ZSCALER AND CROWDSTRIKE DEPLOYMENT GUIDE

OCTOBER 2023, VERSION 1.13

BUSINESS DEVELOPMENT GUIDE

Contents

Terms and Acronyms

The following terms and acronyms are used in this document. When applicable, a Request for Change (RFC) is included in the Definition column for your reference.

Acronym	Definition
API	Application Programming Interface
CA	Central Authority (Zscaler)
CLI	Command Line Interface
CSV	Comma-Separated Values
CVE	Common Vulnerabilities and Exposures
DLP	Data Loss Prevention
DNS	Domain Name Service
DPD	Dead Peer Detection (RFC 3706)
EDR	Endpoint Detection and Response
GRE	Generic Routing Encapsulation (RFC2890)
HEC	HTML Ethernet Channel
IKE	Internet Key Exchange (RFC2409)
IoC	Indicator of Compromise
IPS	Intrusion Prevention System
IPSec	Internet Protocol Security (RFC2411)
NGAV	Next-Generation Antivirus
NSS	Nanolog Streaming Service
PFS	Perfect Forward Secrecy
PSK	Pre-Shared Key
RFM	Reduced Functionality Mode (CrowdStrike)
SaaS	Software as a Service
SIEM	Security Information and Event Management
SSL	Secure Socket Layer (RFC6101)
XFF	X-Forwarded-For (RFC7239)
ZIA	Zscaler Internet Access (Zscaler)
ZPA	Zscaler Private Access (Zscaler)
ZTA	Zero Trust Assessment

Trademark Notice

About This Document

This document provides information on how to configure Zscaler and CrowdStrike for deployment.

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. Zscaler Internet Access (ZIA) and Zscaler Private Access (ZPA) 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 the Zscaler website or follow us on Twitter @zscaler.

CrowdStrike Overview

CrowdStrike (NASDAQ: CRWD), is a leading cybersecurity company protecting customers from all cyber threats by leveraging its Security Cloud to stop breaches. From its inception in 2011, driven by George Kurtz's vision, CrowdStrike was created as a different kind of cybersecurity company. Cloud-native, CrowdStrike immediately brought a threat perspective, effectiveness, scalability, and flexibility never seen before in the industry—seamlessly aligning People, Technology, and Processes. The CrowdStrike Falcon platform has revolutionized enterprise security for the cloud era. Its single lightweight-agent architecture leverages artificial intelligence (AI) and offers real-time protection and visibility across the enterprise, preventing attacks on endpoints and workloads on or off the network. To learn more, refer to CrowdStrike'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, refer to:

- Appendix A: Requesting Zscaler Support
- Zscaler Resources
- CrowdStrike Resources

Software Versions

This document was authored using ZIA and ZPA (with Zscaler Client Connector) along with CrowdStrike Falcon Agent 6.18.13211 on Windows 10.

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 CrowdStrike Introduction

Overviews of the Zscaler and CrowdStrike 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, please contact your Zscaler Account Representative.

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:

- Set up a tunnel (GRE or IPSec) to the closest Zscaler data center (for offices).
- Forward 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 Browser 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 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 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.
ZPA and CrowdStrike ZTA integration	Blog on the benefits of ZPA and CrowdStrike integration.

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 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.
ZPA and CrowdStrike ZTA integration	Blog on the benefits of ZPA and CrowdStrike integration.

CrowdStrike Falcon Endpoint Protection Enterprise Platform Overview

CrowdStrike Falcon Endpoint Protection Enterprise Platform sets the new standard with the first cloud-native security platform that delivers an endpoint breach prevention solution. It is the only endpoint breach prevention solution that unifies NGAV, EDR, managed threat hunting, and threat intelligence automation in a single cloud-delivered agent.

CrowdStrike Zero Trust Assessment Overview

CrowdStrike Zero Trust Assessment (ZTA) delivers real-time security posture assessments across all endpoints regardless of location, network, and user. CrowdStrike ZTA enables enforcement of dynamic conditional access based on device health and compliance checks that mitigate the risk to users and the organization. Every endpoint is granted least-privileged access and is assessed before gaining access to sensitive data and corporate assets—ensuring Zero Trust enforcement across all endpoints. By expanding Zero Trust beyond authentication and including device security, CrowdStrike ZTA helps organizations maintain a holistic cybersecurity approach that protects their data and users from the sophisticated tactics of cyber adversaries.

CrowdStrike Resources

The following table contains links to CrowdStrike support resources.

Name	Definition
CrowdStrike Falcon Admin Portal	Link to CrowdStrike administration portal.
CrowdStrike Support Portal	CrowdStrike support portal for submitting requests and issues.
CrowdStrike Documentation	Link to all CrowdStrike online documentation.
CrowdStrike ZTA Documentation	Documentation for CrowdStrike ZTA.
CrowdStrike ZTA Demo	Link to a video demonstration of CrowdStrike ZTA.
CrowdStrike Falcon IoC Sharing integration with ZIA	GitHub repository of examples that show CrowdStrike and Zscaler integration.

Use Case 1: ZPA Posture Check Integration with CrowdStrike ZTA

In this use case:

- CrowdStrike calculates a ZTA security score from 1 to 100 for each host. A higher score indicates a better security posture for the host. Security scores are derived from two distinct assessment sources:
  - OS settings: Settings that track native OS security options, firmware availability, and Common Vulnerabilities and Exposures (CVE) mitigations.
  - CrowdStrike Falcon sensor settings (Windows and macOS): CrowdStrike Falcon sensor configurations that track Reduced Functionality Mode (RFM) status as well as prevention and Real-Time Response policies.
- ZPA uses CrowdStrike's ZTA score (also known as a device posture score) and allows only compliant endpoints to access selected applications. ZPA checks for any changes to the CrowdStrike's device posture score because the score can change over time. ZTA check is supported currently by CrowdStrike for Windows and macOS endpoints.
- ZPA achieves conditional access by evaluating ZPA access policies that, in turn, reference device level posture check profiles. ZPA administrators can specify that a minimum ZTA score is needed for the endpoint to grant access to internal applications that are referenced in the ZPA access policy. The end-device's ZTA score must be greater than or equal to the threshold referenced in ZPA access policy (via posture check profile). Otherwise, ZPA blocks the application access from that host.

Currently, customers who want to use this integration must contact the CrowdStrike Support team to turn on a feature flag in their CrowdStrike tenant. When enabled on the CrowdStrike back end, ZPA accesses and uses the per-device ZTA score.

You can reach the CrowdStrike Support team at support@crowdstrike.com. Any device trying to access applications over ZPA must be running Zscaler Client Connector version 3.4 or later for this integration to work. The CrowdStrike sensor must be version 6.20 or later.

The following diagram shows a conceptualization of the integration.

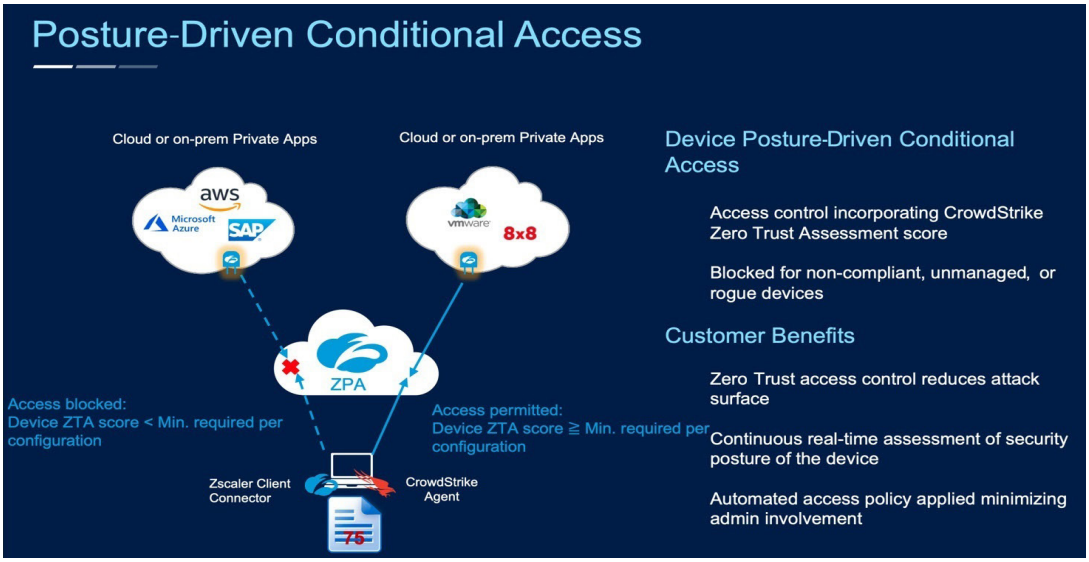

Figure 1. High-Level overview

If interoperability issues arise, update your CrowdStrike policy's sensor visibility exclusions to account for Zscaler Client Connector-related details and re-test. To learn more, see Zscaler Client Connector Processes to Allowlist.

Configuring CrowdStrike ZTA Integration in the CrowdStrike Tenant

Currently, to integrate with Zscaler, you must contact the CrowdStrike Support team to turn on a ZTA feature flag in your CrowdStrike tenant.

When enabled on the CrowdStrike back end, ZPA can access and use the per-device ZTA score.

Before proceeding further, reach out to support@crowdstrike.com to turn on this flag.

Also, ensure that you use Zscaler Client Connector version 3.4 or later on the end host from which you are testing. CrowdStrike sensor version must be 6.20 or later.

Configuring ZPA

This guide assumes that you have a working ZPA set up and provides instructions to integrate ZTA-based conditional access into your existing ZPA deployment.

Log into ZPA Admin Portal

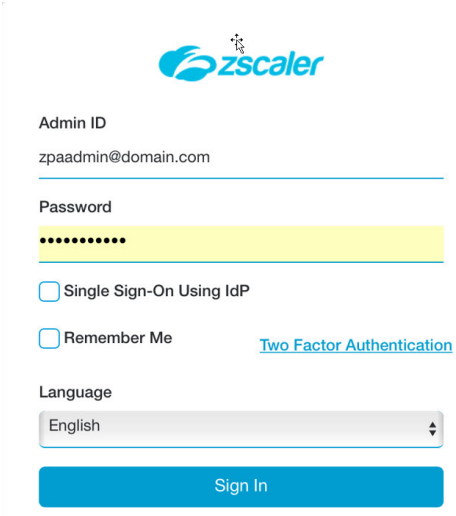

Figure 2. Log in to ZPA Admin Portal

Go to Zscaler Client Connector

Click the **Client Connector** icon to access the Zscaler Client Connector.

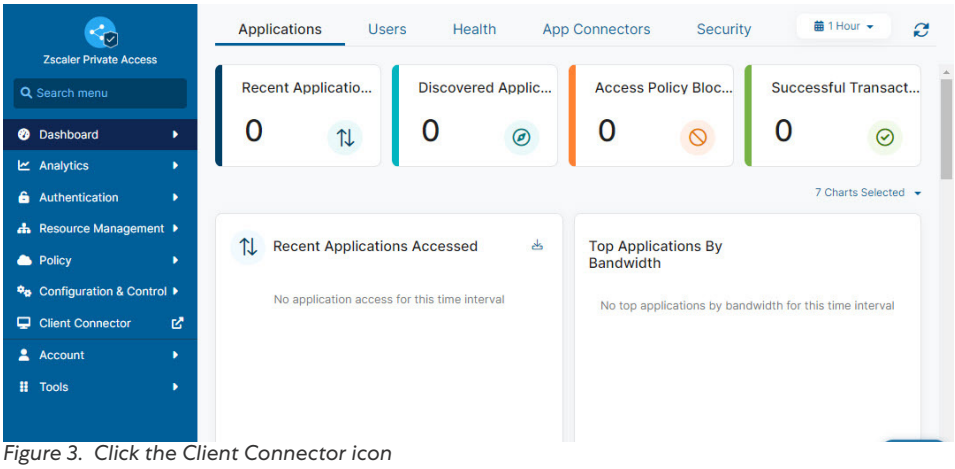

Figure 3. Click the Client Connector icon

Create New Posture Profile

Log into the Zscaler Client Connector and go to **Administration** > **Device Posture**. Then click **Add Device Posture Profile**.

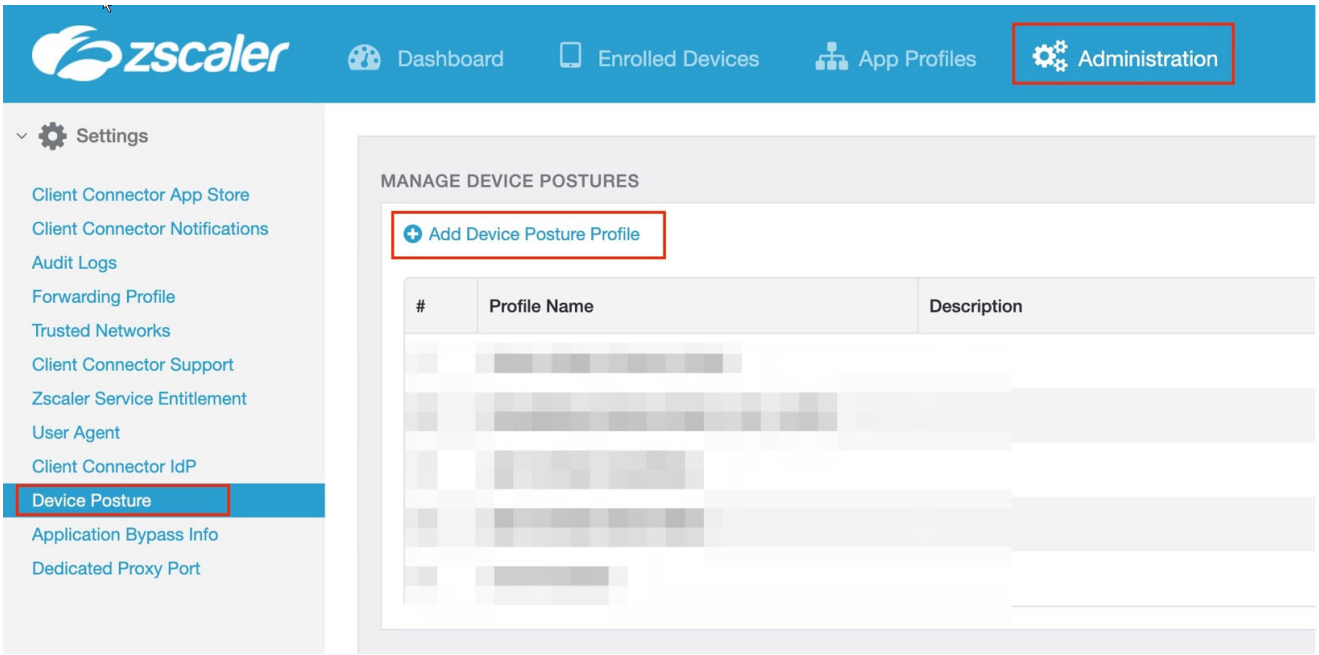

Figure 4. Add a device posture profile

Add a New CrowdStrike ZTA Posture Profile

Complete the following steps:

1. Select only **Windows** and **macOS**.
2. Click the **Posture Type** drop-down menu.
3. Select **CrowdStrike ZTA Score**.
4. Name this policy.
5. Provide the minimum value for a ZTA score.
6. Click **Save**.

ZPA passes a posture check if the end device's ZTA score (calculated by CrowdStrike) is greater than or equal to the value configured in this procedure. This posture profile is referenced in a ZPA access policy. You can set up access policies to allow or deny application access based on whether the posture check passes or fails.

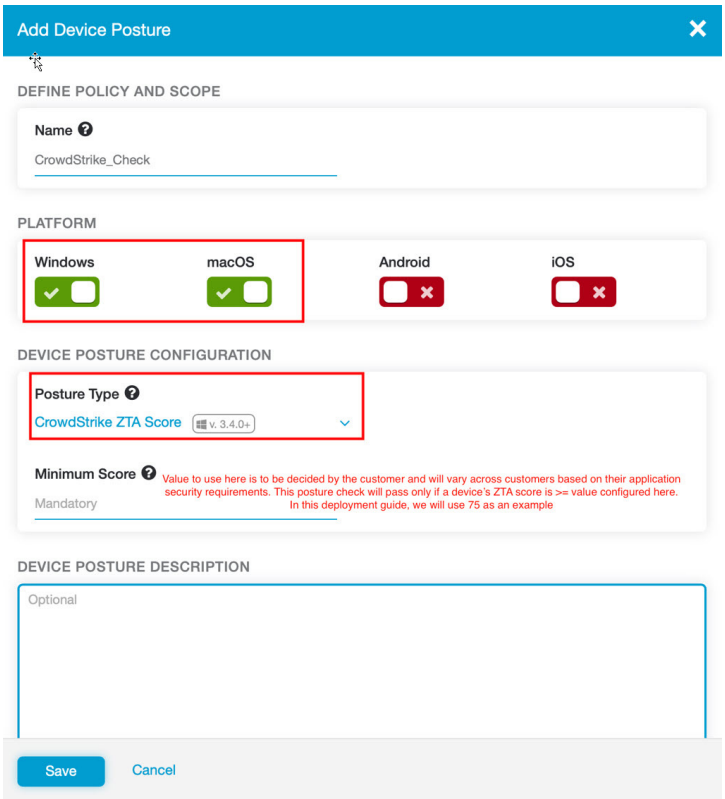

Figure 5. Add a CrowdStrike ZTA posture profile

Decide Which Applications Need Conditional Access Based on ZTA

From the ZPA Admin Portal, go to **Resource Management** > **Application Management** > **Application Segments**. This page lists applications that are accessed over ZPA. Select one of these applications and reference it in an access policy so that access to it is granted conditionally based on an end device's ZTA score.

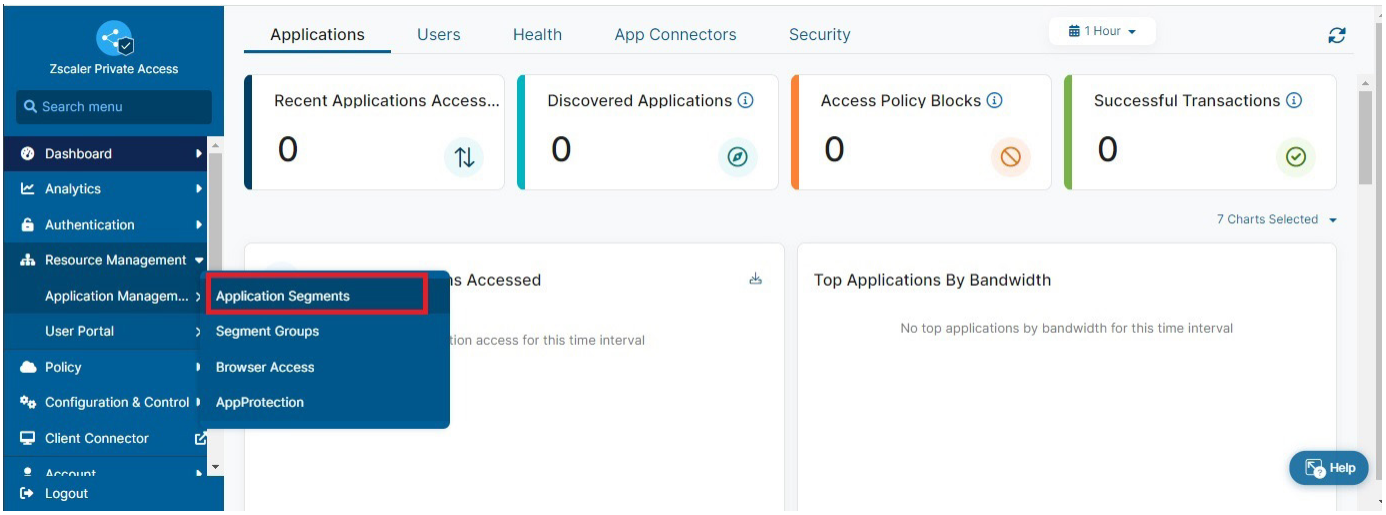

Figure 6. Go to Application Segments

In the following example, applications hosted under the domain *.preview-dev-company.net are accessed over ZPA (and allowed conditional access based on the ZTA score of the end device).

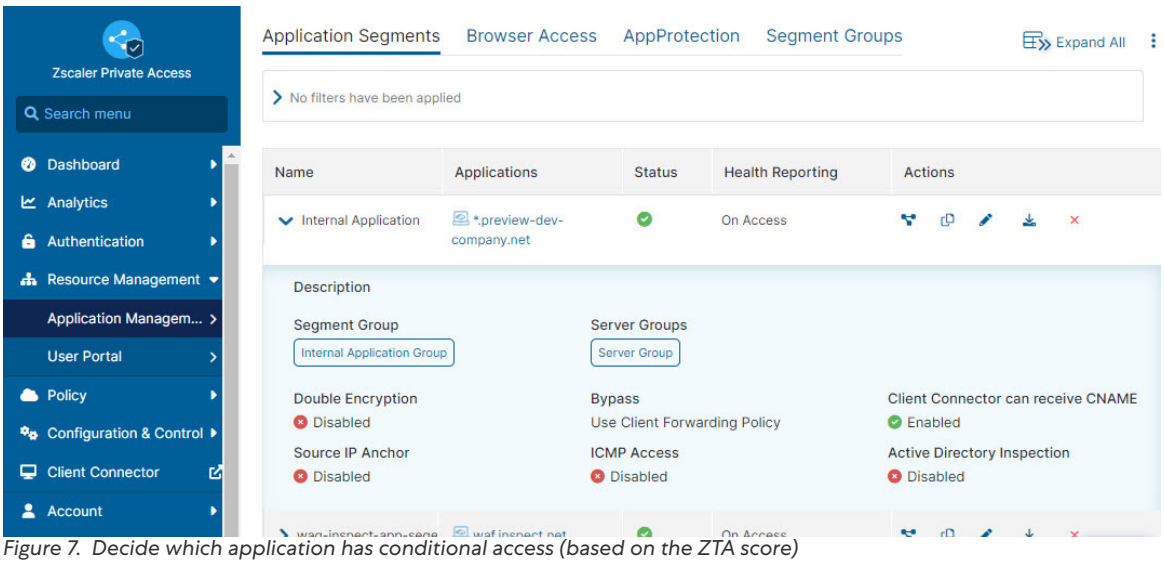

Figure 7. Decide which application has conditional access (based on the ZTA score)

Set Up an Access Policy

From the ZPA Admin Portal, go to **Policy** > **Access Policy**.

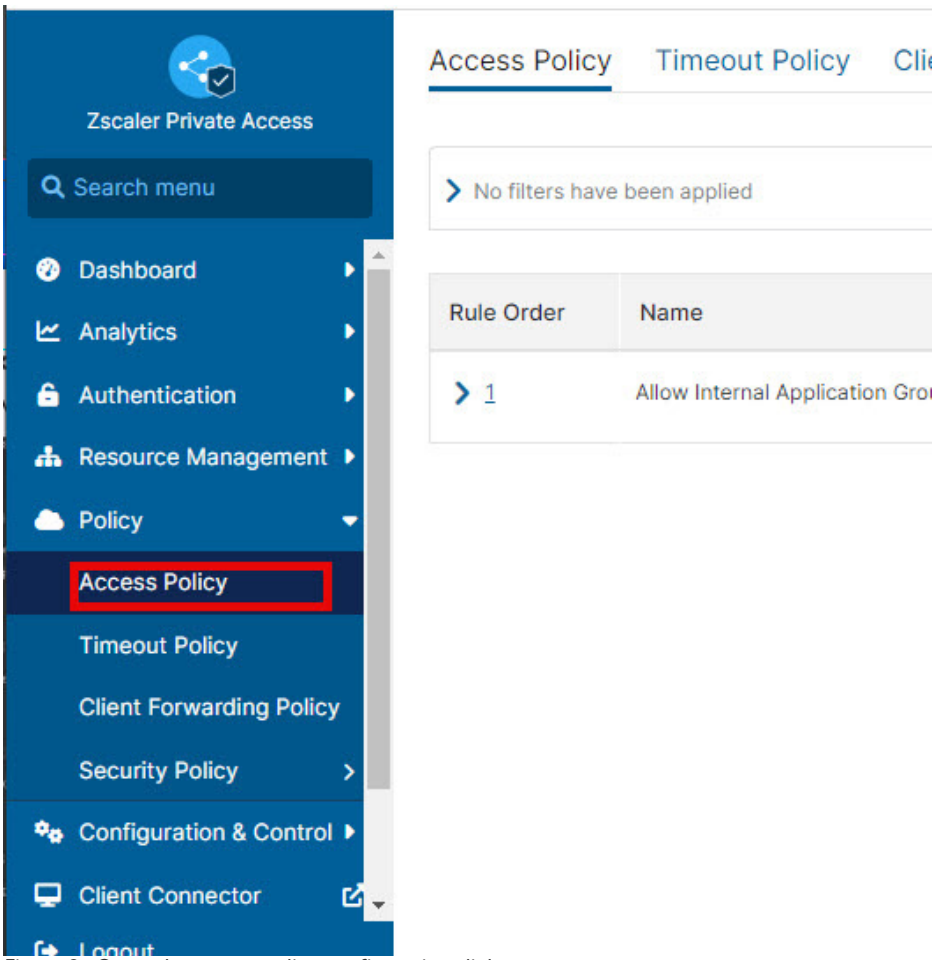

Figure 8. Open the access policy configuration dialog

Tie the Posture Profile to the Access Policy

On the **Access Policy** tab, click **Add Rule** and reference the previously-created posture profile.

You can set up different access policies to protect different internal applications. These access policies, in turn, can reference different ZTA posture check profiles based on the ZTA score requirement. A customizable (and optional) notification message is shown to the end users when application access is allowed or denied, informing them about policy evaluation.

In the following example, an access policy is added to block user access to an application if the ZTA posture check fails (Rule#1). If the end device's ZTA score is greater than or equal to the configured threshold, then Rule#1 fails and the policy evaluation proceeds to Rule#2 (which grants application access).

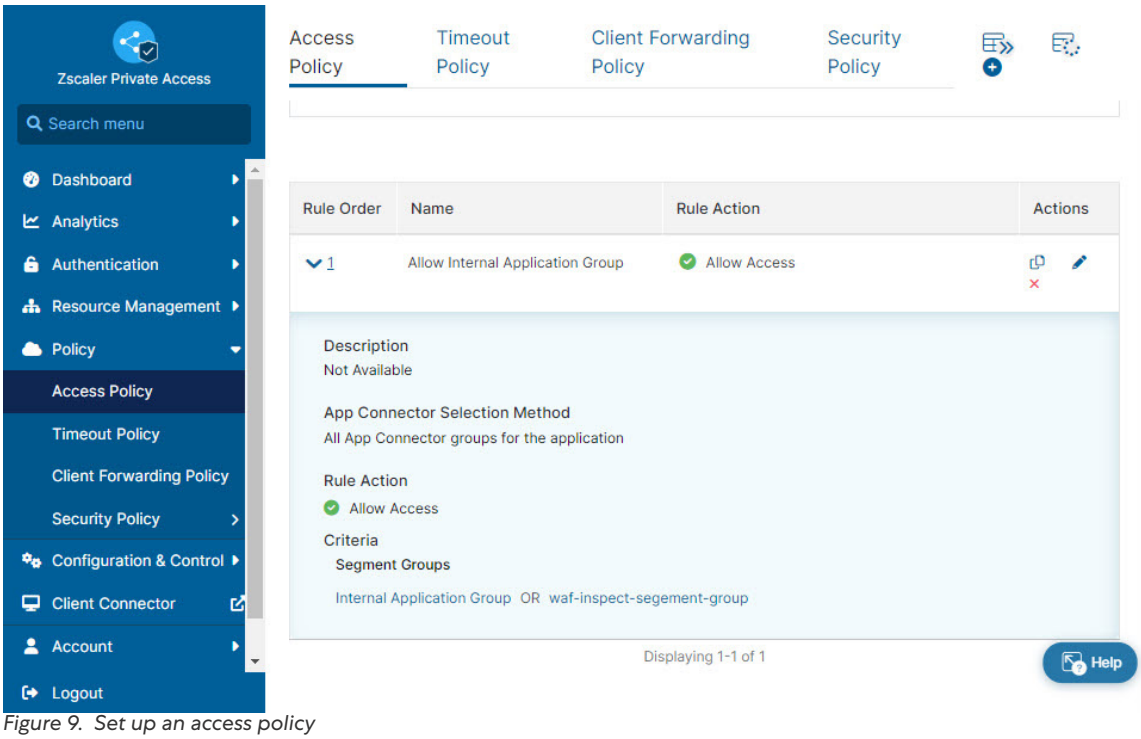

Figure 9. Set up an access policy

Verify ZTA-Based Conditional Access from an Endpoint

Ensure that you are logged into ZPA with Zscaler Client Connector version 3.4 or later, and access the application referenced in the previous step's access policy. The app is accessible from the endpoint if the device's CrowdStrike calculated ZTA score is greater than or equal to the value configured in the posture profile. Otherwise, the access is blocked by ZPA.

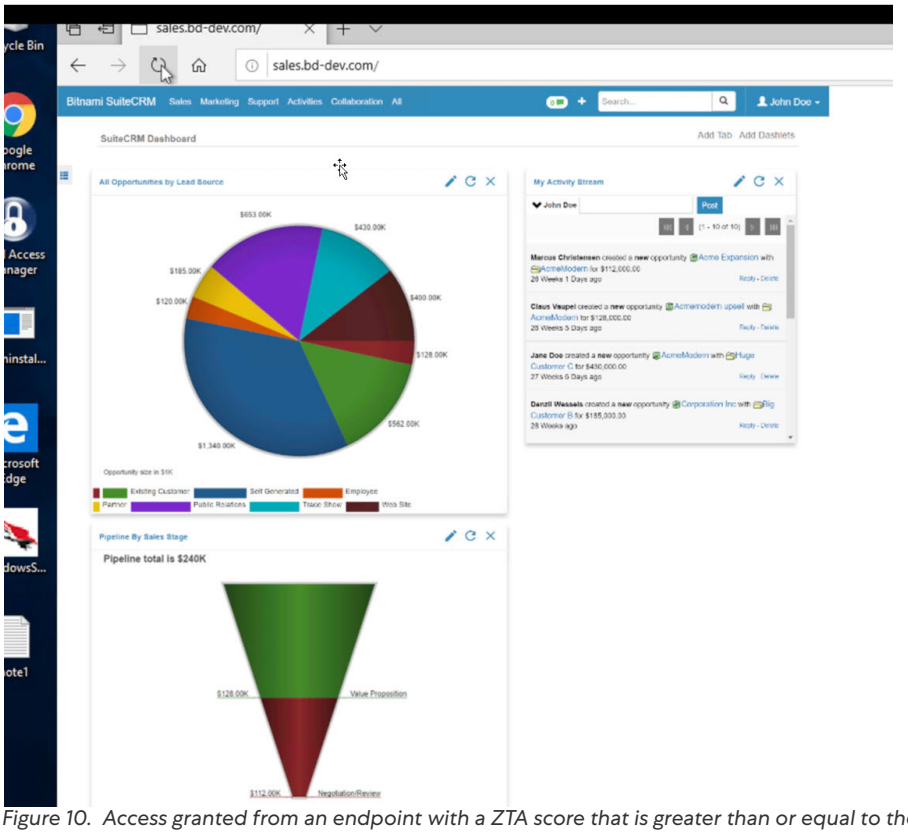

Figure 10. Access granted from an endpoint with a ZTA score that is greater than or equal to the value configured in the posture profile

If access is blocked, the following message is displayed.

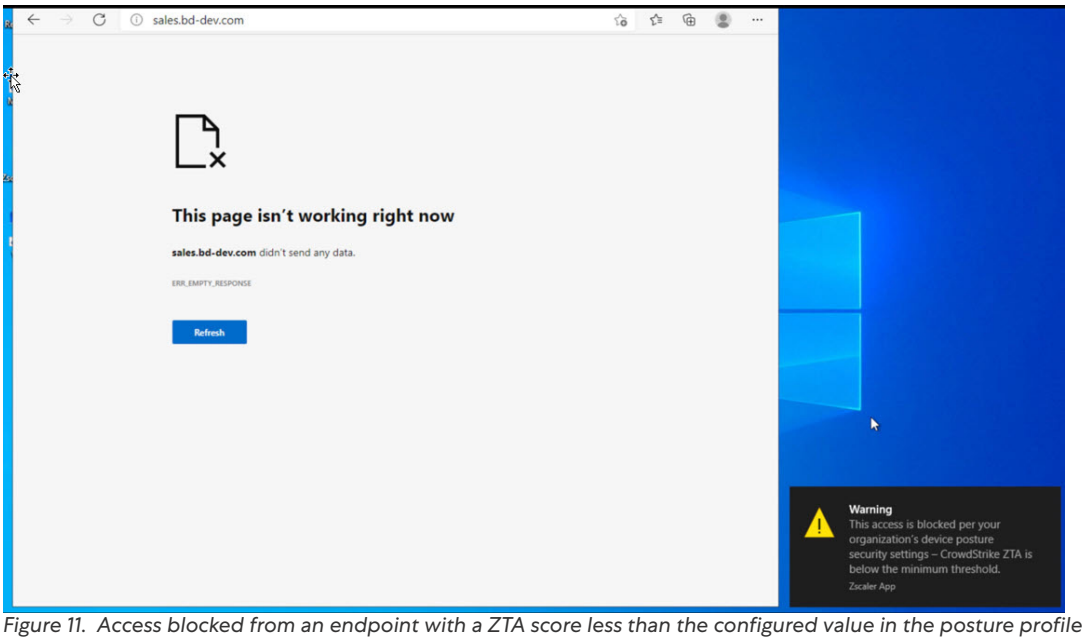

Figure 11. Access blocked from an endpoint with a ZTA score less than the configured value in the posture profile

Use Case 2: ZIA Posture Check Integration with CrowdStrike ZTA

In this use case:

- CrowdStrike calculates a ZTA security score from 1 to 100 for each host. A higher score indicates a better security posture for the host.
- Security scores are derived from two distinct assessment sources:
  - OS settings: Settings that track native OS security options, firmware availability, and CVE mitigations.
  - CrowdStrike Falcon sensor settings (Windows and macOS): CrowdStrike Falcon sensor configurations that track RFM status as well as prevention and Real-Time Response policies.
- ZIA uses CrowdStrike's ZTA score (also known as a device posture score) and allows only compliant endpoints to access selected applications. ZIA checks for any changes to the CrowdStrike's device posture score because the score can change over time. ZTA check integration is supported currently for Windows and macOS endpoints.
- ZIA achieves conditional access based on Device Trust level. Device Trust levels are determined by ZIA posture profiles that, in turn, reference device level posture check profiles. ZIA administrators can simply specify that a minimum ZTA score is needed for the endpoint to grant access to resources that are referenced via various ZIA policies (such as URL filtering policy, Firewall policy, File Control policy, Cloud apps policy, or SSL inspection policy). You can combine the ZTA score threshold requirement with other device posture policies to allow only compliant devices to access resources through ZIA.

The following diagram shows a conceptualization of the integration.

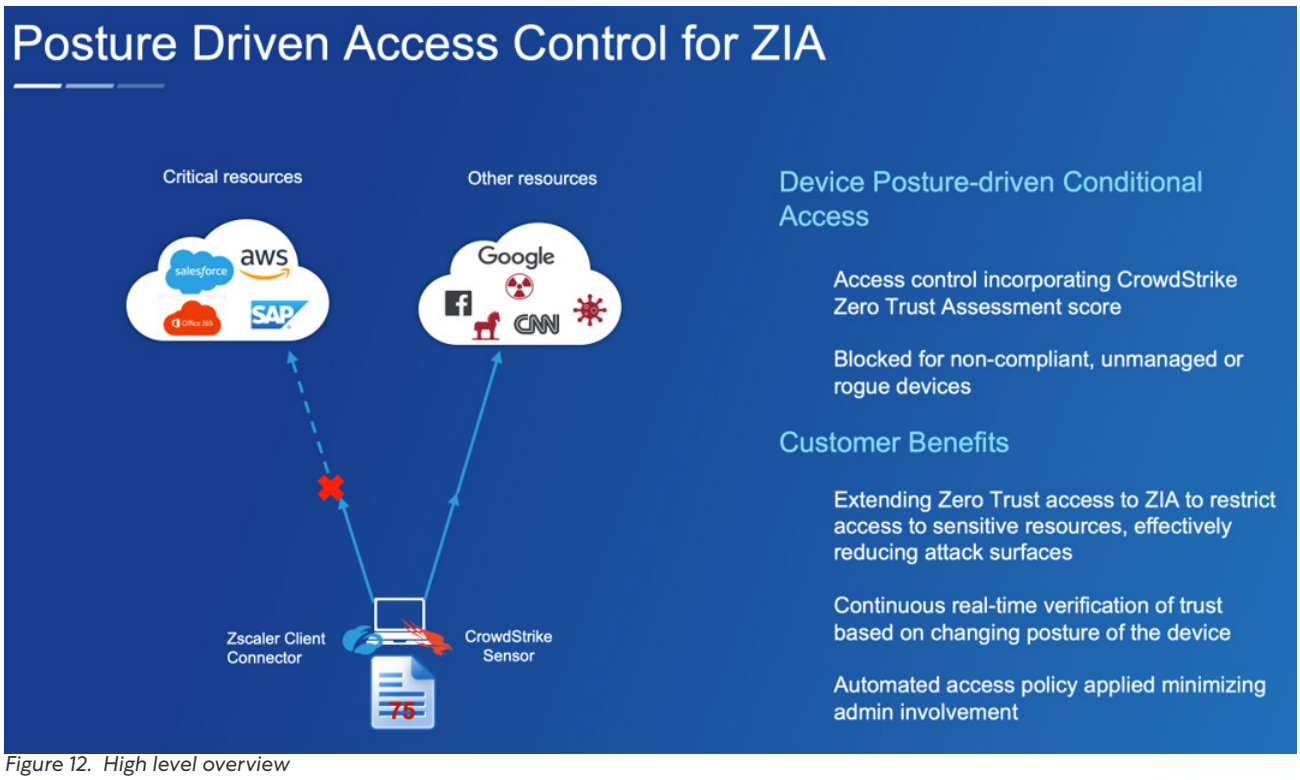

Figure 12. High level overview

If interoperability issues arise, update your CrowdStrike policy's sensor visibility exclusions to account for Zscaler Client Connector-related details and re-test. To learn more, see Zscaler Client Connector Processes to Allowlist. (government agencies, see Zscaler Client Connector Processes to Allowlist).

Configuring CrowdStrike ZTA Integration in the CrowdStrike Tenant

Currently, to integrate with Zscaler, you must contact the CrowdStrike Support team to turn on a ZTA feature flag in their CrowdStrike tenant.

When enabled on the CrowdStrike back end, ZIA can access and use the per-device ZTA score.

Before proceeding further, reach out to support@crowdstrike.com to turn on this flag.

Also, ensure that you use Zscaler Client Connector version 3.4 or later on the end host from which you are testing. CrowdStrike sensor version must be 6.20 or later.

Configuring ZIA

This guide assumes that you have a working ZIA set up and provides instructions to integrate ZTA-based conditional access into your existing ZIA deployment.

Log into ZIA Admin Portal

Log in to the ZIA Admin Portal as an administrator.

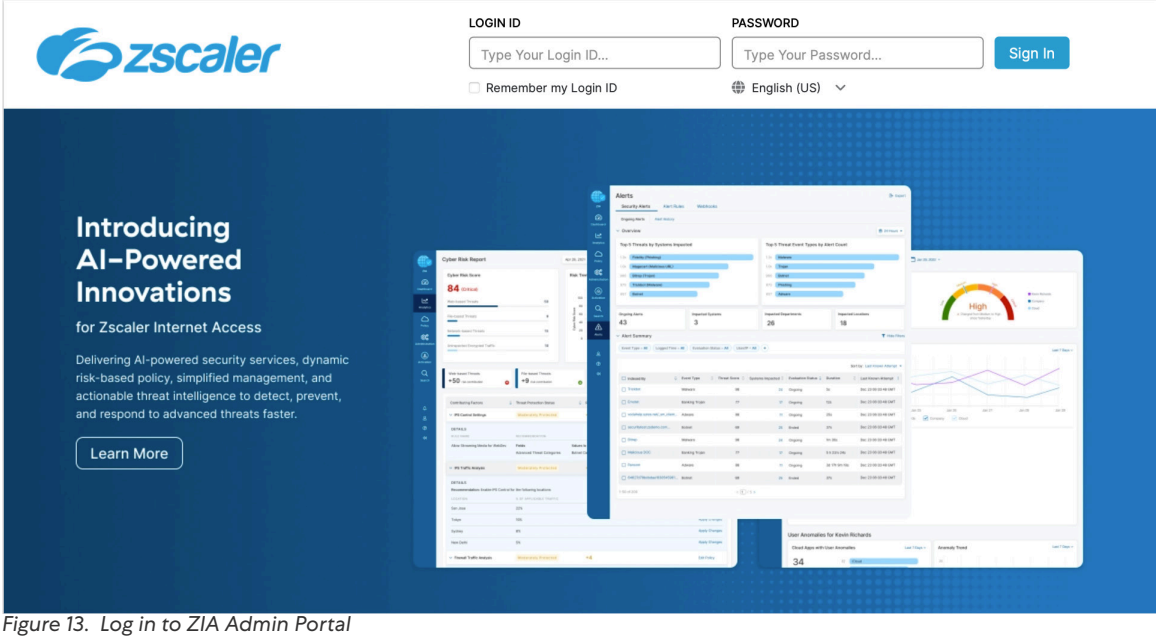

Figure 13. Log in to ZIA Admin Portal

Go to Zscaler Client Connector

Click the **Zscaler Client Connector Portal** link under the **Policy** section to access the Zscaler Client Connector Portal.

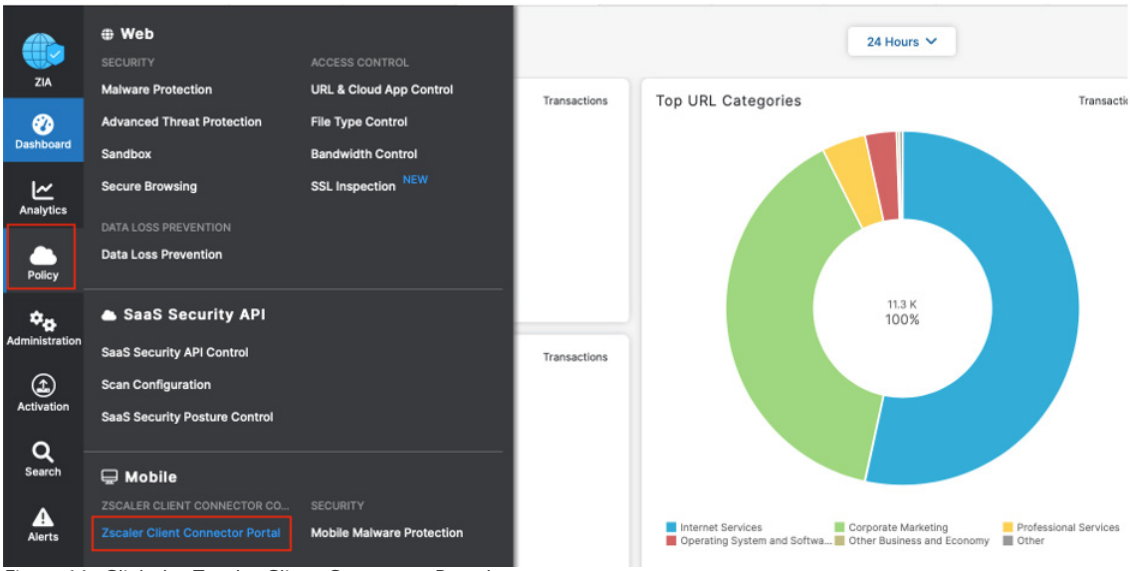

Figure 14. Click the Zscaler Client Connector Portal

Create New Posture Profile

Log in to the Zscaler Client Connector and go to **Administration** > **Device Posture** > **Add Device Posture Profile**.

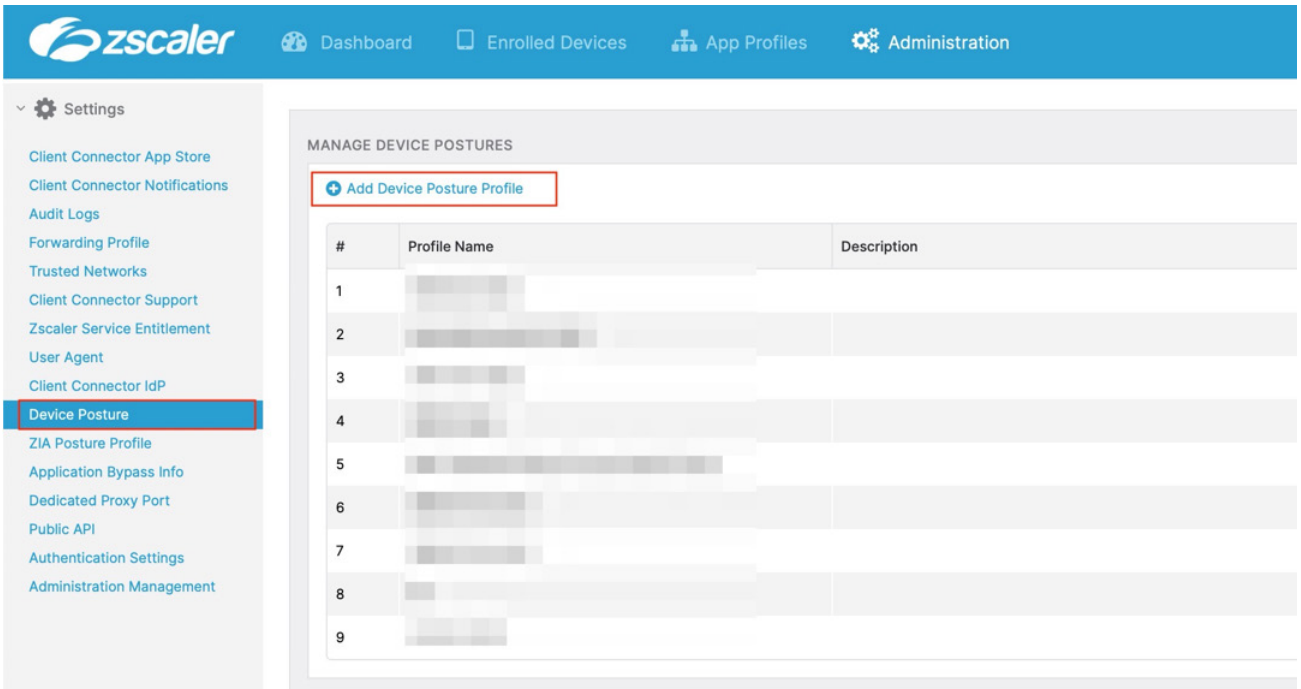

Figure 15. Add a device posture profile

Add a New CrowdStrike ZTA Posture Profile

To add a new CrowdStrike ZTA posture profile:

1. Select only **Windows**, **macOS**, or both.
2. Select **CrowdStrike ZTA Score** under the **Posture Type** drop-down menu.
3. **Name** this policy.
4. Provide the minimum value for the **CrowdStrike ZTA Score**.
5. Click **Save**.

The device posture check passes if the end device's ZTA score (calculated by CrowdStrike) is greater than or equal to the value configured in this procedure. This device posture profile is, in turn, referenced in a ZIA posture profile to determine the trust bucket or level that a device falls in. You can set up conditional access policies to allow or deny access through ZIA based on device trust levels.

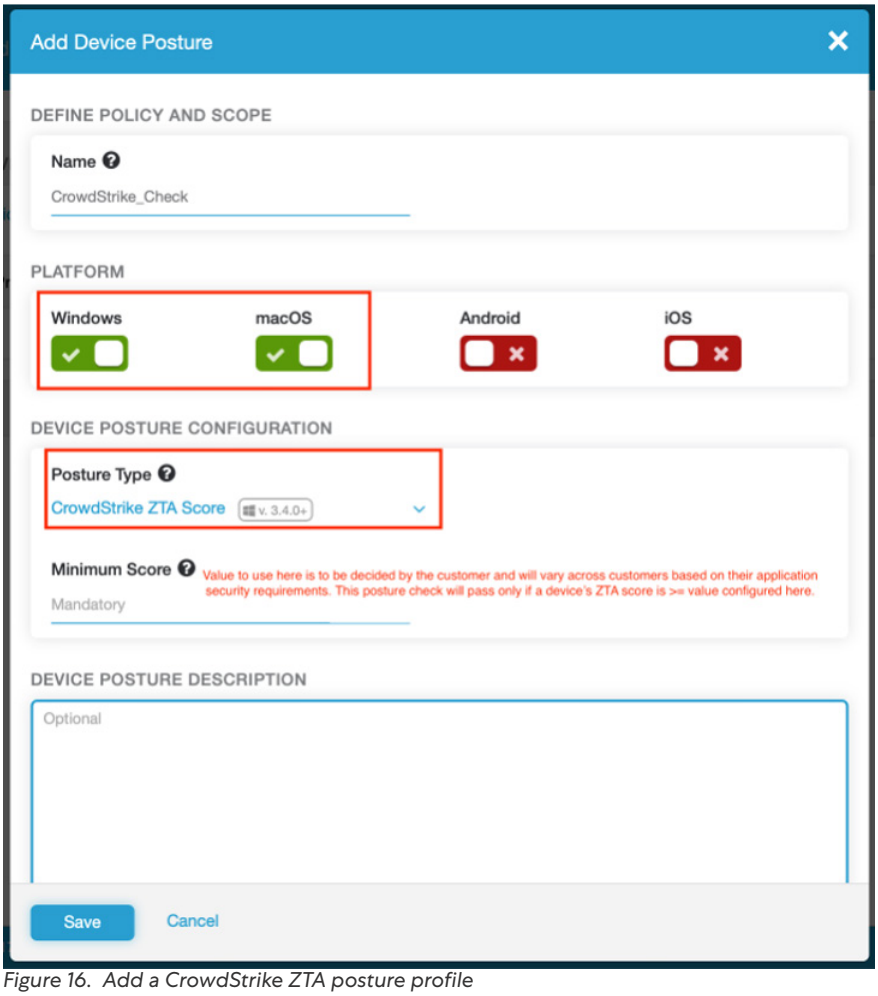

Figure 16. Add a CrowdStrike ZTA posture profile

Multiple Device Posture Profiles Defined Referencing Different ZTA Scores

You can define multiple ZTA-based device posture templates to correspond to different Device Trust levels. As an example, you can set up multiple device posture profiles referencing different ZTA scores.

The following example defines three different device posture profiles, which set the minimum device ZTA threshold scores to be 75, 65, and 50 respectively. These are used for qualifying a device to fall in the High, Medium, Low trust level (respectively) via ZIA posture profile configuration.

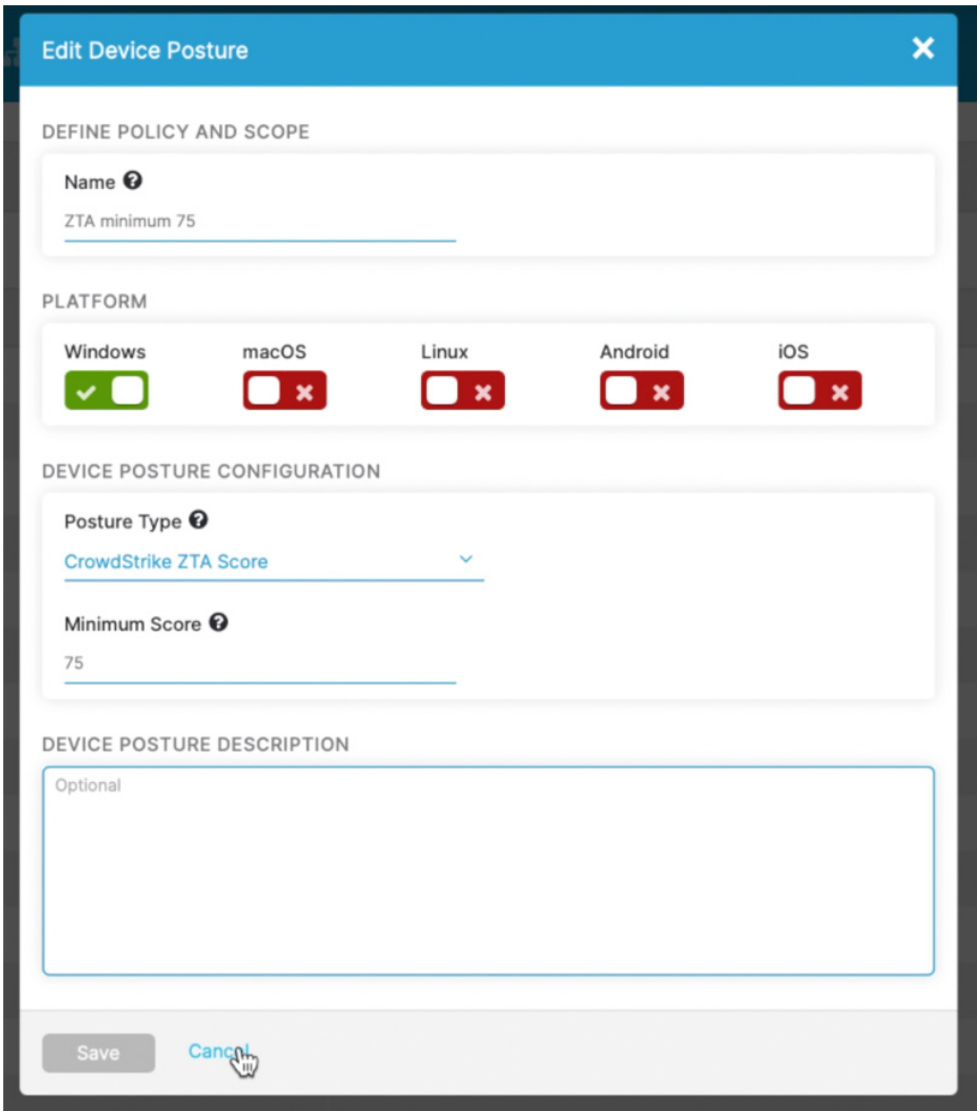

Figure 17. Device posture policy based on 75 ZTA score

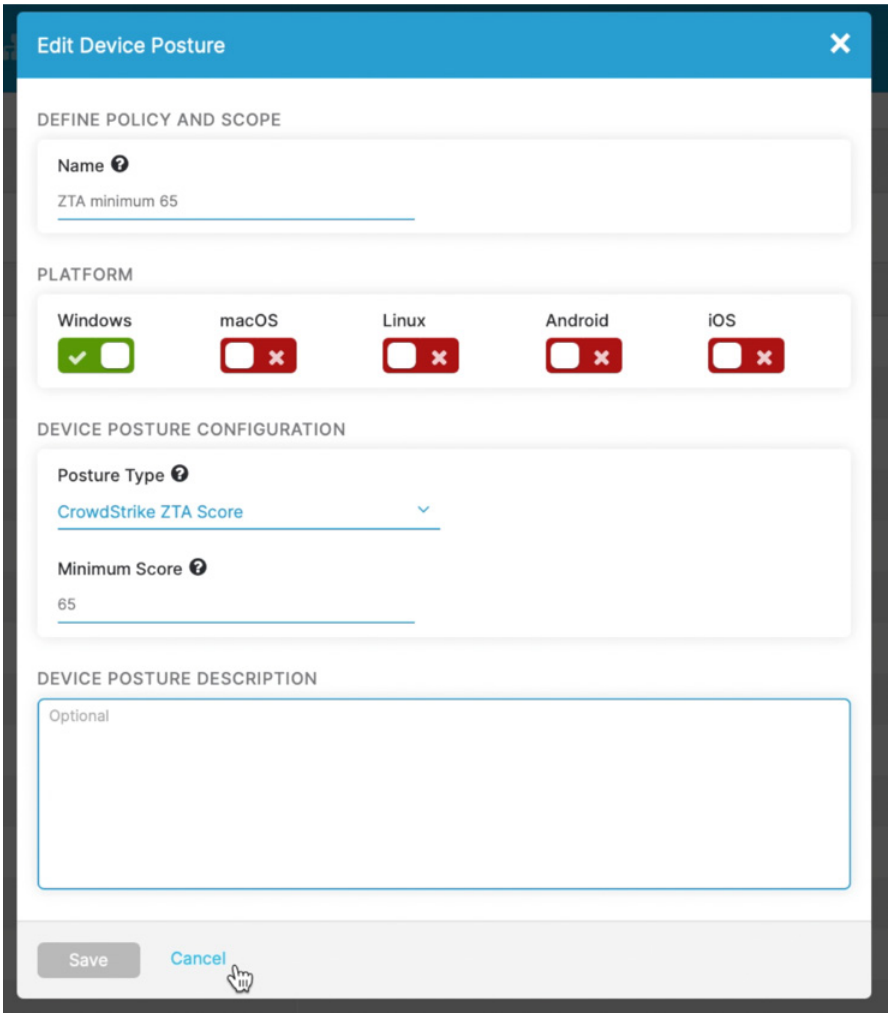

Figure 18. Device posture policy based on 65 ZTA score

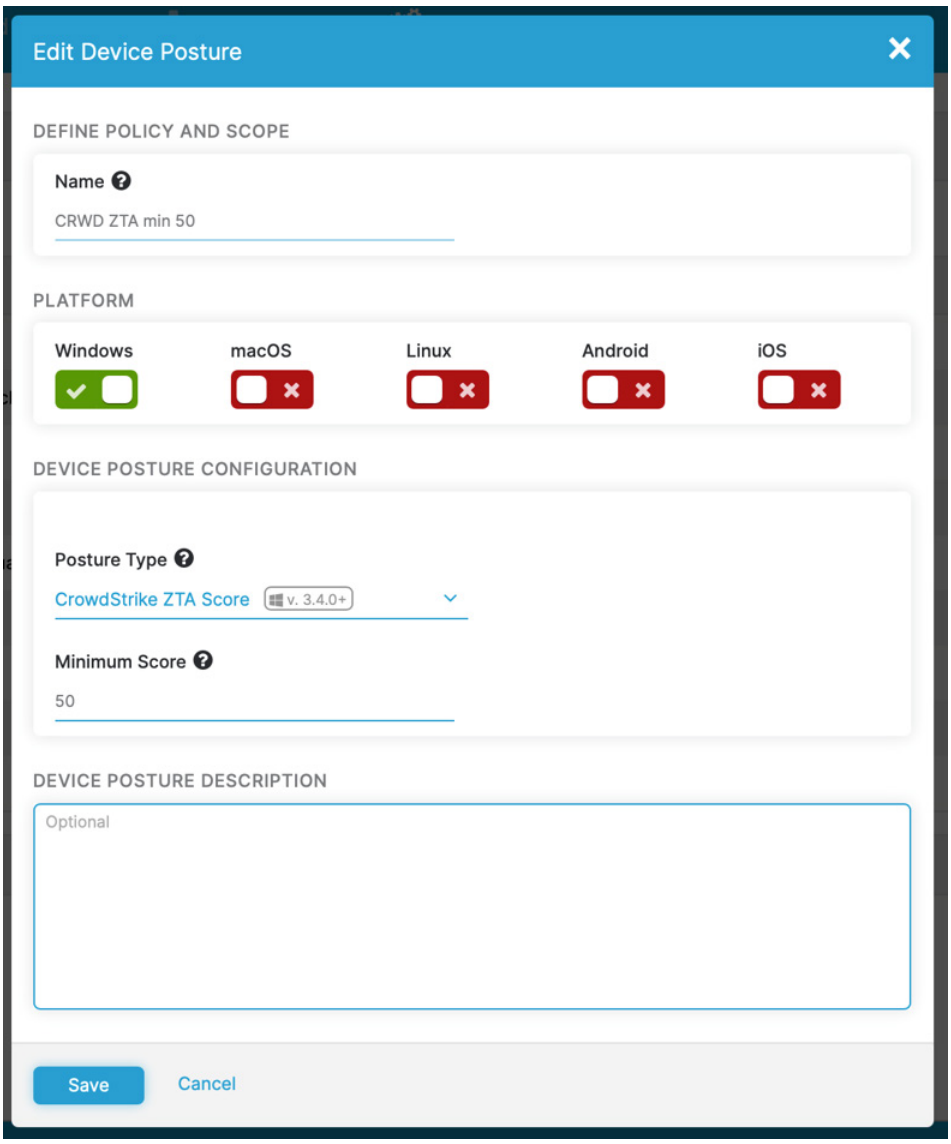

Figure 19. Device posture policy based on 50 ZTA score

Reference the Device Posture Profile in ZIA Posture Profile

In the Zscaler Client Connector **Administration** section, go to **ZIA Posture Profile** > **Add ZIA Posture Profile** option.

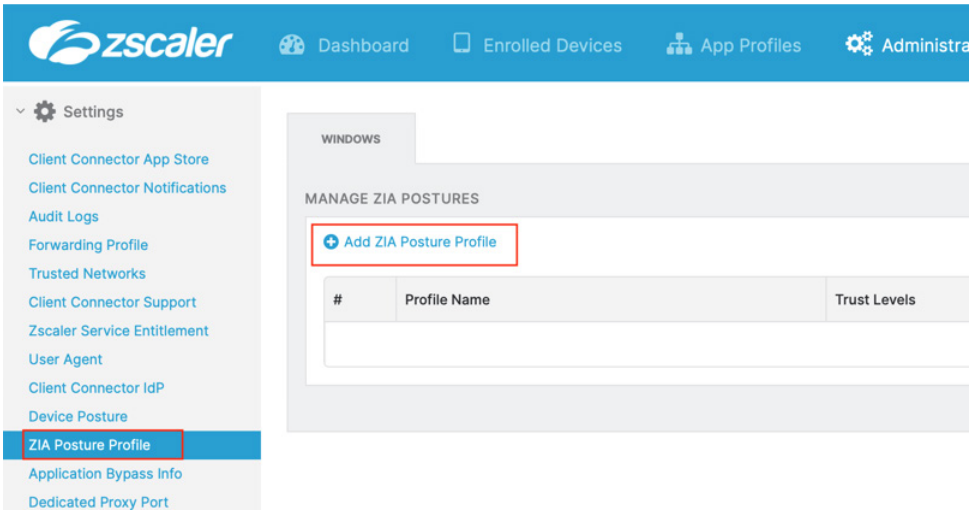

Figure 20. Add a new ZIA Posture Profile

In this new ZIA Posture Profile, you reference the previously created device posture profiles to establish Device Trust criteria.

In the following example, as long as the device's ZTA score is greater than or equal to 75, it falls into the High trust level. A device with a ZTA score greater than or equal to 65 (but less than 75) falls into the Medium trust level. A device with a ZTA score greater than or equal to 50 (but less than 65) falls into the Low trust bucket.

Trust criteria are evaluated in the top-down order, and the evaluation stops at the first match.

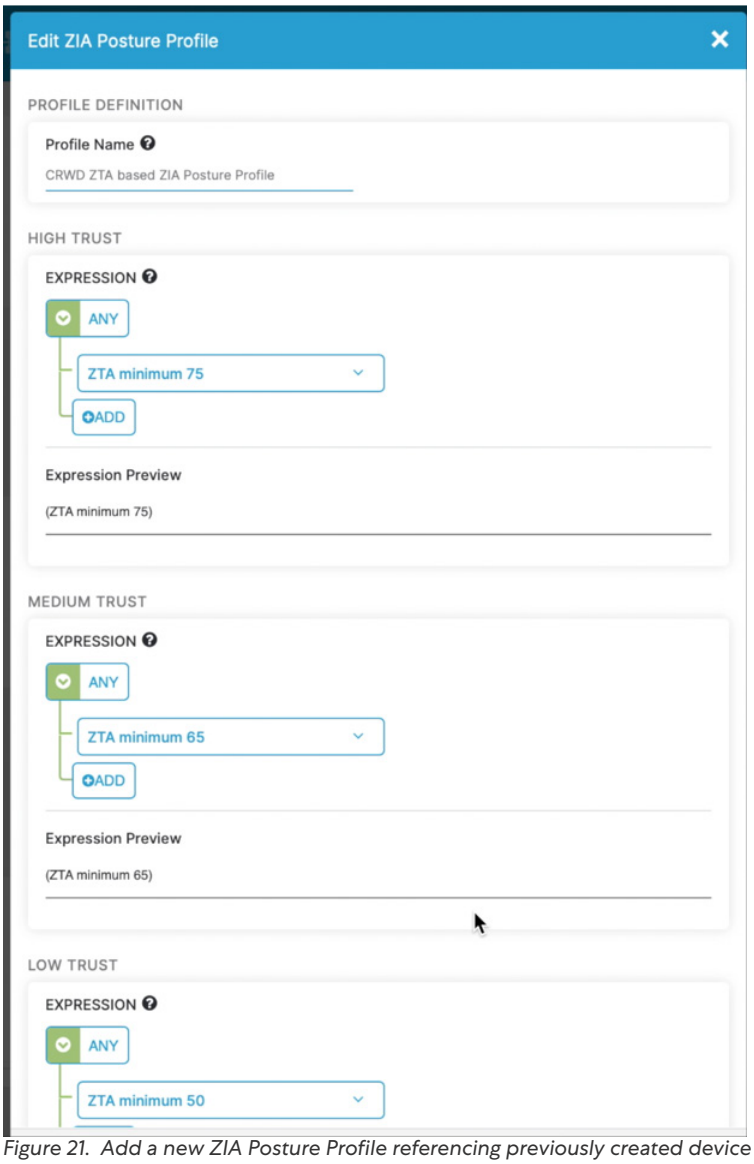

Figure 21. Add a new ZIA Posture Profile referencing previously created device posture profiles

Create an App Profile to Reference the ZIA Posture Profile

Go to the App Profile section and create a new App Profile. The following example creates an App Profile for Windows machines. The previously created ZIA Posture Profile is referenced in this App Profile.

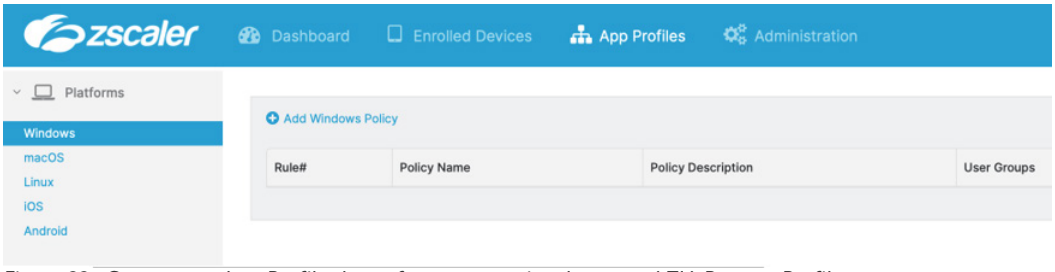

Figure 22. Create new App Profile that references previously created ZIA Posture Profile

Configure the App Profile.

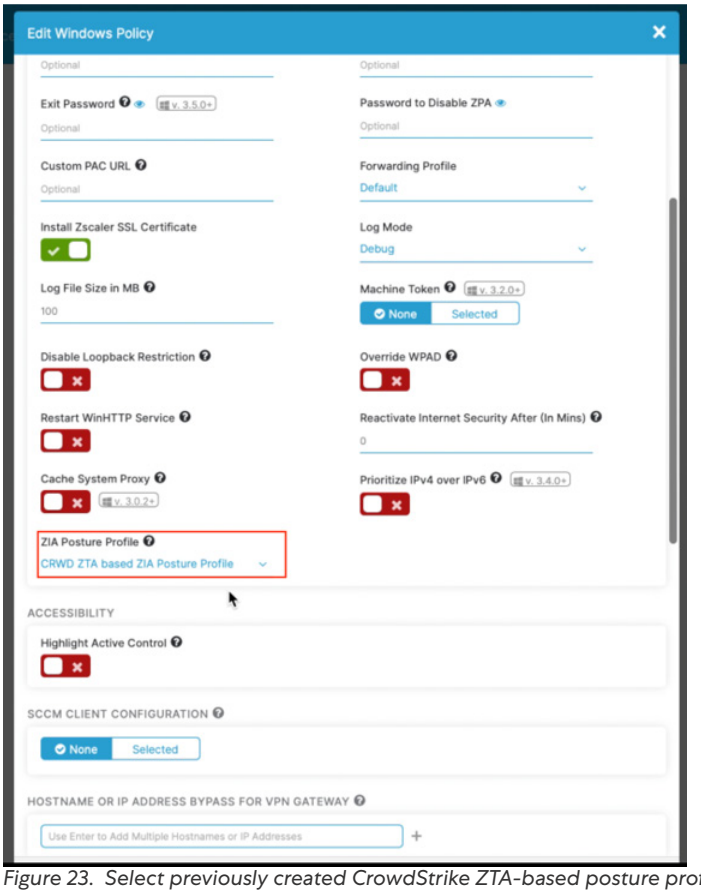

Figure 23. Select previously created CrowdStrike ZTA-based posture profile

Create a Traffic Enforcement Policy Using Device Trust Level as Criteria

The following example creates a URL and App Control policy that uses Device Trust level as a criterion for conditional access that allows access, shows a caution page, or blocks access through ZIA based on Device Trust level. You can set up such policies for File type control, SSL inspection, Firewall control, etc.

1. Go to **Policy** > **URL & Cloud App Control**.

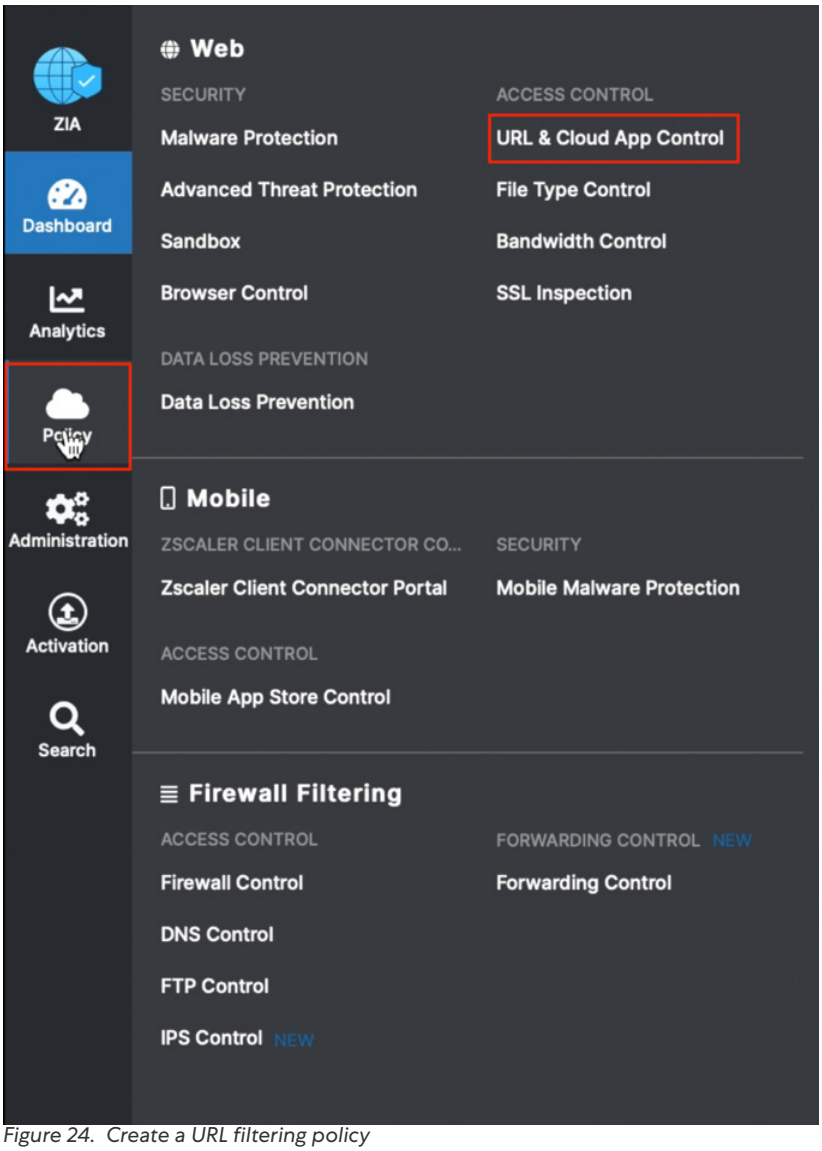

Figure 24. Create a URL filtering policy

2. Use the **Device Trust** level as a criterion for policy enforcement.

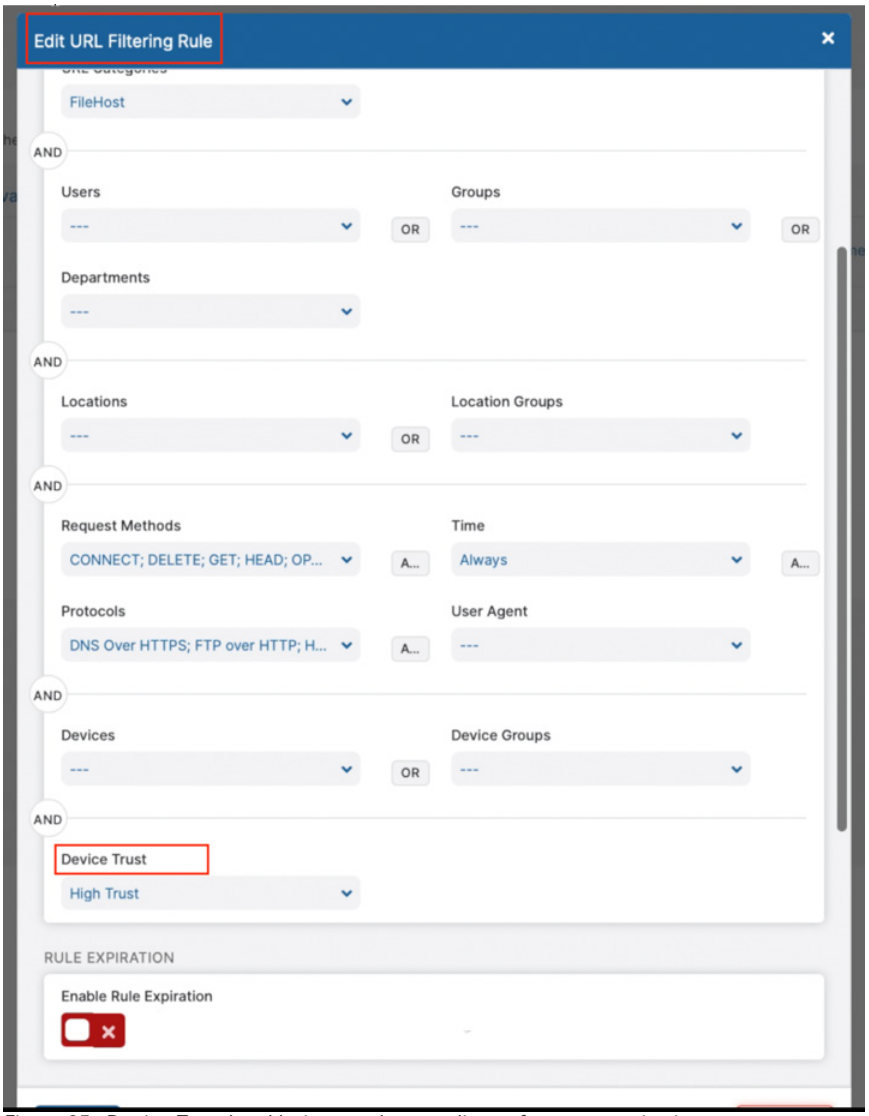

Figure 25. Device Trust level being used as a policy enforcement criterion

The following image shows different actions associated in URL filtering policy based on the Device Trust level.

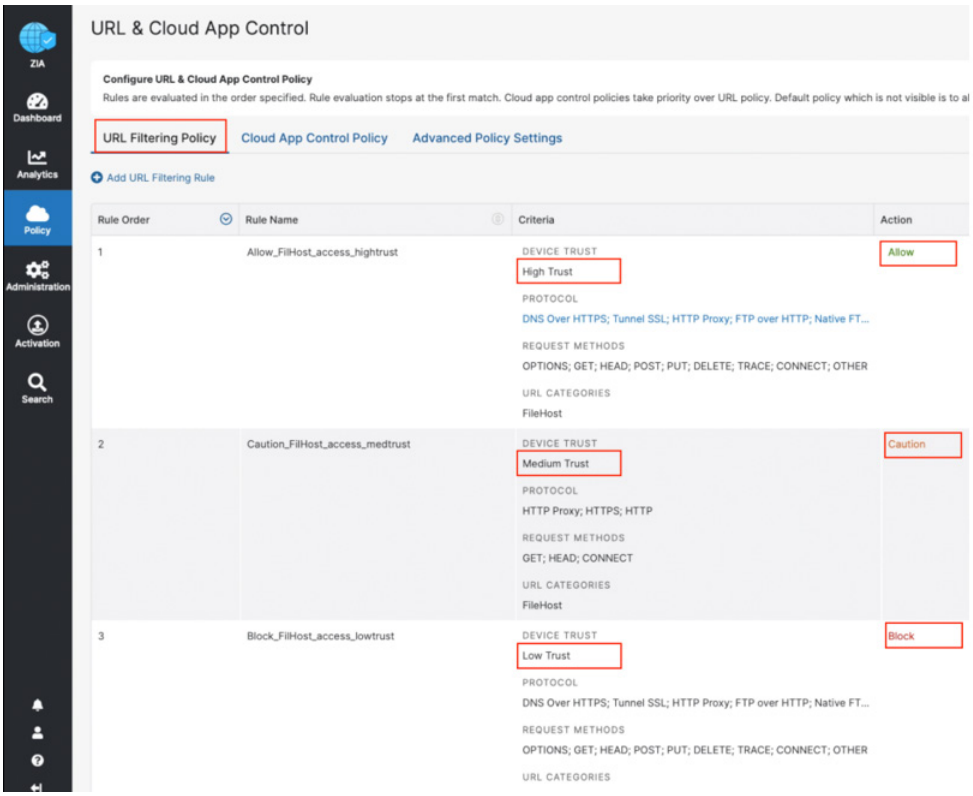

Figure 26. URL filtering policy based on Device Trust criterion

The same user logged into Zscaler Client Connector on three different machines sees pertinent URL filtering policies enforced based on Device Trust level.

In this case, you see the same user accessing a file hosting website from three different machines that fall under High (in the following example, a ZTA greater than or equal to 75), Medium (a ZTA greater than or equal to 65, but less than 75), and Low (greater than or equal to 50, but less than 65) trust levels due to those machines' CrowdStrike ZTA scores.

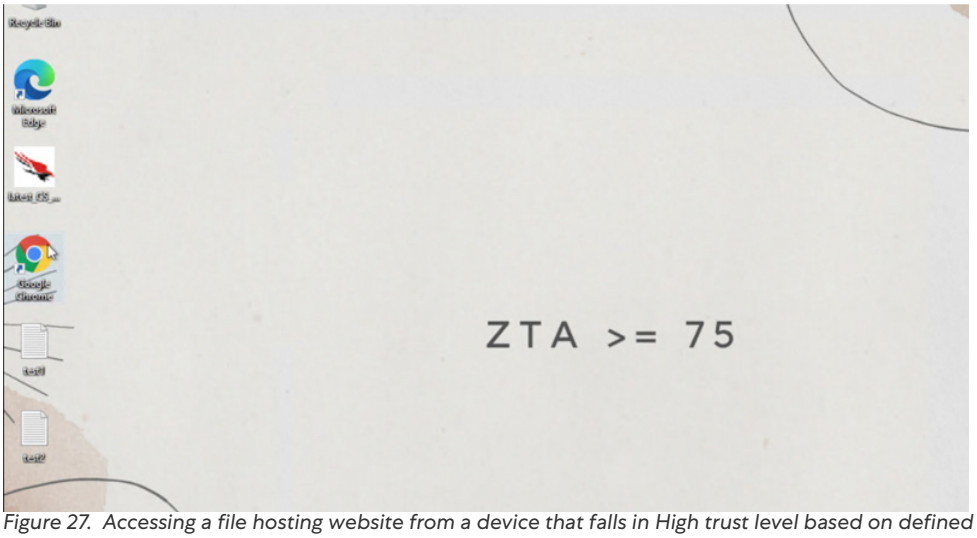

Figure 27. Accessing a file hosting website from a device that falls in High trust level based on defined ZIA posture policies

Access is allowed by ZIA per URL filtering policy:

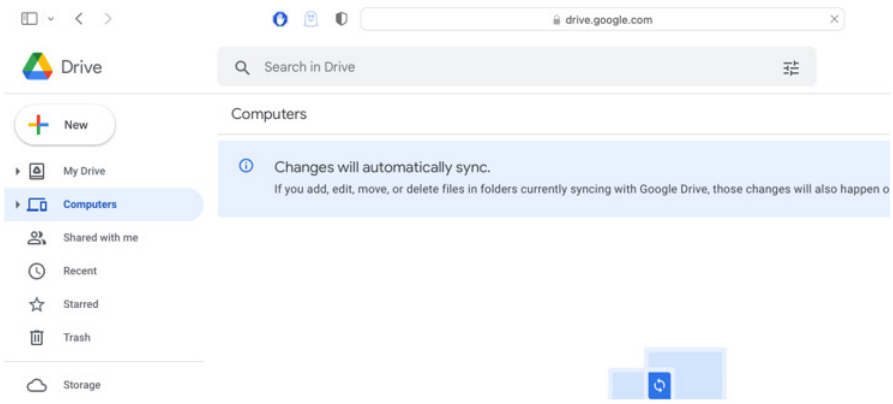

Figure 28. Access allowed by ZIA per the URL filtering policy due to Device Trust level

Devices that fall in the Medium trust level:

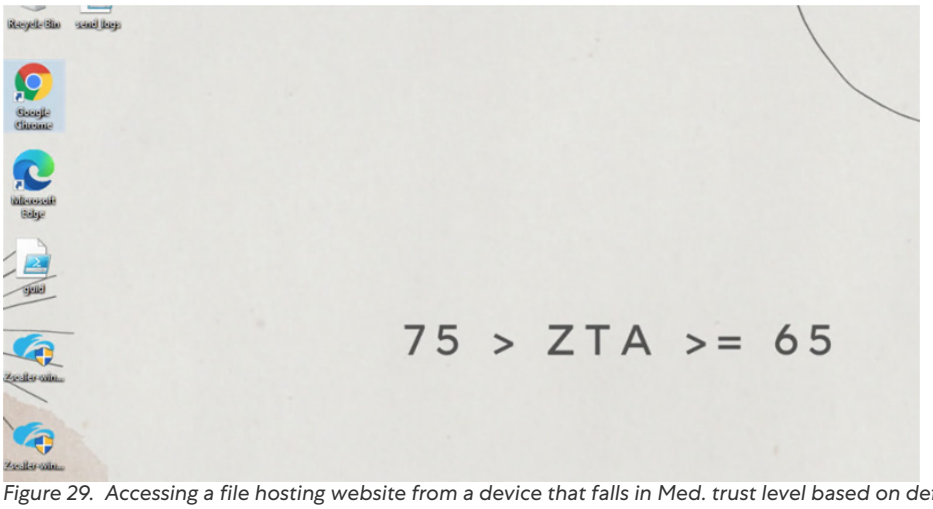

Figure 29. Accessing a file hosting website from a device that falls in Med. trust level based on defined ZIA posture policies

Caution warning from a URL filtering policy based on trust level:

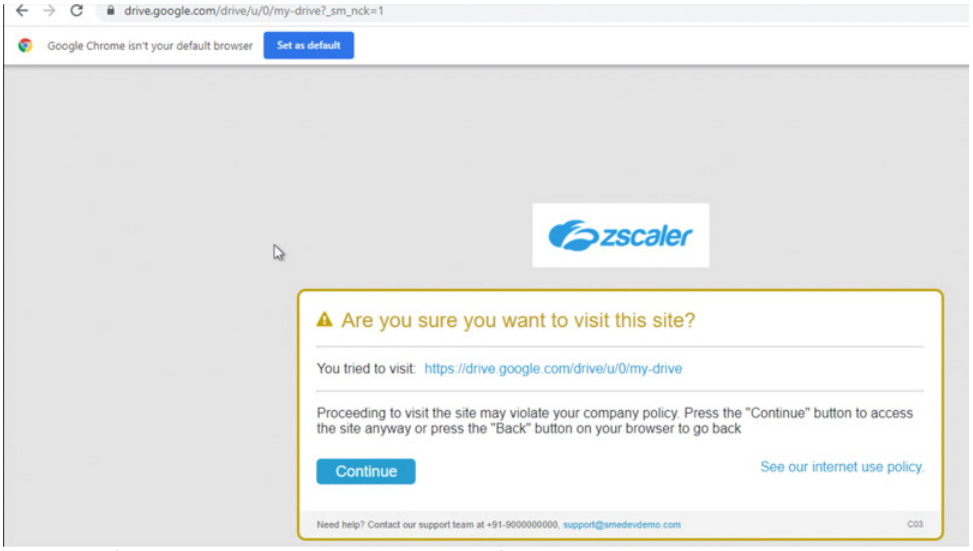

Figure 30. Caution page shown by ZIA per the URL filtering policy due to Device Trust level

Low Trust level based on posture policy:

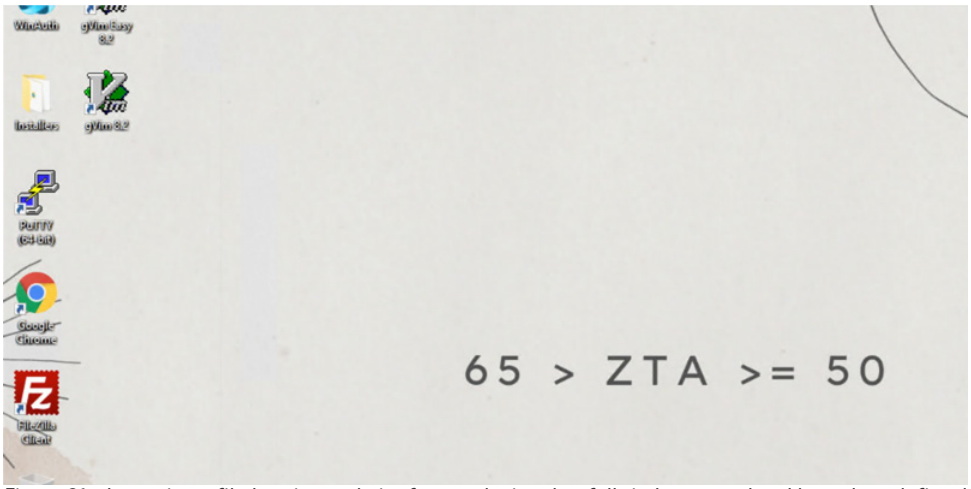

Figure 31. Accessing a file hosting website from a device that falls in Low trust level based on defined ZIA posture policies

Access blocked due to Device Trust level:

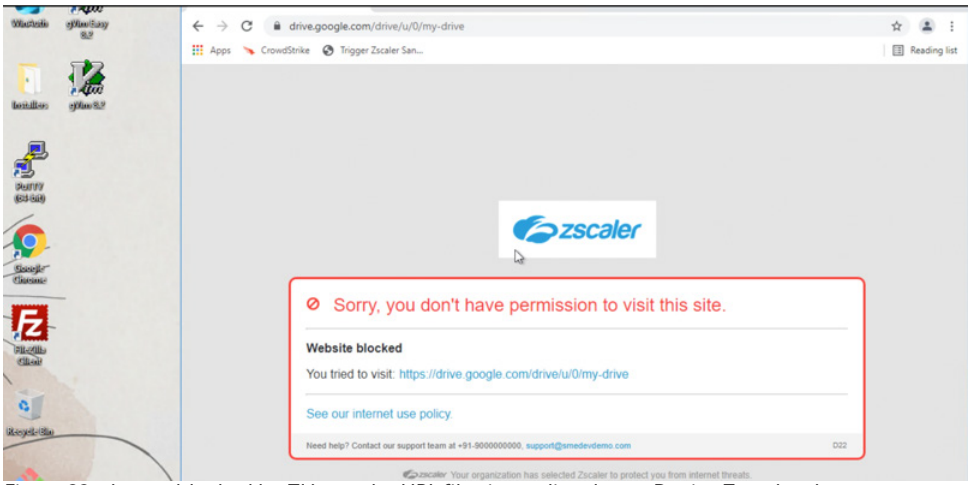

Figure 32. Access blocked by ZIA per the URL filtering policy due to Device Trust level

The same user logged into different machines under different trust levels gets pertinent enforcement policies applied by ZIA.

Use Case 3: ZPA Posture Check Integration with CrowdStrike

In this use case:

- ZPA verifies the presence of a running CrowdStrike Falcon process on the endpoint as an assessment of end device posture. ZPA is configured to allow only compliant endpoints (ones that pass the posture check) to access selected applications.
- ZPA evaluates ZPA access policies for conditional access. The policies, in turn, reference device level posture check profiles. The ZPA administrator specifies (for Windows and macOS workstations) that a CrowdStrike Falcon agent is installed and running on the endpoint so that the endpoint is granted access to internal applications referenced via the ZPA access policy.

This ZPA integration was implemented before ZTA functionality was available from CrowdStrike.

ZTA-based posture check (Use Case 1 in this guide) is an enhancement to this use case and is preferred due to its nuanced posture checking abilities.

See the following conceptual diagram for an overview of the integration.

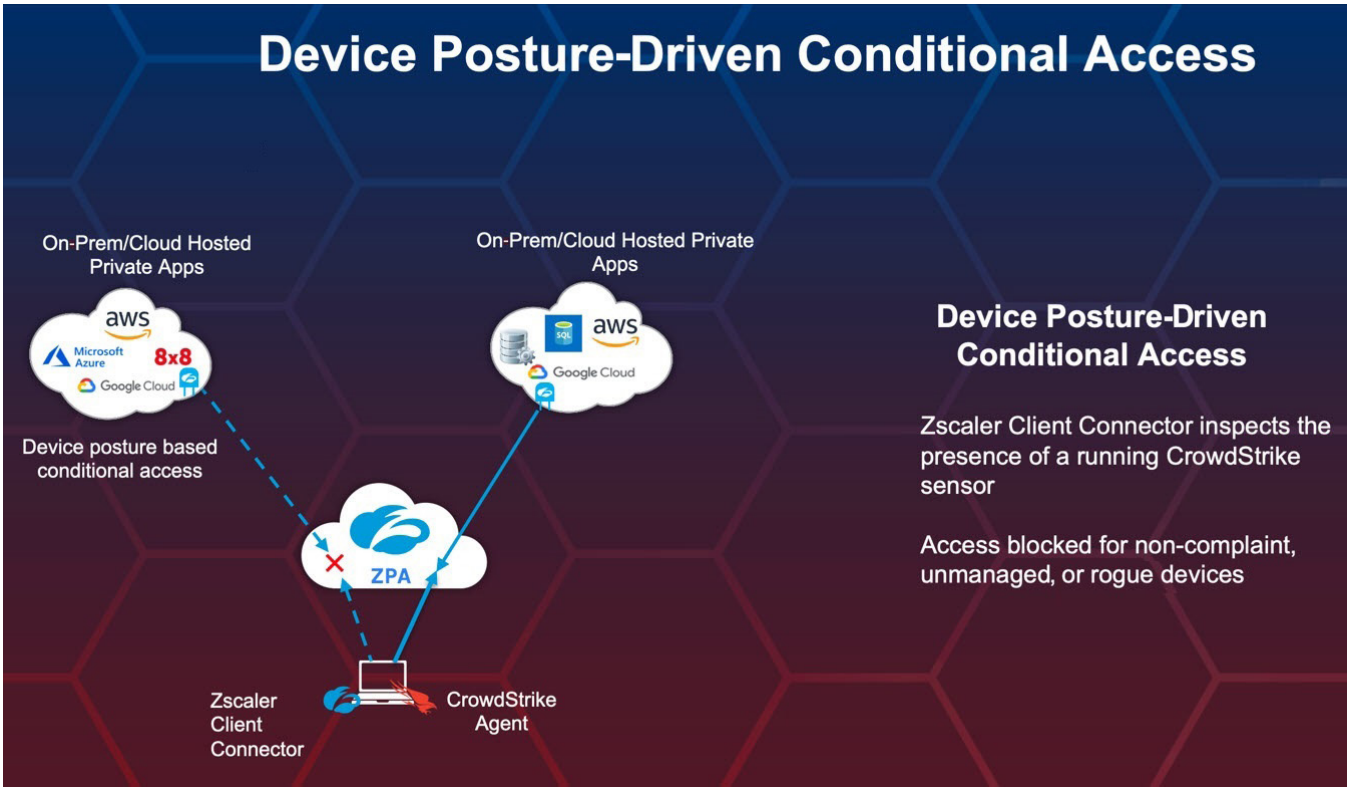

Figure 33. High-Level overview

Configuring ZPA

This guide assumes that you have a working ZPA set up and provides instructions to integrate posture-based conditional access as part of your existing ZPA deployment.

Log into ZPA Admin Portal

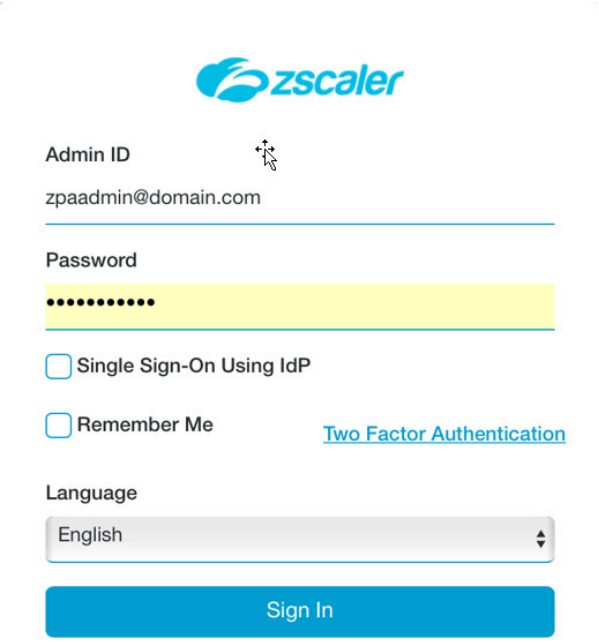

Figure 34. Log in to ZPA Admin Portal

Go to the Zscaler Client Connector

Click the **Client Connector** icon to open the Client Connector window.

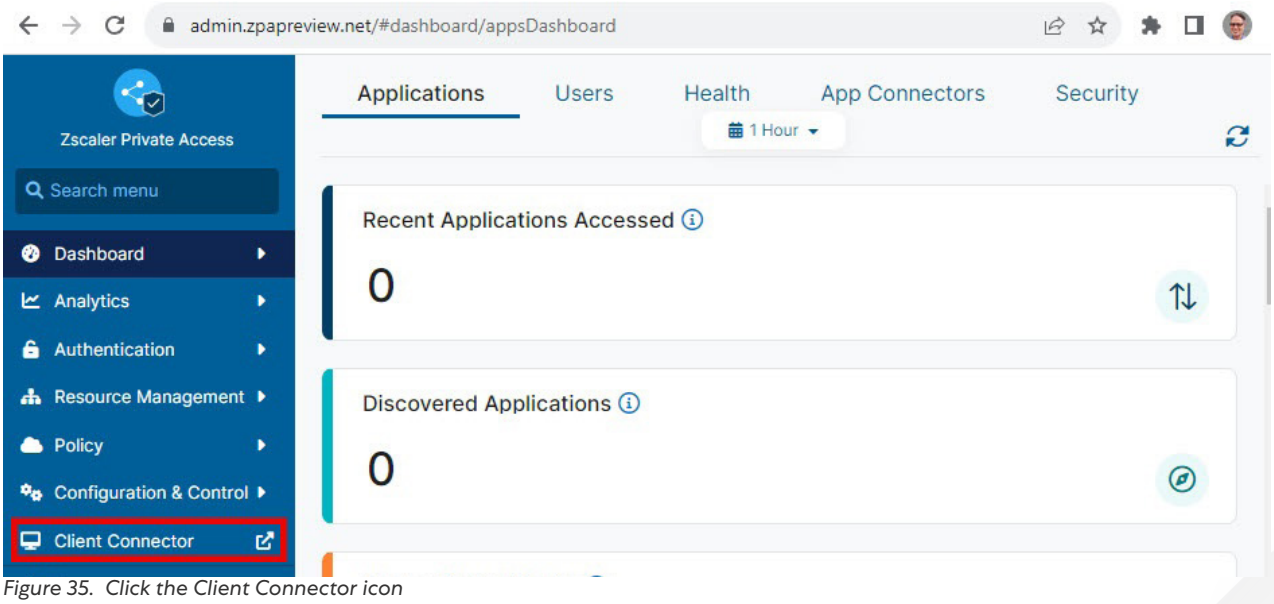

Figure 35. Click the Client Connector icon

Create a New Posture Profile

Log in to the Zscaler Client Connector and go to **Administration** > **Device Posture**. Then click **Add Device Posture Profile**.

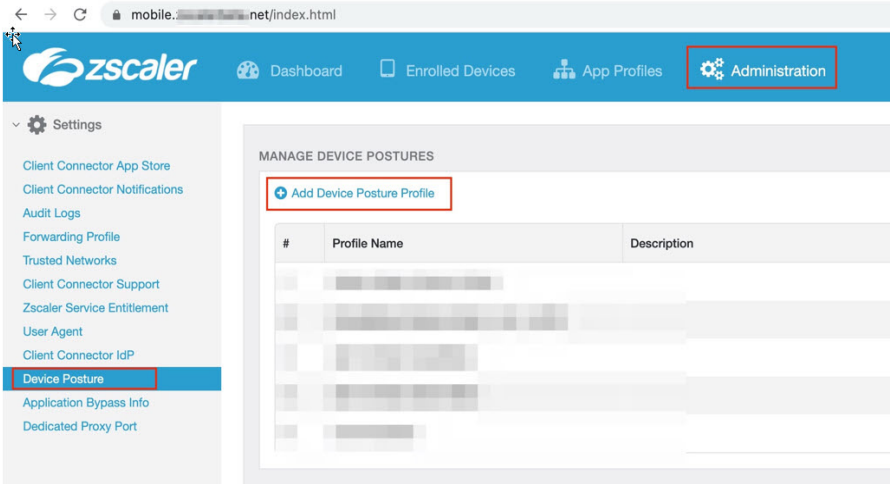

Figure 36. Add a Device Posture Profile in the ZPA Admin Portal

Add a New CrowdStrike Posture Profile

Complete the following steps:

1. Select only **Windows** and **macOS**.
2. Click the **Posture Type** drop-down menu.
3. Select **Detect CrowdStrike**.
4. Name this policy and click **Save**.

   This posture profile is referenced in a ZPA access policy. You can set up access policies to allow or deny application access based on whether the posture check passes or fails.

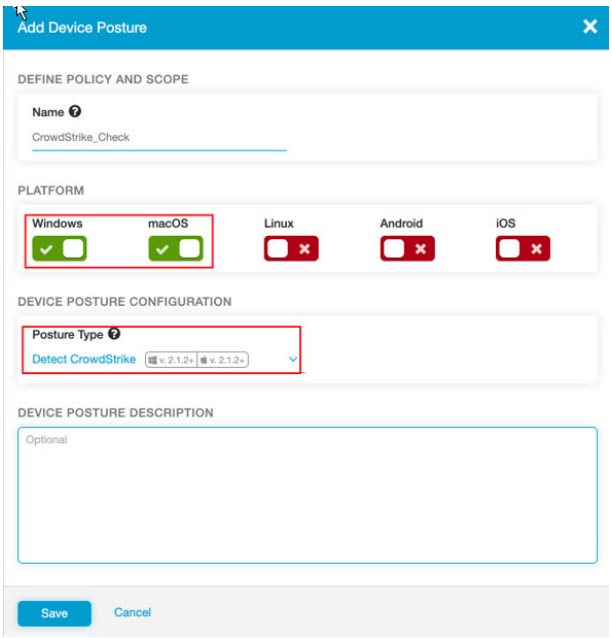

Figure 37. Add a detect CrowdStrike posture profile

Decide Which Applications Need Conditional Access

From the ZPA Admin Portal, go to **Resource Management** > **Application Management** > **Application Segments**.

This page lists which applications are accessed by ZPA. Select one of these applications and reference it in an access policy so that access to the application is granted based on the end device's posture.

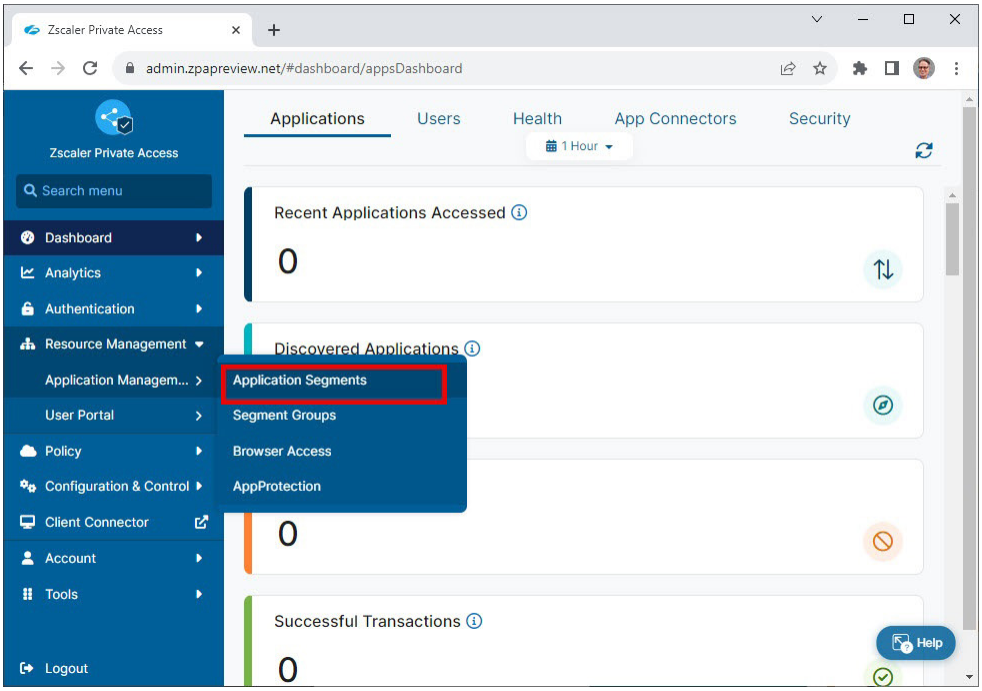

Figure 38. Go to Application Segments in the ZPA Admin Portal

In the following example, ZPA can access applications hosted under the domain *.bd-dev.com, based on the posture of the end device.

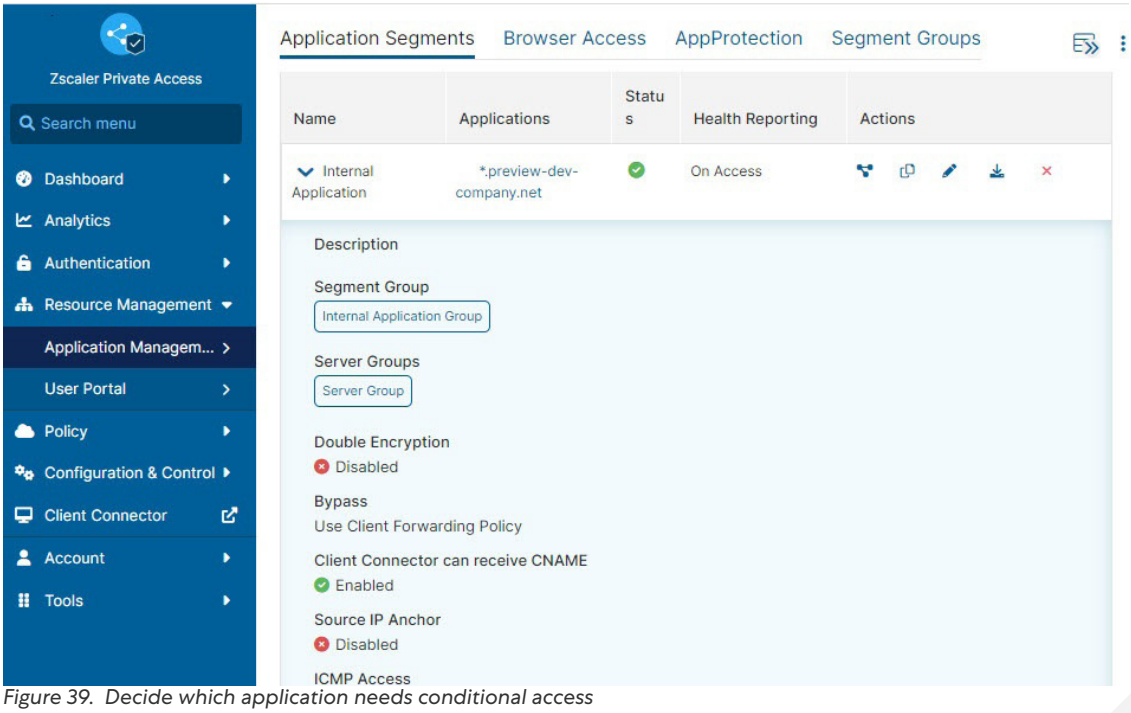

Figure 39. Decide which application needs conditional access

Set Up an Access Policy

From the ZPA Admin Portal, go to **Policy** > **Access Policy**.

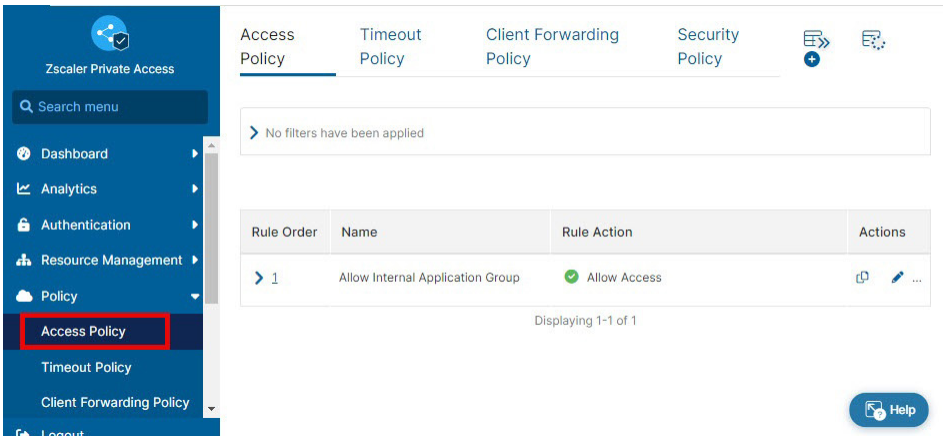

Figure 40. Open access policy configuration dialog

Tie the Posture Profile to this Access Policy

Create a new access policy by clicking **Add Rule** and reference the previously-created posture profile. You can set up different access policies to protect different internal applications. You can show a customizable (and optional) notification message to the end users when application access is allowed or denied, informing them about the policy evaluation.

In the following example, an access policy is added to block user access to the application if the CrowdStrike posture check fails (Rule#1). If CrowdStrike is not running on the endpoint, Rule#1 is marked true and access is blocked. Otherwise, the policy evaluation proceeds to Rule#2 (which grants application access).

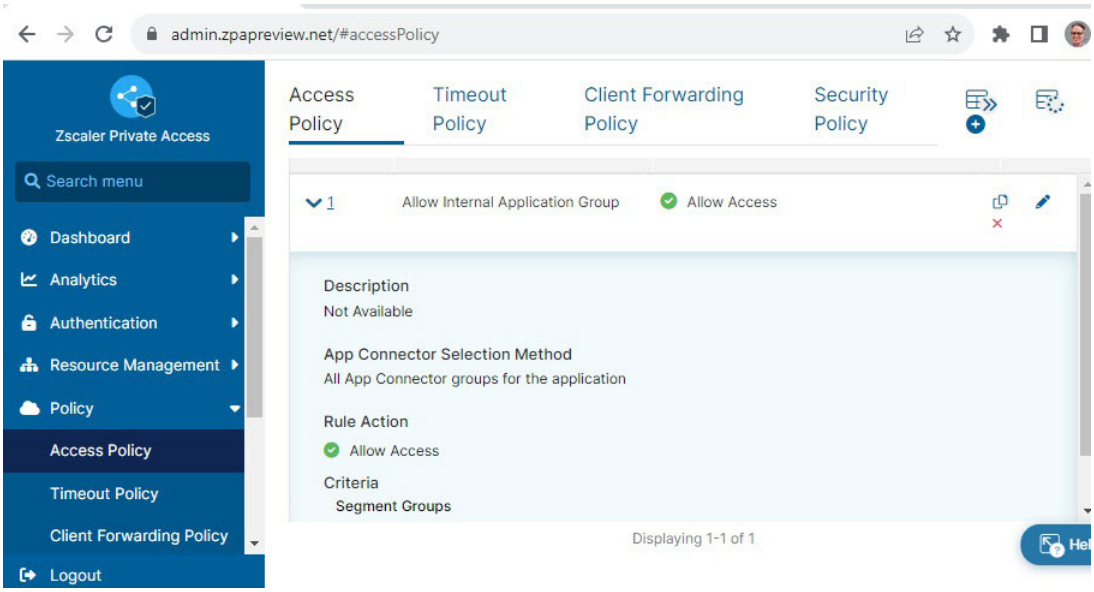

Figure 41. Set up an access policy

Verify Conditional Access from an Endpoint

The endpoint accesses the application if the endpoint device has a CrowdStrike agent installed and running. Otherwise, the access is blocked by ZPA.

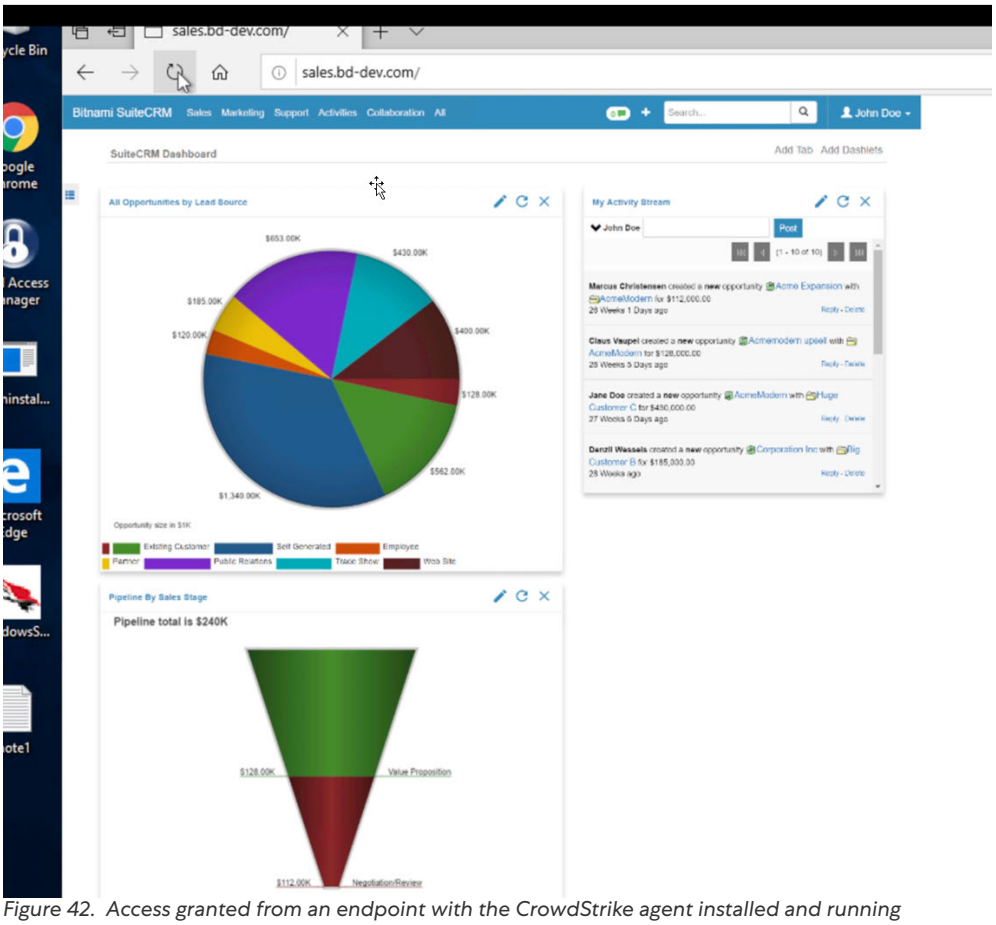

Figure 42. Access granted from an endpoint with the CrowdStrike agent installed and running

The following message is displayed if access is blocked because CrowdStrike isn't running:

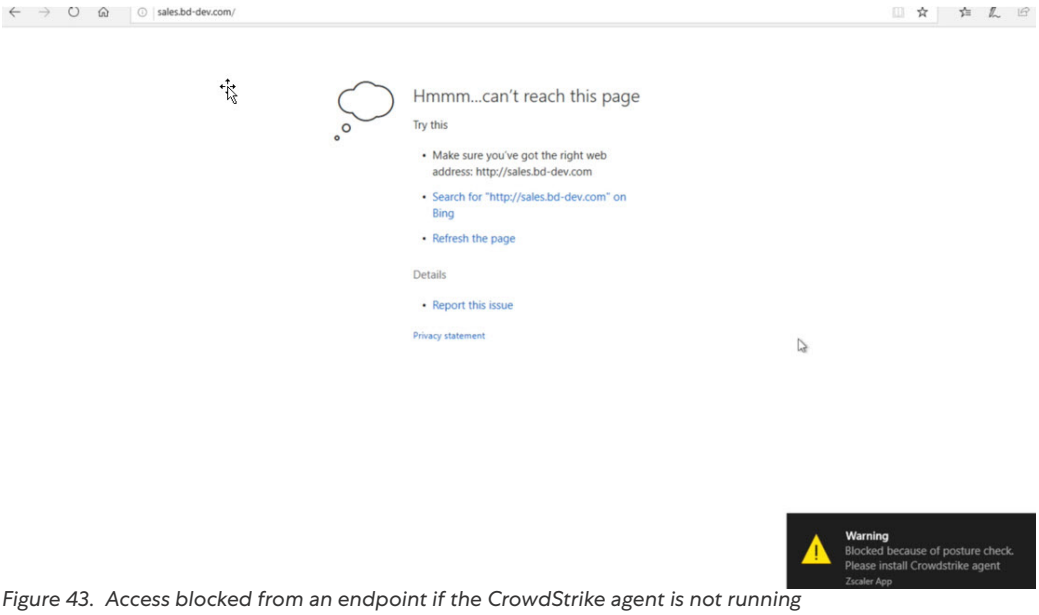

Figure 43. Access blocked from an endpoint if the CrowdStrike agent is not running

Use Case 4: ZIA Sandbox Integration with CrowdStrike

In this use case:

- ZIA Cloud Sandbox detects zero-day malicious files via Zscaler Cloud Sandbox and produces an insight log about the file hash. In the same report, you receive relevant CrowdStrike endpoint telemetry data. The endpoint data is retrieved dynamically via an API session established by a one-time setup process in the ZIA Admin Portal.
- The same report also includes a contain or quarantine action button, which enables the administrator to trigger a network contain or quarantine request to CrowdStrike Falcon platform. A network contained or quarantined host can talk to only CrowdStrike backend IPs and IPs explicitly placed on the allowlist by the CrowdStrike admin. All other network access is suspended.
- Alternatively, an administrator can click the CrowdStrike Agent ID within the insight log to access the Falcon Console. Then, in the Falcon Console, the administrator can further investigate and mitigate operations for that Agent ID

The following diagram shows a conceptualization of the integration.

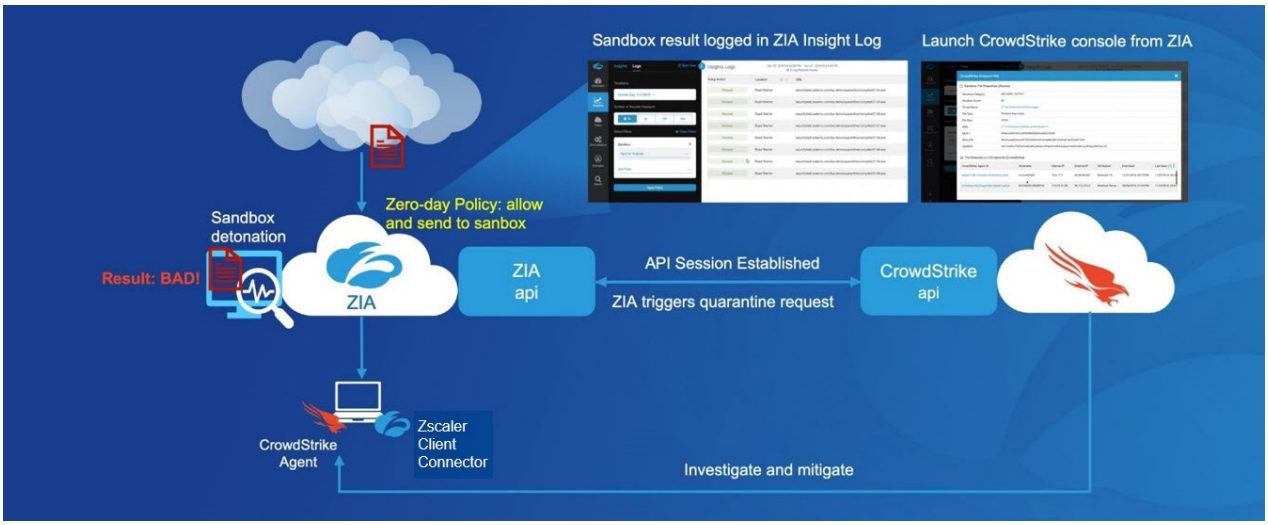

Figure 44. High level overview

Configuring CrowdStrike for ZIA

To establish the API connection between CrowdStrike and Zscaler, you must first generate an OAuth 2.0 token from the Falcon Console and then copy it to the ZIA Admin Portal.

Zscaler requires the following values to establish the API connection. You can get the values from the Falcon Console.

- CrowdStrike API Auth URL
- Client ID
- Secret
- Customer ID

The following steps assume that the CrowdStrike Falcon platform and CrowdStrike sensors are deployed and properly configured. Refer to the CrowdStrike documentation on how to deploy and configure CrowdStrike components before proceeding.

Log in to CrowdStrike

Log in to CrowdStrike using your administrator account. If you are unable to log in using your administrator account, contact CrowdStrike support.

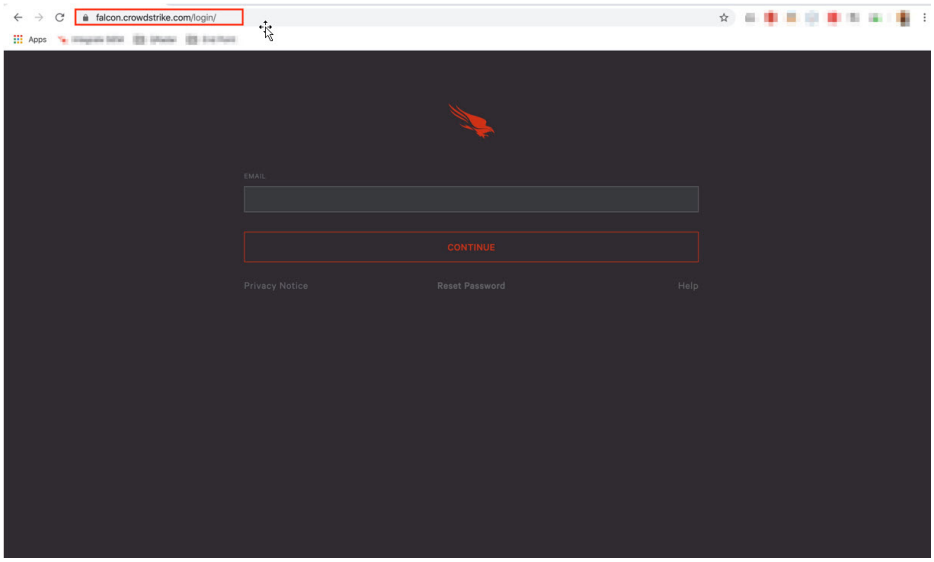

Figure 45. Log in to CrowdStrike

Access Your CrowdStrike Customer ID

After logging in to Falcon Console, click the User icon to access your Customer ID.

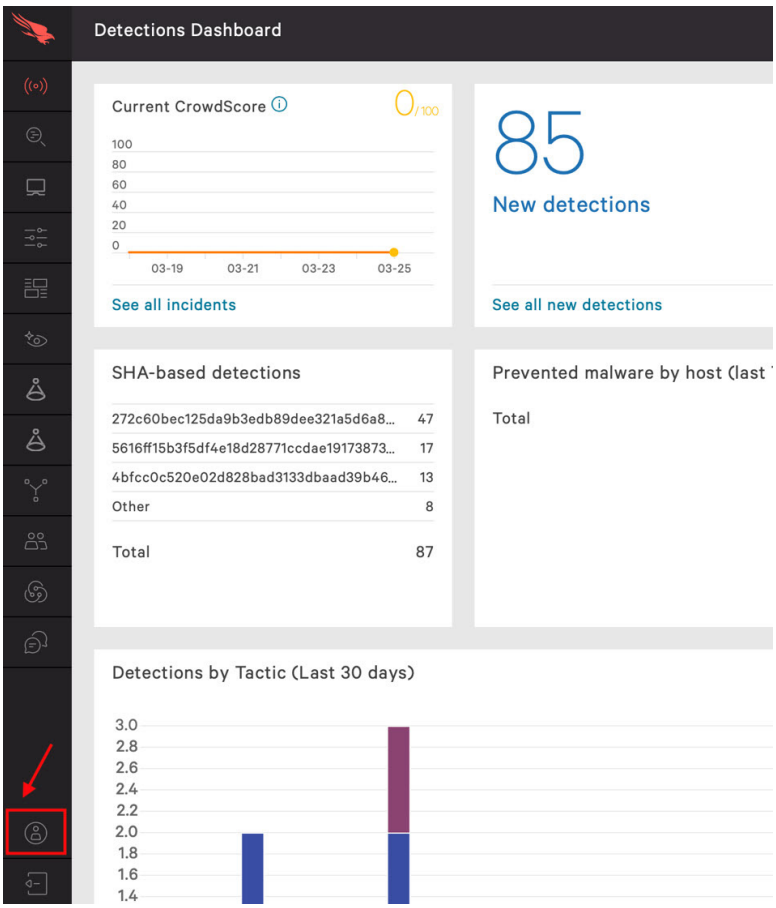

Figure 46. Click the User icon to access your CrowdStrike Customer ID

Note Your CrowdStrike Customer ID

You see your CrowdStrike Customer ID. You enter this ID into the ZIA Admin Portal later.

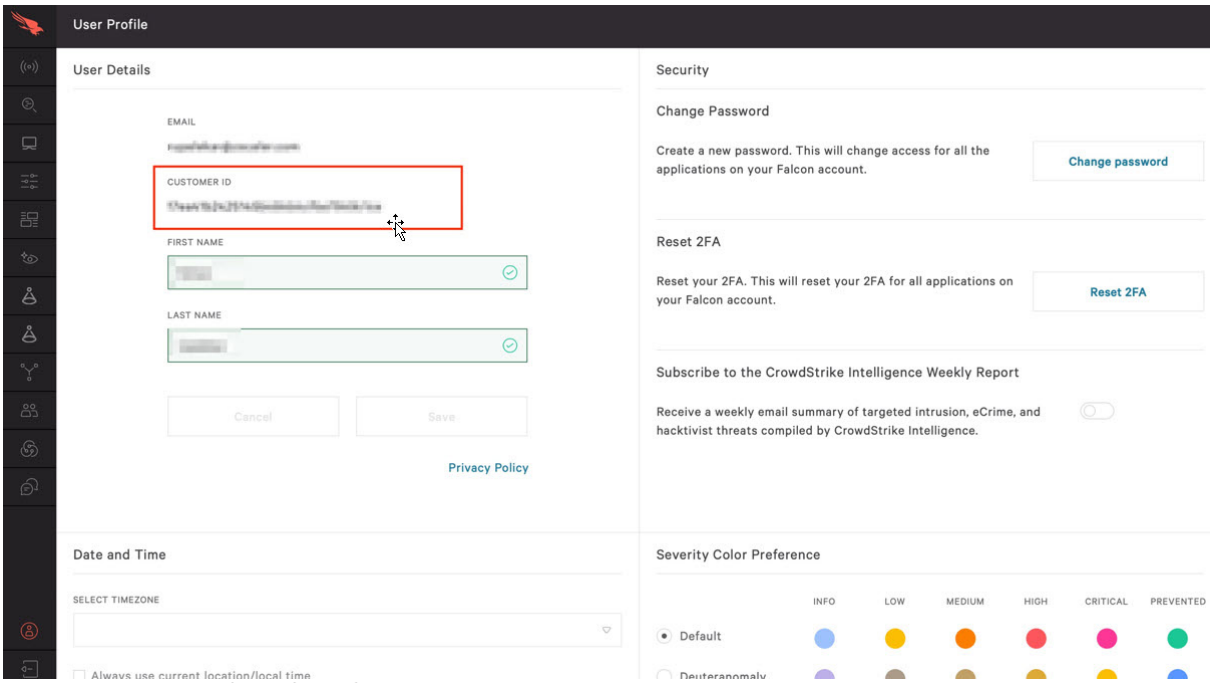

Figure 47. Access your CrowdStrike Customer ID

Go to the API Section

While still in the Falcon Console, go to **Support** > **API Clients and Keys**.

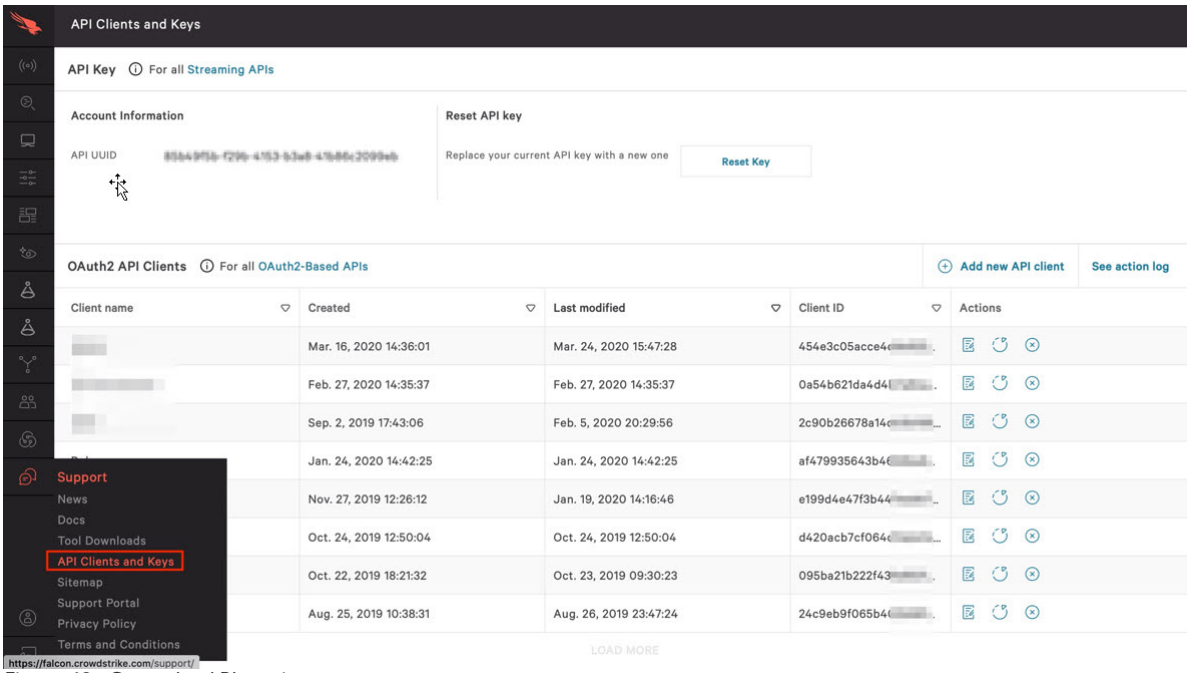

Figure 48. Go to the API section

Add a New API Client

In this use case, you create a new API client with specific permissions required for the use case. This is a one-time setup.

Click **Add new API client**.

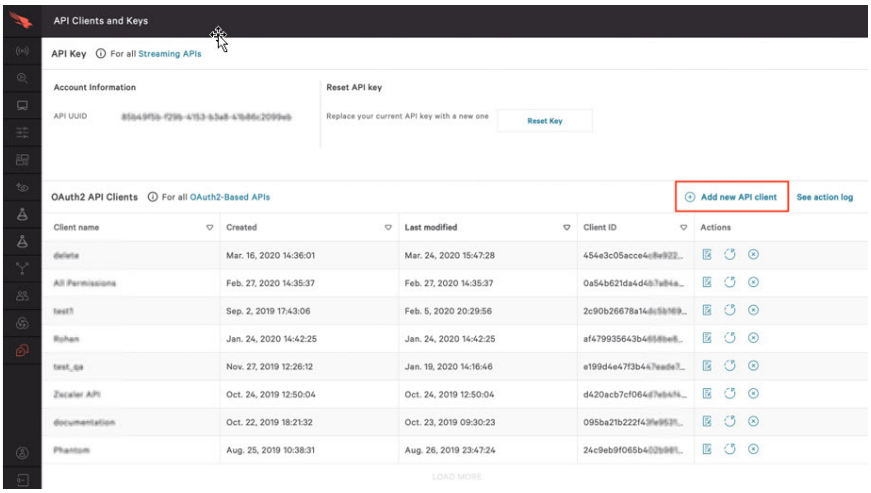

Figure 49. Add new API client

Create API Client for ZIA

Create an API client with following settings:

- Read-Write permission for Hosts (write permission is required for containment action)
- Read only permission for indicator of compromise (IoC)
- Read only permission for detections

When complete, click **Save**.

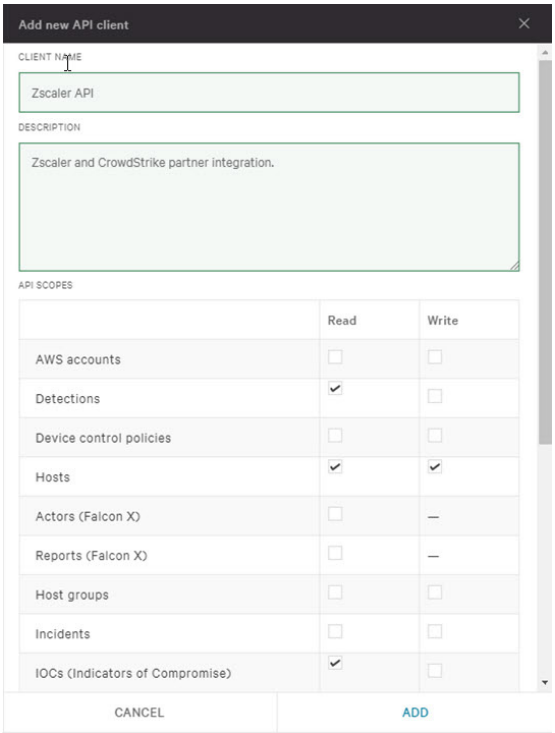

Figure 50. Create and save API client

Make a Note of the API Credentials

After the API Client is created, you can access a **Client ID** and a **Secret**. Note your **Secret** value and then click **Done**. You must provide the ID and secret in Zscaler.

You cannot re-access the Secret after you click **Done**. If you lose the **Secret**, you must reset the CrowdStrike API credentials.

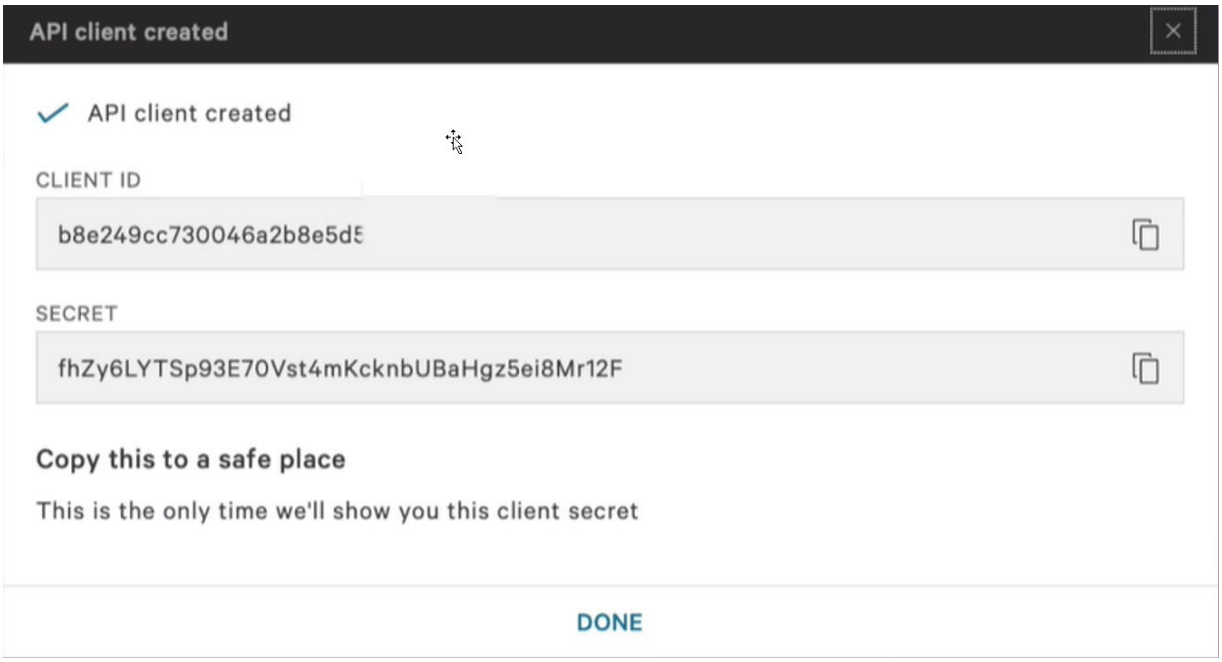

Figure 51. Note the API credentials

Allowlist Zscaler Hub IP Range Used to Make API Calls to CrowdStrike (Optional)

You can lock down which public IPs are allowed to make API calls to your CrowdStrike tenant. This step is optional. If IP-based restriction is in place, you must allowlist Zscaler hub IP ranges (since ZIA initiates API calls to CrowdStrike for this integration) based on which Zscaler cloud you use. The Zscaler cloud IP range must be added to the IP Allowlist Management under the Host Setup and Management section in the CrowdStrike tenant.

Zscaler hub IP ranges are found at https://config.zscaler.com/zscaler.net/hubs (government agencies, see https://config.zscaler.us/zscalergov.net/hubs). Filter for the cloud where your ZIA tenant resides. Allowlist IPs in the Recommended column for API call purposes.

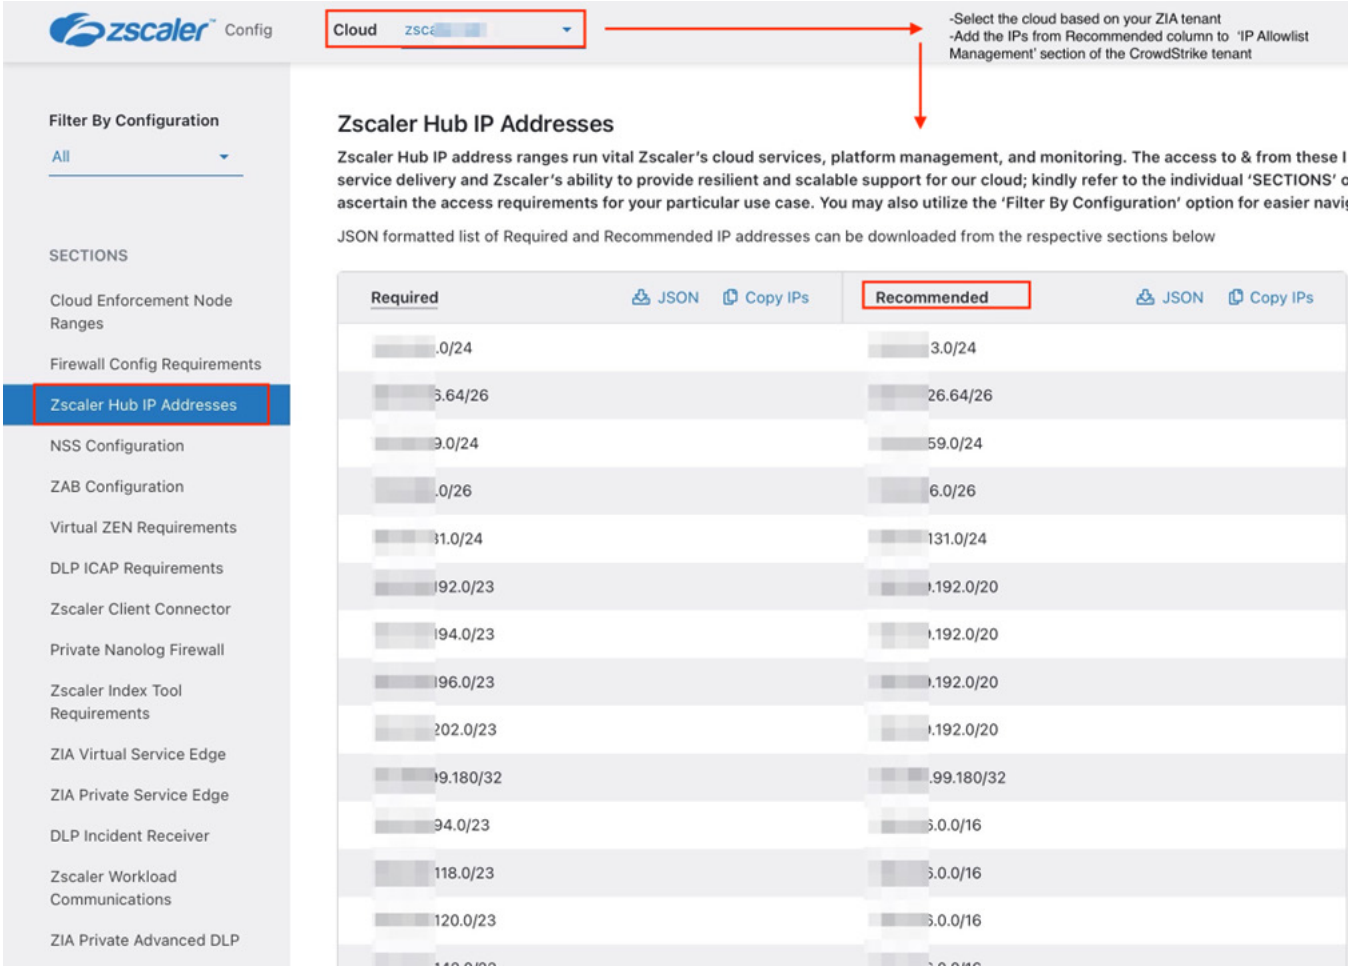

Figure 52. Zscaler IP range for API calls

Configuring ZIA for CrowdStrike

Endpoint telemetry data from the CrowdStrike Falcon Platform is passed to the ZIA Admin Portal via an API integration. Correlating the endpoint data enables the ZIA Admin Portal to display the Sandbox report. In addition, you see information about the originating endpoint device and other infected endpoints in the environment, including the following:

- CrowdStrike Agent ID.
- Host Name.
- Timestamps that capture when malicious files appear on the endpoint (e.g., an infection via a different attack surface, such as via a USB thumb drive).

This automatic correlation of malware detection with an endpoint device reduces the time and effort needed for investigation and remediation.

In this section, you configure the ZIA Admin Portal with the ID and Key generated in the previous section.

Log in to ZIA Admin Portal

Log in to ZIA Admin Portal using your admin account. If you are unable to log in using your administrator account, contact Zscaler Support (government agencies, see Zscaler Support).

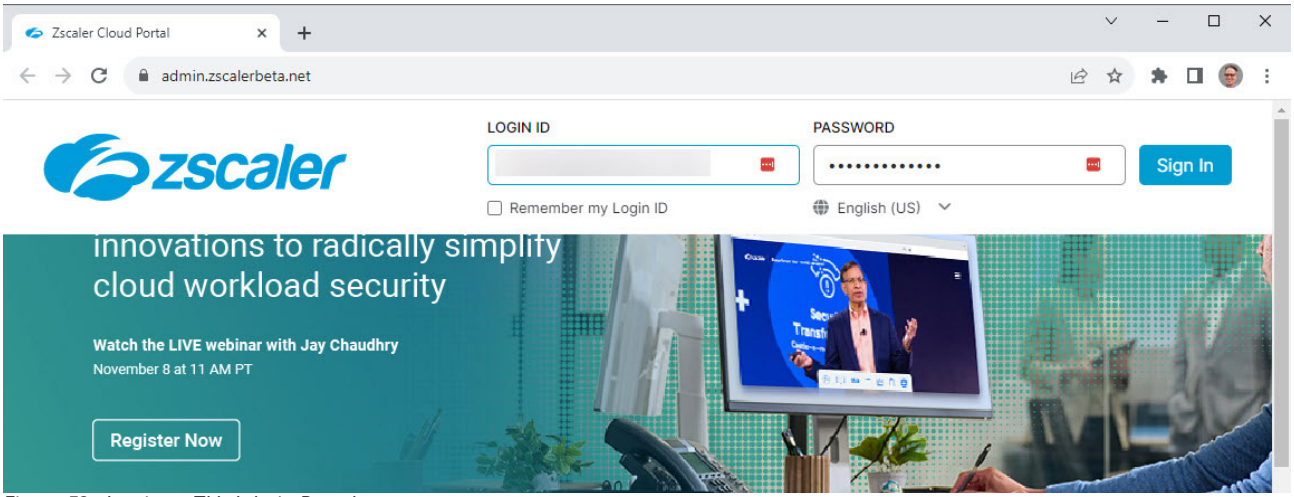

Figure 53. Log in to ZIA Admin Portal

Configure Partner Integration

From the ZIA Admin Portal, go to **Administration** > **Partner Integrations**. The **API Auth FQDN** URL depends on which CrowdStrike cloud you use.

Complete the following steps:

1. Paste or enter your CrowdStrike API credentials (**Client ID**, **Secret,** and **Customer ID**) in the appropriate fields.
2. Click **Save**. Then wait a few seconds for a status message.

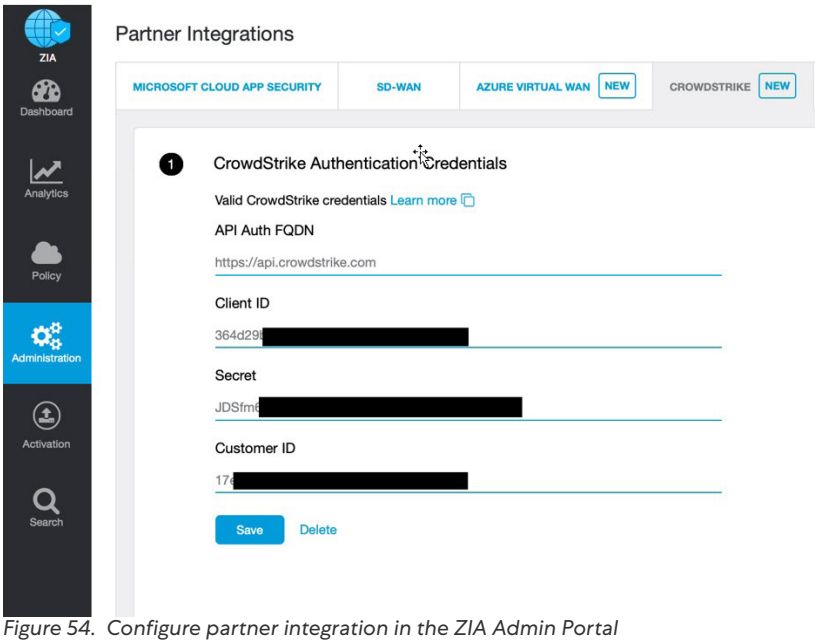

Figure 54. Configure partner integration in the ZIA Admin Portal

Verify the Partner Integration

When you see the message **Valid API token(s). The configuration is complete** in green background, you have successfully configured the API connection for the ZIA integration.

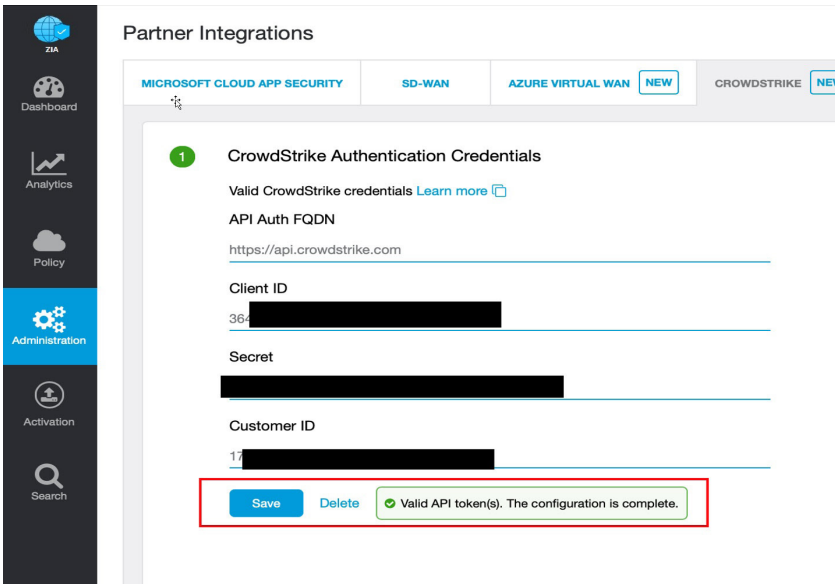

Figure 55. Verify partner integration

Activate Pending ZIA Configuration

Any time you make a change in ZIA, you see a number displayed over the **Activation** icon on the left-side navigation. This lets you know that you have changes pending in the queue for activation.

When you are ready to commit all changes in the queue, hover your cursor over the **Activation** icon and click **Activate**.

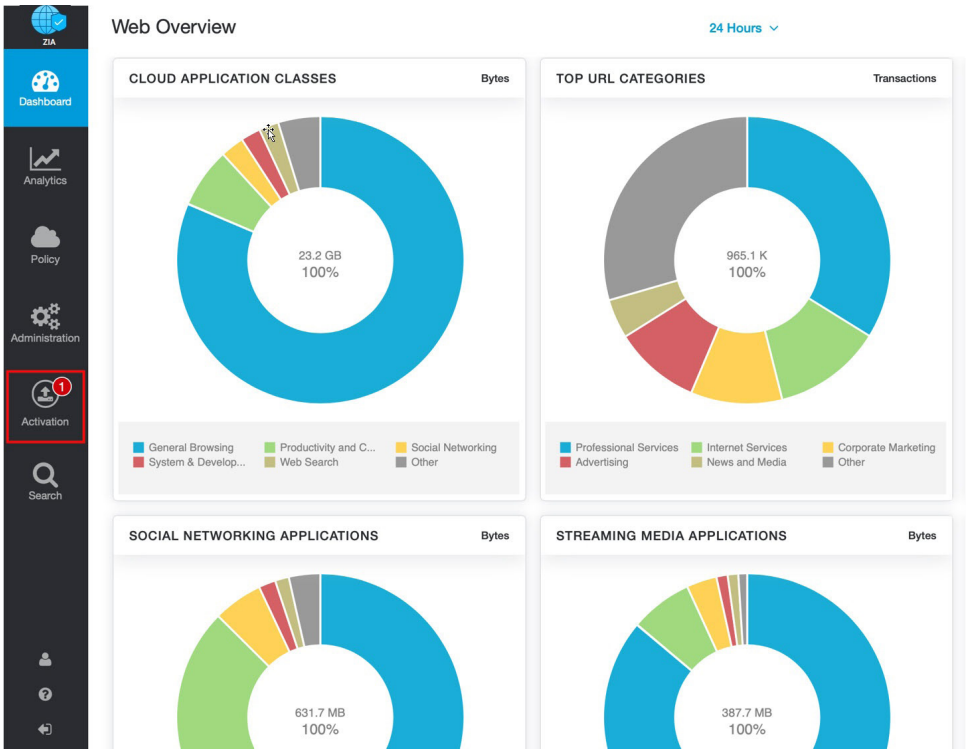

Figure 56. Activate pending ZIA configuration

Viewing CrowdStrike Endpoint Hits

Thanks to this integration, a file detonated by Zscaler Cloud Sandbox is automatically correlated with CrowdStrike endpoint device information within the ZIA Admin Portal.

Go to Web Insights

From the ZIA Admin Portal, complete the following steps:

1. Select the **Analytics** icon.
2. Click **Web Insights**.

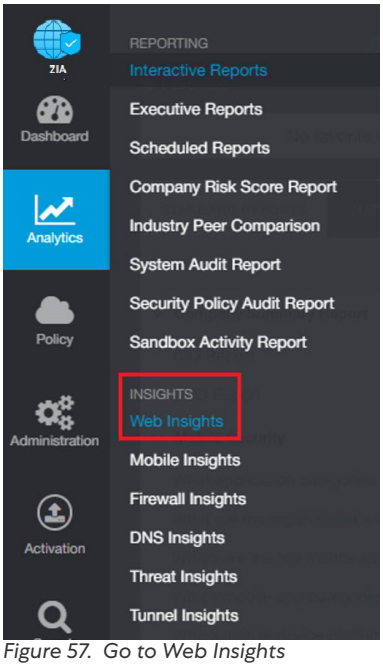

Figure 57. Go to Web Insights

Select Logs

On the **Logs** tab, click **Add Filter**.

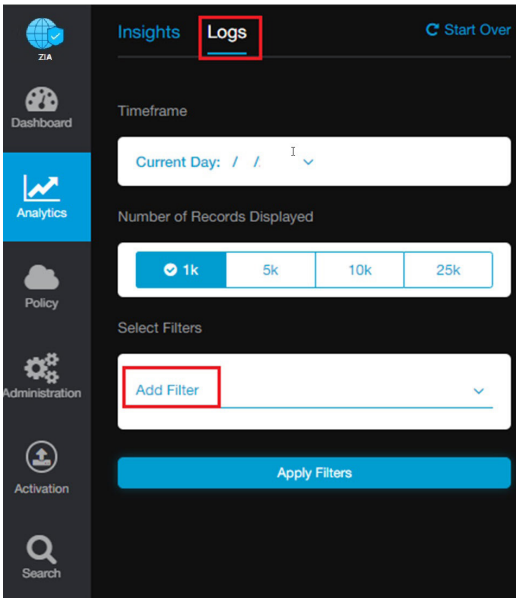

Figure 58. Select Logs

Filter on Sandbox Related Logs

Select **Sandbox** as the **Threat Class** and click **Apply Filters**.

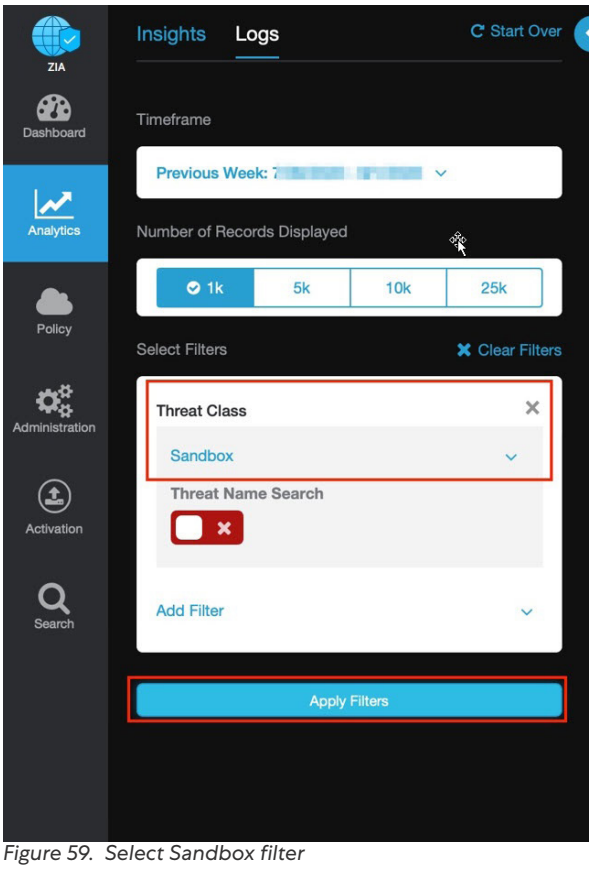

Figure 59. Select Sandbox filter

Confirm the File was Sent to the Sandbox

After clicking **Apply Filters**, if the file in question was detonated or is currently being detonated by the Cloud Sandbox, corresponding log entries are listed on right side of the portal.

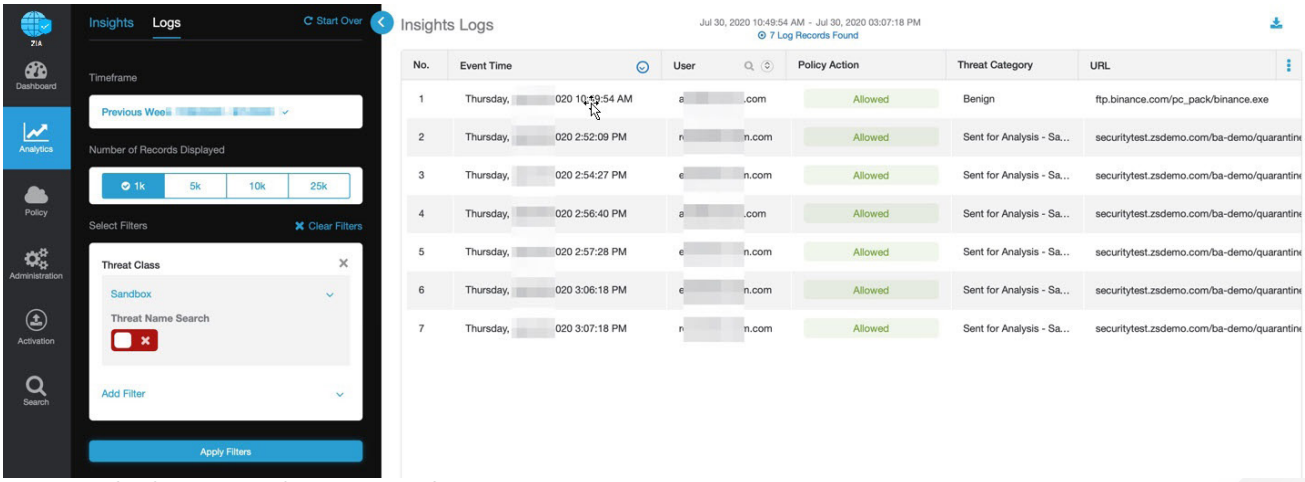

Figure 60. Confirm whether file was sent to Sandbox

Access Sandbox Report

Within the list of log entries, select an MD5 hash, and right-click the entry to access the drop-down menu. Select **View Sandbox Detail Report**.

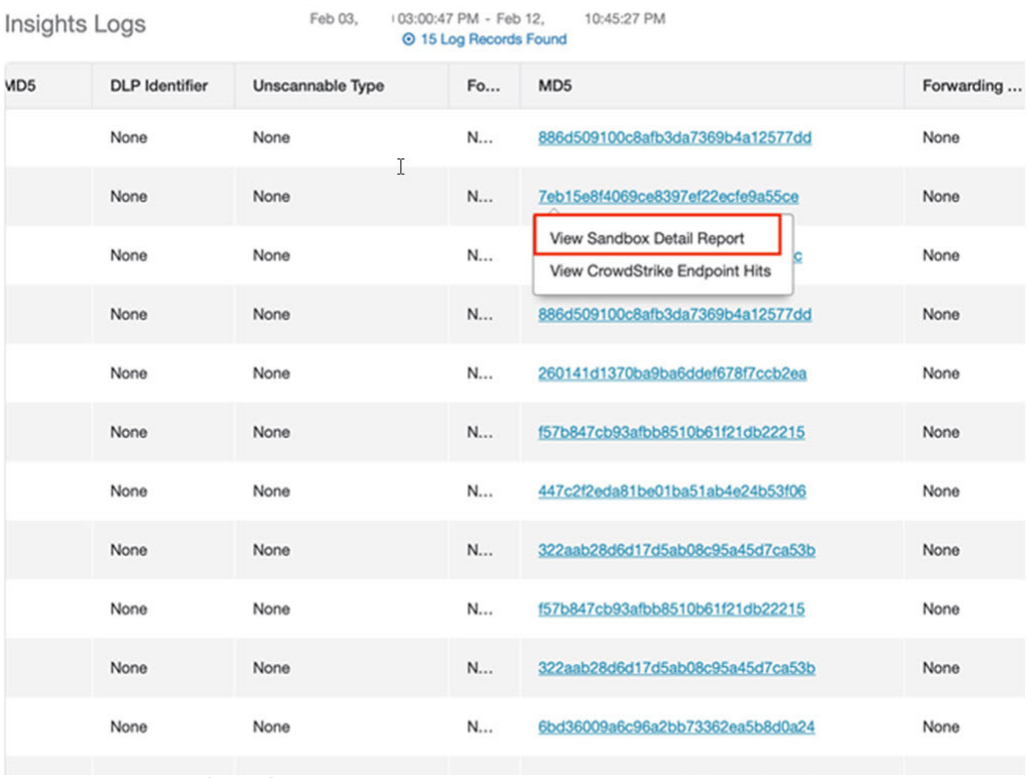

Figure 61. Access the Cloud Sandbox report

Sandbox Detail Report

The Sandbox Detail report provides detailed information regarding file detonation results.

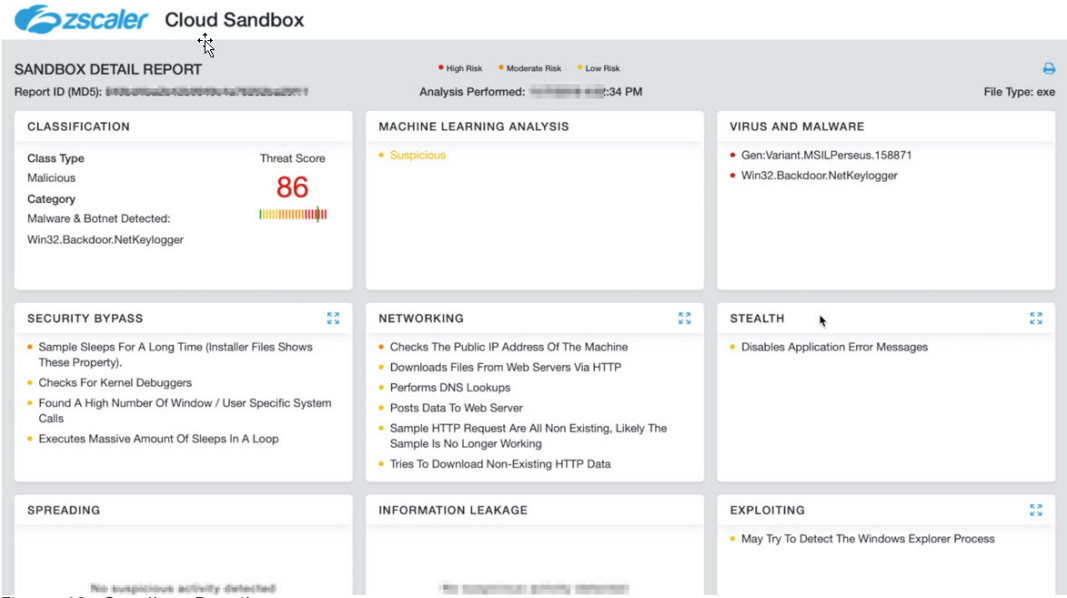

Figure 62. Sandbox Detail report

Access the CrowdStrike Endpoint Hits Report

Within the list of log entries, select the same or a different MDS entry, and click the entry to access the drop-down menu. Select **View CrowdStrike Endpoint**.

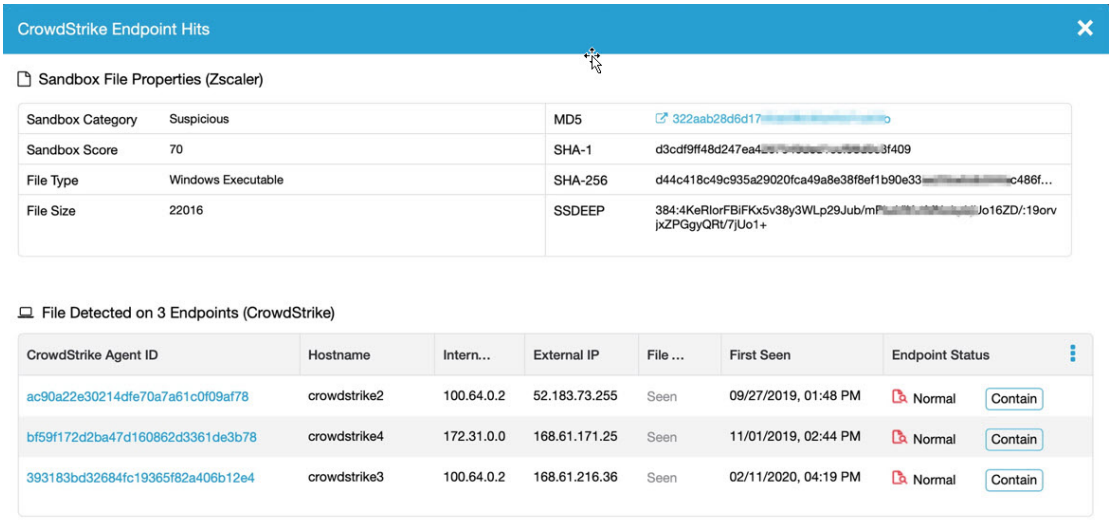

Figure 63. CrowdStrike Endpoint Hits report

Network Contain an Endpoint

Click **Contain** to trigger an API call to CrowdStrike. When you trigger the API to contain the endpoint, CrowdStrike prevents the endpoint from connecting to the network. This isolates the endpoint and prevents it from allowing malicious software on the endpoint from accessing the rest of the network via the network connection.

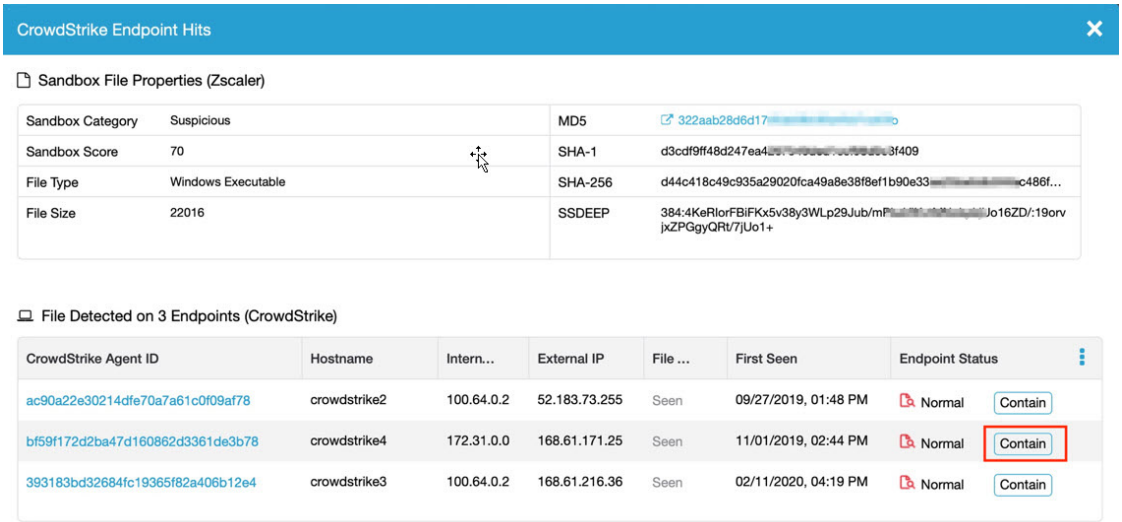

Figure 64. Contain an endpoint

Confirm Endpoint Quarantine Status

Access the CrowdStrike Endpoint Hits report again to confirm the containment status.

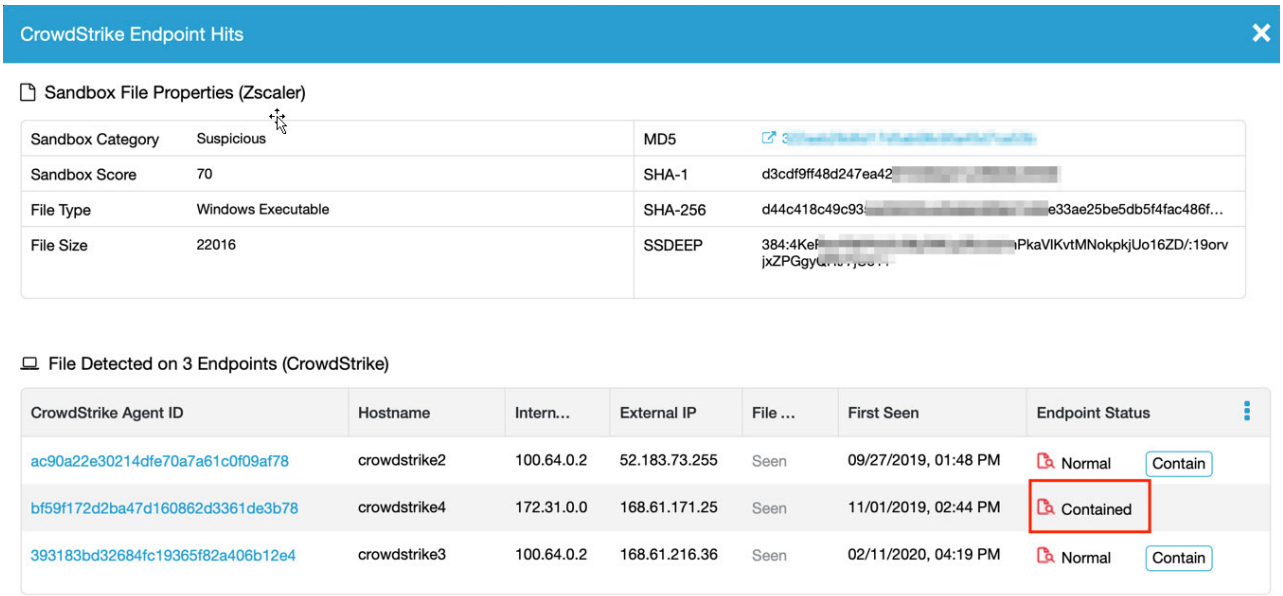

Figure 65. Confirm containment status

Use Case 5: Threat Intelligence Sharing—CrowdStrike Falcon and ZIA

When CrowdStrike's Falcon threat intelligence data is shared with Zscaler Zero Trust Exchange, seamless usage integrations provide stronger protection and increased visibility.

When a mutual customer of Zscaler and CrowdStrike activates this integration, the integration fetches malicious IPs or URLs from CrowdStrike's Intel platform and pushes them to that customer's ZIA custom URL list. This custom URL list is then referenced by ZIA URL policies to block the end-user access. Before the push happens from CrowdStrike to ZIA, all the indicators of compromise (IoCs) are checked against Zscaler's global IoC database and only IoCs that ZIA doesn't currently qualify as malicious are pushed into the ZIA tenant.

You must download and run Python code (hosted on GitHub) in order for the integration to work.

The integration is maintained by CrowdStrike. Updates and improvements are added to the GitHub page and the GitHub repository is the authoritative and latest resource.

Overview

Use Case:

- ZIA maintains a global database of malicious IPs, Domains, or URLs (i.e., IoCs) and blocks these threats inline in all ZIA customer tenants if pertinent security engines are enabled by ZIA admins. ZIA also maintains per-tenant custom URL lists. You can bring in your own custom threat feeds and populate these URL lists. You can then reference these custom URL lists in ZIA URL policies for granularly controlling end-user access within that ZIA tenant.
- CrowdStrike Falcon expands your defenses with real-time access to global IoCs delivered by CrowdStrike. An existing CrowdStrike Falcon intelligence and ZIA customer can set up this integration to continually push high value threats from the Falcon Platform into their ZIA tenant.

The following diagram shows a conceptualization of the integration.

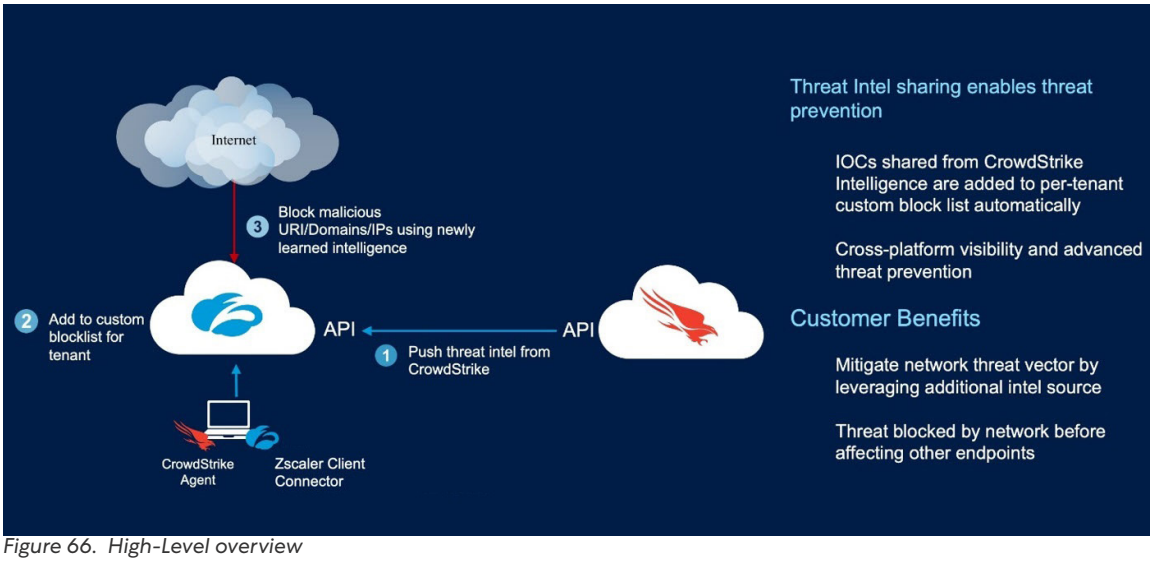

Figure 66. High-Level overview

The new version of the CrowdStrike Falcon ZIA Intel Bridge features a codebase overhaul rewritten for legibility, ease of modification, and simplicity in troubleshooting. CrowdStrike Falcon programmatic operations provide a more modularized source code and comprehensible modification process.

Furthermore, improved logging capabilities and error handling simplifies troubleshooting and support. Any user challenged with excessive runtime troubles can find detailed logs in the /logs directory. Including the log information in a support case ensures efficient, effective remediation.

CrowdStrike Falcon approaches this use case slightly differently. Originally, the CrowdStrike Falcon Indicator database included 175,000 high-confidence, malicious URLs, which was compatible with Zscaler's 275K custom URL limit. Over the past year and a half, the number of high-confidence malicious URLs in the CrowdStrike Falcon Indicators database has grown to several million. A different approach to maintaining the Intel Bridge was required.

The integration version ensures that multiple indicators are present in the ZIA custom URL category. You can configure the number of indicators to your environment's unique limit. ZIA users get the latest indicators from CrowdStrike's intelligence team. By the time those indicators start to appear in other third-party feeds, the Intel Bridge already updates the ZIA URL category with the most recent indicators available.

Prerequisites

- License for CrowdStrike Falcon intelligence
- Admin access to Falcon Console
- Admin access to ZIA Admin Portal
- Python 3.7+ environment

You can push the default number, 25K, of custom IoCs into a ZIA tenant.

You can increase the IoC quota (with a maximum up to 275K) by contacting your Zscaler Account team and purchasing licenses for additional custom URLs.

The integration scripts consider the `urllookup` API rate limit on the ZIA side (~40K lookups per hour) and throttles the lookups to avoid running into any rate limit issues.

Create Zscaler URL Category

To create a Zscaler URL category:

1. Log in to your ZIA Admin Portal.
2. Go to **Administration** > **URL-Categories**.
3. Add a new URL category with the name **CrowdStrike**.
4. In the **URL Super Category** drop-down menu, select **User-Defined**.
5. Click **Save**.

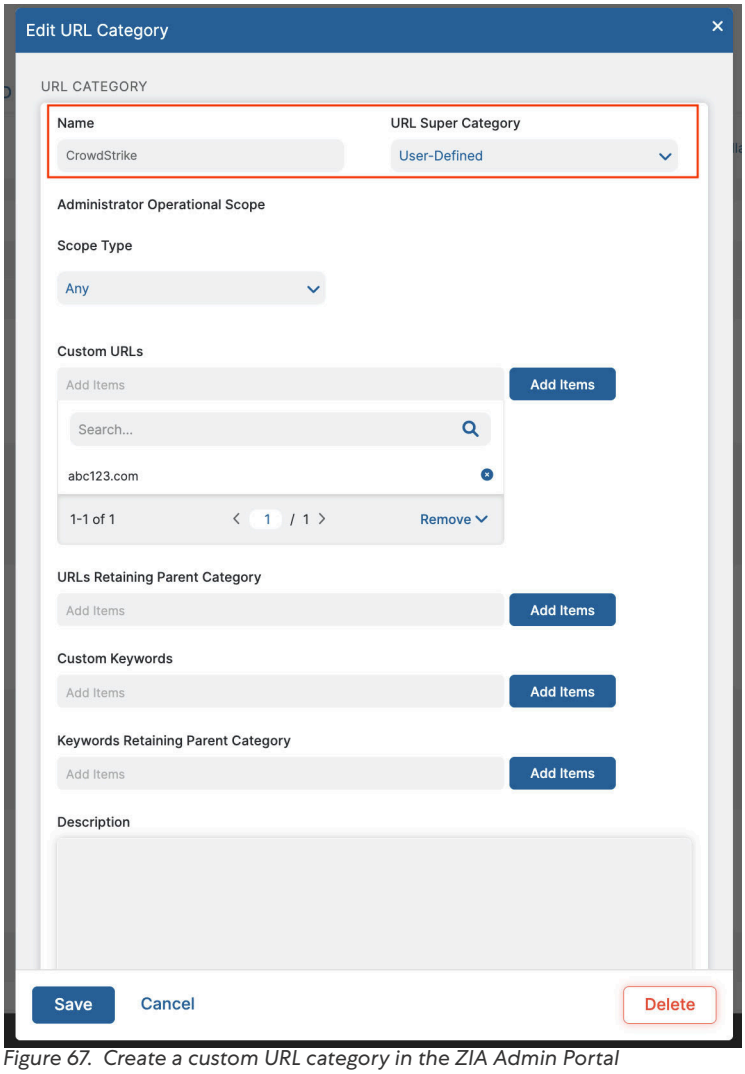

Figure 67. Create a custom URL category in the ZIA Admin Portal

CrowdStrike OAuth 2.0 Token Scope

In the CrowdStrike Falcon Console:

1. Go to **API Clients and Keys**.
2. Click **Add a New API Client**.
3. Create a client with READ permissions for **Indicators (Falcon X)**.
4. Save the resulting values. You need them to run the integration.

Download Repository

```
git clone https://github.com/CrowdStrike/zscaler-FalconX-integration.git
```

cd zscaler-FalconX-integration

Install Dependencies with pip3

```
pip3 install -r requirements.txt
```

Configure

Input your configurations in config.ini. Do not use quotes or ticks for any of these values.

Most of the fields are self-explanatory, but be sure to put some thought into the LIMIT field. This field determines how many malicious URLs the Intel Bridge maintains in your ZIA tenant. Zscaler offers different subscription tiers with varying maximum custom URLs (from 25K to 275K). Consider this and your existing custom URL categories when you choose a value, as going over the limit causes runtime errors.

For example, if you have a limit of 25K, and are already using 10K in another URL category, a value like 14000 won't exceed the limit and leaves some overhead.

```
[CROWDSTRIKE]

client=Your Falcon API Client ID

secret=Your Falcon API Client Secret

base_url=Your Falcon API Base URL (ex: https://api.crowdstrike.com)

limit=Number of indicators to maintain (Max: 275,000 Default 10,000)

[ZSCALER]

hostname=Your zscaler Hostname (Hostname only requires the base URL (i.e. https://zsa-
pi.zscalerthree.net))

username=Your ZIA Username

password=Your ZIA Passsword

token=Your ZIA API token

[CHRON]

disable_loop=Change this value to 1 if you are running the Intel Bridge via Chron job.
This will force the program to quit after running. (Default 0, looping enabled)

[LOG]

log_indicators=Change this value to 1 for indicators to be logged in logs/data_log as
they are deleted and loaded.
```

Running the Integration

With Python 3.7+ installed:

```
python3 intelbridge
```


Further Details About the Integration

The integration is maintained by CrowdStrike. Updates and improvements are added to the GitHub page. The GitHub repository is the authoritative and latest resource.

Find instructions on how to get the ZIA API key in the ZIA Cloud Service API Developers Guide (government agencies, see ZIA Cloud Service API Developers Guide) section of the online documentation.

Use Case 6: CrowdStrike Humio Essential Configuration (Using Cloud-to-Cloud logging - HTTPS POST)

This section details the steps required to stream ZIA logs to a CrowdStrike Humio cloud from Zscaler, using Humio's HTTP/HTTPS API-based log ingestion functionality.

Cloud NSS is Zscaler's cloud-to-cloud log streaming service that allows you to stream logs directly from the ZIA cloud into a supported SIEM, without the need to deploy, manage and monitor an NSS VM for web or Firewall. The service supports all ZIA log types: web, SaaS security, tunnel, Firewall, and DNS.

For the CrowdStrike Humio cloud, the log ingestion API is the HEC input (`/api/v1/ingest/hec/raw`).

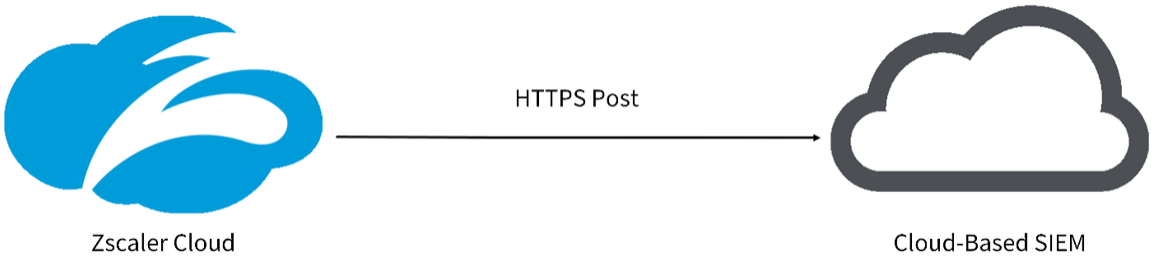

Figure 68. High-level overview of cloud-to-cloud logging

You can subscribe to Cloud NSS, which allows direct cloud-to-cloud log streaming for all types of ZIA logs into a CrowdStrike Humio instance.

To learn more, see:

- About Nanolog Streaming Service (government agencies, see About Nanolog Streaming Service).
- About Cloud NSS Feeds (government agencies, see About Cloud NSS Feeds).
- Adding Cloud NSS Feed for Web Logs (government agencies, see Adding Cloud NSS Feed for Web Logs).

Configure CrowdStrike Humio Cloud to Ingest ZIA Logs over HEC Input

This section requires that you have admin access to a working instance of CrowdStrike Humio cloud.

Installation of the package is straightforward, and the installation deploys the parser, saved queries, and dashboards directly into the repository that you select. You must add five ingest tokens, one for each NSS log source assigned to its corresponding parser.

Then create a new cloud NSS feed in the ZIA Admin Portal dashboard for the sources included in this package. Configure these feeds with your Humio API URL and ingest token.

Log into CrowdStrike Humio Cloud Tenant

Log in to CrowdStrike Humio tenant via the online portal.

1. Go to humio.com and select **Cloud login** > **Enterprise Login**.

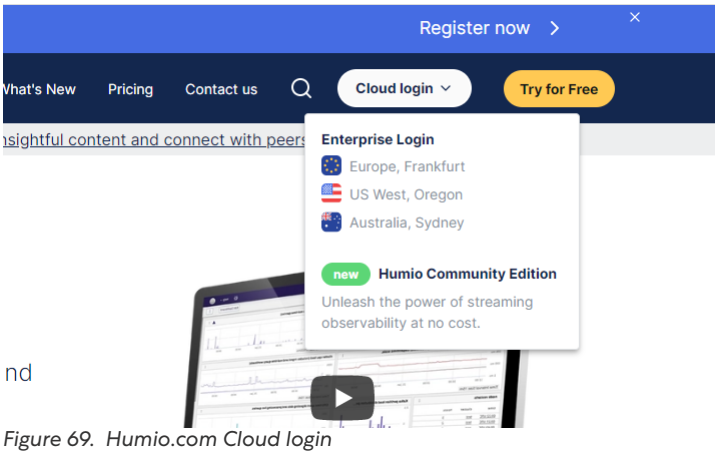

Figure 69. Humio.com Cloud login

2. Select the **Region** from the drop-down menu, then select your login option.

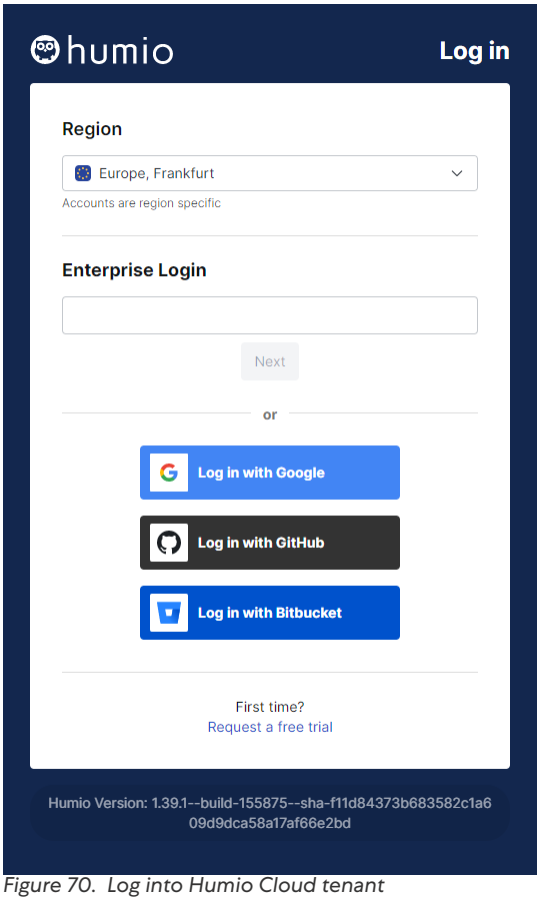

Figure 70. Log into Humio Cloud tenant

3. Log in to the cloud with admin credentials.

Install Zscaler Package in Your Cloud Tenant

After logging in:

1. Go to **Settings** > **Packages** > **Marketplace**.

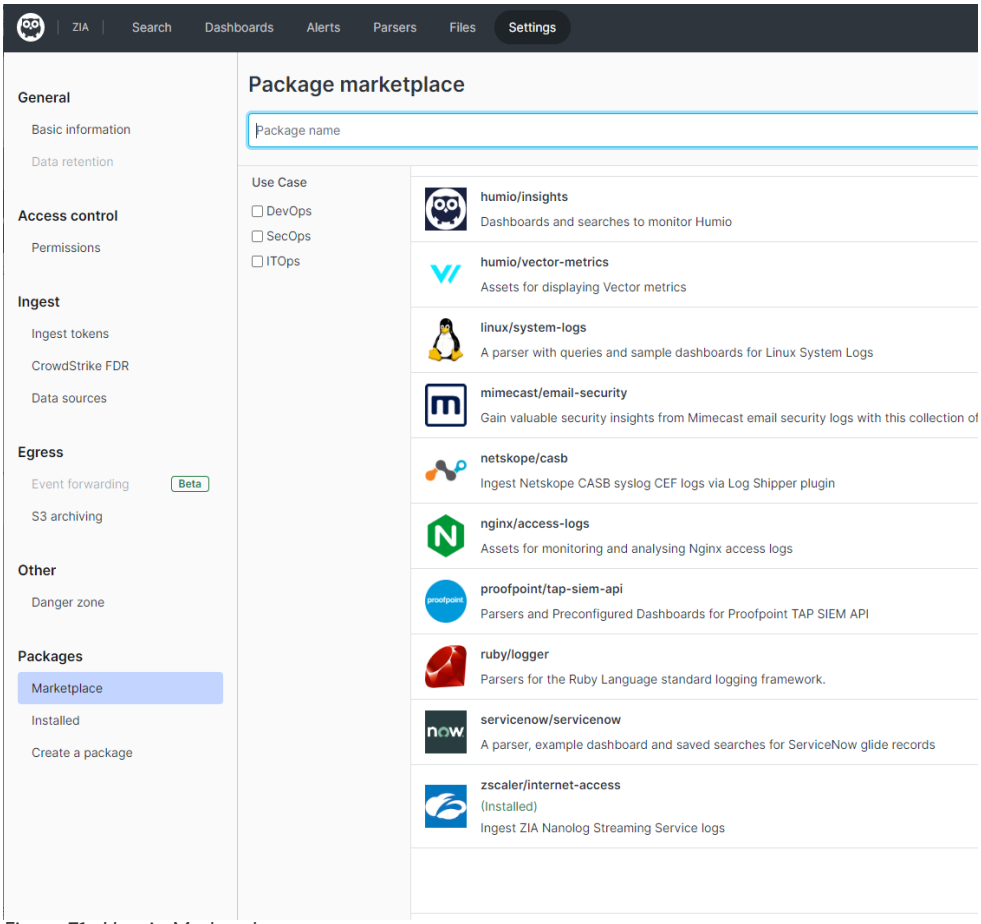

Figure 71. Humio Marketplace

2. Select **zscaler/internet-access** from the list of apps.

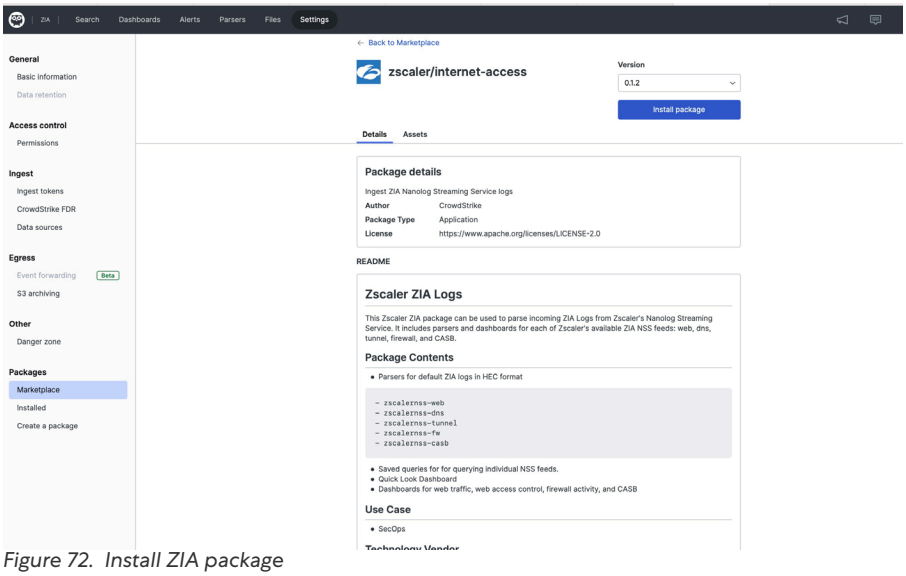

Figure 72. Install ZIA package

3. Click **Install package**. A list of app features is displayed.

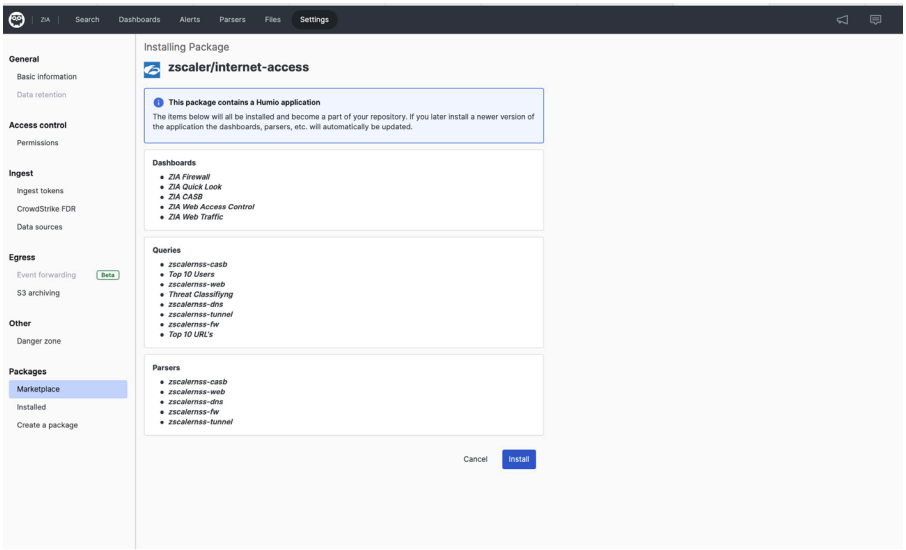

Figure 73. List of ZIA package features

4. Click **Install**.

Create and Add Ingest Tokens in CrowdStrike Humio

After installing the ZIA package:

1. Go to **Settings** > **Ingest** > **Ingest tokens**.
2. Click **+ Add token**.

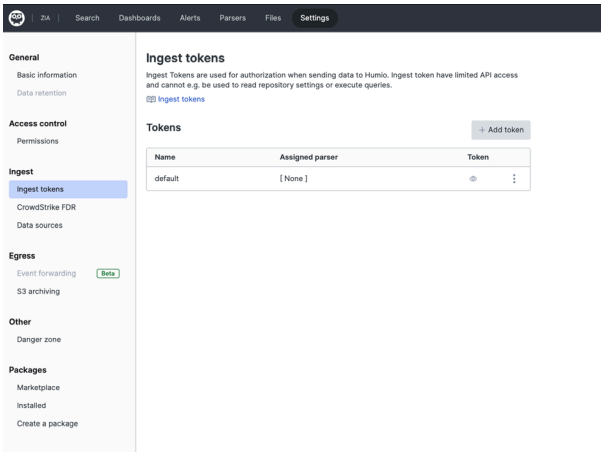

Figure 74. Ingest tokens window

3. In the **New token** window:
   - **Token name**: Enter an intuitive name for the token.
   - **Assigned parser**: Select the corresponding parser from the drop-down menu depending on which logs you want ZIA to send to Humio:
     - zscalernss-web
     - zscalernss-dns
     - zscalernss-tunnel
     - zscalernss-fw
     - zscalernss-casb

4. Click **Save**. The **Token** dialog displays with the token value.

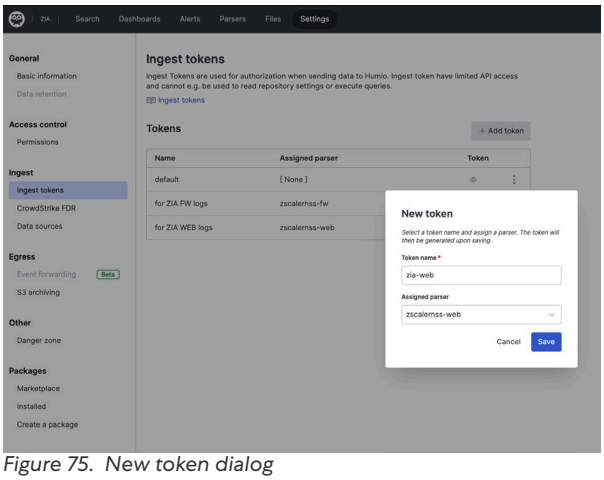

Figure 75. New token dialog

5. Copy this token value to paste into Zscaler when creating the cloud NSS feed for Humio.

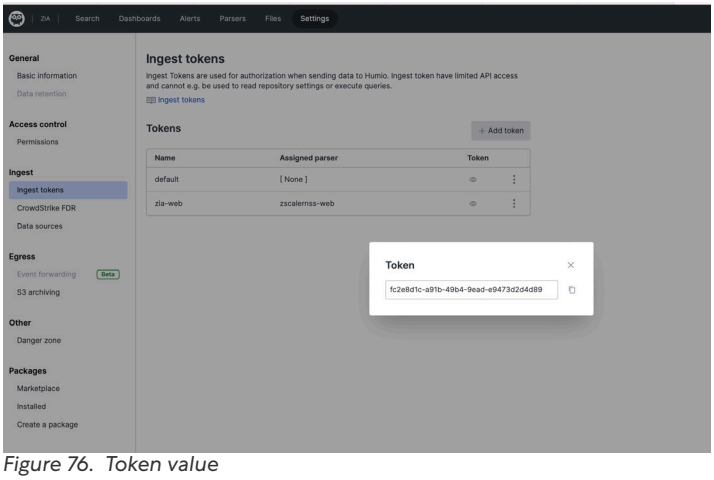

Figure 76. Token value

Configure Zscaler for Cloud-to-Cloud Logging

You can subscribe to Cloud NSS, which allows direct cloud-to-cloud log streaming for all types of ZIA logs into a Humio instance. Rather than deploying, managing, and monitoring on-premises NSS VMs, you can configure an HTTP/HTTPS API feed that pushes logs from the Zscaler cloud service into an HTTP/HTTPS API endpoint on the SIEM. The following steps show how to set up the log feed for web logs. Repeat these steps to set up other Zscaler log types (e.g., Firewall or DNS logs).

Go to Cloud-to-Cloud Logging Section in ZIA Portal

From the ZIA Admin Portal, go to **Administration** > **Nanolog Streaming Service** > **Cloud NSS Feeds** > **Add Cloud NSS Feed**.

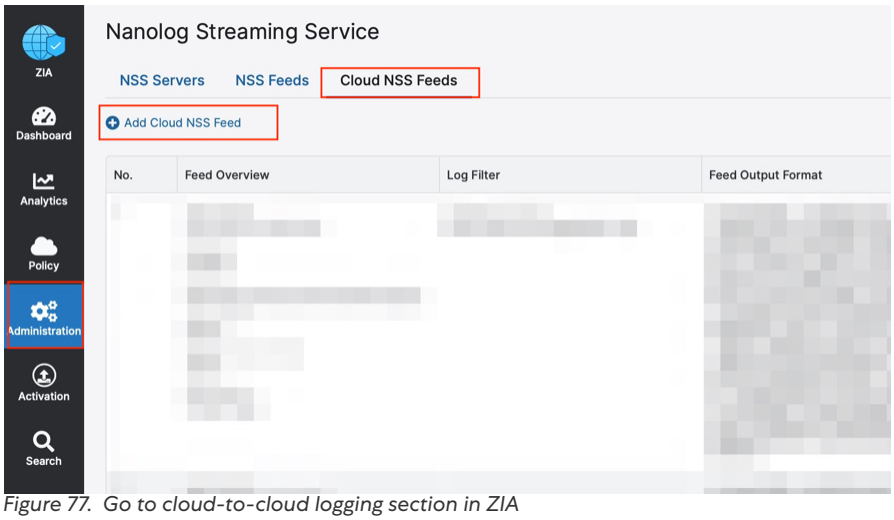

Figure 77. Go to cloud-to-cloud logging section in ZIA

Set Up the Cloud NSS Log Feed (Web)

1. Select **Other** as the **SIEM type** from the drop-down menu. The **API URL** is a Humio URL, which is dependent on the customer's Humio cloud location. The authorization header contains the relevant Humio HEC token created in previous steps.
2. In the **Add Cloud NSS Feed** dialog, **Key1** is "Authorization". **Value1** is the Humio HEC token in the format "Bearer XXX-XXX-XXX" (replace XXX with actual HEC token value).
3. **Feed Output Type** is **JSON** from the drop-down menu. Save the configuration after providing the required parameters. Add **",** \ (double quote, comma, backslash) to the **Feed Escape Character** list.

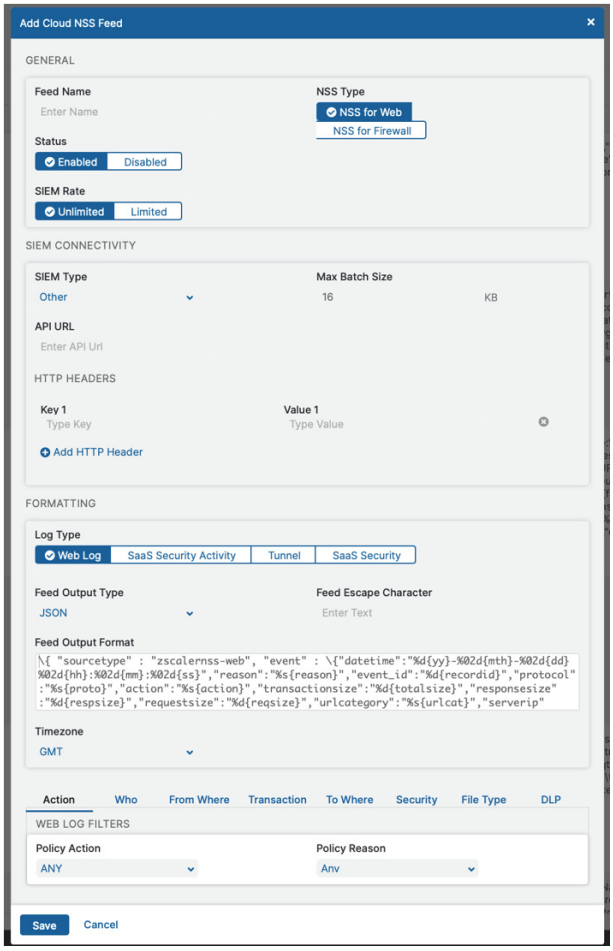

Figure 78. Configure cloud NSS feed

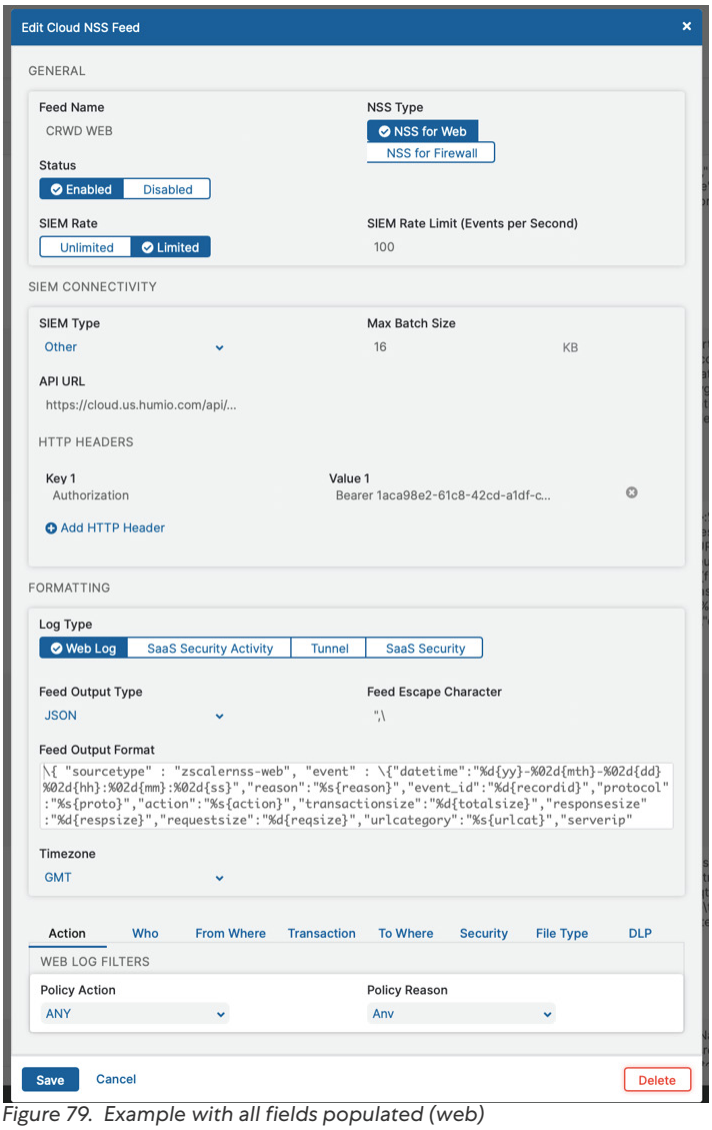

Figure 79. Example with all fields populated (web)

Set Up the Cloud NSS Log Feed (Firewall)

To set up the Cloud NSS log feed:

1. Select **Other** as the **SIEM Type** from the drop-down menu. The **API URL** is a Humio URL, which is dependent on the customer's Humio cloud location.
2. Repeat the steps from Create and Add Ingest Tokens in CrowdStrike Humio section of this document to create a new ingest token for use with the Firewall log type. Each log type (web, Firewall, DNS, etc.) requires its own unique ingest token on the Humio side. The authorization header contains the relevant Humio HEC token.
3. In the **Add Cloud NSS Feed** dialog, **Key1** is "Authorization". **Value1** is the HEC token in format "Bearer XXX-XXX-XXX" (replace XXX with actual HEC token value).
4. Select the **Feed Output Type** of **JSON** from the drop-down menu.
5. Provide the required parameters and then click **Save**.

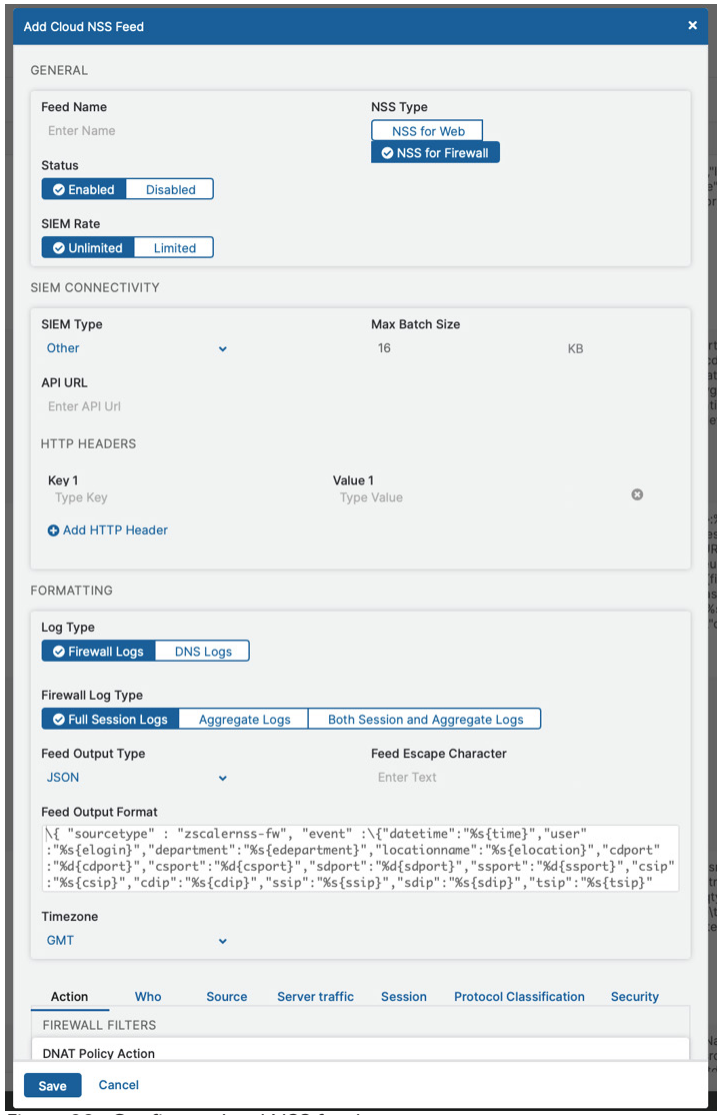

Figure 80. Configure cloud NSS feed

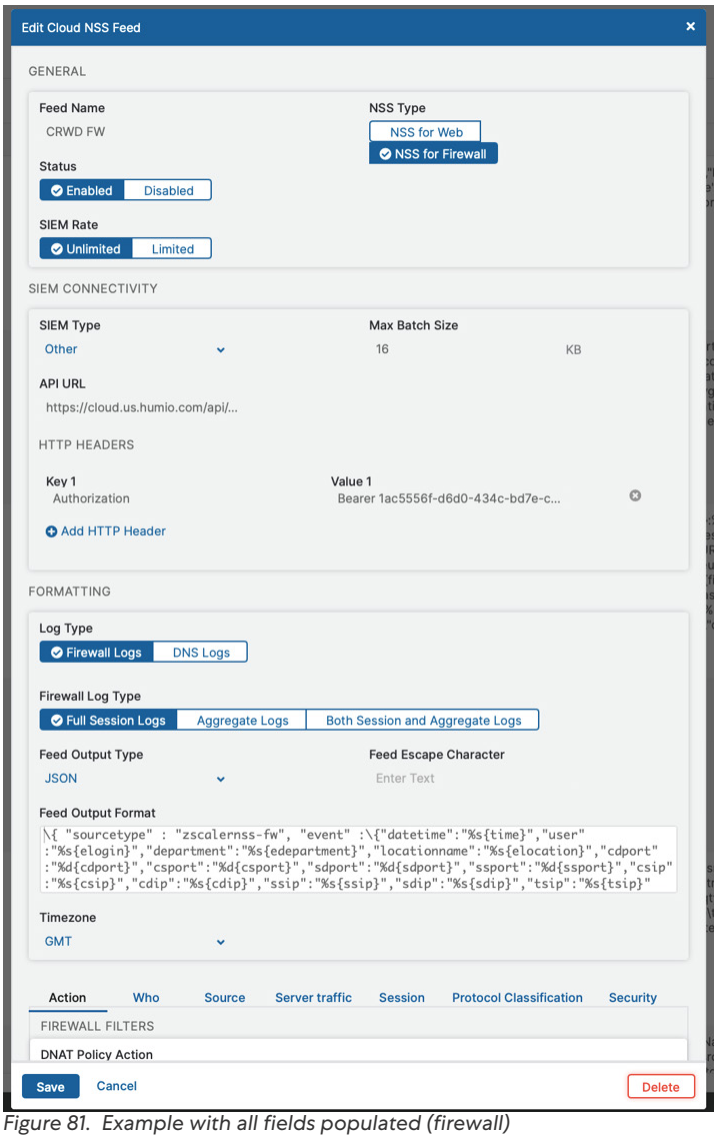

Figure 81. Example with all fields populated (firewall)

Add Other Log Sourcetypes

Repeat the previous steps to add other log source types (e.g., DNS logs, tunnel logs, etc.).

Validate NSS Cloud Configuration

After you save the configuration, click the Verify icon shown in the following screenshot to verify connectivity from ZIA cloud to the Humio cloud. This sends a sample or test log message from the ZIA cloud to Humio. Cloud-to-cloud connectivity is verified if Humio sends the expected response.

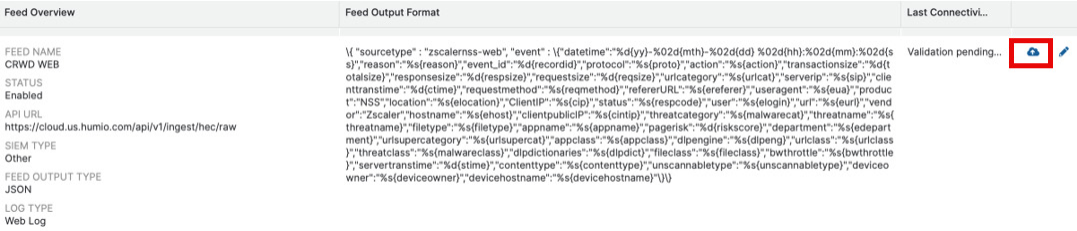

Figure 82. Verify connectivity to Humio cloud

When the connectivity is verified, the **Last Connectivity Test** column changes from **Last Validation Pending** to **Last Validation Successful**.

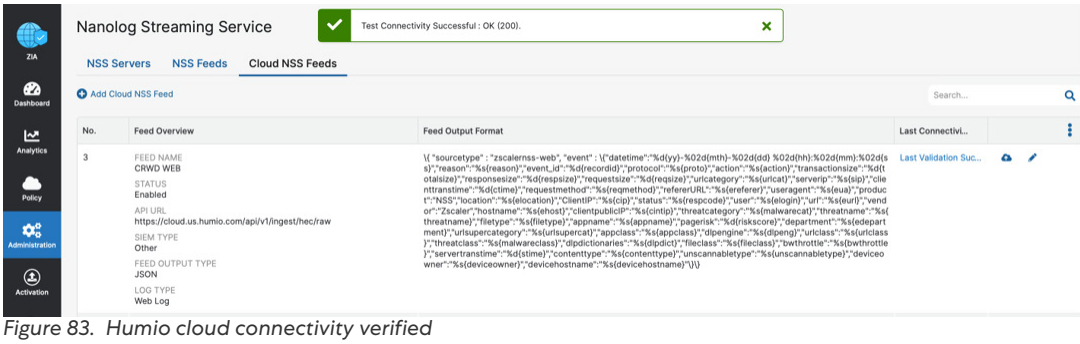

Figure 83. Humio cloud connectivity verified

Verify Log Flow Using Humio's Zscaler Package

Log back into your Humio cloud tenant and go to the **Dashboards** section. Then select the appropriate Zscaler dashboard depending on the type of logs being sent from ZIA to Humio.

The dashboard is populated with incoming Zscaler log data.

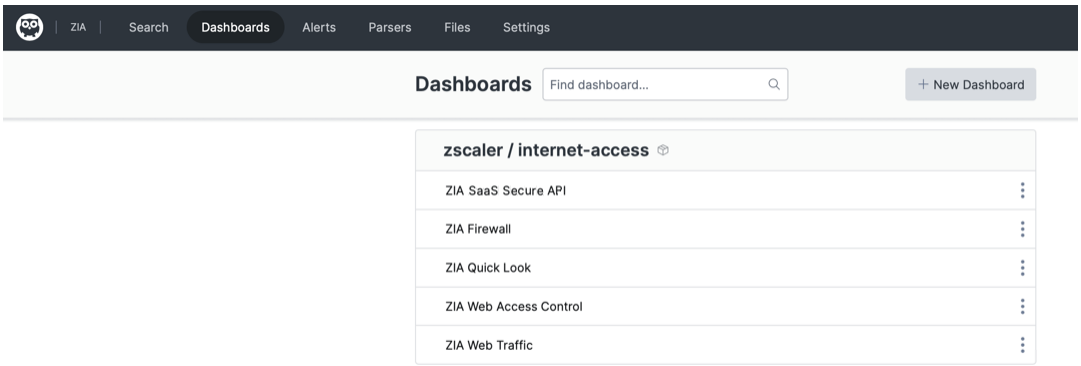

Figure 84. Check Humio Zscaler dashboards

If you see that a particular widget is not populated, click the More icon (ellipsis) next to it to see the query that the panel is running. Understanding the query helps you troubleshoot.

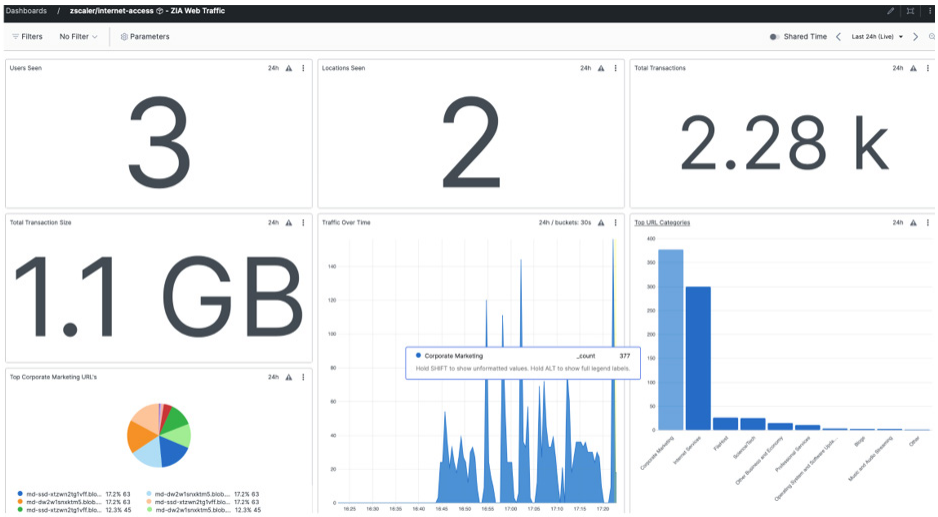

Figure 85. Widgets within Humio Zscaler dashboards

Use Case 7: Extended Detection and Response (XDR)

Falcon's XDR telemetry combined with Zscaler extends detection and response capabilities beyond the endpoint. The combination of Falcon and Zscaler automatically correlates suspicious activity across attack surfaces and leverages powerful analysis tools that help your team accelerate security operations, reduce risk, and improve threat visibility and detection across the enterprise.

Using Falcon XDR telemetry and Zscaler, you can:

- Ingest security data from supported third-party vendors.
- Search cross-domain security data in the Falcon Console.
- Create XDR scheduled searches and reports.
- Create custom XDR detections from queries.
- Triage and investigate XDR detections generated by the Falcon platform and supported third-party data sources from a unified interface.
- Perform targeted response actions from within an XDR detection or the graph view. Extend XDR response with supported third-party integrations.
- Configure response actions that execute automatically when a triggering condition occurs with Falcon Fusion workflows.

Requirements

Subscriptions:

- Falcon Insight XDR with one or more XDR connector packs. To leverage XDR capabilities, you must have at least one additional supported Falcon subscription and the associated Falcon platform XDR connector pack, or a third-party XDR connector pack. Supported Falcon subscriptions:
  - Falcon Identity Protection. To enable identity response actions, you must have Falcon Identity Protection and the associated Falcon platform XDR connector pack.
  - Falcon Cloud Workload Protection (CWP).
  - Falcon for Mobile.

Roles:

- Falcon Administrator
- XDR Administrator
- XDR Security Lead
- XDR Analyst
- XDR Analyst (Read Only)

Clouds:

- US-1
- US-2
- EU-1

Other requirements:

- For Zscaler response actions:
  - ZIA
  - Enable SCIM API integration for ZIA. To learn more, see the Configuring SCIM (government agencies, see Configuring SCIM).
  - Ensure user's machines are (Active Directory) domain-joined.

Understanding XDR

Using XDR, you can ingest third-party telemetry to unify detection and response across your security stack. You can hunt, detect, and investigate adversary activity across attack surfaces from a unified interface in the Falcon Console:

- XDR detection monitoring: Falcon monitors activity across your specified domains and data sources. Falcon correlates and analyzes data in real-time upon ingestion, and automatically generates XDR detections when suspicious activity is detected.
- XDR event search: From the Falcon Console, search against cross-domain data to hunt for suspicious activity or further investigate detections.
- XDR scheduled searches: Create XDR event searches that run automatically and recur on a schedule that you set. You can download and share the search results, and your specified recipients can receive notifications each time a scheduled search completes. Configure notifications to be sent when a search produces results, when a search produces no results, or both.
- Custom XDR detections from queries: For scheduled searches that indicate suspicious activity, you can optionally generate an XDR detection each time the search returns results. These custom XDR detections support MITRE ATT&CK framework mappings, severity, and descriptions. They appear in a consolidated view alongside CrowdStrike-generated XDR detections to facilitate monitoring and triage.
- Response actions: Perform targeted response actions, either manually from within an XDR detection or automatically through a Falcon Fusion workflow. Optionally, extend XDR response capabilities with supported third-party integrations.

Data Sources

You can extend Falcon platform telemetry by ingesting security telemetry from supported third-party vendors (such as Zscaler). Upon ingestion, data is parsed and mapped into the CrowdStrike XDR data schema, which provides a common language for correlation and analysis of data from different sources.

Supported third-party vendor data sources:

- CrowdStrike supports third-party data from CrowdXDR Alliance partners. To learn more, refer to CrowdXDR Alliance. Work with your CrowdStrike onboarding team for more information about vendor-specific support.
- CrowdStrike anticipates adding more third-party data sources over time. Request support for additional data sources through CrowdStrike Ideas (CrowdStrike support portal).

XDR Data Ingestion

The data ingestion setup process varies depending on the specific vendors you work with, and your unique environment and configuration. During XDR onboarding, your onboarding team provides vendor-specific instructions for each of the partners you work with.

How Data is Processed and Stored

Upon ingestion, data is automatically parsed and mapped into the CrowdStrike XDR data schema, which provides a common language for correlation and analysis of data from different sources. The XDR data schema includes XDR events and XDR indicators. Vendor-specific telemetry is also preserved and stored because it can contain relevant information for investigation and response.

XDR Telemetry

The following table defines the XDR telemetry values.

Item	Description
XDR detection	A confirmation that activity is malicious or otherwise warrants further investigation. An XDR detection can be composed of one indicator or multiple correlated indicators.
XDR indicator	A signal of activity that's interesting from a security perspective. XDR indicators are composed of one or more events and associated telemetry that represent unusual activity or behavior known to be used by threat actors. Examples: • Ten failed logins in a row. • Metasploit traffic. • Suspicious web proxy activity. • Suspicious email attachment executed. An indicator on its own isn't necessarily confirmation of suspicious activity and might not warrant further investigation. XDR indicators are the building blocks that form XDR detections.
XDR event	A signal that an activity occurred. Examples: • An email was received. • A login attempt failed. • A file was written. View events in XDR search, or in the context of indicators or correlated detections.

Setup

To set up XDR, work with your CrowdStrike onboarding team to configure ingestion of third-party data. Your onboarding team provides vendor-specific instructions during the onboarding process.

Create XDR scheduled searches and detections as needed. To learn more, refer to How to Leverage Scheduled Searches.

Configure automated response actions, including third-party responses, and Falcon Fusion workflows as needed. To learn more, refer to XDR response actions.

XDR Detection Monitoring

Falcon monitors activity across your specified Falcon platform and third-party domains and data sources. It surfaces suspicious collections of signals and indicators in the form of correlated XDR detections. XDR detections include output that's generated automatically by CrowdStrike and output that's generated by any custom XDR detections that you created from scheduled searches.

XDR detections appear in an XDR-specific interface in the Falcon Console, where you can triage, investigate, and respond to cross-domain alerts.

- Triage: Filter XDR detections by data domain, data source, status, severity, and more. Assign detections, update detection status, and comment on XDR detections.
- Investigate: Review detection details in the detection summary, including information about the indicators that triggered the detection. Pivot to the XDR graph explorer to visualize the detection and explore connections between events. Pivot from the summary into XDR event search to view event data associated with the alert.
- Respond: In the XDR detection summary, add users to the Falcon Identity Protection watchlist or perform a supported third-party response action.

View XDR detections at **Endpoint security** > **Monitor** > **XDR detections** in the Falcon Console.

XDR Event Investigation

With XDR event search in the Falcon Console, search against cross-domain data to hunt for suspicious activity or further investigate detections. XDR event searches are written in the LogScale Query Language. To learn more, refer to XDR queries.

Search XDR events at **Endpoint security** > **XDR search** > **Search** in the Falcon Console.

XDR Scheduled Searches and Custom XDR Detections

Create XDR event and indicator searches that run automatically and recur on a schedule that you set. You can download and share the search results, and your specified recipients can receive notifications each time a scheduled search is completed. Configure notifications to be sent when a search produces results, when a search produces no results, or both.

For recurring searches that indicate suspicious activity, you can opt to generate an XDR detection each time that the search returns results. These custom XDR detections appear alongside CrowdStrike-generated XDR detections to facilitate monitoring and triage.

The XDR scheduled search activity log shows the history of your completed scheduled searches, including any searches that generated errors. The XDR scheduled search audit log shows the history of changes to your configured XDR scheduled searches and custom XDR detection queries.

EDR scheduled searches are managed separately from XDR scheduled searches. To learn more, refer to Scheduled Searches for EDR.

Frequency and Timing

After you save an XDR scheduled search, the first instance of that search runs at your specified start time and generates results for the specified time interval. For example, if you configured the search to run every two hours, the first run of that search generates results for the previous two hours.

If more than five searches are scheduled to run at the same time, they're placed in a queue and then run as resources become available. If a given search isn't complete by the time the next search is scheduled to start, the next instance of the search fails.

If a running search hasn't completed after 60 minutes, it times out. If a specific scheduled search is timing out frequently, consider editing the scheduled search to refine its query syntax. To learn more, refer to Edit an XDR scheduled search or custom XDR detection.

You can retry a failed search manually. To learn more, refer to Retry an XDR search.

When you create a scheduled search without detections, you can configure an optional offset window. Specifying an offset window skips a period immediately before the search and looks at older, more complete data. Configure an offset time that's at least as long as the time between scheduled log uploads.

When you create a custom XDR detection query, you specify a search frequency and a search window. The search frequency determines how often Falcon runs the search for the detection. The search window determines the period of time that's searched. The first time a detection is generated for a given query, a time window is established between the first and last events. When the query runs again, a new detection is generated only if a new event is included in the result and is outside the time window from the previous result.

To ensure that all relevant XDR data is searched for detections, configure a search window that's at least as long as the search frequency. Overlapping searches help to ensure that all XDR data is included, regardless of when logs are uploaded. However, with overlapping searches, a given XDR event or indicator might appear in more than one custom detection.

Limitations

Falcon stores a maximum of 10,000 XDR events per scheduled search. Each generated search result set is retained and available for download for 30 days, after which it's permanently deleted.

Your customer ID (CID) can have a maximum of 200 active XDR searches, and an unlimited number of inactive searches. If your CID already has 200 active searches, any new searches created are inactive by default. For information about activating and deactivating scheduled searches, refer to Activate or deactivate an XDR scheduled search.

Deactivating a search stops the search from running any new queries, but the configured scheduled search remains available for use for 30 days. After 30 days, the deactivated scheduled search is deleted permanently.

If the user associated with a scheduled search is removed or has their permissions revoked, the search is deactivated. Additionally, a notification explaining the deactivation is sent to all notification recipients for that scheduled search. If a different user reactivates the search, that user becomes the new owner of the search.

XDR detections appear in the Falcon Console for 90 days after they're generated.

Notifications

You can set up automatic delivery of notifications to alert members of your team each time an XDR scheduled search has run, or whenever an XDR detection results in an error.

A notification contains a summary and, if applicable, a link to download the search results through the Falcon Console. For XDR scheduled searches, you can specify when to send notifications in any combination:

- When matching search results are returned.
- When no matching search results are returned.
- When errors or warnings are generated during a search.

For custom XDR detections, you can configure notifications to be sent if errors or warnings are generated.

Recipients cannot unsubscribe from scheduled search notifications. Only the scheduled search creator, an XDR Administrator, or a Falcon Administrator can remove recipients from the scheduled search notification settings.

Notification Delivery Options

Send scheduled search notifications to individual users by email or to groups of users through Slack, PagerDuty, Microsoft Teams, or webhook integrations. For each scheduled search, you can configure one or more delivery methods.

Before you can configure notifications through Slack, PagerDuty, Microsoft Teams, or webhook integration, you must set up the relevant app integrations through the CrowdStrike Store.

Email

Send scheduled search notifications to a specified list of email addresses in your approved domains. You can designate up to ten email notification recipients per scheduled search, including people who are not Falcon users.

You can configure email notifications to include the full search result set as a JSON file attachment. Note that attaching results sends data out of Falcon to systems that might have different security standards or terms and conditions.

Slack

Send notifications to one or more channels in your integrated Slack account. A Slack integration through the CrowdStrike Store is required. To learn more, refer to CrowdStrike Store App Integrations.

PagerDuty

Configure notifications to automatically open an incident in PagerDuty and alert relevant user groups. Select the PagerDuty source and severity from your connected service to configure notification delivery. A PagerDuty integration through the CrowdStrike Store is required. To learn more, refer to CrowdStrike Store App Integrations.

Microsoft Teams

Send notifications to one or more channels in your integrated Microsoft Teams account. A Microsoft Teams integration through the CrowdStrike Store is required. To learn more, refer to CrowdStrike Store App Integrations.

Webhook

Distribute notifications to other applications through a webhook. A webhook integration through the CrowdStrike Store is required. To learn more, refer to CrowdStrike Store App Integrations.

XDR Response Actions

Perform targeted Falcon platform response actions from within XDR detection details or the graph view. Optionally, extend XDR response with supported third-party integrations. You can also automate response actions through Falcon Fusion workflows.

Response actions can take up to ten minutes to complete.

Falcon Identity Protection Response

Perform targeted response actions by adding users to the Falcon Identity Protection watchlist from within an XDR detection or the graph view. You can configure the watchlist to enforce additional security controls such as requiring multifactor authentication.

To learn more, refer to Identity Protection Overview.

Third-Party Response

Extend XDR response with integrations from supported third-party vendors such as Zscaler.

ZIA Response

Integrate ZIA with Falcon so that you can perform your configured ZIA response actions from within an XDR detection or the graph view. Examples of ZIA response actions that you can configure:

- Add a user to custom-defined restricted user group.
- Remove a user from custom-defined restricted user group

This option is available only after you add the user to the group.

You can also extend your ZIA response capabilities by creating a Falcon Fusion workflow. For example, automatically add a user to a custom-defined restricted user group when a triggering condition occurs.

To learn more, refer to ZIA response management.

Falcon Fusion Workflows from XDR Detections

Create a Falcon Fusion workflow that initiates a specified action when Falcon generates an XDR detection:

- Send notifications by email, Slack, or webhook.
- Generate a ServiceNow or PagerDuty incident ticket.
- Change status, specify an assignee, or add a comment to the detection.

Create a Falcon Fusion workflow based on an XDR detection by selecting a Trigger category of Alert and a Subcategory of XDR detection. To learn more, refer to Workflows.

XDR Scheduled Search and Custom XDR Detection Management

Your CID can have a maximum of 200 active XDR searches, including searches that generate custom XDR detections, and an unlimited number of inactive searches. If your CID already has 200 active XDR searches, any new searches created are inactive by default. To learn more, refer to Activate or deactivate an XDR scheduled search.

Get to XDR Scheduled Searches and Custom XDR Detections

The XDR Scheduled search and detections page is where you can view, create, edit, and delete your XDR scheduled searches and custom XDR detection queries. From this page, you can also view XDR search activity, download XDR search results, and view audit logs.

Go to **Endpoint security** > **XDR search** > **Scheduled search and detections** in the Falcon Console.

Create an XDR Scheduled Search Without Detections

To create an XDR scheduled search:

1. On the **XDR Scheduled search and detections** page, write and run the search query that you want to schedule. To learn more, refer to XDR queries.
2. Click **Schedule search**. The search details are prepopulated with the query information that you provided.
3. Specify a name for the scheduled search.
4. (Optional) Add a descriptive comment about the scheduled search.
5. In the **XDR query** field, confirm or modify the search query.
6. (Optional) Test the query in **XDR Search**, and then modify the query details as needed.
7. Click **Scheduled search report**, and then click **Next**.
8. Configure the schedule:
    a. **Start date** and **Start time**: Specify when the search begins running.
    b. (Optional) **End date** and **End time**: If you want the scheduled search to run for a finite period of time, specify when to stop running the search. If no end date is specified, the scheduled search runs indefinitely according to the configured frequency.
    c. **Search frequency**: Specify how often to run the search.
    d. (Recommended) **Search offset**: Configure an offset time that's at least as long as the time between scheduled log uploads. Specifying an offset window skips a period immediately before the search and looks at older, more complete data. To learn more, refer to Frequency and timing.
9. Click **Next**, and then configure notification settings as needed. To learn more, refer to XDR notification configuration fields.
10. Click **Schedule search**.

Create a Custom XDR Detection from Scheduled Search Results

Configure Falcon to generate an XDR detection each time a given scheduled search returns results.

On the XDR Scheduled search and detections page, write and run the search query that you want to base the detection on. To learn more, refer to XDR queries.

1. Click **Schedule with detections**. The detection details are prepopulated with the query information that you provided.
2. Specify a name for the detection.
3. (Optional) Add a descriptive comment about the detection.
4. In the **XDR query** field, confirm or modify the search query.
5. (Optional) Test the query in **XDR Search**, and then modify the query details as needed.
6. Click **XDR detection**.
7. Select a severity for the detection.
8. (Optional) Assign a MITRE ATT&CK tactic and technique to the detection.
9. Click **Next**, and then configure a schedule for the detection:
    a. **Start date** and **Start time**: Specify when the search for this detection begins running. The start time must be at least 15 minutes from the current time.
    b. (Optional) **End date** and **End time**: If you want the search for this detection to run for a finite period of time, specify the last day on which to run the search. If no end date is specified, the scheduled search runs indefinitely according to the configured frequency.
    c. **Search frequency**: Specify how often to run the search for this detection.
    d. **Search window**: Specify the period of time that is searched. To help ensure that all relevant XDR data is searched for detections, configure a **Search window** value that's at least as long as the **Search frequency** value. To learn more, refer to Frequency and timing.
10. Click **Next**, and then configure notification settings as needed. To learn more, refer to XDR notification configuration fields.
11. Click **Schedule**.

Edit an XDR Scheduled Search or Custom XDR Detection

To edit an exiting XDR scheduled search or custom XDR detection:

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**.
2. For the XDR scheduled search or custom XDR detection query that you want to edit, click **Edit** from the **Ellipsis** icon.
3. Modify the settings as needed. To learn more, refer to Create an XDR scheduled search without detections or Create a custom XDR detection from scheduled search results.
4. (Optional) Test the query in **XDR Search**, and then modify the query details as needed.
5. Click **Update search**.

Deactivate an XDR Scheduled Search

To stop running a scheduled search, deactivate it. If a scheduled search is deactivated, its **Status value** changes to Inactive and all further scheduled searches stop. You can resume the scheduled search by reactivating it.

Deactivating a search stops the search from running any new queries, but the configured scheduled search remains available for use for 30 days. After 30 days, the deactivated scheduled search is deleted permanently. However, any previously generated detections continue to appear in the list of detections.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**.
2. For the scheduled search that you want to activate or deactivate, click **Activate** or **Deactivate** from the **Ellipsis** icon.

Delete an XDR Scheduled Search or Custom XDR Detection

Permanently delete a scheduled search that you no longer need. Previously generated search results remain available to download for 30 days. Delete scheduled searches with caution. Deleting a scheduled search removes all references to the search.

As an alternative to permanent deletion, you can deactivate a scheduled search. Deactivating a search stops the search from running any new queries, but the configured scheduled search remains available for use for 30 days. After 30 days, the deactivated scheduled search is deleted permanently. To learn more, refer to Activate or deactivate an XDR scheduled search.

If you delete a custom XDR detection query, any previously generated detections continue to appear in the list of detections.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**.
2. For the XDR scheduled search or custom XDR detection query that you want to delete, click **Delete** from the **Ellipsis** icon.

View XDR Scheduled Searches and Custom XDR Detections

View your configured XDR scheduled searches and custom XDR detection queries, and any XDR searches that are currently running or queued. Refine the results through sorting, filtering, or specifying which columns are visible.

If the user associated with a scheduled search has been removed or has had their permissions revoked, the search is deactivated and the User field shows a value of None. To learn more, refer to Limitations and Activate or deactivate an XDR scheduled search.

For information about viewing search activity and search results, refer to View XDR search activity.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**. A full list of your configured XDR scheduled searches and custom XDR detection queries appears.
2. Refine the list of results as needed:
   - Apply filters:
     - Click a filter at the top of the list.
     - Select or clear the filter-specific metadata options, and then click **Apply**.
   - Specify which columns are visible:
     - Click the **Configure table columns** icon.
     - Select the checkboxes for the columns that you want to see.
3. Click any scheduled search to see additional details.

View the XDR Scheduled Search and Custom Detection Audit Log

View the history of changes to your configured scheduled searches and custom XDR detection queries.

If a scheduled search has been deleted, the Name field shows a numerical unique identifier for the scheduled search instead of its configured name. To avoid this, you can deactivate a scheduled search instead of deleting it. To learn more, refer to Activate or deactivate an XDR scheduled search and Delete an XDR scheduled search or custom XDR detection.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**, and then click **Audit log**.
2. Adjust your view by filtering or sorting the log entries.
3. Click any audit log to see additional details.

View XDR Search Activity

View the activity history for your completed XDR scheduled searches, including searches that generated errors. View all of your generated scheduled search results, or refine the results through sorting and filtering. To learn more, refer to Download XDR search results.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**, and then click **Search log**.
2. Adjust your view by filtering or sorting the log entries.
3. Click any activity event to see more details.

Download XDR Search Results

View generated XDR search results by downloading them as a JSON file.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**, and then click **Search log**.
2. Adjust your view by filtering or sorting the log entries.
3. Click any entry to see more details.
4. Click **Download results**.

Retry an XDR Search

Retry a failed search manually. To learn more, refer to Frequency and timing.

1. Go to **Endpoint security** > **XDR search** > **Scheduled search and detections**, and then click **Search log**.
2. For the search that you want to retry, click **Retry**.

The **Retry** option appears only if a search has failed.

XDR Notification Configuration Fields

During configuration, set up notifications if you want to alert others each time that a scheduled search has run or a custom XDR detection query resulted in an error. To learn more, refer to Notifications.

To learn more, refer to Create an XDR scheduled search without detections, Create a custom XDR detection from scheduled search results, and Edit an XDR scheduled search or custom detection.

For each XDR scheduled search, configure one or more notification methods, or disable notifications. To set up multiple delivery methods, click **Add another notification after adding a notification method**.

Notification Method	Description and Options
Send email	Send an email notification to the specified recipients each time the scheduled search has run. In the **Recipients** field, type an email address and then press Enter. You can enter up to ten email addresses. To attach the full search result set as a JSON file, select the **Send results as an attachment** checkbox. Note that attaching results sends data out of Falcon to systems that might have different security standards or terms and conditions.
Send Slack message	Send a notification to the specified Slack channel each time the scheduled search has run. From the **Name** list, select the applicable channel.
Create a PagerDuty incident	Create a PagerDuty incident each time the scheduled search has run. From the **Name** list, select the applicable configuration.
Send Microsoft Teams message	Send a notification to the specified Microsoft Teams channel each time the scheduled search has run. From the **Name** list, select the applicable channel.
Send webhook notification	Send a webhook notification each time the scheduled search has run. From the **Name** list, select the applicable webhook.
None	Don't send notifications to any recipients for this scheduled search.

When configuring a scheduled search, for each notification method, indicate when to send notifications. Select any combination of these options.

Field	Description
No results are found	Sends a notification each time the scheduled search runs and generates no matching results.
Results are found	Sends a notification each time the scheduled search runs and generates matching results.
Errors or warnings occur	Sends a notification each time the scheduled search runs, or attempted to run, and resulted in errors or warnings.

ZIA Response Management

Configuring ZIA response actions is a multi-step process. The exact steps are determined by your ZIA response goals.

> Many of these steps are performed in third-party products. The CrowdStrike Falcon platform integrates the relevant settings as you configure them. However, Falcon does not validate any third-party configurations. Perform the following steps with care, and validate your configurations before finalizing ZIA response actions in Falcon.

In this summary of the configuration process, the example goal is to enable the ability to add users to a customer-defined restricted user group:

- In your identity provider (IdP), create a user group for restricted users. See Step 1: Create a User Group for ZIA Response.
- Push the user group from your IdP to ZIA. See Step 2: Push the User Group to ZIA.
- In ZIA, create a response policy that restricts user access, and associate the policy with the user group. See Step 3: Create a ZIA Policy.
- In ZIA, enable SCIM support for your IdP. See Step 4: Enable SCIM Support.
- In the CrowdStrike Store, configure ZIA response actions for XDR. See Step 5: Configure ZIA Response Actions for XDR. After you complete this step, ZIA integration is configured, and you can perform manual response actions from within an XDR detection or the graph view.
- (Optional) Create a Falcon Fusion workflow that automatically initiates a specified ZIA action when a triggering condition occurs. To learn more, see Step 6: Create a Falcon Fusion Workflow for ZIA Response.

Step 1: Create a User Group for ZIA Response

In your IdP, create a user group for ZIA response purposes. You can add users to this group from a relevant XDR detection.

> CrowdStrike recommends that this group be used for ZIA response purposes only. After you create this user group, do not modify or use it for any other purpose.

These steps are performed in your IdP's administration interface. For more detailed information, refer to the applicable product documentation for your IdP.

In your IdP, create a user group. Make a note of the group name because you'll need it during subsequent steps.

Step 2: Push the User Group to ZIA

Sync the group to ZIA.

These steps are performed in your IdP's administration interface. For more detailed information, refer to the applicable product documentation for your IdP.

In your IdP, push the user group to ZIA.

Step 3: Create a ZIA Policy

Create a ZIA policy for the user group that you created. The type of policy that you create depends on your ZIA response goals. For example, the policy could restrict user access or define quarantine thresholds.

These steps are performed in the ZIA Admin Portal. For more information, see the applicable ZIA product documentation (government agencies, see ZIA product documentation).

1. Create a policy that supports your ZIA response goals.
2. Associate the new policy with the new user group that you created in an earlier step.

Step 4: Enable SCIM Support

In ZIA, enable SCIM support for your IdP. The SCIM protocol offers a generic interface for IdPs.

These steps are performed in the ZIA Admin Portal. For more information, see the applicable ZIA product documentation (government agencies, see ZIA product documentation).

1. Enable SCIM support for your IdP.
2. Make a note of these values because you'll need them for the next step:
   a. **Base URL**
   b. **Bearer token**

Step 5: Configure ZIA Response Actions for XDR

In the CrowdStrike Store, configure the ZIA-based response actions that you want to enable.

1. Go to the CrowdStrike Store.
2. Select **Zscaler Internet Access Response Actions for Falcon Insight XDR**.

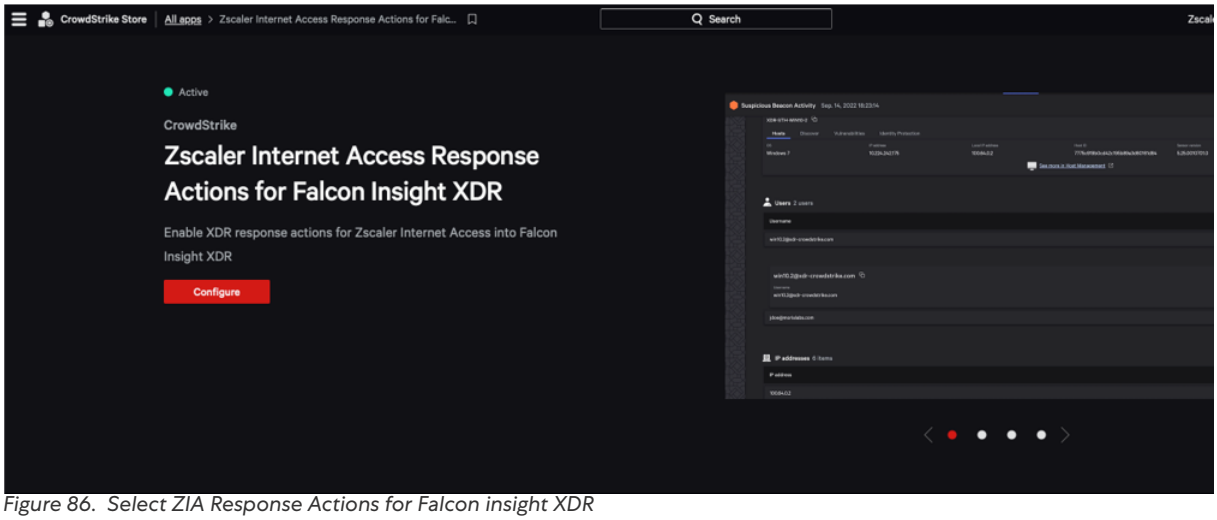

Figure 86. Select ZIA Response Actions for Falcon insight XDR

3. Click **Configure**.
4. Specify a name for the configuration. This name appears in the Falcon Fusion workflow configuration interface.
5. Type the base URL and bearer token values exactly as they appear in your identity provider (IdP), but omit the https:// protocol identifier.

6. Type the name of the user group that you created earlier. To learn more, see Step 1: Create a User Group for ZIA Response.

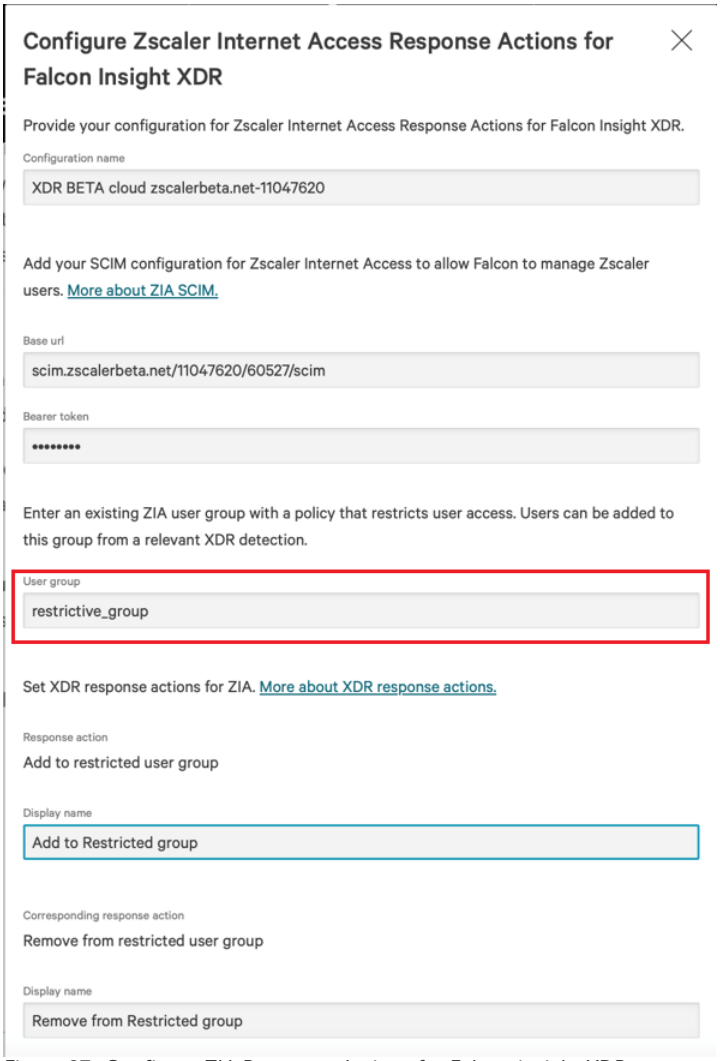

Figure 87. Configure ZIA Response Actions for Falcon insight XDR

Type the group name exactly as it appears in Zscaler. Falcon cannot validate the value that you enter in this field.

7. In the two **Display name** fields, enter a short, descriptive name for each response action. For example, you might create two corresponding actions of **Add** to restricted user group and **Remove** from restricted user group. These names appear as selectable ZIA response actions in Falcon Console menus.
8. Click **Save configuration**.

Step 6: Create a Falcon Fusion Workflow for ZIA Response

Optionally, create a Falcon Fusion workflow that initiates a specified ZIA response action when a triggering condition occurs. For example, automatically add a user to a customer-defined restricted user group when the severity of an XDR detection is greater than or equal to Medium.

Create a Falcon Fusion workflow for ZIA response by selecting these values.

Falcon Fusion Setting	Value
Trigger	Alert > XDR Detection.
Condition	(Optional) Customer preference.
Action	Containment > Zscaler response action that you want to take.
Configuration	Select the name of the Zscaler configuration that you created in Step 5: Configure ZIA Response Actions for XDR.

To learn more, refer to Workflows (Falcon Fusion).

Use Case 8: Prevent Lateral Movement—Zscaler Deception and CrowdStrike

The following steps demonstrate how to prevent malicious software from moving laterally using Zscaler Deception.

Step 1: Ensure the Prerequisites are Met

Ensure there is network connectivity from the Zscaler Deception Admin Portal to CrowdStrike Falcon Insight on HTTPS port 443.

Step 2: Create a Client and Secret Key in CrowdStrike Falcon Insight

First, log in to the CrowdStrike Falcon Insight platform:

1. Go to **Support** > **API Clients and Keys**.
2. Click **Add new API client**.
3. In the **Add new API client** window:
    a. Enter the **Client Name** of the new API.
    b. Enter the **Description**.
    c. Under **API Scopes**:
        - For **Detections**, select **Read**.
        - For **Hosts**, select **Read and Write**.
        - For **Custom IoA Rules**, select **Read and Write**.
        - For **IoCs (Indicators of Compromise)**, select **Read and Write**.

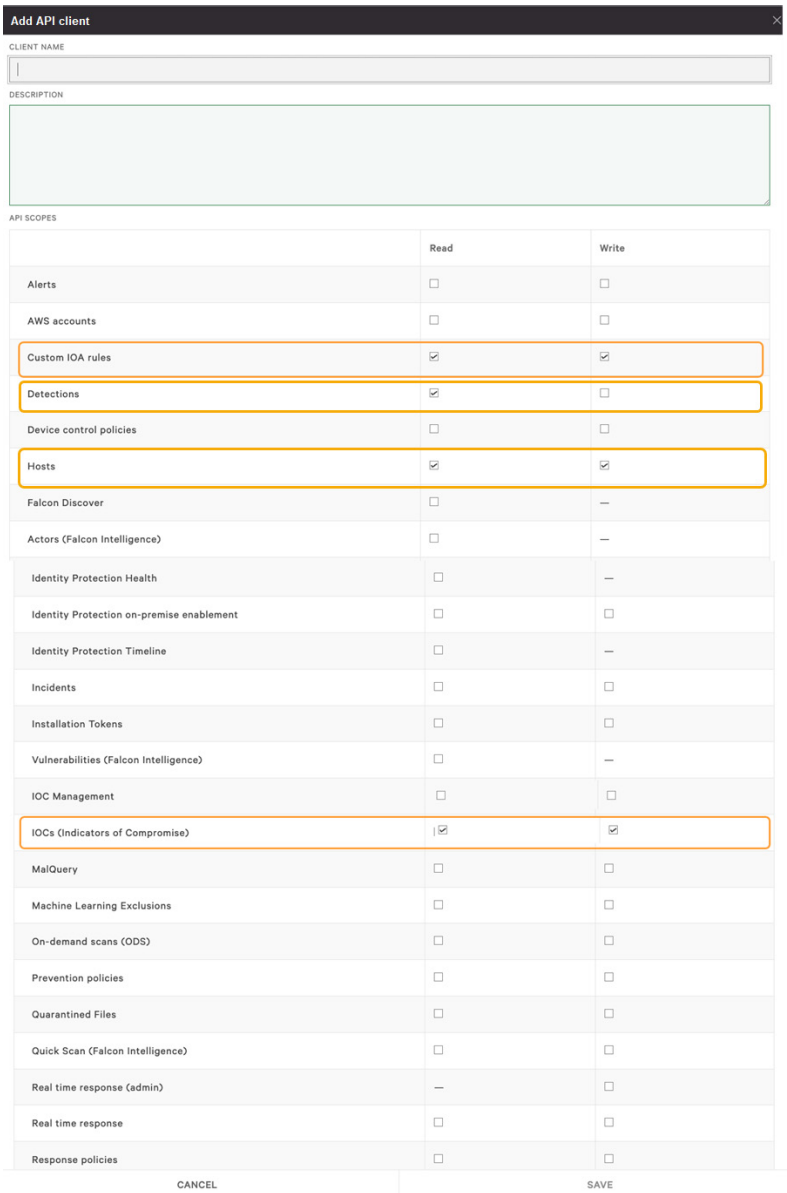

Figure 88. Read and write options for API scopes

4. Click **Save**. The **API client created** window appears.

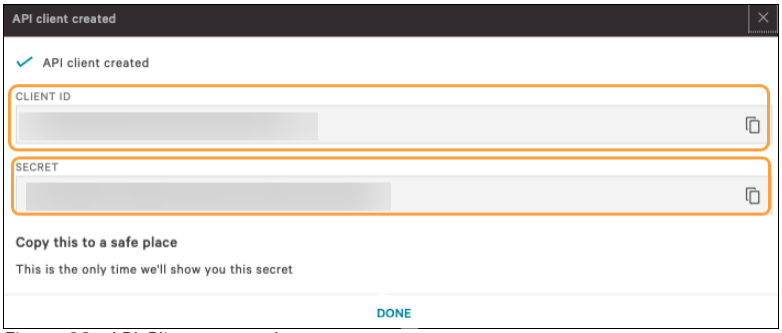

Figure 89. API Client created

5. Copy the client ID and secret key.
6. Click **Done**.

Step 3: Configure the Containment Integration Between Deception and CrowdStrike Falcon Insight

From the Zscaler Deception Admin Portal:

1. Go to **Orchestrate** > **Containment**.
2. In the table, locate CrowdStrike and click the **Edit** icon under the **Actions** column.

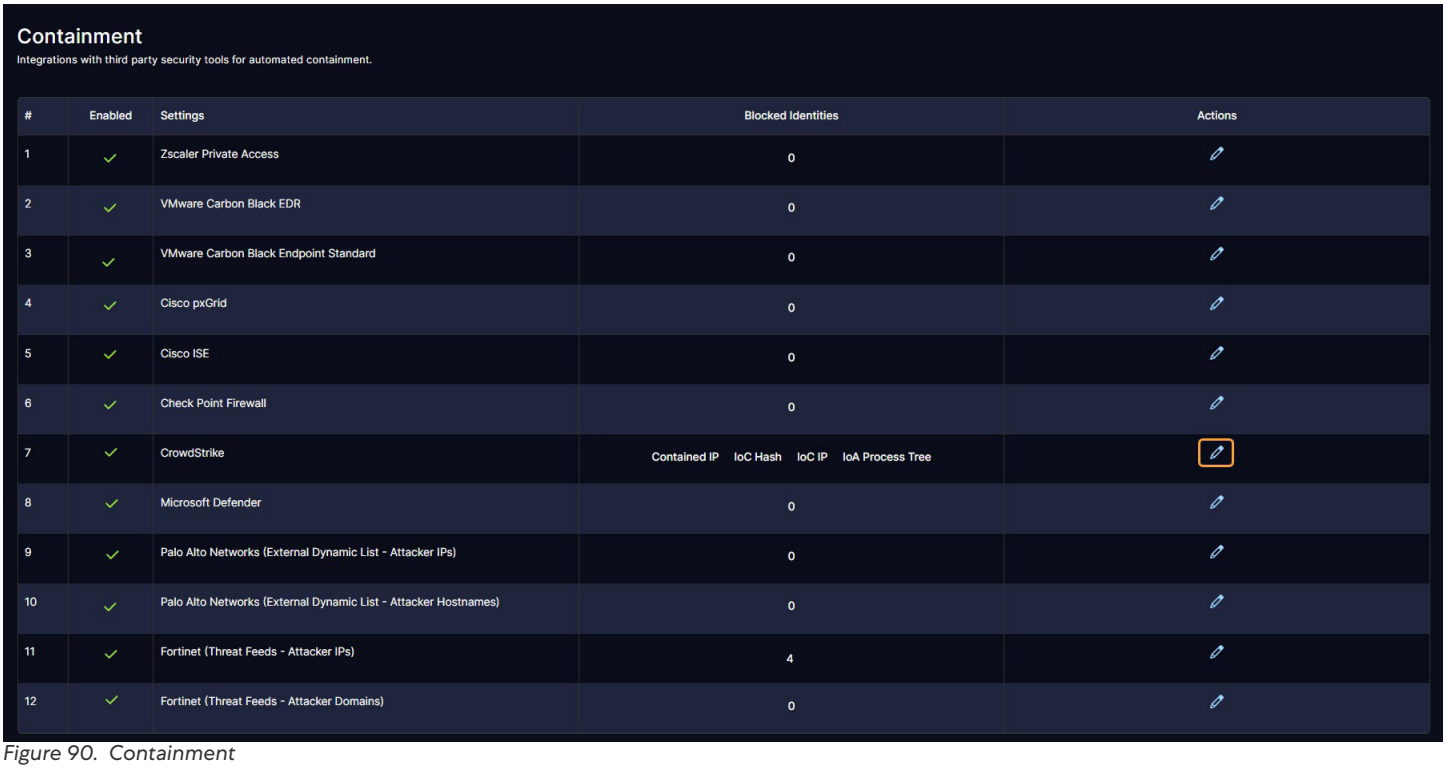

Containment

Integrations with third party security tools for automated containment.

#	Enabled	Settings	Blocked Identities	Actions
1	✓	Zscaler Private Access	0	
2	✓	VMware Carbon Black EDR	0	
3	✓	VMware Carbon Black Endpoint Standard	0	
4	✓	Cisco pxGrid	0	
5	✓	Cisco ISE	0	
6	✓	Check Point Firewall	0	
7	✓	CrowdStrike	Contained IP IoC Hash IoC IP IoA Process Tree	
8	✓	Microsoft Defender	0	
9	✓	Palo Alto Networks (External Dynamic List - Attacker IPs)	0	
10	✓	Palo Alto Networks (External Dynamic List - Attacker Hostnames)	0	
11	✓	Fortinet (Threat Feeds - Attacker IPs)	4	
12	✓	Fortinet (Threat Feeds - Attacker Domains)	0	

Figure 90. Containment

3. Perform steps 4 and 5 in the **CrowdStrike Falcon Insight configuration** window.
4. **Enabled**: Select to enable the containment.

5. **URL**: Enter the CrowdStrike API endpoint URL. For example, enter `https://api.crowdstrike.com`.
   a. **Client ID**: Enter the client ID that you copied in Step 2.
   b. **Client Secret**: Enter the client secret that you copied in Step 2.
   c. **Marking IoA Safe Processes**: (Optional) To avoid sharing legitimate process trees with CrowdStrike for IoA actions, follow these steps to designate legitimate processes as safe processes:
      i. Click **IoA Safe Processes**.
      ii. In the **Safe Processes** window, click **Add Safe Process**.
      iii. In the **Create Safe Process** window, provide the following details:
         - **Name**: Enter a name for the safe process.
         - **Process Path**: Enter the path of the executable file from which the process originates.
      iv. Click **Create Safe Process**.

   Processes are marked as safe only if all elements of the process tree have an entry in the safe processes list.

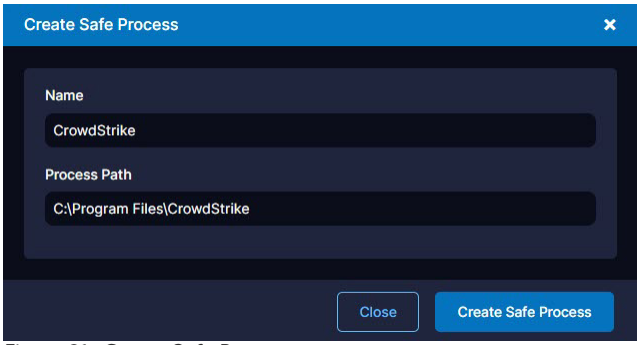

Figure 91. Create Safe Process

6. Click **Save**.
7. Click **Test** to verify the reachability of the CrowdStrike Falcon Insight platform.

Figure 92. CrowdStrike Falcon Insight

Step 4: Configure Orchestration Rule or Take Action to Contain Detected Attackers

You can contain detected attackers automatically by creating an orchestrated rule, or manually by taking action from the Investigate page.

1. Create a rule.
   a. Go to **Orchestrate** > **Rules**.
   b. Click **Add Rule**.

Figure 93. Rules

2. In the **Rule Details** window:
   a. Enter the name of the rule.
   b. Select **Enabled**.
   c. Create a rule using queries and conditions.

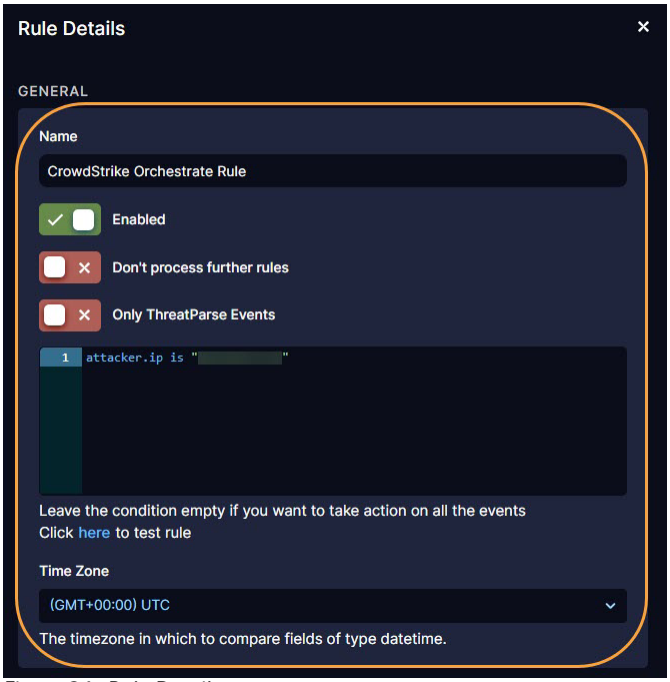

Figure 94. Rule Details

d. Under **Respond**, locate the **CROWDSTRIKE** section, and configure the following options:

- **Falcon Insight Containment**: Enable to contain an endpoint using CrowdStrike when an event matching the configured rule occurs.
- **IoC Hash**: Enable to share file hashes that are considered indicators of compromise with CrowdStrike and select an appropriate IoC Hash Action and IoC Hash Severity.
- **IoC Hash Action**: Choose the action for CrowdStrike to perform when IoC hashes are shared:
  - **Detect Only**: Show the indicator as a detection and take no other action.
  - **Block**: Add the indicator to the blocklist and show it as a detection.
  - **Block, Hide Detection**: Add the indicator to the blocklist and do not detect it.
  - **Allow**: Add the indicator to the allowlist and do not detect it.
  - **No Action**: Save the indicator for future use and take no action.
- **IoC Hash Severity**: Designate an appropriate severity level for indicators to share with CrowdStrike for IoC hashes. The available levels are **Informational**, **Low**, **Medium**, **High**, and **Critical**.
- **IoC IP**: Enable to share IPs that are considered indicators of compromise with CrowdStrike and select an appropriate IoC IP Action and IoC IP Severity.
- **IoC IP Action**: Choose the action for CrowdStrike to perform when IPs are shared:
  - **Detect Only**: Show the indicator as a detection and take no other action.
  - **No Action**: Save the indicator for future use and take no action.
- **IoC IP Severity**: Designate an appropriate severity level for indicators to share with CrowdStrike for IoC IPs. The available levels are **Informational**, **Low**, **Medium**, **High**, and **Critical**.
- **IoA Process Tree**: Enable to share process trees that are considered indicators of attack with CrowdStrike and select an appropriate IoA Action and IoA Severity.
- **IoA Action**: Choose the action for CrowdStrike to perform when IoA process trees are shared:
  - **Detect**: Show the indicator as a detection and take no other action.
  - **Monitor**: Use the indicator for informational use only and take no action, nor show it as a detection.
  - **Block Execution**: Terminate the process and show the indicator as detection.

  Exercise caution when choosing the Block Execution action as CrowdStrike terminates all processes that match with the process tree across all endpoints in your network. To avoid terminating legitimate processes, make sure that you have configured IoA safe processes.

- **IoA severity**: Designate an appropriate severity level for indicators to share with CrowdStrike for IoC hashes and IPs. The available levels are **Informational**, **Low**, **Medium**, **High**, and **Critical**.

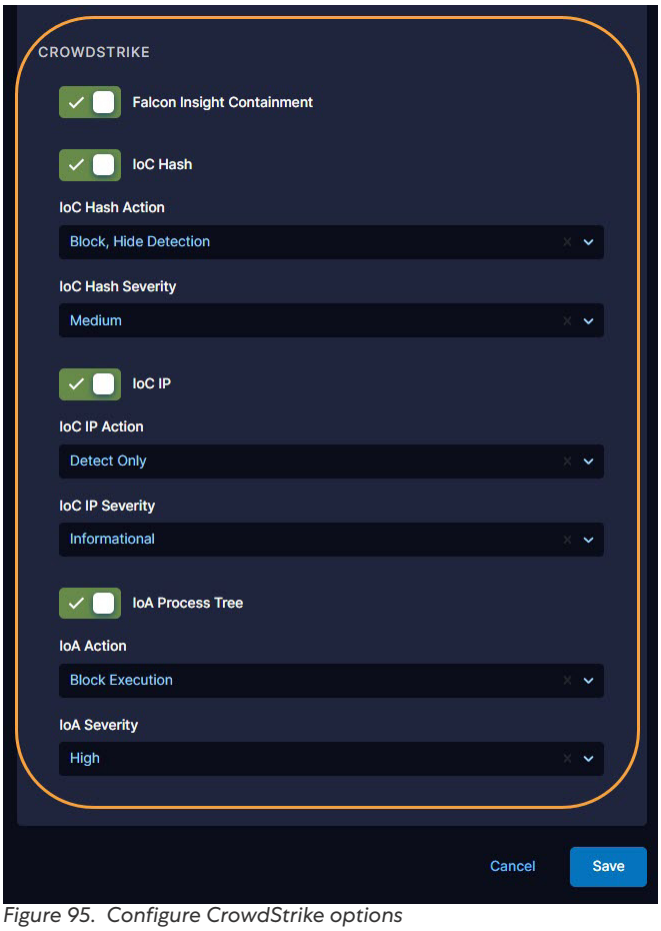

Figure 95. Configure CrowdStrike options

e. Click **Save**.

3. Take action from the **Investigate** page. For more information, see the [Zscaler Deception and CrowdStrike Deployment Guide](#).
   a. On the **Investigate** page, click an **Attack** icon on the alert graph. The attacker details pane opens.

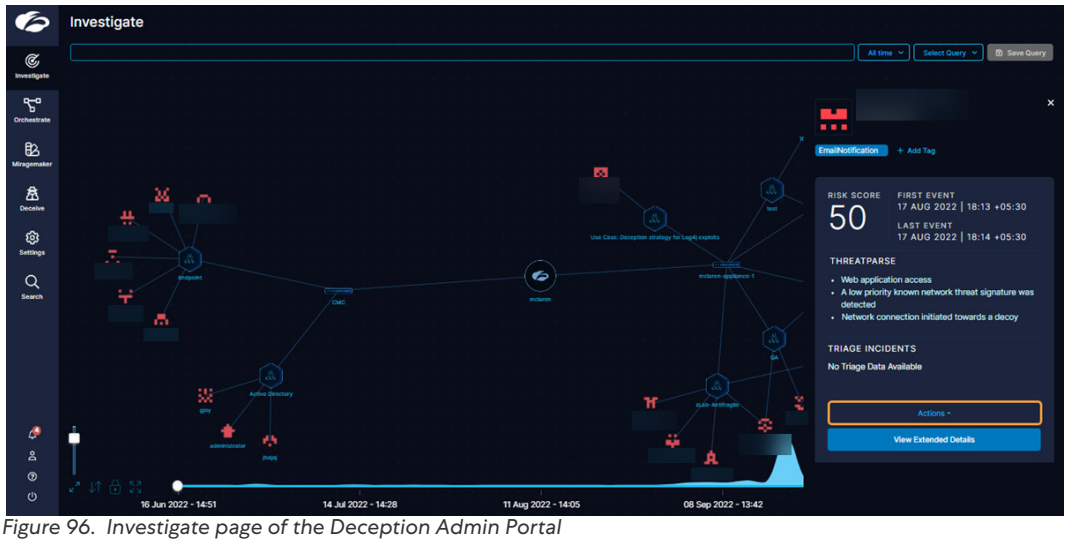

Figure 96. Investigate page of the Deception Admin Portal

   b. Click **Actions**. A list of actions that you can take to remediate the threat appears. If any orchestration rules were set up to automate sharing with CrowdStrike, the IoC, or IoA that has been shared with CrowdStrike appears in the Containment section of the Deception Admin Portal.
   c. (Optional) You can also select one of the following containment actions manually.

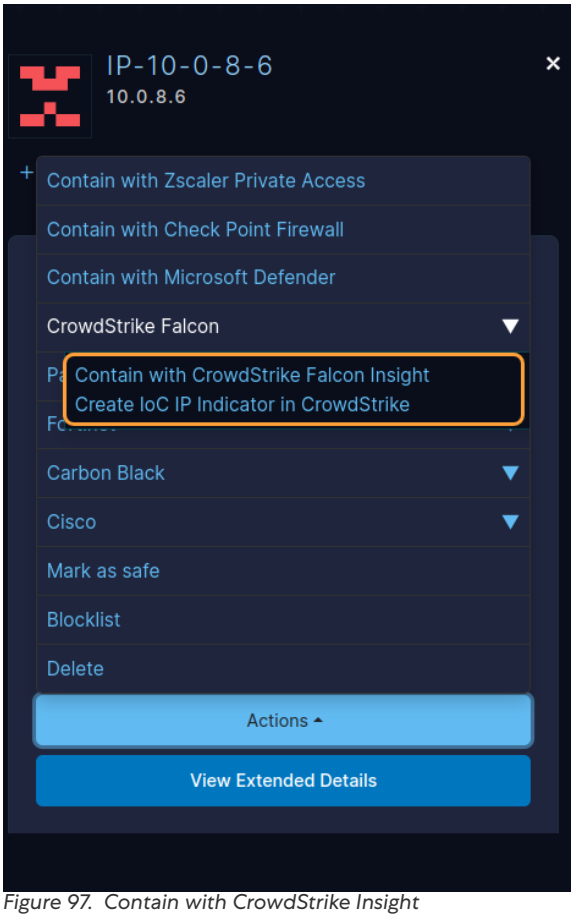

Figure 97. Contain with CrowdStrike Insight

Appendix A: Requesting Zscaler Support

If you need Zscaler Support for provisioning certain services, or troubleshooting configuration and service issues, it is available 24/7/365.

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

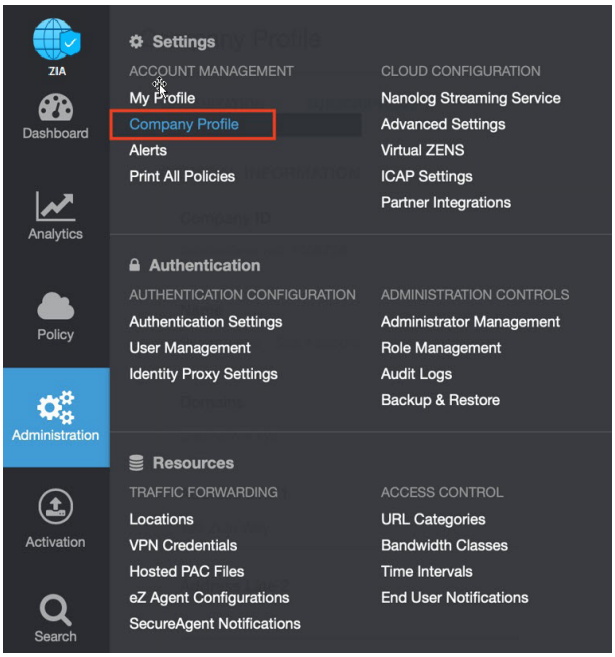

Figure 98. Collecting details to open support case with Zscaler TAC

Save Company ID

Copy your Company ID.

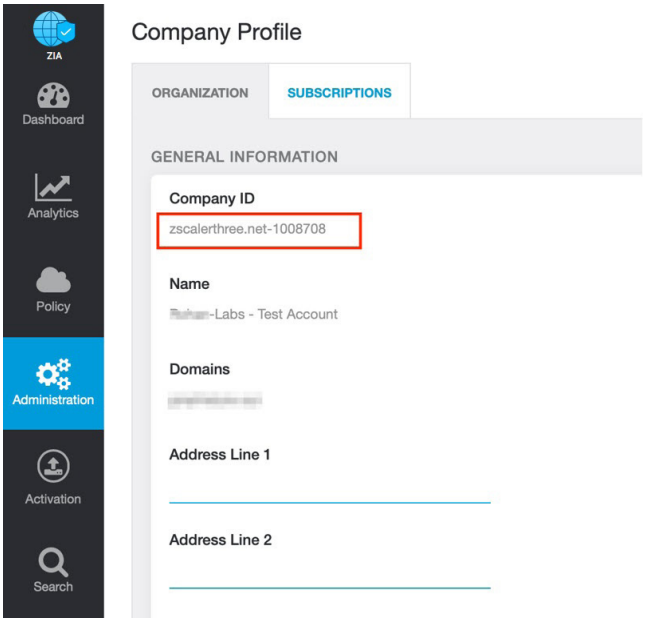

Figure 99. Company ID

Enter Support Section

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

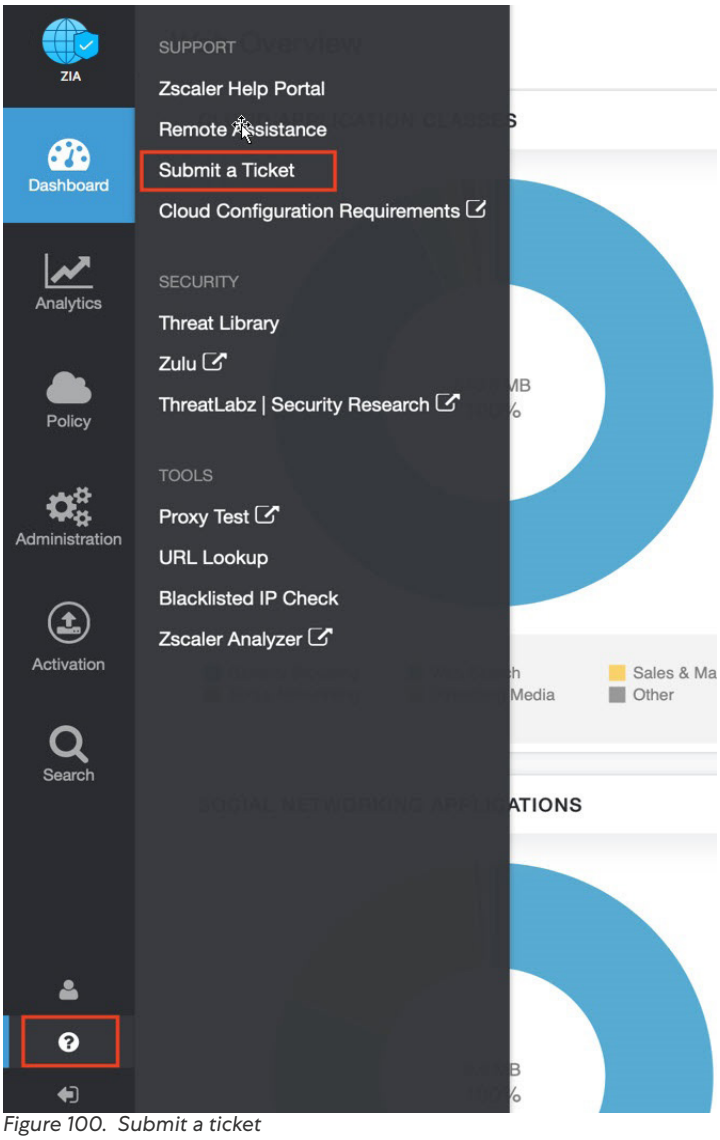

Figure 100. Submit a ticket

