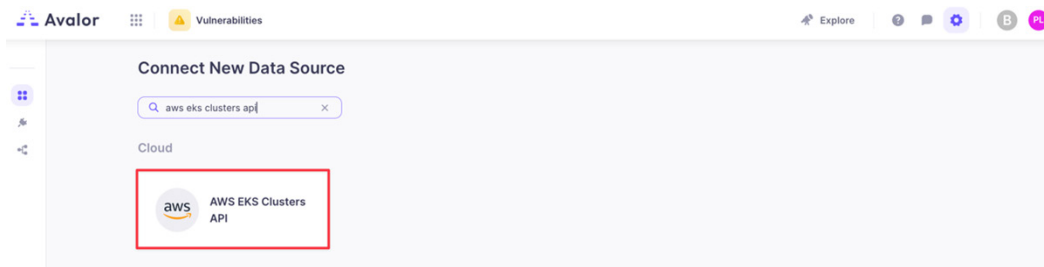


Figure 241. Connect New Data Source

4. Click 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 applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

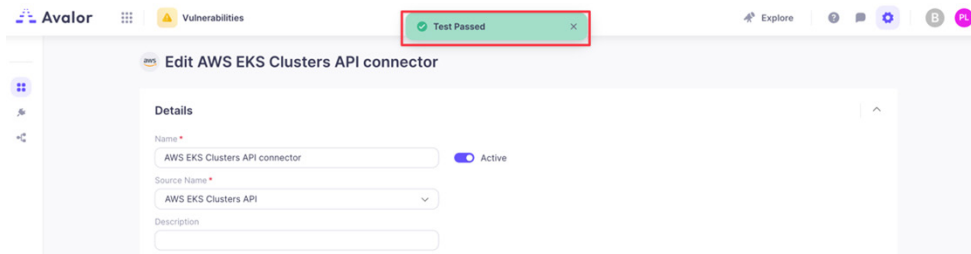


Figure 242. Test Passed

7. Click **Save**.

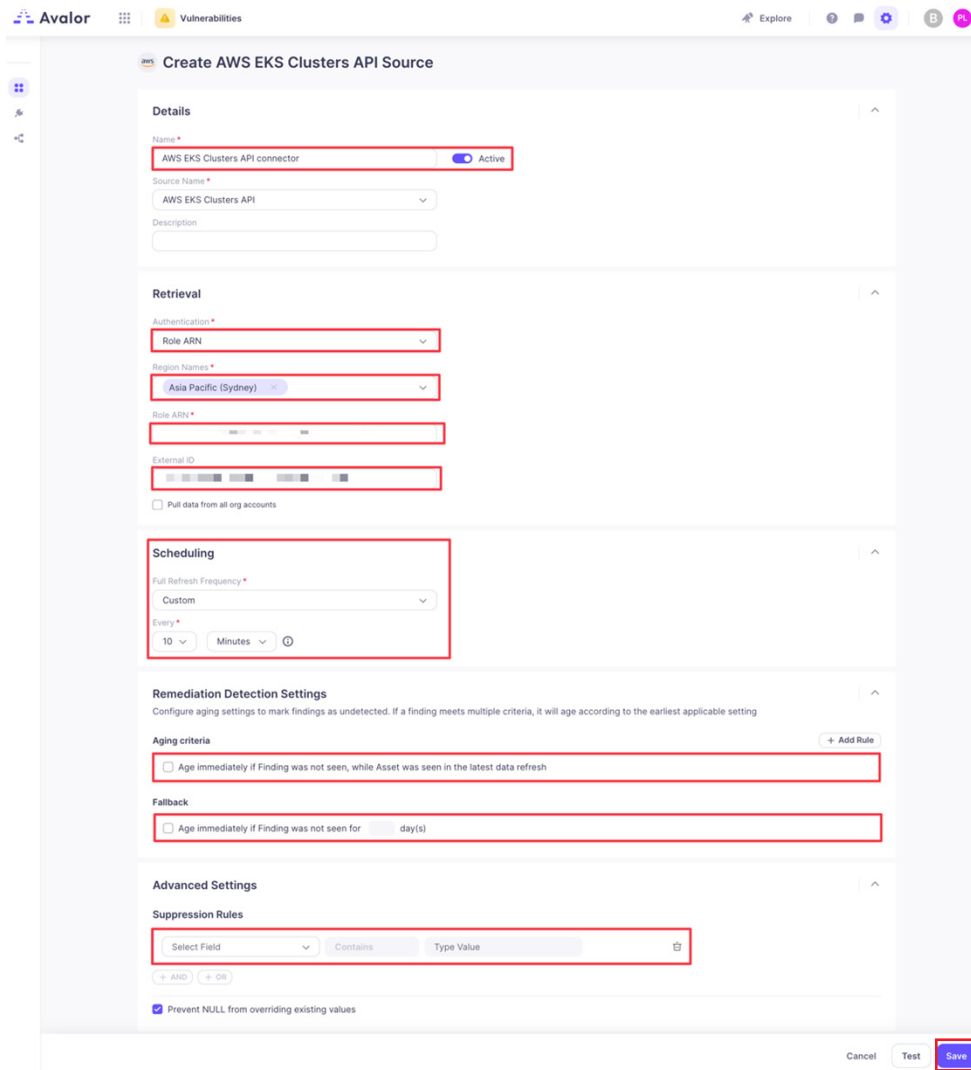


Figure 243. Create AWS EKS Clusters API Source

Configure the AWS Inspector Findings Data Source

To configure the AWS inspector findings data source:

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

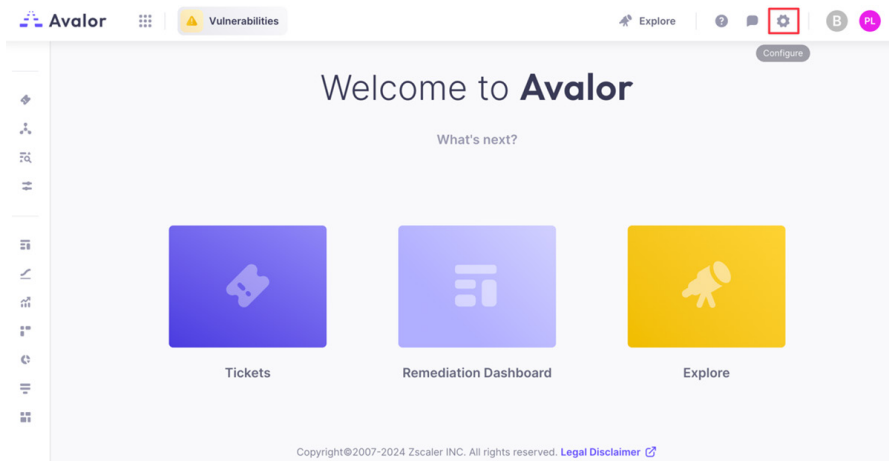


Figure 244. Zscaler UVM Platform

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

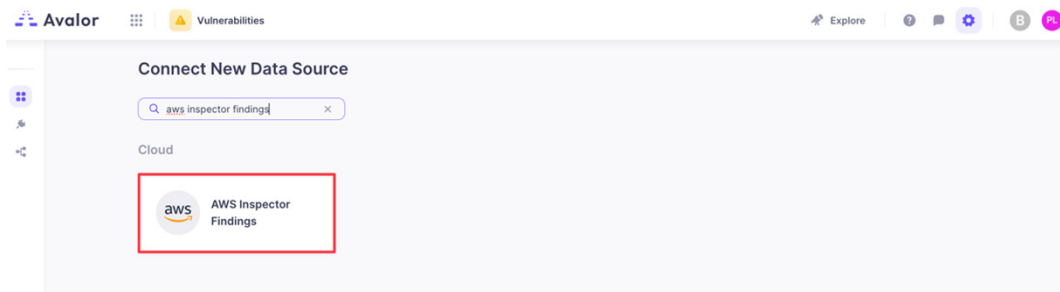


Figure 245. Connect New Data Source

4. Click 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 applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

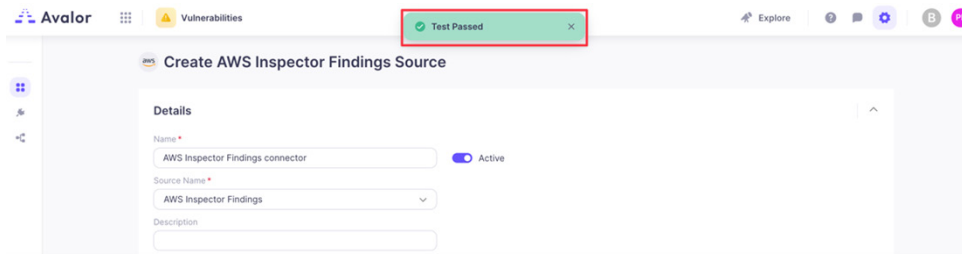


Figure 246. Test Passed

7. Click **Save**.

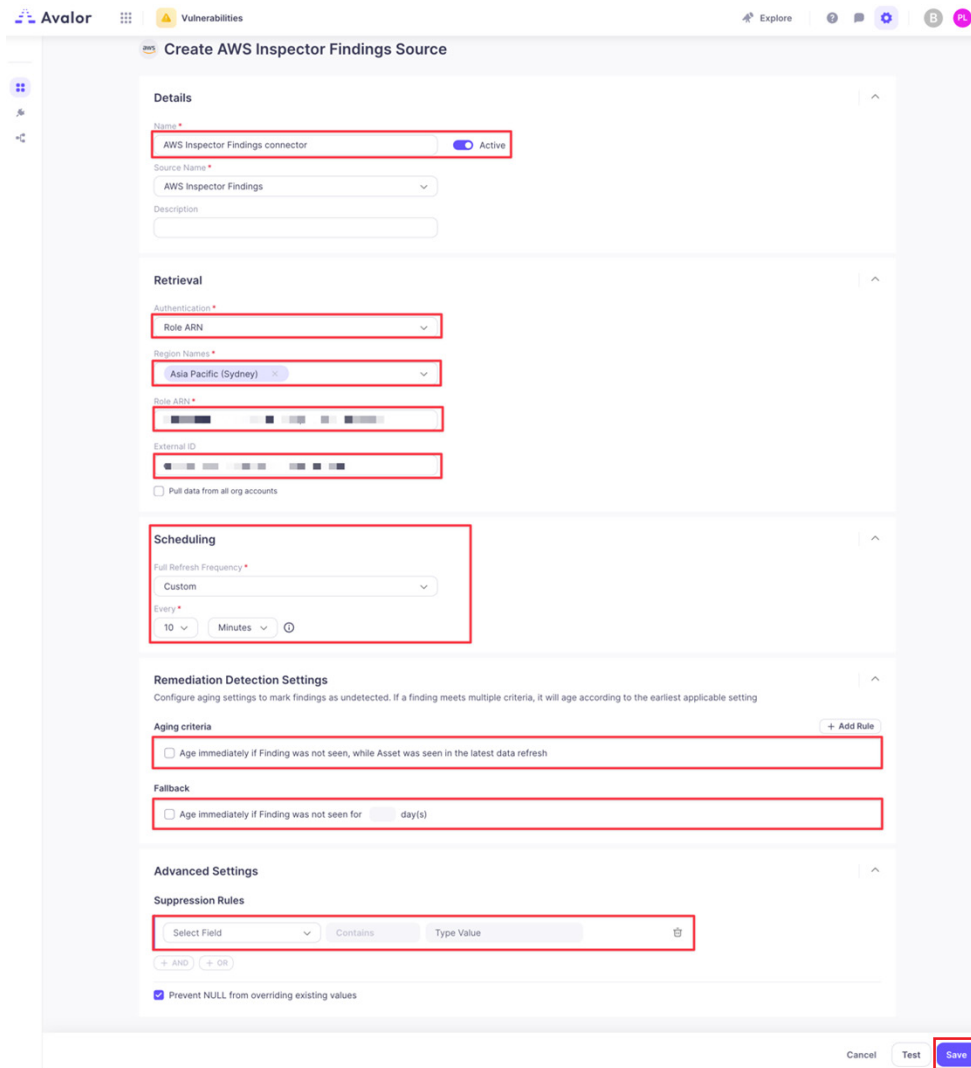


Figure 247. Create AWS Inspector Findings Source

Configure the AWS RDS Data Source

To configure the AWS RDS data source:

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

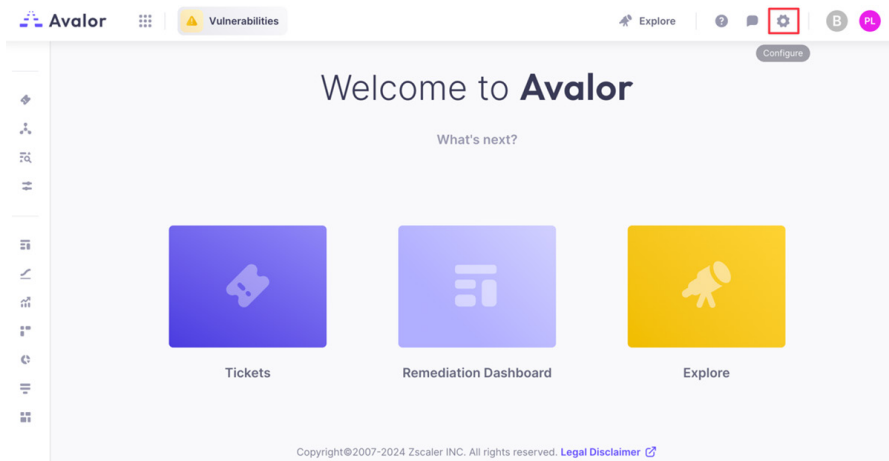


Figure 248. Zscaler UVM Platform

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

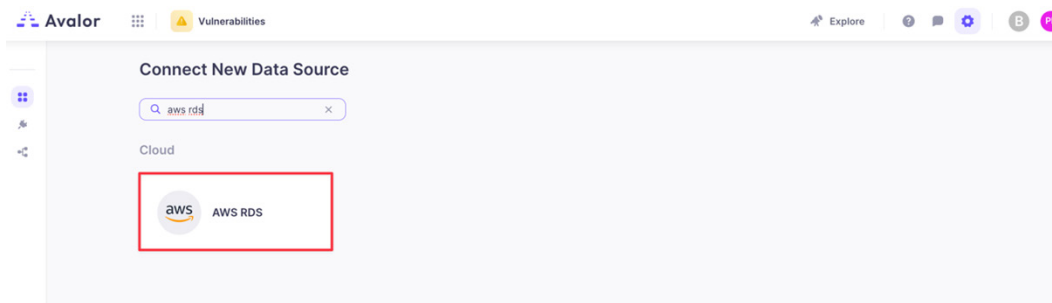


Figure 249. Connect New Data Source

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 this data source applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

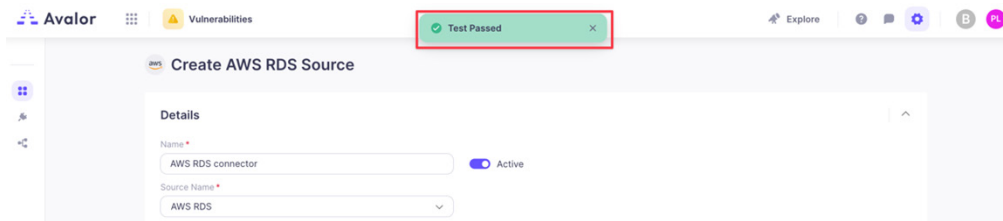


Figure 250. Test Passed

7. Click **Save**.

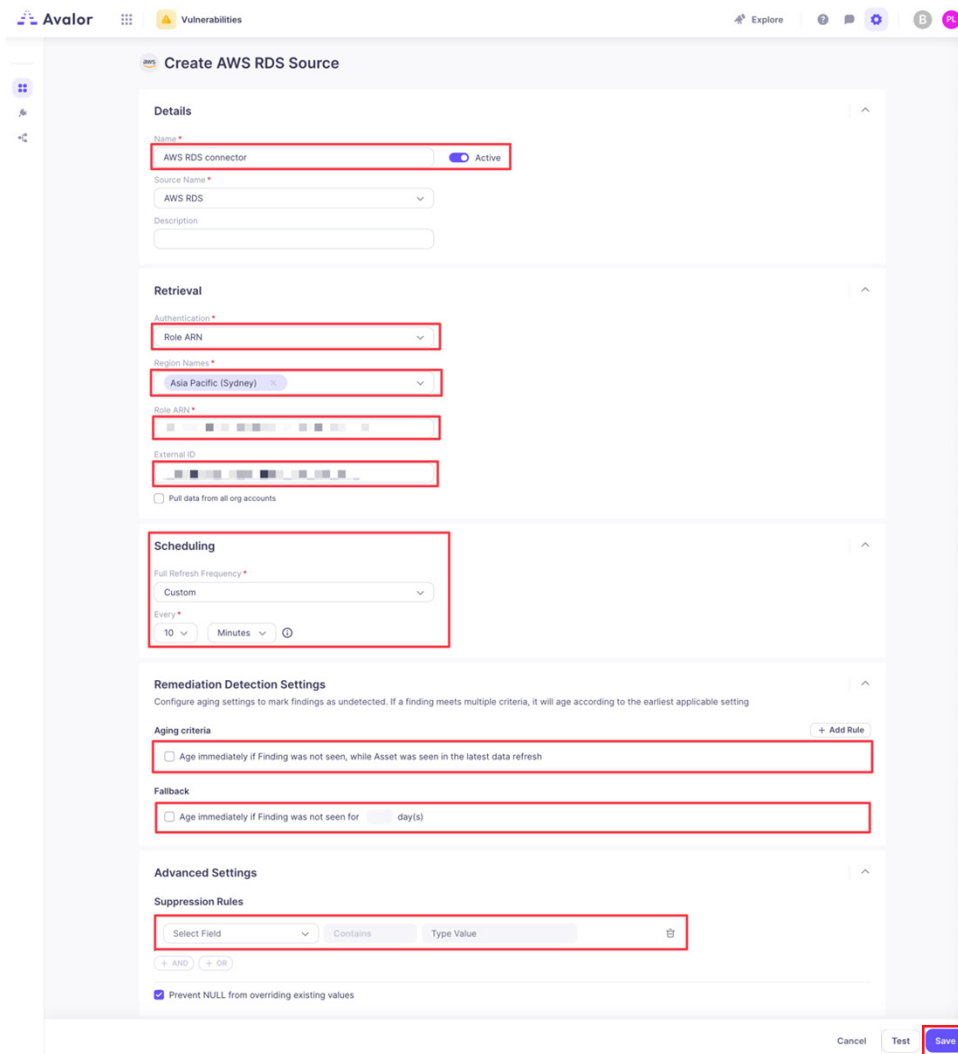


Figure 251. Create AWS RDS Source

Configure the AWS S3 Buckets Data Source

To configure the AWS S3 buckets data source:

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

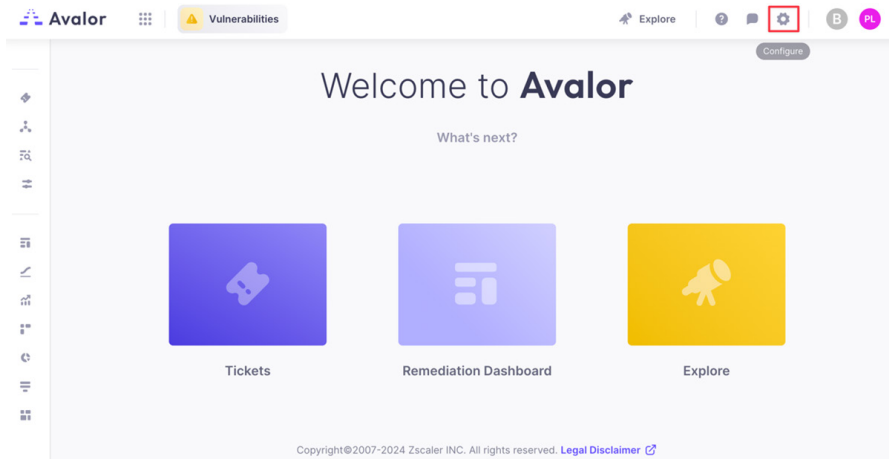


Figure 252. Zscaler UVM Platform

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

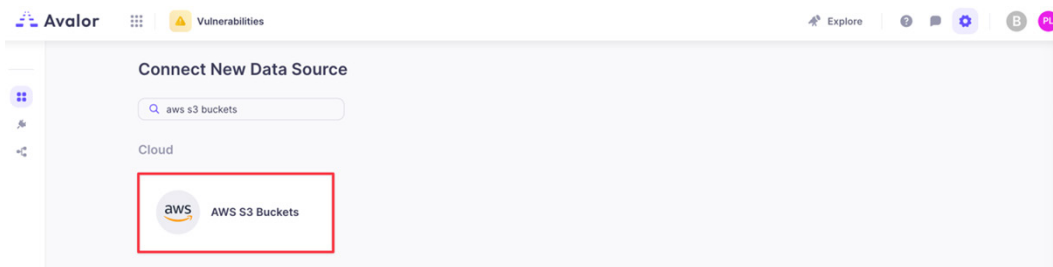


Figure 253. Connect New Data Source

4. Click 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 applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

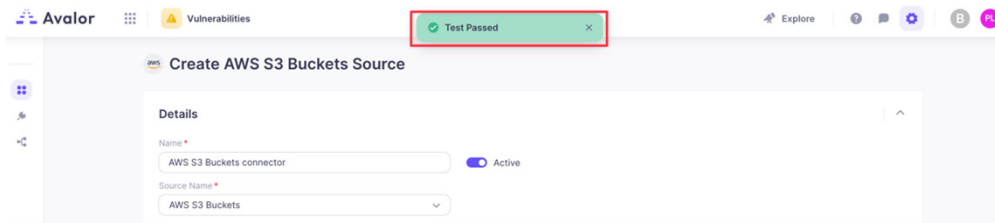


Figure 254. Passed Test

7. Click **Save**.

 This screenshot shows the 'Create AWS S3 Buckets Source' form with several fields highlighted by red boxes. The 'Details' section includes 'Name' (AWS S3 Buckets connector), 'Source Name' (AWS S3 Buckets), and an 'Active' toggle. The 'Retrieval' section includes 'Authentication' (Role ARN), 'Region Names' (Asia Pacific (Sydney)), 'Role ARN', and 'External ID'. The 'Scheduling' section includes 'Full Refresh Frequency' (Custom) and 'Every' (10 Minutes). The 'Remediation Detection Settings' section includes 'Aging criteria' (Age immediately If Finding was not seen, while Asset was seen in the latest data refresh) and 'Fallback' (Age immediately If Finding was not seen for 1 day(s)). The 'Advanced Settings' section includes 'Suppression Rules' (Select Field, Contains, Type Value) and a checkbox for 'Prevent NULL from overriding existing values'. The 'Save' button is highlighted in the bottom right corner.

Figure 255. 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 Zscaler UVM Platform.
2. Click **Configure**.

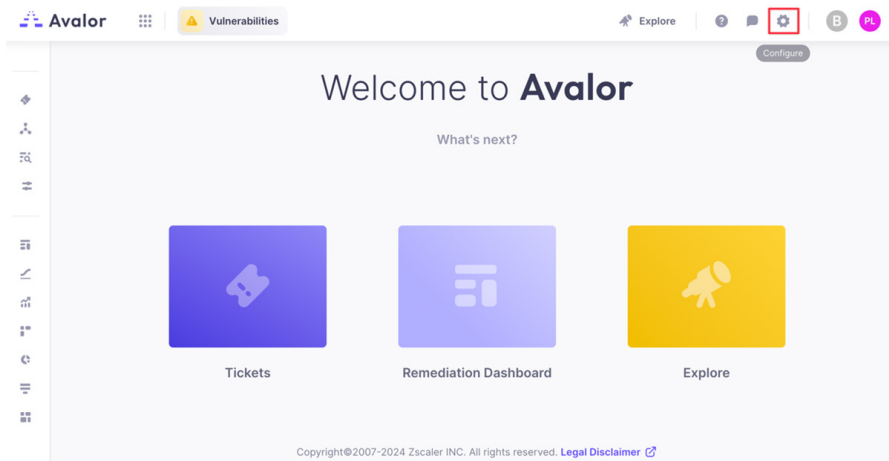


Figure 256. Zscaler UVM Platform

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

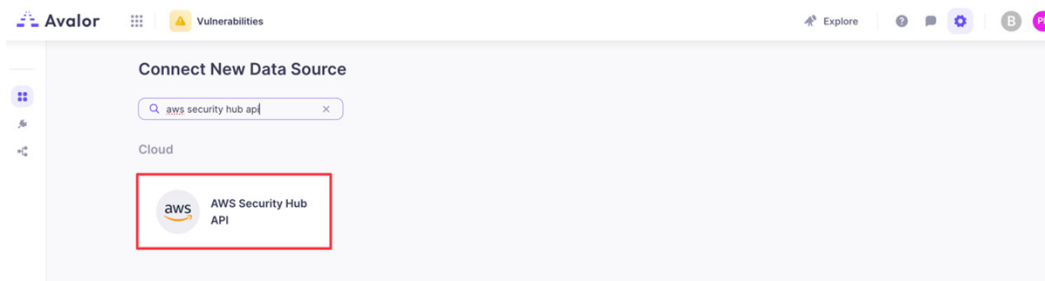


Figure 257. Connect New Data Source

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 this data source applies 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 to undetected. To learn more, see the [Zscaler 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 Zscaler system. To learn more, see the [Zscaler documentation](#).

6. Click **Test**. If the Role ARN and External ID have been entered correctly, the system responds with **Test Passed**.

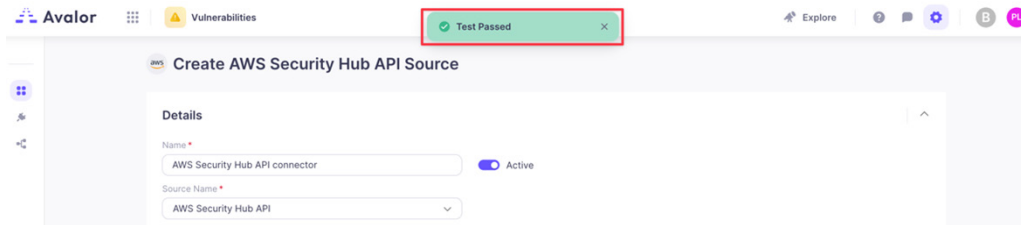


Figure 258. Test Passed

7. Click **Save**.

 The screenshot shows the 'Create AWS Security Hub API Source' form with several sections highlighted by red boxes:

- Details:** 'Name' (AWS Security Hub API connector), 'Source Name' (AWS Security Hub API), and the 'Active' toggle.
- Retrieval:** 'Authentication' (Role ARN), 'Region Names' (Asia Pacific (Sydney)), 'Role ARN' (a long alphanumeric string), and 'External ID' (another long alphanumeric string). There is also a checkbox for 'Pull data from all org accounts'.
- Scheduling:** 'Full Refresh Frequency' (Custom) and 'Every' (10 Minutes).
- Remediation Detection Settings:** 'Aging criteria' (Age immediately if Finding was not seen, while Asset was seen in the latest data refresh) and 'Fallback' (Age immediately if Finding was not seen for 1 day(s)).
- Advanced Settings:** 'Suppression Rules' (Select Field, Contains, Type Value) and a checked checkbox for 'Prevent NULL from overriding existing values'.

 At the bottom right, the 'Save' button is highlighted with a red box.

Figure 259. Create AWS Security Hub API Source

Review and Adjust Data Model Mapping

(Optional) Zscaler UVM automatically maps ingested data to the default Data Model, so analysis can begin immediately. 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 Management Console.
2. Select **Services** > **EC2**.

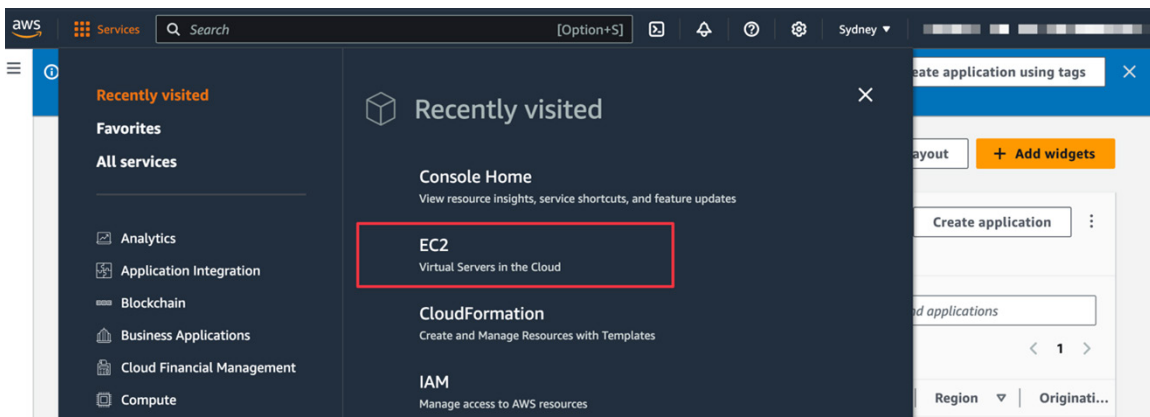


Figure 260. EC2

3. Click **Instances**.

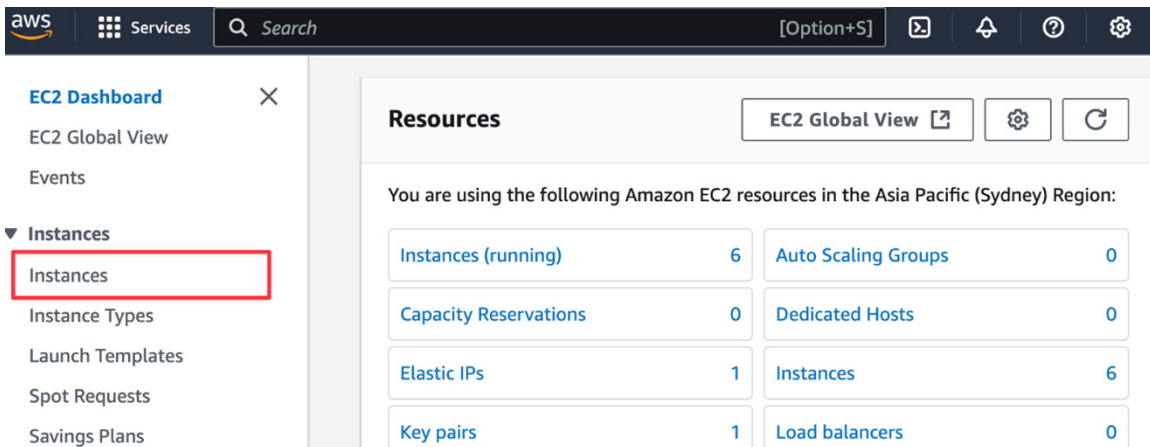


Figure 261. Instances

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

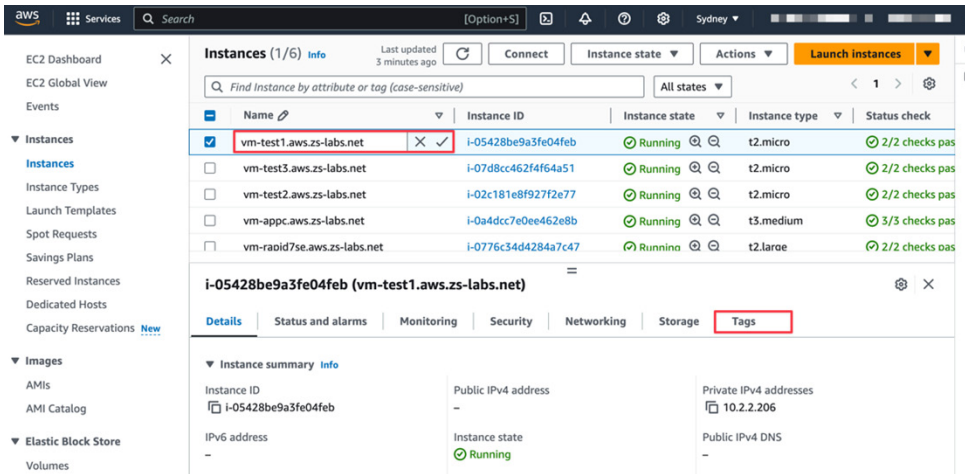


Figure 262. Tags

5. Click **Manage tags**.

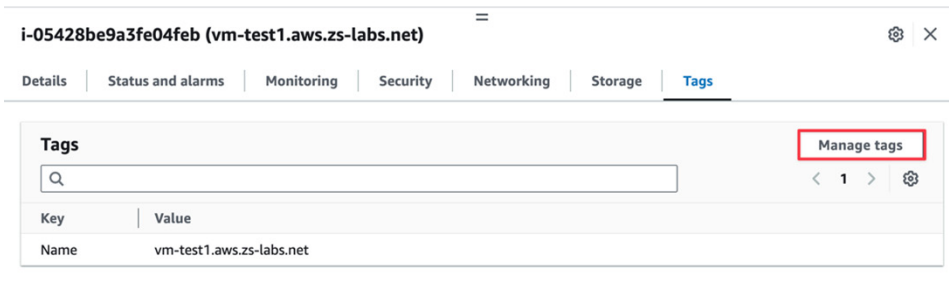


Figure 263. Manage tags

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

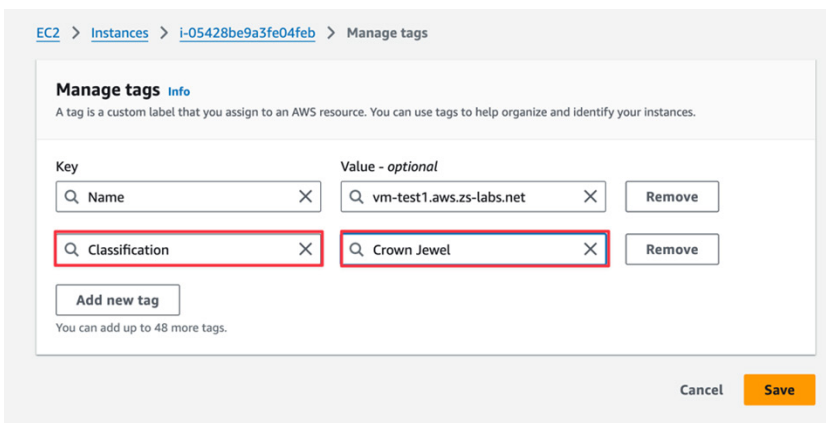


Figure 264. Manage tags

7. Click **Save**.

Map the AWS EC2 Data Source

To map the AWS EC2 data source:

1. Select **Configure** > <the newly created ZCC Devices Connector> > **Map Data**.

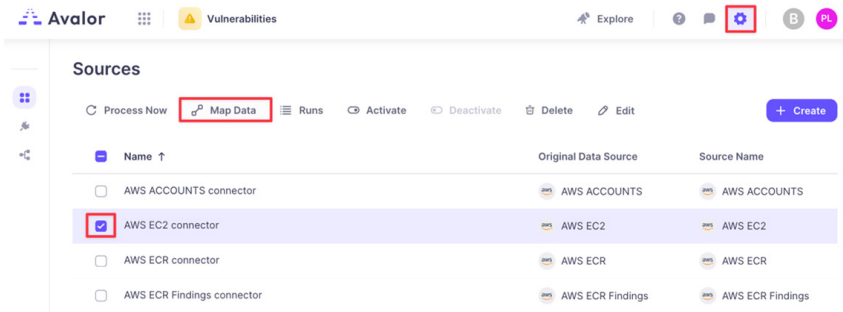


Figure 265. Map Data

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

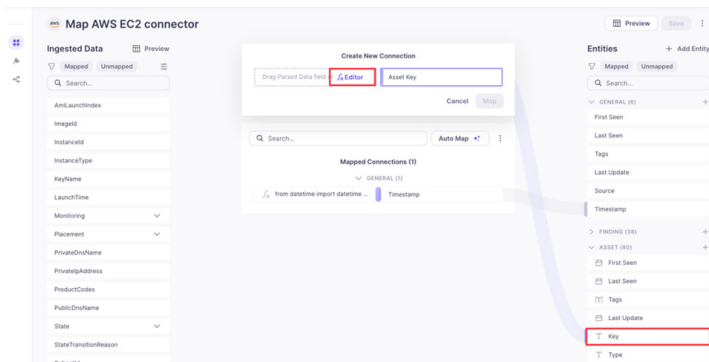


Figure 266. Asset Key

- c. 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)
```

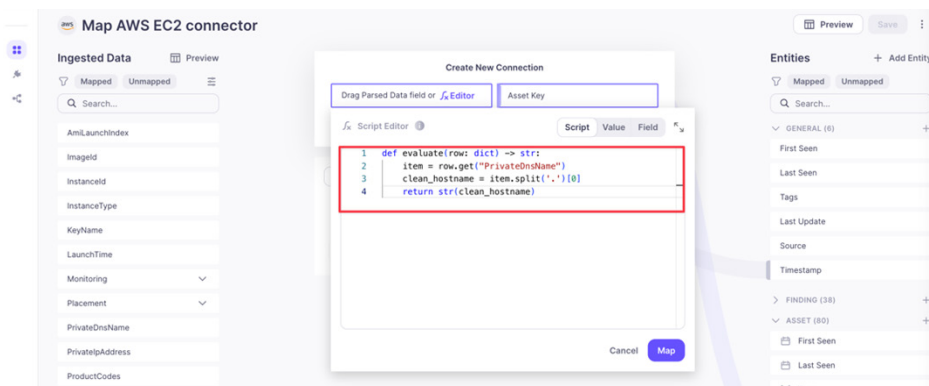


Figure 267. Script field

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

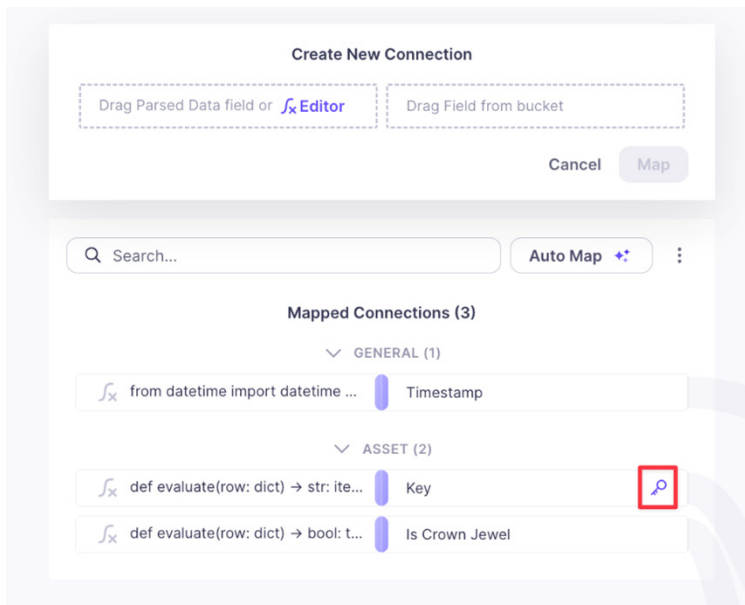


Figure 268. Asset Key

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

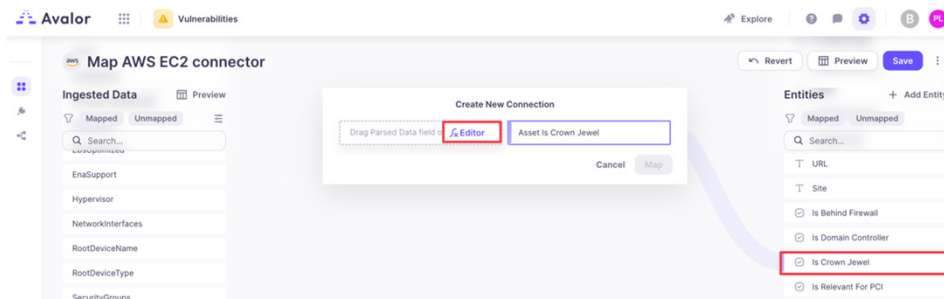


Figure 269. Editor element

- 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
```

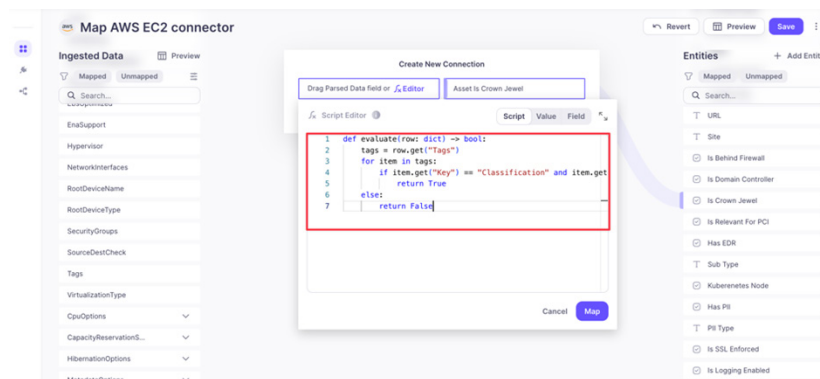


Figure 270. Script field

- Click **Map**.
- f. 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/io.avalor.prot...	true	2024-10-22T00:00:00Z	ip-10-2-2-206	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/io.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-2-8	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/io.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-153	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/io.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-33	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/io.avalor.prot...	false	2024-10-22T00:00:00Z	ip-10-2-1-247	["AWS EC2"]	2024-10-23T05:45:15Z
type.googleapis.com/io.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

Figure 271. Preview

- g. 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, Zscaler 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).

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

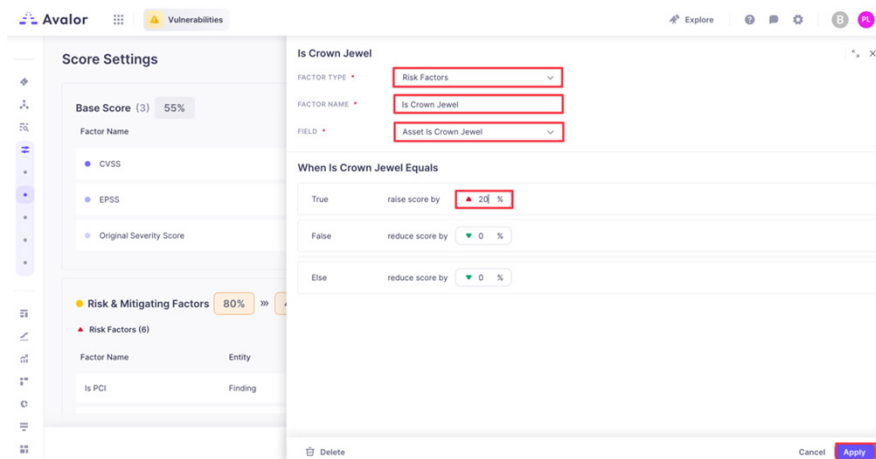


Figure 272. Boolean login section

- e. Click **Apply**, then **Save & Run**.
4. Click **Add Factor** in the **Risk & Mitigating Factors** section.
 5. If **Crown Jewel** is not already a **Risk Factor**, in the **Add new factor** modal, 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**.

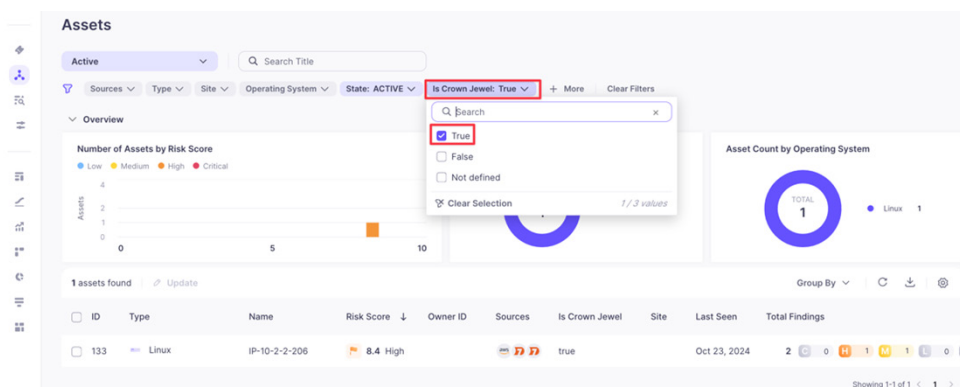


Figure 273. Assets dashboard

- 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).

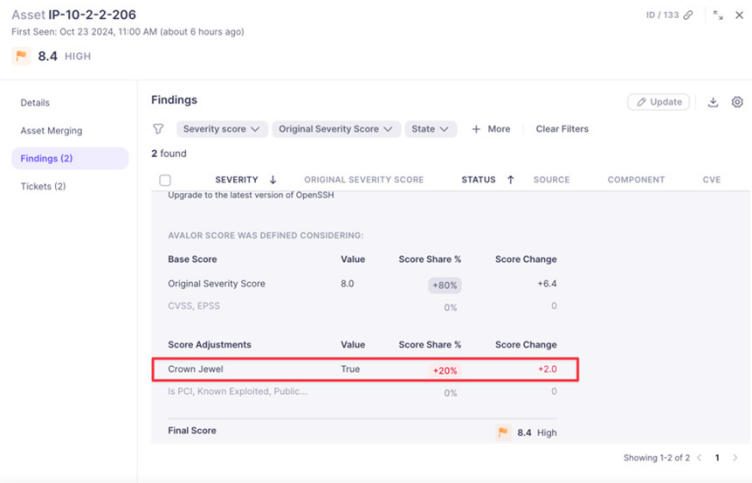


Figure 274. Findings tab

Appendix A: AWS Transit Gateway Lab Environment

The [AWS GitHub page](#) shows a diagram of a Transit Gateway (TGW) lab from an example AWS blog post, which includes a Cloud formation template. You can refer to this page as you test your Site-to-Site VPN Connection. You can find instructions on testing on the [AWS blog post](#).

The followign is a Transit Gateway lab environment diagram:

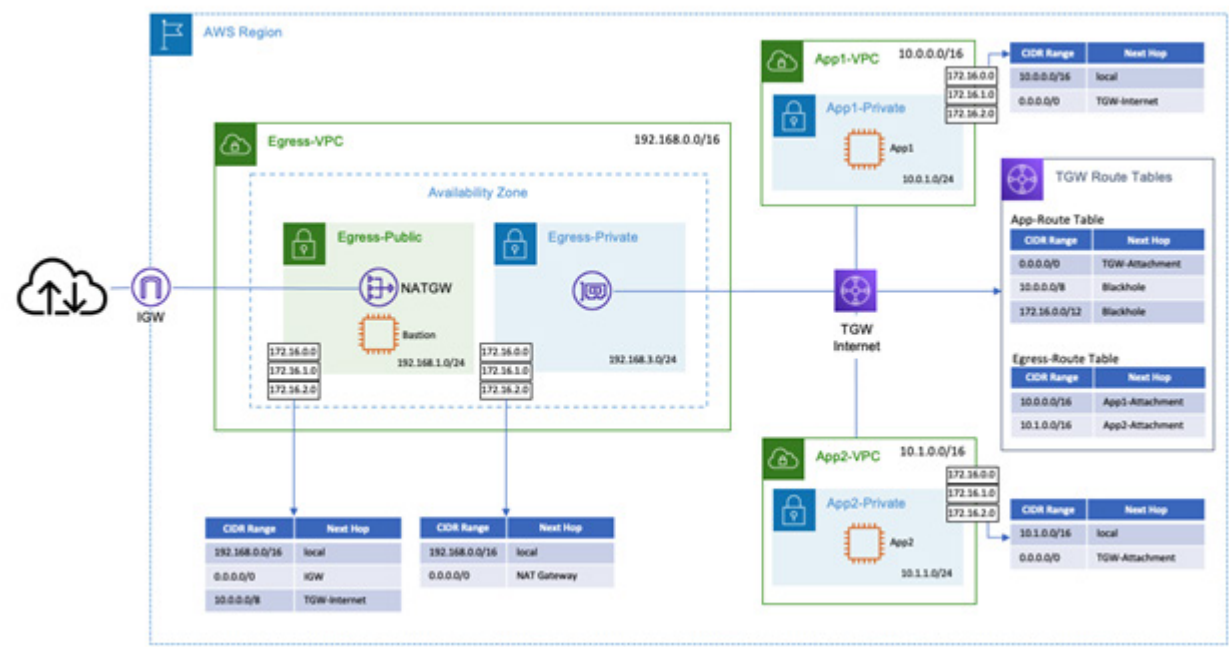


Figure 275. Example transit gateway lab environment diagram

Appendix B: Creating a Trail

To create a trail:

- 1. Go to **Services > CloudTrail > Trails by** and click **Create trail**.
- 2. Enter a name for the trail and either choose an existing S3 bucket to use or create a new S3 bucket. The Log file SSE-KMS encryption option is enabled by default. In this example, it is disabled.
- 3. Click **Next**.

S3-Test-TailDeleteStop logging

General detailsEdit

Trail logging

Logging

Trail log location

s3-test-logs/AWSLogs/008866442200

Log file validation

Enabled

SNS notification delivery

Disabled

Trail name

S3-Test-Tail

Last log file delivered

January 18, 2022, 15:30:09 (UTC-06:00)

Last file validation delivered

-

Last SNS notification

-

Multi-region trail

Yes

Log file SSE-KMS encryption

Not enabled

Apply trail to my organization

Not enabled

Figure 276. CloudTrail general details

- 4. Select the **Events > Event types** that you want to log, and the **Data event > Data event** type to use as the source. In the following example, it is **S3**.

▼ Data event: S3Remove

Data event type

Choose the source of data events to log.

S3▼

Log selector template

Log all events▼

Figure 277. Data event

5. Click **Next**. The following image shows a management event.

Choose log events

Events [Info](#)
Record API activity for individual resources, or for all current and future resources in AWS account. [Additional charges apply](#) [↗](#)


Event type
Choose the type of events that you want to log.

☒ **Management events**
Capture management operations performed on your AWS resources.

☒ **Data events**
Log the resource operations performed on or within a resource.

☒ **Insights events**
Identify unusual activity, errors, or user behavior in your account.

Management events [Info](#)
Management events show information about management operations performed on resources in your AWS account.

 No additional charges apply to log management events on this trail because this is your first copy of management events.

API activity
Choose the activities you want to log.

☒ **Read** ☒ **Write**

☐ Exclude AWS KMS events

☐ Exclude Amazon RDS Data API events

Figure 278. Management events

6. Review your input, then click **Create trail**.

Appendix C: Testing Notes

Configuring the Data at Rest scanning policy is documented in the [Understanding the Data at Rest Scanning Policy](#) (government agencies, see [Understanding the Data at Rest Scanning Policy](#)).

When configuring the Data Loss Prevention and the Malware Detection policy, you must select Public Cloud Storage at the top of each page to create a policy for your S3 SaaS application tenant. The following image shows a DLP scan.

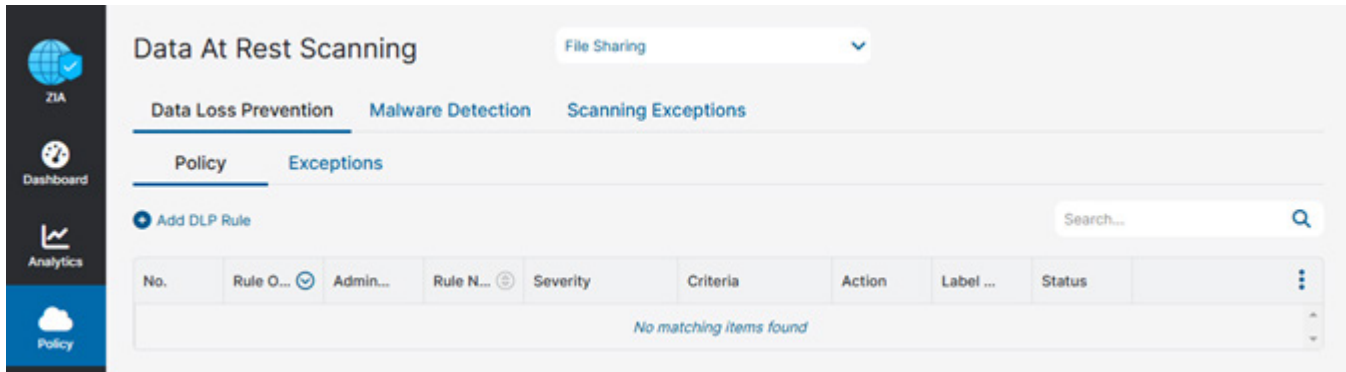


Figure 279. Data at Rest Scanning Policy – DLP

The following image shows a malware scan.

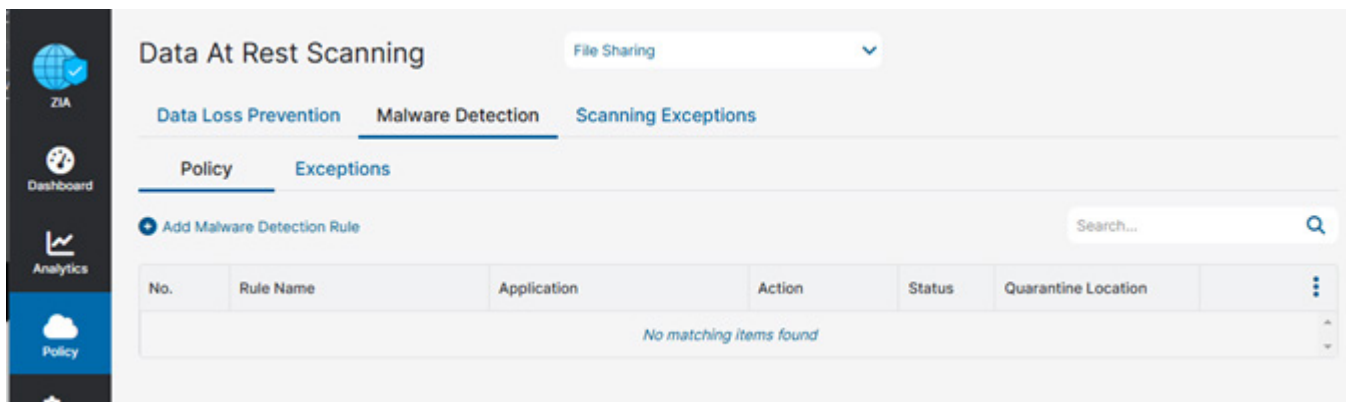


Figure 280. Data at Rest Scanning Policy – Malware Detection

You cannot select specific buckets for each of these policies until you have configured the Scan Schedule and selected all possible buckets to include. Then you can go back into the DLP and Malware policies (select Public Cloud Storage at the top again) to select specific buckets (if multiple buckets were selected in the Scan Schedule).

After you save the Scan Configuration, click Start. This changes the Status to Running.

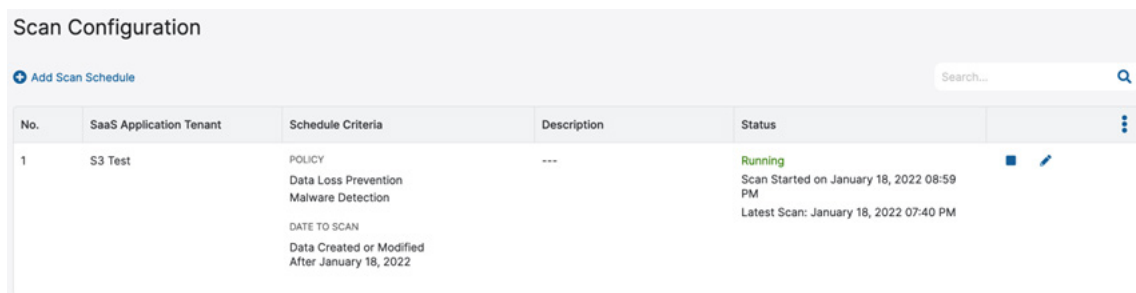


Figure 281. SaaS configuration

You can find information about DLP and Malware incidents in the following locations:

- Analytics > SaaS Assets Summary Report (see the following sample)
- Analytics > SaaS Security Report > Assets
- Analytics > SaaS Security Insights (see the following sample)

The following figure shows an SaaS assets summary report.

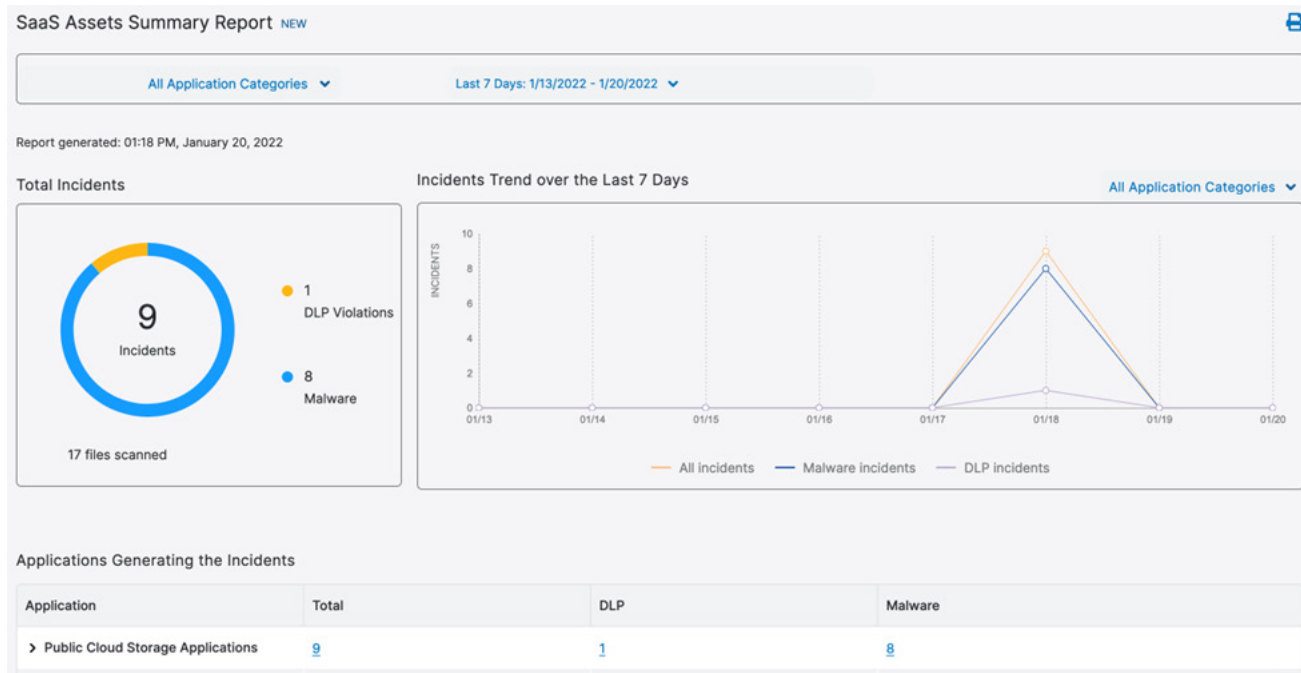


Figure 282. SaaS Assets Summary report

The following figure shows an SaaS Security Insights report.

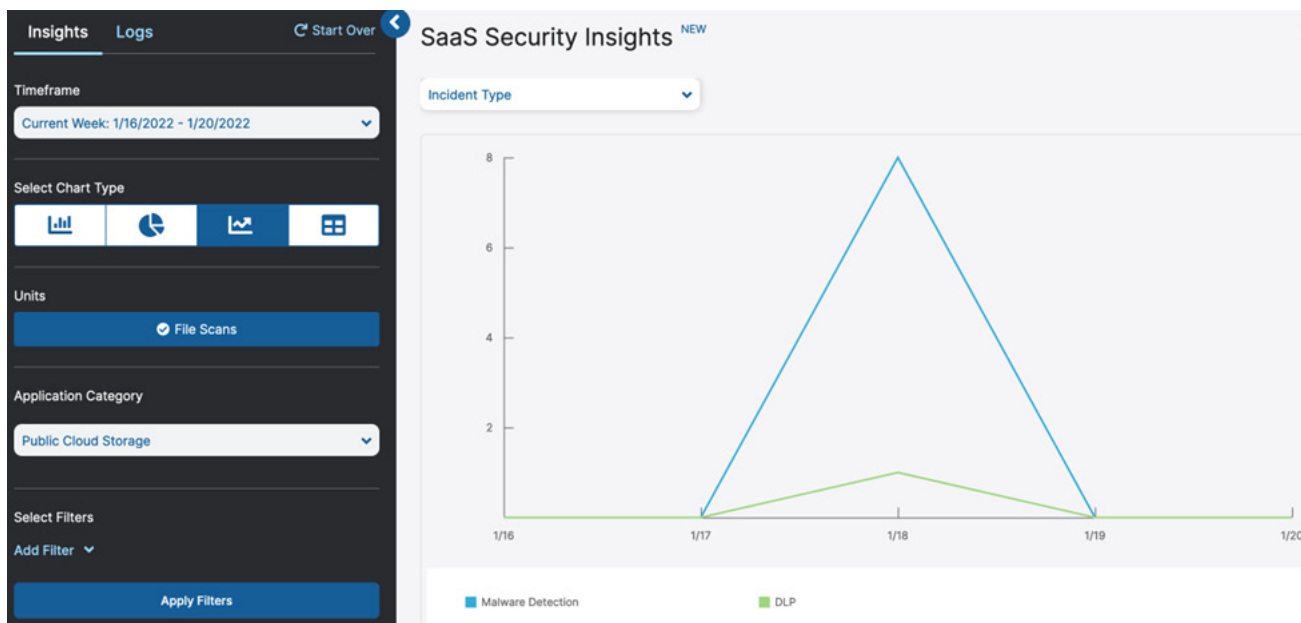


Figure 283. SaaS Security insights

Appendix D: AWS SSM Distributor

Distributor, a tool of AWS Systems Manager (SSM), allows you to package, store, and securely distribute software across AWS SSM-managed nodes, including EC2 instances and on-premises servers connected via hybrid activations. It simplifies the deployment and version management of custom and third-party software by handling packaging, permissions, and distribution workflows natively within AWS. Distributor is an effective method for deploying the Zscaler Microsegmentation agent to AWS EC2 instances, as it ensures secure, scalable, and consistent agent installation across multiple environments without manual intervention.

Prerequisites

Make sure the following prerequisites are met:

- The AWS SSM agent must be installed on the endpoints. For more information, refer to the [AWS documentation](#).
- An S3 bucket is required to host the deployment package. It is advisable to enable Versioning on the bucket. [Refer to this document](#) for more information.
- IAM permissions must be configured to allow SSM to manage EC2 instances. The Default Host Management Configuration option is sufficient. For more information, refer to the [AWS documentation](#).

Configuration

First, create the Installation Package:

1. Unzip the starter-package.zip file that was provided alongside this document.
2. Copy the appropriate Agent Provisioning Key from the ZPA Admin Portal. Paste this value into the provision_key file. It is advisable to do this using a command line text editor to make certain not to introduce file formatting characters.
3. Run both `./build-linux.sh` and `./build-windows.sh`. If operating on Windows, open these files with Notepad and manually perform the equivalent PowerShell commands.
4. Note the SHA256 hash output from each build script. Paste these values into the manifest.json file where indicated with `UPDATE_THIS_VALUE`. There is one value for the Windows zip file and one value for the Linux zip file.

Upload the Installation Package

To upload the installation package:

1. Verify that all prerequisites are met.
2. Upload the following files to the S3 bucket:
 - a. manifest.json
 - b. zscaler-microsegmentation-agent-linux.zip
 - c. zscaler-microsegmentation-agent-windows.zip

Configure AWS SSM Distributor

To configure the AWS SSM distributor:

1. Select **Create Package** in the **AWS SSM Distributor** UI.
2. Select **Advanced**.
3. Provide a **Name** and **Version**.
4. Provide the **S3 Bucket Name** and **S3 Key Prefix**.
 - a. The **Key Prefix** should be the name of the S3 folder where the installation package files are stored.
5. the **Manifest** setting as **Extract From Package**. Select **View Manifest File**. Verify there are no Warnings or Errors listed in the Manifest file viewer.
6. Select **Create Package**.

Deploy the Installation Package

AWS SSM provides two options for deploying a Distributor package: Install One Time and Install on a Schedule.

Install One Time

What follows describes a one-time installation.

1. Select **Install One Time** in the **AWS SSM Distributor Package Details** UI.
2. On the next screen, **Name** and other details should already be populated.
3. Configure the **Target Selection** section.
4. Optionally, specify a location to send logs in the **Output Options** section.
5. Select **Run**.
6. On the next screen, monitor the status and result of the **Run Command**.

Updates

Sometimes it is necessary to update the Installation Package (such as with a new Provisioning Key). To update the provision_key file or to make similar modifications to the contents of the installation package:

1. Make the required changes.
2. Delete the existing zscaler-microsegmentation-agent-*.zip files.
3. Run the build-*.sh scripts.
4. Update the manifest.json file with the new SHA256 values.
5. Upload the new zscaler-microsegmentation-agent-*.zip files and manifest.json file to the S3 bucket, overwriting the previous files.
6. In the **AWS SSM Distributor Package Details** UI, select the submenu option **Versions**.
7. Select **Add Version** and complete the wizard.
8. Select the radio button that matches the new version and then select **Set Default Version**. Subsequent deployments of the installation package use the latest version unless specified otherwise on the **Run Command** window.

Appendix E: Metrics Config Options

The following table shows the possible values for various metrics described in the CloudWatch use case discussed previously.

Detail Level	Metrics Included
Basic	<p>Mem: mem_used_percent</p> <p>Disk: disk_used_percent</p> <p>The disk metrics such as disk_used_percent have a dimension for Partition, which means that the number of custom metrics generated is dependent on the number of partitions associated with your instance. The number of disk partitions you have depends on which AMI you are using and the number of Amazon EBS volumes you attach to the server.</p>
Standard	<p>CPU: cpu_usage_idle, cpu_usage_iowait, cpu_usage_user, cpu_usage_system</p> <p>Disk: disk_used_percent, disk_inodes_free</p> <p>Diskio: diskio_io_time</p> <p>Mem: mem_used_percent</p> <p>Swap: swap_used_percent</p>
Advanced	<p>CPU: cpu_usage_idle, cpu_usage_iowait, cpu_usage_user, cpu_usage_system</p> <p>Disk: disk_used_percent, disk_inodes_free</p> <p>Diskio: diskio_io_time, diskio_write_bytes, diskio_read_bytes, diskio_writes, diskio_reads</p> <p>Mem: mem_used_percent</p> <p>Netstat: netstat_tcp_established, netstat_tcp_time_wait</p> <p>Swap: swap_used_percent</p>

Appendix F: ZPA and ZIA Configuration for Private AI Data Protection

Private AI data protection is only required if you plan to have AI frontends accessible in private networks, such as a private VPC in AWS. This requires you to properly configure and run a Source IP Anchor and ensure the following conditions are met:

- Configure App Connectors with a private IP address and NAT'd to a public IP address.
- Do not expose the App Connector by configuring a public IP address directly on the App Connector interface.
- App Connector's public IP address is the anchored source IP address.
- Ensure the firewall allows the App Connector to reach the destination server.

Source IP Anchoring uses ZIA forwarding policies and ZPA App Connectors to selectively forward the application traffic to the appropriate destination servers. You can configure forwarding rules in the ZIA Admin Portal to forward Source IP Anchored traffic to ZPA through ZIA threat and data protection engines.

The following diagram shows a typical design pattern for protecting private Generative AI and Zscaler configured and deployed to protect Gen AI.

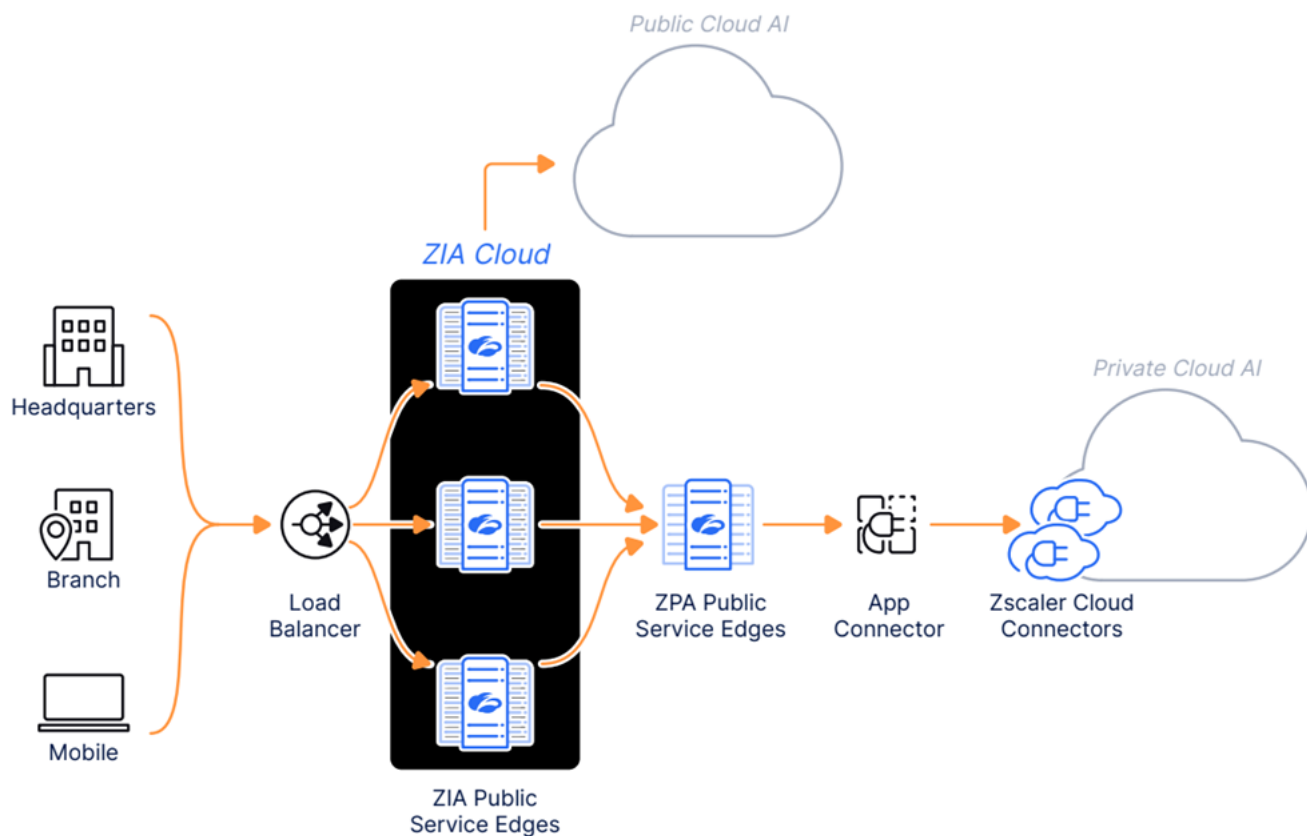
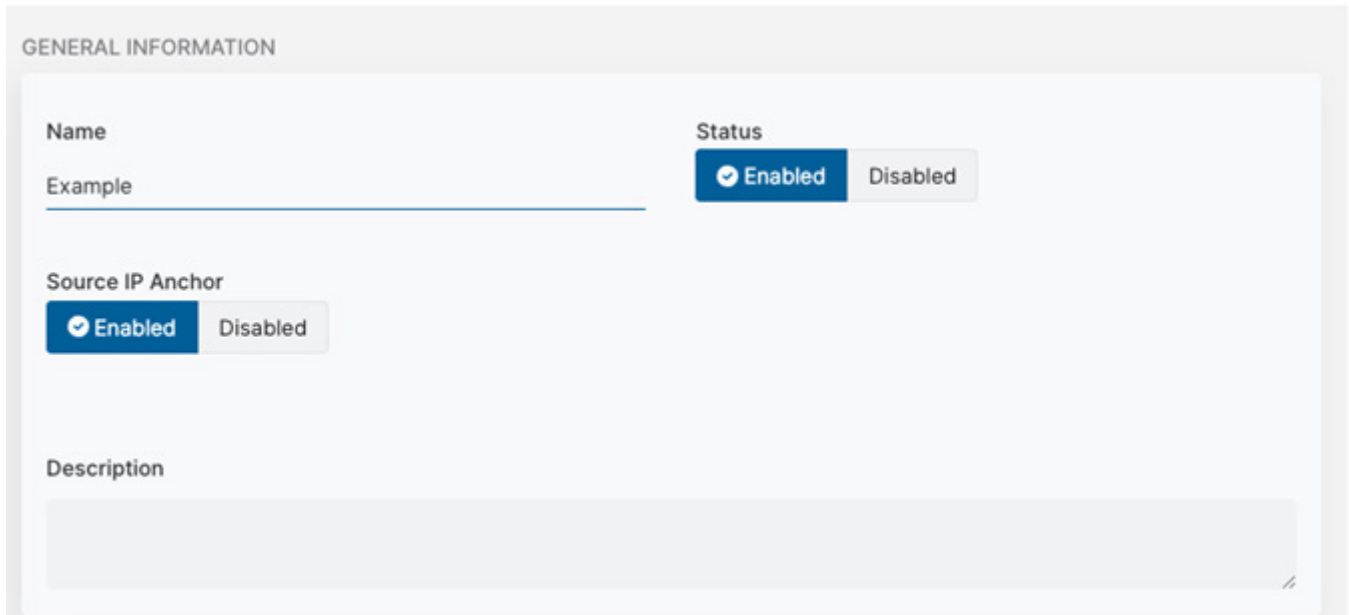


Figure 284. ZIA and ZPA Data Protection

Configure Application Segment

Create and configure an application segment that uses Source IP Anchoring. Ensure you enable the Source IP Anchor option and select Use Client Forwarding Policy under the Bypass field while configuring the application segment.

1. Log in to the ZPA Admin Portal.
2. Go to **Resource Management > Application Management > Application Segment**.
3. Click the **Ellipsis** in the right-side of the window and click **Add Application Segment**.
4. Enter a **Segment Name** in the **Name** field.
5. Verify **Status** is set to **Enabled**.
6. Set **SourceIP Anchor** to **Enabled**.



GENERAL INFORMATION

Name
Example

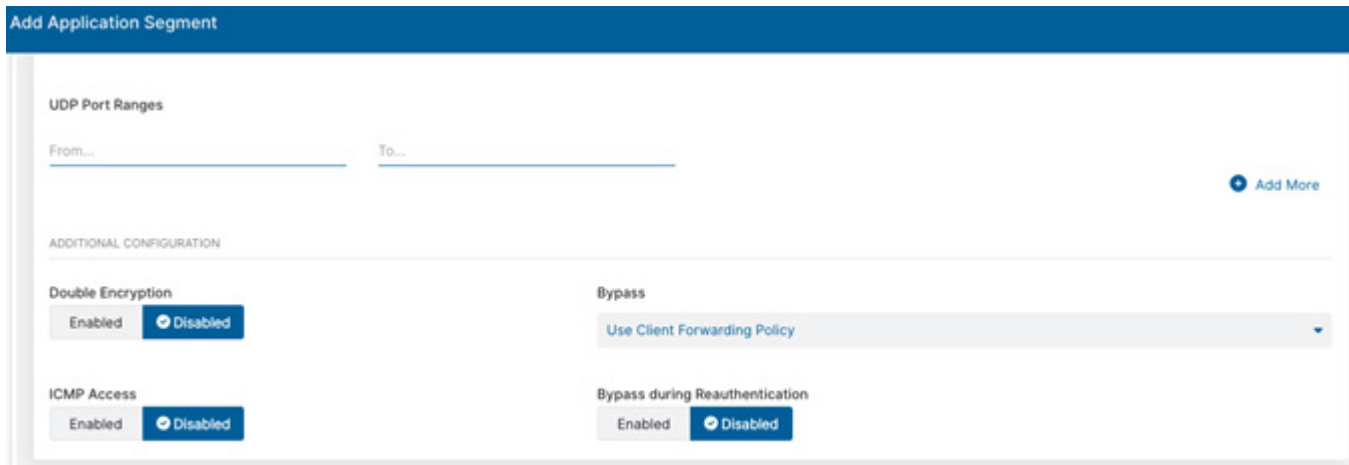
Status
☒ Enabled
 ☐ Disabled

Source IP Anchor
☒ Enabled
 ☐ Disabled

Description

Figure 285. Add Application Segment Source IP Anchor

7. In **Additional Configuration** set **Bypass** to **Use Client Forwarding Policy**.
8. Complete the remaining configuration steps for **Application Segment** configuration by following [Configuring Defined Application Segments](#) (government agencies, see [Configuring Defined Application Segments](#)).
9. Click **Save**.



Add Application Segment

UDP Port Ranges
 From... To...
 + Add More

ADDITIONAL CONFIGURATION

Double Encryption
☐ Enabled
 ☒ Disabled

Bypass
 Use Client Forwarding Policy

ICMP Access
☐ Enabled
 ☒ Disabled

Bypass during Reauthentication
☐ Enabled
 ☒ Disabled

Figure 286. Add Application Segment Bypass

Configure ZPA Client Forwarding Policy

Configure a client forwarding policy for the application segment. Create separate client forwarding policy rules for IP address-based and domain-based applications.

For IP address-based applications, select the Only Forward Allowed Applications rule action for Source IP Anchoring application segments. For IP address-based applications, configure the following rule:

1. In the ZPA Admin Portal, go to **Policy > Client Forwarding Policy**.
2. Click **Add Rule**.
3. Enter the **Client Forwarding Policy Name** in the **Name** field.
4. (Optional) Enter the **Client Forwarding Policy Description** in the **Description** field.
5. In **Rule Action**, select **Only Forward Allowed Applications**.
6. Click **Add Criteria**:
 - a. Select **Applications**.
 - b. In the **Application Segments** drop-down menu, select the **SIPA Application Segment** created in [Configure Application Segment](#).
7. Click **Save**.

Edit Client Forwarding Policy [X]

Name
[Text Field]

Description
[Text Field]

ACTION

Rule Action

Forward to ZPA **Only Forward Allowed Applications** Bypass ZPA

CRITERIA

Application Segments [Dropdown Menu] [Add Criteria]

OR

Segment Groups
[Select one or more segment groups]

Save Cancel

Figure 287. Add Client Forwarding Policy—Only Forward Allowed Applications

For domain-based applications, configure the following two rules:

- **Rule 1:** Select the **Bypass ZPA** rule action for **Source IP Anchoring Segment Groups** and **Client Types > Client Connector** as described in [Rule 1: Enable the Bypass ZPA Rule Action](#).
- **Rule 2:** Select the **Forward to ZPA** rule action for **Source IP Anchoring Segment Groups** and **Client Types > ZIA Public Service Edge** as described in [Rule 2: Enable the Forward to ZPA Rule Action](#).

Rule 1: Enable the Bypass ZPA Rule Action

1. Go to **Policy > Client Forwarding Policy**.
2. Click **Add Rule**.
3. Enter the **Client Forwarding Policy Name** in the **Name** field.
4. (Optional) Enter the **Client Forwarding Policy Description** in the **Description** field.
5. In **Rule Action**, select **Bypass ZPA**.
6. Click **Add Criteria**.
 - a. Select **Applications**.
 - b. In the **Application Segments** drop-down menu, select the **SIPA Application Segment** created in **Configure Application Segment** earlier.
7. Click **Add Criteria**:
 - a. Select **Client Types**.
 - b. In the **Client Types** drop-down menu, select **Client Connector**.
8. Click **Save**.

Figure 288. Add Client Forwarding Policy—Bypass ZPA

Rule 2: Enable the Forward to ZPA Rule Action

1. Go to **Policy > Client Forwarding Policy**.
2. Click **Add Rule**.
3. Enter the **Client Forwarding Policy Name** in the **Name** field.
4. (Optional) Enter the **Client Forwarding Policy Description** in the **Description** field.
5. In **Rule Action**, select **Forward to ZPA**.
6. Click **Add Criteria**.
 - a. Select **Applications**.
 - b. In the **Application Segments** drop-down menu, select the **SIPA Application Segment** created in [Configure Application Segment](#).
7. Click **Add Criteria**.
 - a. Select **Client Types**.
 - b. In the **Client Types** drop-down menu, select **ZPA Service Edge**.
8. Click **Save**.

Edit Client Forwarding Policy [X]

Name

Description

ACTION

Rule Action

Forward to ZPA | Only Forward Allowed Applications | Bypass ZPA

CRITERIA

Application Segments [X] [Add Criteria]

OR

Segment Groups

Select one or more segment groups

AND

Client Types [X]

ZPA Service Edge

Save Cancel

Figure 289. Add Client Forwarding Policy—Forward to ZPA

Configure ZPA Access Policy

The following steps create and configure an access policy for the application segment. Create separate access policy rules for IP address-based and domain-based applications.

For IP address-based applications, configure the following rule:

1. Select the **Allow Access** rule action and add only the **ZIA Public Service Edge** client type for the application segments.
2. For domain-based applications, allow the **Source IP Anchoring** client (**ZIA Public Service Edge** client type) to access the applications.

For IP Address-Based Applications

Configure the following rules. Allow Access rule action and add only the ZIA Public Service Edge.

1. Go to **Policy > Access Policy**.
2. Click **Add Rule**.
3. Enter the **Access Policy Name** in the **Name** field.
4. (Optional) Enter the **Access Policy Description** in the **Description** field.
5. In **Rule Action**, select **Allow Access**.
6. Click **Add Criteria**:
 - a. Select **Applications**.
 - b. In the **Application Segments** drop-down menu, select the **SIPA Application Segment** created in [Configure Application Segment](#).
7. Click **Add Criteria**.
8. Select **Client Types**.
9. In the **Client Types** drop-down menu, select **ZIA Public Service Edge**.
10. Click **Save**.

Figure 290. Add Access Policy—Allow Access Public Service Edge Only

For Domain-Based Applications

Configure the following rule:

1. Go to **Policy > Access Policy**.
2. Click **Add Rule**.
3. Enter the **Access Policy Name** in the **Name** field.
4. (Optional) Enter the **Access Policy Description** in the **Description** field.
5. In **Rule Action**, select **Allow Access**.
6. Click **Add Criteria**:
 - a. Select **Applications**.
 - b. In the **Application Segments** drop-down menu, select the **SIPA Application Segment** created in [Configure Application Segment](#).
7. Click **Add Criteria**:
 - a. Select **Client Types**.
 - b. In the **Client Types** drop-down menu, select **ZPA Service Edge**.

8. Click **Save**.

Figure 291. Add Access Policy—Allow Access Public Service Edge Only Domain Access policy



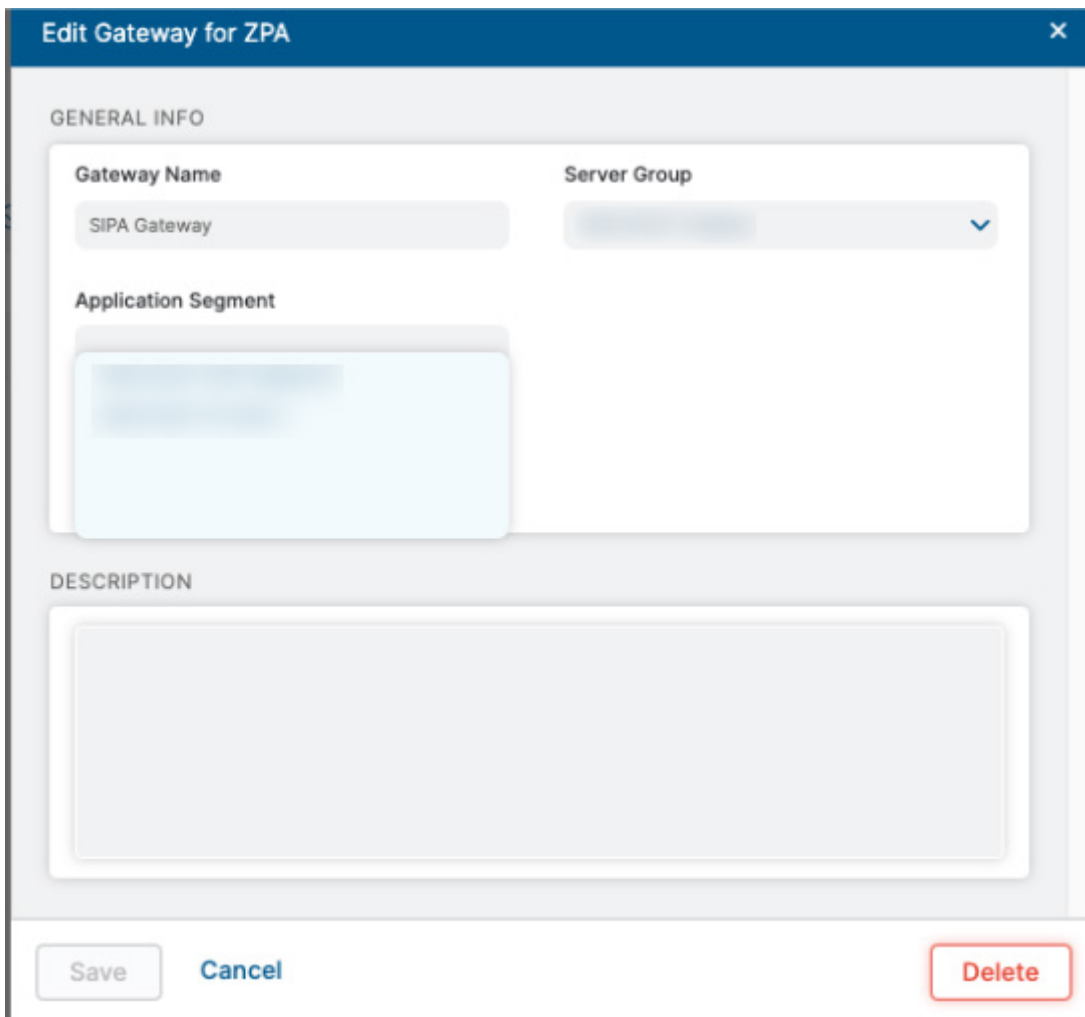
To learn more, see [Configuring Client Forwarding Policies](#) (government agencies, see [Configuring Client Forwarding Policies](#)).

If configuration is required for source IP direct for disaster recovery mode, see [Understanding Source IP Anchoring Direct](#) (government agencies, see [Understanding Source IP Anchoring Direct](#)).

Configure ZPA Gateway

Configure ZPA gateways on the ZIA Admin Portal to map it to the ZPA server groups and the associated application segments that require Source IP Anchoring.

1. Log in to the ZIA Admin Portal.
2. Go to **Administration > Zscaler Private Access**.
3. Click **Add Gateway for ZPA**. The **Add Gateway for ZPA** window appears.
4. From the **Server Group** drop-down menu, select the server group that you configured in ZPA for Source IP Anchoring. All the application segments that are associated with the selected server group for which Source IP Anchoring is enabled appear in the **Application Segment** field.
5. Click **Save** and **Activate** the changes.



The screenshot shows the 'Edit Gateway for ZPA' window. The 'GENERAL INFO' section contains the following fields:

- Gateway Name:** SIPA Gateway
- Server Group:** A dropdown menu with a blue arrow icon.
- Application Segment:** A list of application segments, with one segment highlighted in blue.

The 'DESCRIPTION' section contains a large text area for additional information.

At the bottom of the window, there are three buttons: 'Save', 'Cancel', and 'Delete'.

Figure 292. Add Gateway for ZPA



To learn more, see [Saving and Activating Changes in the ZIA Admin Portal](#) (government agencies, see [Saving and Activating Changes in the ZIA Admin Portal](#)).

Configure Forwarding Policy for ZPA

Zscaler uses forwarding control policies to forward selective Zscaler traffic to specific endpoints. For example, if you want to forward web traffic to a third-party proxy service or if you want to forward application traffic to a ZPA App Connector, you can configure your forwarding policy with appropriate rules.

The following steps configure forwarding policies to forward ZIA traffic for Source IP Anchoring. Zscaler provides a predefined forwarding rule, ZIA Inspected ZPA Apps, enabled by default. This rule forwards all ZPA application segment traffic for ZIA inspection that has the Inspect Traffic with ZIA field enabled in the ZPA Admin Portal. You cannot edit this rule.

1. Go to **Policy > Forwarding Control**.
2. Click **Add Forwarding Rule**. The **Add Forwarding Rule** window appears.
3. Under the **Forwarding Rule Section**, configure the following attributes.
 - a. **Rule Order**: Enter the order of the rule. Policy rules are evaluated in ascending numerical order (Rule 1 before Rule 2, and so on), and the Rule Order reflects this rule's place in the order. You can change the value based on your requirements. However, if you've enabled Admin Rank, your assigned admin rank determines the Rule Order values you can select.
 - b. **Rule Name**: Enter a user-friendly name for the rule. The **Forwarding Control** automatically creates a rule name, which you can change. The maximum length is 31 characters.
 - c. **Forwarding Method**: Select **ZPA**.
 - d. Under **Action Forward to ZPA Gateway** select the gateway created in Configure ZPA Gateway.
 - e. Under **Criteria select Destination and Application Segment** select the application segment created in Configure Application Segment.
4. Click **Save** and **Activate** the changes.

Edit Forwarding Rule

FORWARDING RULE

Rule Order: 1

Rule Name: SIPA Forwarding

Rule Status: Enabled

Rule Label: ---

Forwarding Method: ZPA

CRITERIA

General Applications Source **Destination**

Application Segment: ...

ACTION

Forward to ZPA Gateway: SIPA Gateway

DESCRIPTION

Save Cancel Delete

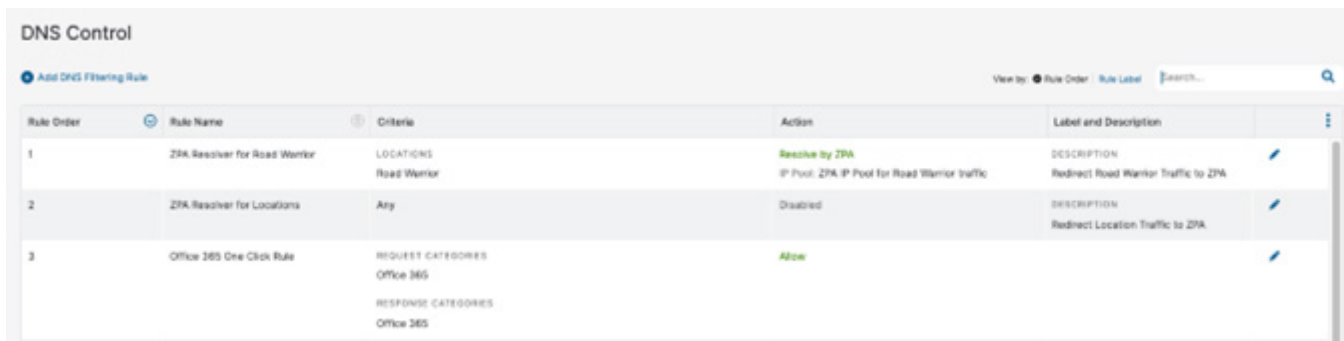
Figure 293. Forwarding Control

Configure DNS Control

Configure Source IP Anchoring for all traffic forwarded to the ZIA Admin Portal, enable the appropriate preconfigured DNS filtering rule.

1. Go to **Policy > DNS Control**.
2. For location users, enable the **ZPA Resolver for Locations** rule.
3. For remote users, enable the **ZPA Resolver for Road Warrior** rule.

Ensure that these DNS rules are the top rules (i.e., Rule 1 and Rule 2) to configure Source IP Anchoring. The DNS rules are associated with the respective preconfigured IP pools under Administration > IP & FQDN Groups > IP Pool. You can edit the IP pools based on your needs. To learn more, see [About IP Pool](#) (government agencies, see [About IP Pool](#)). Any change in the IP pool is reflected in the Action column of the respective DNS rule when the rule is enabled.



The screenshot shows the 'DNS Control' interface with a table of DNS filtering rules. The table has columns for Rule Order, Rule Name, Criteria, Action, and Label and Description. There are three rules listed:

Rule Order	Rule Name	Criteria	Action	Label and Description
1	ZPA Resolver for Road Warrior	LOCATIONS Road Warrior	Resolve by ZPA IP Pool: ZPA IP Pool for Road Warrior traffic	DESCRIPTION Redirect Road Warrior Traffic to ZPA
2	ZPA Resolver for Locations	Any	Disabled	DESCRIPTION Redirect Location Traffic to ZPA
3	Office 365 One Click Rule	REQUEST CATEGORIES Office 365 RESPONSE CATEGORIES Office 365	Allow	

Figure 294. DNS Control



When the ZPA Resolver for Road Warrior rule is disabled, the remote user traffic automatically falls under the ZPA Resolver for Locations rule instead of blocking the traffic. Therefore, Zscaler does not recommend disabling the ZPA Resolver for the Road Warrior rule.

To support Source IP Anchoring for Zscaler Tunnel (Z-Tunnel) 1.0 traffic, you must enable the Enable Firewall for Z-Tunnel 1.0 and PAC Road Warriors option under Administration > Advanced Settings.

Zscaler also recommends having open firewall rules for the Source IP Anchoring pools while sending DNS traffic to the Zscaler service for the Source IP Anchoring domains (i.e., set the Action column on the Firewall Filtering policy to Allow for the Source IP Anchoring pools).

Appendix G: 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.

To contact Zscaler Support:

1. Go to **Administration > Settings > Company Profile**.

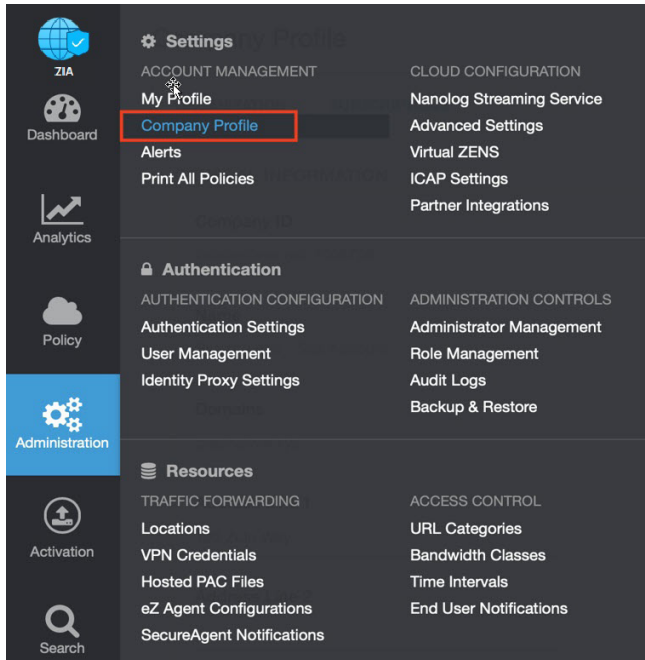


Figure 295. Collecting details to open support case with Zscaler TAC

2. Copy your **Company ID**.

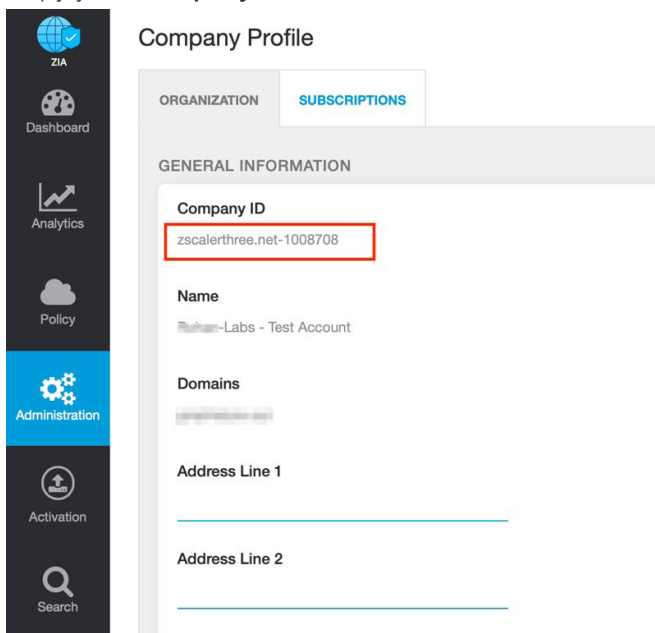


Figure 296. Company ID

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

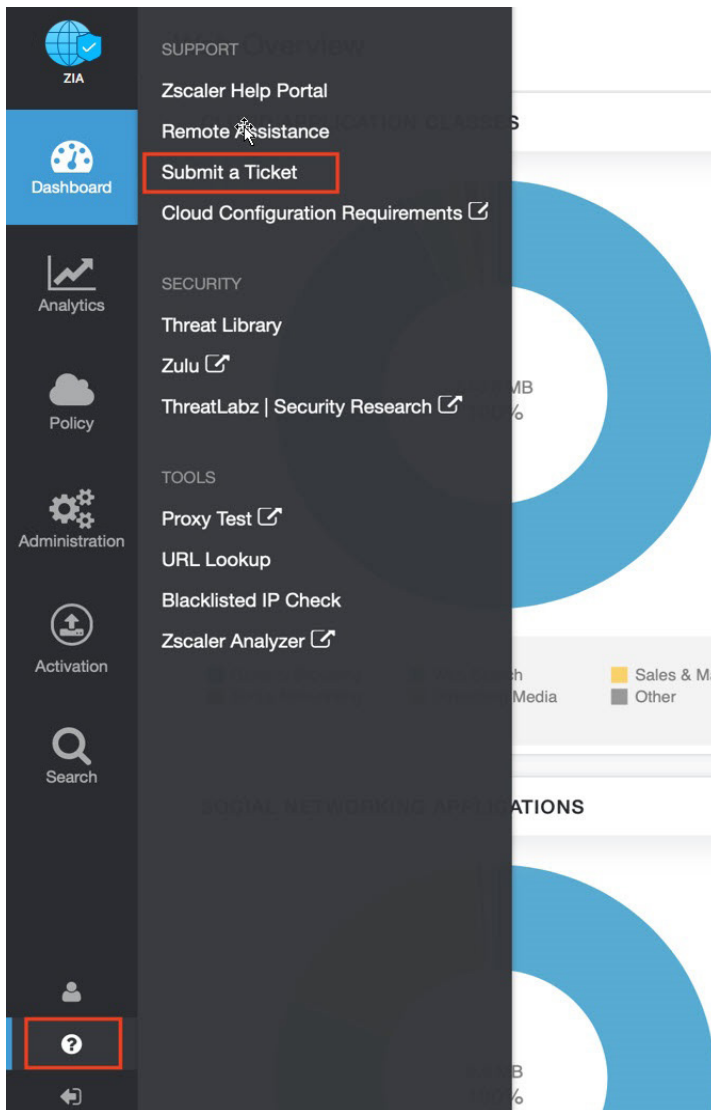


Figure 297. Submit a ticket